THE LABOR DEMAND AND LABOR SUPPLY CHANNELS OF MONETARY POLICY

Sebastian Graves
Christopher K. Huckfeldt
Eric T. Swanson

The Labor Demand and Labor Supply Channels of Monetary Policy
Sebastian Graves, Christopher K. Huckfeldt, and Eric T. Swanson
NBER Working Paper No. 31770
October 2023, Revised February 2026
JEL No. E32, E52, J22, J23, J64

ABSTRACT

Monetary policy is conventionally understood to influence labor demand, with little effect on labor supply. We estimate the response of labor market flows to high-frequency changes in interest rates around FOMC announcements and Fed Chair speeches and find evidence that, in contrast to the consensus view, a contractionary monetary policy shock leads to a significant increase in labor supply: workers reduce the rate at which they quit jobs to nonemployment, and non-employed individuals increase their job-seeking behavior. These effects are quantitatively important: holding supply-driven labor market flows constant, the decline in employment from a contractionary monetary policy shock would be twice as large. To interpret our findings, we estimate a heterogeneous agent model with frictional labor markets and an active labor supply margin. The model rationalizes existing estimates of small labor supply responses to idiosyncratic transfers with our new evidence of a large labor supply response to an aggregate shock.

Sebastian Graves
University of Cambridge
sebgraves@googlemail.com

Christopher K. Huckfeldt
Federal Reserve Board of Governors
chris.huckfeldt@gmail.com

Eric T. Swanson
University of California, Irvine
Department of Economics
and NBER
eric.swanson@uci.edu

An online appendix is available at http://www.nber.org/data-appendix/w31770

THE LABOR DEMAND AND LABOR SUPPLY CHANNELS OF MONETARY POLICY

SEBASTIAN GRAVES, CHRISTOPHER HUCKFELDT, AND ERIC T. SWANSON

ABSTRACT. Monetary policy is conventionally understood to influence labor demand, with little effect on labor supply. We estimate the response of labor market flows to high-frequency changes in interest rates around FOMC announcements and Fed Chair speeches and find evidence that, in contrast to the consensus view, a contractionary monetary policy shock leads to a significant increase in labor supply: workers reduce the rate at which they quit jobs to nonemployment, and non-employed individuals increase their job-seeking behavior. These effects are quantitatively important: holding supply-driven labor market flows constant, the decline in employment from a contractionary monetary policy shock would be twice as large. To interpret our findings, we estimate a heterogeneous agent model with frictional labor markets and an active labor supply margin. The model rationalizes existing estimates of small labor supply responses to idiosyncratic transfers with our new evidence of a large labor supply response to an aggregate shock.

1. INTRODUCTION

"Policies to support labor supply are not the domain of the Fed: Our tools work principally on demand." –Federal Reserve Chairman Jerome Powell, November 30, 2022

Monetary policy is traditionally viewed as affecting labor demand and having little effect on labor supply, as reflected in the quote by Fed Chair Powell, above. This conventional wisdom is also embodied in the original Keynesian IS-LM framework, as discussed by Galí (2013), and in statements by other monetary policymakers around the world. It is further reflected in some parts of the New Keynesian (NK) literature, where various researchers have argued that estimated models with sticky wages often leave little quantitative role for labor supply in employment's response to monetary policy shocks, as we discuss below.

In contrast to this view, we offer new empirical evidence consistent with a substantial labor supply response to monetary policy. We begin by identifying labor market flows (and components of flows) that are plausibly more reflective of labor supply considerations insofar as they are initiated by the worker.[1] Thus, we classify flows from unemployment (U) to nonparticipation (N) and vice versa as supply-driven, given that such flows occur when an individual decides to stop or start searching actively for work. Similarly, we classify quits to nonemployment as supply-driven, given that these separations are initiated by the worker. One contribution of our paper is to provide a new decomposition of flows between employment (E) and nonparticipation (N) into quits and layoffs, which we use to show that a large and procyclical component of flows from E-to-N is due to quits.

We estimate the response of labor market flows to exogenous monetary policy shocks by extending a standard structural monetary policy vector autoregression (VAR) to include those flows. Following Stock and Watson (2012), Gertler and Karadi (2015), and others, we identify the effects of monetary policy using high-frequency changes in interest rate futures around FOMC announcements as an external instrument. Crucially, we also employ the recent methodology of Bauer and Swanson (2023b) to improve the relevance and exogeneity of our instrument, in part by exploiting additional interest rate variation around Fed Chair speeches. We are thus able to obtain substantially more accurate estimates of the response of labor market flows to monetary policy shocks than are available in the existing literature.

Consistent with the view described above, our VAR analysis shows that flows from E-to-U increase following a monetary policy tightening, and flows from U-to-E decrease, in line with an interpretation of lower labor demand amidst a weakening economy.[2] However, we also show that flows from N-to-U significantly *increase* following the monetary policy tightening, and flows from U-to-N *decrease*, consistent with heightened job search from nonemployment. We further identify a significant reduction in quits from employment to nonparticipation. Separately, we estimate a similar co-movement of labor market flows from impulse responses to the demand-like Main Business Cycle shock à la Angeletos, Collard and Dellas (2020), indicating that the co-movement of supply-driven flows we document is not unique to monetary policy shocks.

Importantly, we verify that cyclical changes in the composition of workers within labor market states plays only a limited role in explaining our estimated responses of supply-driven labor market flows to a monetary policy shock. The response of these flows thus seem to be largely driven by variation at the individual level. This finding, however, does not preclude

---

[1]We formalize this approach in Section 6 using a structural model where extensive-margin labor supply decisions—which respond to changes in job-finding rates, layoff rates, and prices—map into supply-driven labor market flows.

[2]We use the terminology "flows" and "transition probabilities" interchangeably throughout.

different labor market responses across different subgroups of workers: indeed, we document evidence consistent with larger increases in labor supply among lower-educated workers.

We quantify the importance of the response of supply-driven flows using the methods of Shimer (2012) and Elsby, Hobijn and Şahin (2015). We construct hypothetical impulse responses of employment holding candidate labor market flows constant at their average values, allowing us to quantify the contribution of such flows to the total employment response. Holding the response of supply-driven labor market flows fixed, we find the response of employment to a contractionary monetary policy shock would be roughly twice as large—a quantitatively significant effect.

To formalize our economic interpretation of supply-driven flows in the data and understand the implications of our new empirical findings, we study a model of frictional labor markets with an active extensive margin of labor supply following the seminal contribution of Krusell et al. (2017), where agents vary in their assets holdings, labor productivity, and disutility of searching for work. We consider the effects of a contractionary monetary policy shock in the model by feeding in our VAR estimates for the response of the job-finding rate, layoff rate, real interest rate, and real wages. We then study the labor supply response of agents in the model and estimate the model's key parameters to best match the impulse response functions for the labor market transition rates between employment, unemployment and nonparticipation.

The model closely matches our empirical estimates of the responses of these flows to a monetary policy shock, while also being consistent with estimates from the literature of a relatively modest marginal propensity to earn (MPE) out of idiosyncratic transfers. We show that the model matches our estimated impulse responses through a broad-based increase in labor supply: the decline in employment in the model is roughly 80 percent larger if we simulate the model holding labor supply policy functions fixed at steady-state. Consequently, our model establishes that a modest labor supply response to an idiosyncratic transfer (i.e., a modest MPE) does not rule out a large change in the propensity to work following an aggregate shock.

We further use the estimated model to decompose and explain the response of labor market flows and stocks following a monetary policy shock. We show that a crucial driver of individual labor supply decisions in the model is the decline in the job-finding rate. When the job-finding rate falls, workers anticipate longer spells of nonemployment and lower consumption. This discourages quits and increases the likelihood of job search and acceptance. The model generates stronger effects for less productive workers, mirroring our estimates of heterogeneous responses in the data.

Our empirical estimates and model indicate an important role for labor supply in explaining the response of employment to a monetary policy shock. Notably, several authors

have argued that common calibrations of sticky-wage New Keynesian models allow only a negligible short-run role for labor supply.[3] In contrast, our paper highlights a potentially important role for labor supply in the New Keynesian framework.

We believe our evidence yields insights beyond improving our general understanding of the monetary transmission mechanism: for example, labor supply may have taken on particular importance for the post-pandemic economy, where large fiscal transfers to households were followed by an increase in quits to nonparticipation, a slow recovery of labor force participation, and an increase in inflation. Our findings offer a window into the possibly important role of labor supply during this episode.

After surveying the literature, the remainder of our paper proceeds as follows. In Section 2, we review the standard empirical measures of labor market stocks and flows, we introduce our decompositions of E-to-U and E-to-N flows, and we describe our empirical VAR analysis. In Section 3, we report our baseline estimates of how the labor market responds to a monetary policy shock. In Section 4, we explore the role of composition and document heterogeneity in responses for different education groups. In Section 5, we compute hypothetical responses of employment when shutting down the response of various labor market flows. In Section 6, we introduce and estimate our model of frictional labor markets with an active labor supply margin and use it to decompose the employment response to a monetary policy shock. Section 7 concludes and discusses directions for future research.

Related Literature. Our paper is related to a few recent working papers that also study the conditional responses of labor market flows to monetary policy shocks (e.g., White, 2018; Broer, Kramer and Mitman, 2021; Coglianese, Olsson and Patterson, 2023; Faia, Shabalina and Wiczer, 2023). Our use of the Bauer and Swanson (2023a,b) methodology—and in the Appendix, the Aruoba and Drechsel (2026) shocks—allows us to obtain more precise and plausibly less biased estimates of the response of labor market flows to a monetary policy shock than previously available for the subset of this literature studying U.S. data.

Our paper is also related to the broader empirical literature studying labor market flows and their implications for aggregate labor market variables such as employment, unemployment and labor force participation (e.g., Davis, Faberman and Haltiwanger, 2006; Shimer, 2012; Elsby, Hobijn and Şahin, 2015). A distinctive contribution of our paper to this literature is to develop a methodology to measure quits and layoffs from employment to nonparticipation in the CPS, which we use to show that quits constitute a large component of the total flow of workers from employment to nonemployment. Beyond using this new data to document the unconditional cyclical behavior of quits to nonemployment, we also estimate

---

[3]See, e.g., Christiano (2011, p. 371), Wolf (2023b, p. 2254), and Wolf (2023a, p. 77).

that quits to nonparticipation decrease in response to a surprise monetary policy contraction, and we document that this response plays a particularly important role in shaping the response of the employment-population ratio to a monetary policy shock.[4]

While we estimate a response of labor market flows consistent with an increase in labor supply after a contractionary monetary policy shock, we also estimate a slight (and sluggish) decline in the labor force participation rate, which some authors would interpret as evidence of a decline in labor supply (e.g., Galí, Smets and Wouters, 2012; Christiano, Trabandt and Walentin, 2021). However, we show that the response of supply-driven flows significantly dampens a decline in labor force participation that would have been *much larger* in the absence of that response. We find that labor force participation declines following a contractionary monetary policy shock primarily because unemployment increases, consistent with Hobijn and Şahin (2021), who find similar results based on the unconditional cyclical dynamics of participation. Our quantitative structural model in Section 6 confirms this interpretation: labor force participation declines following a monetary contraction, even as the labor supply decisions of agents in the model reflect a greater willingness to work.

Our paper is complementary to the contemporaneous work of Alves and Violante (2025), who extend a framework similar to that of Krusell et al. (2017) and Heathcote, Perri and Violante (2020) into a rich HANK model to study how alternative monetary policy rules influence inequality and inflation. Instead, we provide new empirical evidence in support of such models and establish a minimal heterogeneous agent modeling environment (with labor market frictions and an extensive margin of labor supply) necessary to interpret our new estimates. Our paper also relates to Blanco et al. (2024), who study inefficient separations through quits and layoffs in an analytic model. Our estimates of the conditional and unconditional responses of quits and layoffs to aggregate disturbances offer validation for their theoretical findings.

Labor supply decisions of lower-productivity workers in our model—whose asset holdings typically place them closer to a borrowing constraint—are more responsive to changes in the aggregate job-finding rate, consistent with an insurance role for labor supply. Thus, our paper relates to the literature exploring the role of labor supply as self-insurance against wage and employment risk (e.g., Parker, Belghitar and Barmby, 2005; Pijoan-Mas, 2006; Eeckhout and Sepahsalari, 2023).

Finally, our results complement those of Cantore et al. (2023), who study the response of labor supply focusing on the intensive margin of average hours per worker. Their finding of an increase in labor supply for lower-wage workers following a contractionary monetary

---

[4]Michaels (2024) applies our methodology to show that quits to nonparticipation accounted for a disproportionate share of the increase in quits during the "Great Resignation" that followed the COVID recession. Subsequent work by Ellieroth and Michaud (2026) also examines the decomposition of E-to-N flows into quits and layoffs using a methodology that closely parallels our own, and their findings on the unconditional cyclical properties of quits and layoffs appear broadly consistent with our estimates.

TABLE 1. Cyclicality of Labor Market Stocks

	Employment-Population Ratio	Unemployment Rate	Participation Rate
mean(x)	61.14	6.19	65.16
std(x)/std(Y)	0.72	8.25	0.23
corr(x, Y)	0.83	-0.85	0.35

Note: x denotes the variable in each column, Y denotes HP-filtered log real GDP. Standard deviations and correlations are computed for HP-filtered and logged quarterly averages. The sample is 1978-2019.

policy shock parallels our findings of a particularly large response of quits to nonemployment for less-educated workers.

2. DATA AND METHODOLOGY

We begin by describing the labor market flows data and its relationship to aggregate labor market variables such as employment and unemployment. We then identify labor market flows (and components of flows) that plausibly reflect labor supply considerations. Finally, we describe how to estimate the responses of labor market flows to exogenous variation in monetary policy by extending a standard structural monetary policy VAR with high-frequency identification.

2.1. **Labor Market Stocks and Flows.** We study the cyclical behavior of aggregate labor market stocks and flows. Our primary data source for gross worker flows is the longitudinally linked data from the monthly Current Population Survey (CPS) from 1978 to 2019. We organize our discussion of labor market stocks and flows in terms of three distinct labor market states: employment (E), unemployment (U), and nonparticipation (N).

Table 1 presents summary statistics for three standard labor market stock measures: the employment-to-population ratio, E/(E+U+N), the unemployment rate, U/(E+U), and the labor force participation rate, (E+U)/(E+U+N). The cyclical properties of these labor market aggregates have been widely documented: the employment-population ratio is procyclical but not very volatile, the unemployment rate is countercyclical and highly volatile, and the labor force participation rate is only modestly procyclical and has very low volatility.

The dynamic behavior of the labor market stocks E, U, and N can be understood by the flows of workers between these three states. Labor markets exhibit considerable churn, with positive gross flows in both directions between any two states. Let p_{XY} denote the fraction of workers in labor market state X moving to state Y. Labor market stocks and flows are

TABLE 2. Cyclicality of Labor Market Flows

	E-to-U	E-to-N	U-to-E	U-to-N	N-to-E	N-to-U
mean(x)	0.014	0.029	0.254	0.226	0.045	0.025
std(x)/std(Y)	5.40	2.35	5.74	4.15	2.84	5.13
corr(x, Y)	-0.81	0.47	0.77	0.70	0.66	-0.67

Note: x denotes the variable in each column, Y denotes HP-filtered log real GDP. Standard deviations and correlations are computed for HP-filtered and logged quarterly averages. The sample is 1978-2019.

then related by the Markov chain

$$
\begin{bmatrix} E \\ U \\ N \end{bmatrix}_{t+1} = \begin{bmatrix} 1 - p_{EU} - p_{EN} & p_{UE} & p_{NE} \\ p_{EU} & 1 - p_{UE} - p_{UN} & p_{NU} \\ p_{EN} & p_{UN} & 1 - p_{NE} - p_{NU} \end{bmatrix}_{t+1} \begin{bmatrix} E \\ U \\ N \end{bmatrix}_{t}. \tag{1}
$$

Equation (1) can be extended to study the dynamics of labor market stocks across longer time periods. Let P_{t+1} denote the transition matrix in equation (1). Given the vector $[E, U, N]'_t$ and a time series of transition matrices $\{P_{t+j}\}_{j=1}^{k}$, we can express labor market stocks at $t + k$ as

$$
\begin{bmatrix} E \\ U \\ N \end{bmatrix}_{t+k} = \left(\prod_{j=1}^{k} P_{t+j} \right) \begin{bmatrix} E \\ U \\ N \end{bmatrix}_{t}. \tag{2}
$$

Thus, given an initial condition, we can understand the dynamic properties of labor market stocks through the time series of labor market flows. In Section 5, we use this relationship to help understand how shifts in supply-driven labor market flows account for the response of labor market stocks to monetary policy surprises.

Table 2 summarizes the average level and cyclical properties of each of the off-diagonal transition probabilities of P_t over the period 1978–2019.[5] Given our subsequent focus on quits and layoffs to nonemployment, we do not adjust for time aggregation bias. Our results are robust to corrections for time aggregation, where such corrections are possible (see Appendix Figure C.11). Appendix Figure A.1a plots the time series of each transition probability for our sample. The properties of these transition probabilities have been well documented in the literature (e.g., Shimer, 2012; Elsby et al., 2015; Krusell et al., 2017). Here we simply note that we label flows between nonparticipation and unemployment as supply-driven, given that such flows are initiated by workers. The procyclicality of U-to-N flows and countercyclicality of N-to-U flows can be interpreted as evidence of greater job-seeking behavior among the

---

[5]We seasonally adjust each flow using the X-13ARIMA-SEATS software provided by the Census Bureau. We also use X-13ARIMA-SEATS to impute flows over the small number of months when a longitudinal link across surveys is not available (see the discussion in Shimer, 2012).

TABLE 3. Components of E-to-U and E-to-N Flows

	E-to-U Flows				E-to-N Flows			
	Total	Quits	Layoffs	Other	Total	Quits	Layoffs	Other
mean(x)	0.014	0.002	0.008	0.004	0.029	0.012	0.003	0.015
std(x)/std(Y)	5.40	8.18	8.10	5.43	2.35	5.84	14.58	4.71
corr(x,Y)	-0.81	0.59	-0.83	-0.53	0.47	0.51	-0.45	0.24

Note: The process for decomposing E-to-U and E-to-N flows into quits, layoffs and other separations is described in Appendix B.1.1. x denotes the variable in each column, Y denotes HP-filtered log real GDP. Standard deviations and correlations are computed for HP-filtered and logged quarterly averages. The sample is 1978-2019.

nonemployed during downturns and account for around one-third of cyclical variation in the unemployment rate (Elsby et al., 2015). Finally, note that the average U-to-E probability is more than five times greater than the average N-to-E probability, consistent with a higher job-finding probability from unemployment compared to nonparticipation.

2.2. **Decomposing Separations into Quits and Layoffs.** To investigate the extent to which E-to-U and E-to-N transitions are driven by labor supply choices, we decompose E-to-U and E-to-N flows into "quits", "layoffs", and "other separations" using additional detail from the CPS.[6] Given that quits (by definition) are initiated by the worker, we classify quits from employment to nonemployment as supply-driven.

Although many authors have studied the cyclicality and composition of E-to-U flows, far less attention has been paid to E-to-N flows, despite the fact that they are roughly twice as large. To the best of our knowledge, we are the first to provide a decomposition of monthly E-to-N flows into quits and layoffs (as well as to study their unconditional cyclical properties). This decomposition is more complicated to construct than that for E-to-U flows, as nonparticipants are only asked their reason for leaving their previous job if they are in the outgoing rotation group of the CPS. Additionally, the possible answers to this question have changed over time. In Appendix B, we offer a detailed discussion of our new methodology for classifying E-to-N transitions into quits and layoffs, including evidence showing that the subsequent labor market transition probabilities of individuals who quit to nonemployment differ significantly from those of individuals who are laid off. We also report, for each possible labor market transition, the average number of observations per month.

The left panel of Table 3 summarizes the size and cyclical properties of the subcomponents of E-to-U flows. About 60% of E-to-U flows are due to layoffs, and these flows are highly countercyclical and volatile. Another 10–15% are due to quits, and although these flows are

---

[6]For example, if a worker transitioning from E-to-U lists the reason for unemployment in the CPS as being a "job leaver", then we classify that transition as a quit, while if they report being a "job loser/on layoff", we classify that transition as a layoff. See Appendix B for additional details.

similarly volatile, they are procyclical. The remaining 25–30% of E-to-U flows that cannot be categorized as either layoffs or quits are less volatile and countercyclical.

The right panel of Table 3 reports the size and cyclical properties of the various components of E-to-N flows. Layoffs from E-to-N are countercyclical and quits from E-to-N are procyclical, as is the case for E-to-U flows. However, in contrast to E-to-U flows, quits represent a much larger share of E-to-N flows than layoffs, implying a much more important role for both the magnitude and cyclicality of quits to nonemployment than has been previously recognized. Indeed, the portion of E-to-N flows that can be identified as quits is of similar magnitude to the entirety of E-to-U flows. This is arguably a conservative estimate, as we do not classify roughly half of E-to-N transitions as either quits or layoffs: a significant fraction of these "other separations" are due to individuals who either report retirement or disability, whom we are thus unable to classify as quits or layoffs, as discussed in Appendix B.1.2. Our finding of a quantitatively significant role for quits to nonparticipation stands in sharp contrast to much of the literature (e.g., Faberman and Justiniano, 2015), which often equates quits with job-to-job transitions.

2.3. **Monetary Policy VARs and High-Frequency Identification.** Several recent papers have used high-frequency interest rate changes around the Federal Reserve's Federal Open Market Committee (FOMC) announcements, or *monetary policy surprises*, to estimate the effects of monetary policy in a VAR (e.g., Cochrane and Piazzesi, 2002; Faust et al., 2003, 2004; Stock and Watson, 2012, 2018; Gertler and Karadi, 2015; Ramey, 2016; Bauer and Swanson, 2023b). Monetary policy surprises are appealing in these applications because their focus on interest rate changes in a narrow window of time around FOMC announcements plausibly rules out reverse causality and other endogeneity problems, as we discuss below.

The core of our VAR includes seven monthly macroeconomic variables: the log of industrial production, the unemployment rate, the labor force participation rate, the log of the number of vacancies, the log of the consumer price index, corporate bond spreads, and the two-year Treasury yield.[7] This specification is very similar to Bauer and Swanson (2023b), except that we include labor force participation and vacancies as additional variables, given our focus on the labor market (and we will also extend this core VAR to include labor market flow variables, below). We stack these seven core variables into a vector Y_t and estimate the

---

[7]Industrial production, the unemployment rate, the labor force participation rate, the CPI, and the two-year Treasury yield are from the Federal Reserve Bank of St. Louis FRED database. Corporate bond spreads are from Gilchrist and Zakrajšek (2012), and vacancy data is from Barnichon (2010). As discussed in Swanson and Williams (2014) and Gertler and Karadi (2015), the two-year Treasury yield was largely unconstrained during the 2009–15 zero lower bound period, making it a better measure of the overall stance of monetary policy than a shorter-term interest rate like the federal funds rate.

reduced-form VAR,

$$Y_t = \alpha + B(L)Y_{t-1} + u_t, \tag{3}$$

where α is a constant, $B(L)$ a matrix polynomial in the lag operator, and u_t is a 7×1 vector of serially uncorrelated regression residuals, with $\text{Var}(u_t) = \Omega$. We estimate regression (3) from January 1978 to December 2019 via ordinary least squares with 6 monthly lags.

We follow standard practice and assume that the economy is driven by a set of serially uncorrelated structural shocks, ε_t, with $\text{Var}(\varepsilon_t) = I$ (see, e.g., Ramey, 2016). Since the dynamics of the economy are determined by $B(L)$, the effects of different structural shocks ε_t on Y_t are completely determined by differences in their impact effects on Y_t in period t, given by

$$u_t = S\varepsilon_t, \tag{4}$$

which we assume are linear, with S a matrix of appropriate dimensions. We assume that one of the structural shocks is a "monetary policy shock", and we order that shock first in ε_t and denote it by ε_t^{mp}.[8] The first column of S, denoted s_1, then describes the impact effect of the structural monetary policy shock ε_t^{mp} on u_t and Y_t.

To identify the impact effect s_1 of the monetary policy shock ε_t^{mp}, we use high-frequency identification: Let z_t denote our set of high-frequency interest rate changes (surprises) around FOMC announcements, minutes releases, and Fed Chair speeches, converted to a monthly series by summing over all the high-frequency surprises within each month.[9] In order for z_t to be a valid instrument for ε_t^{mp}, it must satisfy an instrument *relevance* condition,

$$E[z_t \varepsilon_t^{mp}] \neq 0, \tag{5}$$

and an instrument *exogeneity* condition,

$$E[z_t \varepsilon_t^{-mp}] = 0, \tag{6}$$

where ε_t^{-mp} denotes any element of ε_t other than the first (Stock and Watson, 2012, 2018).

The appeal of high-frequency monetary policy surprises is that they very plausibly satisfy conditions (5)–(6). First, FOMC announcements and Fed Chair speeches are an important

---

[8] If the number of structural shocks in ε_t equals the number of variables in the VAR, and S is nonsingular, then equation (4) implies that the VAR is invertible. However, we do not require invertibility for our analysis and the number of shocks in ε_t is unrestricted. See Bauer and Swanson (2023b) for additional discussion.

[9] High-frequency interest rate changes around FOMC announcements, minutes releases, and Fed Chair speeches are from Swanson and Jayawickrema (2024) and include all 323 FOMC announcements from 1988–2019 (excluding that on September 17, 2001), 184 FOMC minutes releases, and 409 press conferences, speeches, and Congressional testimony by the Fed Chair ("speeches" for brevity) over the same period that had potential implications for monetary policy, according to financial market commentary in the *Wall Street Journal* or *New York Times*. This is larger than the set of speeches in Bauer and Swanson (2023b), who used an earlier version of the data that contained only the 295 most influential Fed Chair speeches. We compute z_t in the same way as Bauer and Swanson, taking the first principal component of the change in the current-quarter and 1-, 2-, and 3-quarter-ahead Eurodollar future rates in a narrow window of time around each announcement, which helps capture changes in forward guidance as well as the federal funds rate.

part of the news about monetary policy each month, so the correlation between z_t and ε_t^{mp} in (5) should be positive and large. Importantly, including Fed Chair speeches provides us with a much more relevant instrument than using FOMC announcements alone, as shown by Bauer and Swanson (2023b). Second, high-frequency monetary policy surprises capture interest rate changes in narrow windows of time around policy announcements. It's therefore unlikely that other structural shocks in ε_t^{-mp} are significantly affecting financial markets at the same time, so that these other shocks should be uncorrelated with z_t, implying (6).[10]

Given our external instrument z_t, we estimate the impact effects s_1 in the SVAR as described in Stock and Watson (2012, 2018), Gertler and Karadi (2015), and Bauer and Swanson (2023b). For concreteness, order the two-year Treasury yield first in Y_t, and denote it by Y_t^{2y}. We then estimate the regression

$$Y_t = \tilde{\alpha} + \tilde{B}(L)Y_{t-1} + s_1 Y_t^{2y} + \tilde{u}_t \tag{7}$$

via two-stage least squares, using z_t as the instrument for Y_t^{2y}.[11] It's straightforward to show that (5)–(6) imply that (7) produces an unbiased and consistent estimate of s_1, with the first element normalized to unity. (In our empirical results below, we rescale s_1 so that the first element has an impact effect of 25 basis points, rather than 1 percentage point.) Once we have estimated s_1, the impulse response functions for each variable follow from the estimated matrix lag polynomial $B(L)$ in (3).[12]

Finally, we follow the prescriptions of Bauer and Swanson (2023a,b) and adjust our FOMC announcement surprises by projecting out any correlation with recent macroeconomic and financial news. As Bauer and Swanson (2023b) show, this purges our estimates of a significant "Fed Response to News" endogeneity bias.

3. ESTIMATES

We present several sets of results. First, we report baseline impulse response functions (IRFs) for the core seven-variable VAR described above. Second, we extend this core VAR to include labor market flow variables and report IRFs for labor market flows. Third, we augment the core VAR to include the quit and layoff subcomponents of E-to-U and E-to-N flows to provide additional evidence of the response of supply-driven flows. Finally, we

---

[10]Swanson and Jayawickrema (2024) use narrow intradaily windows around these announcements and are careful to avoid overlapping with any other macroeconomic data releases.

[11]One can obtain the same point estimates for s_1 by regressing the reduced-form residuals u_t from (3) on u_t^{2y} using z_t as the instrument. Stock and Watson (2018) recommend using (7) to avoid a generated regressor and correctly estimate the first-stage F-statistic of the instrument.

[12]Note that the sample for (7) used to estimate s_1 does not have to be the same as for the reduced-form VAR in (3) used to estimate $B(L)$. Our high-frequency monetary policy surprises are only available from 1988:1–2019:12, while we estimate $B(L)$ over the longer sample 1978:1–2019:12.

FIGURE 1. Response of Aggregate Variables to a Monetary Policy Shock

Robust F-statistic: 16.88

Note: Estimated impulse responses to a 25bp monetary policy tightening shock in the baseline VAR. Solid black lines report impulse response functions, while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. See text for details.

augment our core VAR with additional variables to further understand the response of the labor market to a monetary policy shock.

3.1. **Baseline VAR Impulse Responses to a Monetary Policy Shock.** Estimated IRFs from the core seven-variable monetary policy VAR described above are presented in Figure 1. The solid black line in each panel reports the IRF, while dark and light shaded regions report 68% and 90% confidence intervals, computed using a moving block bootstrap as in Jentsch and Lunsford (2019). We calculate a first-stage F-statistic of 16.9, comfortably above the rule-of-thumb value of 10 for weak instruments described by Stock and Yogo (2005).

The impact effect of a monetary policy shock on the 2-year Treasury yield is normalized to a 25bp tightening. After impact, the 2-year Treasury yield increases slightly and then gradually returns to steady state over the next two years. Corporate bond spreads increase by 5bp on impact and rise for several months before gradually returning to steady state. The three other variables typically considered in a monetary policy VAR—unemployment, industrial production, and the CPI—respond more sluggishly, with essentially no effect on impact. After a few months, industrial production declines and the unemployment rate starts to rise, followed by a decrease in the CPI. The peak effect is a little under 0.2 percentage points for the unemployment rate, -0.75 percent for industrial production, and -0.2 percent for the CPI. These responses are similar to those from monetary policy VARs estimated by other

FIGURE 2. Response of Labor Market Flows to a Monetary Policy Shock

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given labor market flow variable to the baseline VAR from Figure 1, where "E" denotes employment, "U" denotes unemployment, and "N" denotes nonparticipation. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Inset boxes report average transition rates.

authors, such as Bauer and Swanson (2023b), and are consistent with the aggregate economy weakening moderately and inflation falling slightly after a monetary policy tightening.

Given our focus on the effect of monetary policy on the labor market, we also estimate the response of vacancies and the labor force participation rate. Vacancies show virtually no effect on impact, but fall after several months, with a peak effect at around -2.5 percent. Labor force participation shows a slow-moving decline beginning around six months after impact, reaching a peak effect of around -0.04 percentage points after three years. Note that, while a negative response of participation would often be interpreted as reflecting a fall in labor supply, we show in Section 3.4, below, that it represents the net effect of two forces: an increase due to the response of supply-driven flows, but an even larger decrease due to higher unemployment interacting with the high average level of U-to-N relative to E-to-N transitions (22.6 percent vs. 3 percent).

3.2. **Responses of Labor Market Flows to a Monetary Policy Shock.** We next extend our core seven-variable VAR to include labor market flows. Extending the VAR to include all six labor market flows (E-to-N, E-to-U, N-to-E, N-to-U, U-to-E, and U-to-N) at once would introduce too many parameters into the VAR, resulting in poor estimates and overfitting, so we extend the baseline VAR with one labor market flow variable at a time, following the approach used by Gertler and Karadi (2015) to analyze financial market responses to monetary policy shocks.

The results for each labor market flow are reported in Figure 2, where each panel corresponds to a separate eight-variable VAR—the seven variables in the baseline VAR, above, plus the labor market flow variable listed at the top of the panel.[13] Within each panel, we also report the average rate for that flow in the inset box—for example, 1.4 percent of employed workers move to unemployment each month, on average, while 25.4 percent of unemployed individuals move to employment.

In response to a 25bp monetary policy tightening, the labor market flows in Figure 2 respond gradually, with either a small or statistically insignificant effect on impact and a peak effect after about 1.5 years. The response of U-to-E and E-to-U flows is consistent with the conventional narrative of a reduction in labor demand due to a weakening economy: the transition rate from unemployment to employment (U-to-E) in the top middle panel of Figure 2 falls significantly in response to the monetary tightening, consistent with a drop in hiring, while the transition rate from employment to unemployment (E-to-U) in the top left panel increases significantly, consistent with an increase in layoffs. This latter increase may seem small at first glance—a little under 0.03 percentage points at its peak—but it is sizeable relative to the steady-state flow of about 1.4 percent each month.[14] Moreover, the increase in E-to-U flows is highly persistent, especially compared to the more transitory increase in E-to-U flows typically seen at the start of a recession (e.g., Elsby, Michaels and Solon, 2009).

Given the conventional wisdom that monetary policy has little effect on labor supply, the response of the flow from nonparticipation to unemployment (N-to-U) in the lower right panel of Figure 2 is more surprising. Following a monetary policy tightening, the rate at which workers enter the labor force from nonemployment to look for a job *increases* significantly. Simultaneously, the symmetric flow from unemployment to nonparticipation (U-to-N) in the top right panel *declines*. The increase in N-to-U and decrease in U-to-N flows tilts the composition of nonemployment (unemployment + nonparticipation) towards the unemployed, increasing the fraction of active searchers, who accordingly find a job at a higher rate. Such a pattern is consistent with individuals increasing their labor supply in response to a weaker economy, as we formalize with our structural model in Section 6.

The flow from employment to nonparticipation (E-to-N) in the bottom left panel of Figure 2 declines modestly around a year after the shock. We show in the next section that the differential response of quits and layoffs is crucial for explaining why the E-to-N

---

[13]IRFs for the seven flow variables are not reported in Figure 2 in the interest of space, and because they are very similar to those from the baseline VAR in Figure 1. For each VAR in Figure 2, the first-stage F-statistic for the instrument is above the Stock and Yogo (2005) rule-of-thumb value of 10.

[14]Because of the differences in average flows, it can be difficult to compare the relative magnitude in the responses across labor market flows. In Appendix C.7, we apply the procedure of Shimer (2012) and Elsby et al. (2015) to quantify the importance of each flow towards shaping the responses of employment, unemployment, and labor force participation.

rate declines in response to a contractionary shock, while the E-to-U rate rises significantly. Finally, the flow from nonparticipation to employment (N-to-E) in the bottom middle panel also falls, although with a more delayed and smaller response than the decline in the U-to-E flow.[15]

Overall, the labor market flow responses in Figure 2 suggest that monetary policy operates through both labor demand and labor supply channels. Although the E-to-U, U-to-E, and N-to-E flow responses are broadly consistent with the conventional wisdom that contractionary monetary policy leads to lower labor demand, the responses of N-to-U and U-to-N flows—and as we will show in Section 3.3, E-to-N flows as well—provide novel evidence suggestive of a labor supply channel.[16] We formalize this interpretation of our estimates in Section 6, where we present a quantitative structural model that matches our estimates through a broad-based increase in household labor supply.

The few other papers that study the responses of labor market flows to monetary policy shocks in the U.S. show less conclusive estimates, e.g., White's (2018) estimates using Romer and Romer (2004) shocks. In Appendix C.1, we re-estimate our main specifications using previously available instruments, including policy-rule residual shocks from Romer and Romer (2004) and high-frequency identification (HFI) shocks à la Gertler and Karadi (2015). We obtain substantially less precise estimates than our baseline. Note, these instruments are subject to critiques of exogeneity and relevance summarized by Ramey (2016), Bauer and Swanson (2023b), and Aruoba and Drechsel (2026, hereafter "AD").
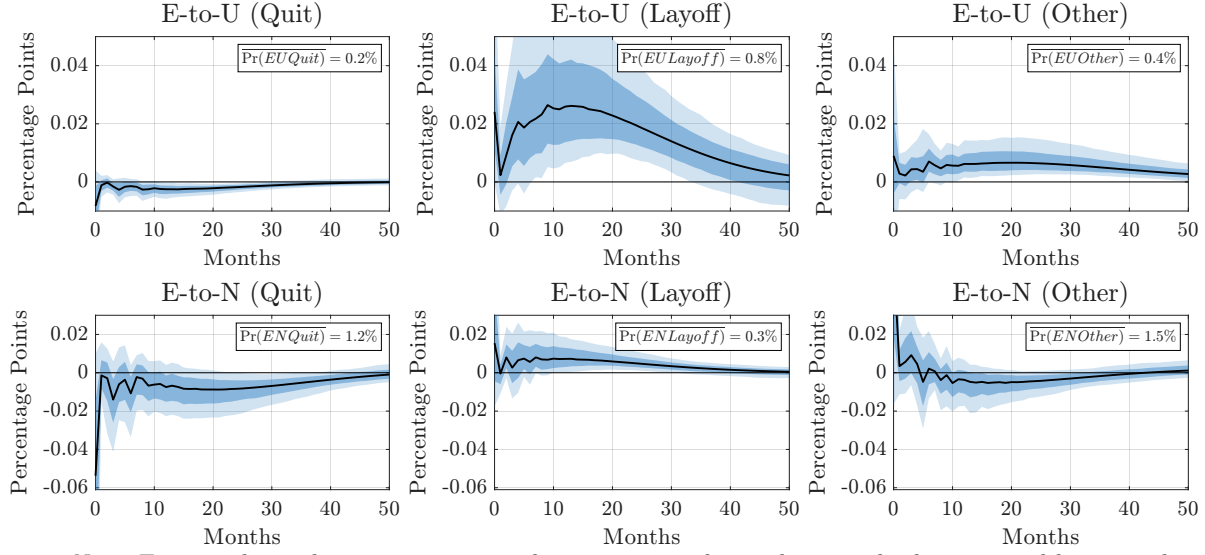
Our baseline approach following Bauer and Swanson (2023a,b) is constructed to address the econometric problems associated with prior HFI surprises. AD develop a complementary approach for policy-rule residual shocks, expanding on Romer and Romer (2004) by incorporating additional information from the text of the Federal Reserve's *Greenbook* as well as Fed staff forecasts. In Appendix C.1, we also estimate impulse responses using AD shocks, as well as a combination of our baseline HFI shocks and AD shocks. The two shock series—each designed independently to confront separate critiques—yield very similar point estimates for the response of labor market flows.

In addition to being robust to different instruments, Appendix C.2 shows that our estimates are robust when a) estimated as local projections, b) estimated in a Bayesian VAR in which all labor market flows and additional macro variables are included simultaneously, and c) the number of lags in the VAR is increased. Finally, in Appendix D, we show that we

---

[15]The structural model we develop in Section 6 also predicts a small rise in the N-to-E rate on impact, through a change in job acceptance behavior.

[16]In Section 4, we show that these results are robust to controlling for cyclical changes in the composition of each employment state. In the Appendix, we quantify the importance of each flow in explaining the response of labor market stocks to a contractionary monetary policy shock, and we show that our estimates are robust to a correction for time-aggregation of labor market flows.

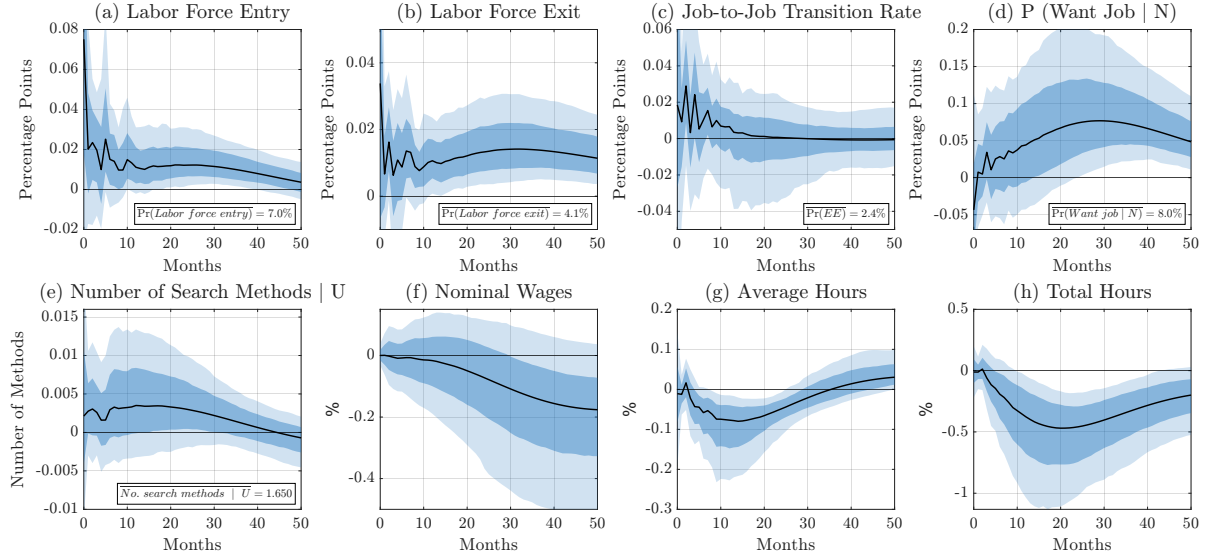FIGURE 3. Decomposition of E-to-U and E-to-N Responses

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given labor market flow variable to the baseline VAR from Figure 1, where "E" denotes employment, "U" denotes unemployment, and "N" denotes nonparticipation. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Inset boxes report average transition rates. Rows share a common scale to facilitate quantitative comparison.

obtain similar results from impulse responses to a Main Business Cycle shock à la Angeletos, Collard and Dellas (2020), a shock with demand-like properties constructed to explain maximal variation of unemployment at business cycle frequencies.

3.3. **Responses of Quits and Layoffs to a Monetary Policy Shock.** We provide further evidence of the response of supply-driven flows by looking at the differential responses of quits and layoffs to a monetary policy shock. Figure 3 reports responses for the quit, layoff and other separation components of both E-to-U and E-to-N flows (defined in Section 2.2) to a 25bp monetary policy tightening. Each of these variables is appended to our core seven-variable VAR one at a time, as in Section 3.2.

We find that layoffs to both unemployment and nonparticipation rise significantly after a monetary policy tightening. Again, this is consistent with the standard narrative of lower labor demand amidst a weakening economy. In contrast, the quit rate to both unemployment and nonparticipation *decreases* after a tightening, consistent with the evidence suggesting an increase in labor supply found in the response of U-to-N and N-to-U flows (as we formalize in the discussion of our model in Section 6). The portion of E-to-U flows that cannot be definitively attributed to layoffs or quits increases modestly, while the unattributed E-to-N

FIGURE 4. Response of Additional Variables to a Monetary Policy Shock

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given variable to the baseline VAR from Figure 1 (except for panel (c), as discussed in the text). Solid black lines report impulse response functions, while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Inset boxes report average values.

flow rate declines slightly with some delay.[17] As layoffs represent a much larger fraction of E-to-U flows than quits, the overall response of E-to-U flows tracks that of the layoffs component. The opposite is true for E-to-N flows: the modest decline in the overall E-to-N rate in response to a contractionary monetary policy shock occurs as the decline in the quit rate to nonparticipation outweighs the rise in layoffs to nonparticipation.

Note that a worker who is laid off has a choice of whether to immediately begin searching (as an E-to-U transition) or enter nonparticipation (as an E-to-N transition). The impulse responses in Figure 3 show a proportionally larger increase in layoffs to U than to N, indicating that the share of laid-off workers immediately searching for work rises following a contractionary monetary policy shock. This finding is also consistent with an increase in labor supply and is replicated in our model.

3.4. **Response of Additional Variables.** Figure 4 plots the estimated impulse response functions of a number of additional labor market variables, again each computed by appending the given variable to the baseline VAR from Section 3.1.

Labor force entry and exit. As shown in panels (a) and (b) of Figure 4, labor force entry and exit rates both increase in response to a contractionary monetary policy shock. Thus,

---

[17]While we do not categorize it as such, this is also consistent with an increase in labor supply. For example, a tightening of monetary policy may lead to a delay in retirement (which constitutes a significant fraction of other separations to nonparticipation).

the overall decline in participation in response to a contractionary monetary policy shock is driven by the increase in labor force exit and mitigated by the increase in labor force entry.

To understand how the responses of labor force exit and entry relate to the responses of the labor market flows discussed in Section 3.2, note that the entry and exit rates satisfy

$$\left(\text{Labor Force Entry Rate}\right)_t = \text{N-to-U}_t + \text{N-to-E}_t, \tag{8}$$

$$\left(\text{Labor Force Exit Rate}\right)_t = u_{t-1} \cdot \text{U-to-N}_t + (1 - u_{t-1}) \cdot \text{E-to-N}_t, \tag{9}$$

where u_{t-1} denotes the unemployment rate. The labor force entry rate is the sum of the flows from nonparticipation to either unemployment or employment, as shown in (8). The increase in labor force entry in response to a contractionary monetary policy shock reflects that the increase in N-to-U flows more than offsets the reduction in N-to-E flows (both shown in previously in Figure 2).

Labor force exit reflects the dynamics of U-to-N and E-to-N flows, as well as those of the unemployment rate u_t itself, as shown in equation (9). Intuitively, an increase in unemployment after a contractionary monetary policy shock tilts the distribution of the labor force towards unemployment, and U-to-N transitions are much more common than E-to-N transitions, about 22.6 percent per month vs. 2.9 percent. This places substantial upward pressure on labor force exits. Recalling from Section 3.2 that both U-to-N and E-to-N flows decrease in response to a monetary contraction, equation (9) implies that the rise in unemployment following a monetary contraction is the sole driver for the increase in labor force exit (and therefore, the entire decline in participation).

Hence, our estimates accord with Hobijn and Şahin (2021), who show that the unconditional dynamics of participation are driven by unemployment. Indeed, as we show in Appendix C.6, the response of supply-driven flows attenuate the overall response of labor force participation to a monetary policy shock. In Section 6, we use our model to further illustrate the disconnect between the participation rate and household labor supply: the estimated model matches the responses of labor market stocks and flows through a broad-based increase in labor supply, yet is also consistent with a decline in the participation rate.

Job-to-Job transitions. An important theoretical literature postulates an important role of job-to-job transitions as drivers of inflation, including Moscarini and Postel-Vinay (2023), Birinci et al. (2024), Alves (2025), Faccini and Melosi (2025),. These papers consider "offer-matching" theories of inflation, whereby competition between firms over workers bids up wages and drives worker mobility across jobs, increasing marginal costs. Such theories imply that the rate of job-to-job changes is a relevant measure of labor market slack: a contractionary monetary policy shock may decrease inflation in part by reducing the rate at which workers meet potential employers, and thus the rate of job-to-job transitions.

To consider the importance of this channel in the monetary transmission mechanism, we estimate the response of the rate of job-to-job transitions to a contractionary monetary policy surprise. We use the measure of job-to-job transitions constructed by Fujita, Moscarini and Postel-Vinay (2024). The estimated IRF is plotted in Panel (c) of Figure 4, showing no significant response of job-to-job transitions to a contractionary monetary policy shock.[18]

Intensive margins of job search. Here, we show that, even within distinct labor market states, workers exhibit behavioral responses to a contractionary monetary policy surprise consistent with an increase in labor supply.

We first look at the fraction of nonparticipants who report wanting a job despite not being engaged in active search. As shown in Table B.2 of the Appendix, such workers are almost four times more likely to move to employment in the following month than nonparticipants who do not want a job, indicating that the stated preference of "wanting a job" is an important indicator that a worker will accept an job offer (and perhaps is more likely to receive one). Panel (d) of Figure 4 shows a robust and persistent increase in the desire to work among nonparticipants following a contractionary monetary policy shock. Thus, at the same time as a larger share of nonparticipants move to unemployment, the remaining nonparticipants are also more likely to want a job.

We then look at the number of job search methods used by workers in unemployment. This metric has been adopted elsewhere in the literature and has been shown to be highly correlated with time spent looking for a job, e.g., Mukoyama, Patterson and Şahin (2018). Panel (e) of Figure 4 shows the response of the number of search methods for unemployed workers: after a contractionary monetary policy surprise, the average number of search methods used by unemployed workers increases modestly.

Wages. Next, we consider the response of nominal wages to a contractionary monetary policy shock, where we adopt the Employment Cost Index (ECI) of the BLS as a measure of nominal wages. As shown in panel (f) of Figure 4, nominal wages do not respond to a contractionary monetary policy shock for approximately a year, after which they begin to decline. As this response is slower than that of the consumer price index (shown in Figure 1), our estimates imply a very modest rise of real wages in the first few years following the shock, before declining back to their steady-state after around four years.

Hours. Finally, we report the response of both average and total hours worked in panels (g) and (h). We find that both average and total hours decline, with the decline in average

---

[18]Note, the series for job-to-job transitions can only be calculated from 1994 (with a correspondingly lower first-stage F-statistic of 7.6). Whereas the original Fujita et al. (2020) series for job-to-job transitions begins in October 1995, we extend the series back to February 1994, when the CPS began collecting information on job-switching for employed workers. As we do for the other flow series, we use the X-13ARIMA-SEATS procedure to impute flows over the period June to September 1995, as described in Footnote 5.

hours explaining around 20% of the decline in total hours. This relatively limited role for changes in average hours explaining changes in total hours in the United States is consistent with previous work studying unconditional variation, and is one motivation for our focus on the extensive margin. Our finding of a modest decline in average hours is consistent with Cantore et al. (2023). In addition, that paper finds notable heterogeneity in the response of hours across the wage distribution, with average hours rising in response to a contractionary shock for workers with the lowest wages—a pattern they interpret as reflecting increased labor supply for this group. This mirrors our results on heterogeneity across education groups that we lay out in the next Section.

In Appendix C.3, we also consider the response of various fiscal variables to a contractionary monetary policy shock. We find no response of federal spending or transfer, which federal tax receipts decline significantly, and federal debt rises. While the response of federal transfers is not significant, we do find a sizeable increase in UI payments, which are largely financed at the state level. The model of Section 6 is consistent with this response.
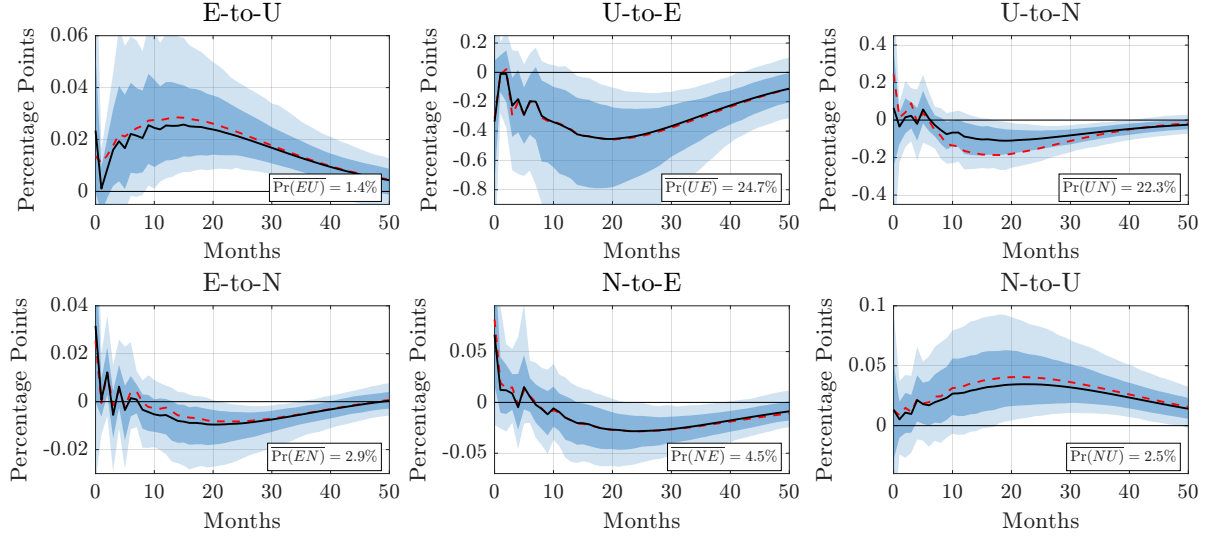
4. Composition and Heterogeneity

The impulse response functions for supply-driven labor market flows in Figures 2 and 3 are suggestive of an increase in labor supply in response to a contractionary monetary policy shock. Here, we establish that these findings are not explained by cyclical changes in the composition of each labor market state, suggesting that our estimated impulse responses reflect true behavioral responses at the individual level. We then explore heterogeneity in the responses across lower- and higher-educated workers.

4.1. **Composition.** Let y_t be an aggregate time series of interest, and $y_{i,t}$ the same time series for a subgroup i with population share $\omega_{i,t}$. Furthermore, denote the time series means of $y_{i,t}$ and $\omega_{i,t}$ as $\bar{y}_i$ and $\bar{\omega}_i$. We can then write

$$y_t = \underbrace{\sum_i y_{i,t} \cdot \bar{\omega}_i}_{\text{variation from } y_{i,t}} + \underbrace{\sum_i \bar{y}_i \cdot \left(\omega_{i,t} - \bar{\omega}_i\right)}_{\text{variation from } \omega_{i,t}} + \underbrace{\sum_i (y_{it} - \bar{y}_i)(\omega_{i,t} - \bar{\omega}_i)}_{\text{covariance}}. \tag{10}$$

The decomposition given by (10) expresses y_t as the sum of three components: a component holding composition fixed, a component allowing composition to vary but holding the variable constant at the group-level, and a final covariance term. Thus, the time series behavior of a variable y_t can be thought of as lying between two extremes: one in which its variation is driven entirely by changes in individual behavior, so that the composition of subgroups remains constant (and only the first term on the right-hand side of (10) is nonzero); and

FIGURE 5. Response of Composition-Adjusted Flows to a Monetary Policy Shock

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given labor market flow variable to the baseline VAR from Figure 1. Solid black lines report impulse response functions for composition-adjusted flows, while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals for composition-adjusted flows. Dashed red lines report impulse responses for unadjusted flows, as in Figure 2.

another in which the time-series variation in y_t is driven entirely by changes in the composition, with individual behavior remaining constant (so that only the second term on the right-hand side of (10) varies over time).

We use this decomposition to estimate the effects of changes in labor force composition on our impulse response functions in Section 3. We follow Elsby et al. (2015) and group individuals according to age (16–24, 25–54, or 55+), gender (male or female), educational attainment (less than high school, high school, some college, or BA+), and reason for unemployment if unemployed (quit, layoff, or other). Thus, we consider 24 subgroups of employed workers, 24 subgroups of nonparticipants, and 72 subgroups of unemployed workers.[19] We construct composition-adjusted labor market flow as the first term on the right-hand side of equation (10), as in Elsby et al. (2015).

As in Section 3.2, we extend our core seven-variable monetary policy VAR to include labor market flows, but now we use the composition-adjusted flows and report the results in Figure 5. Compared to the IRFs for the unadjusted flows (given by the dashed red lines),

---

[19]We differ from Elsby et al. (2015) only in that we do not further classify workers according to their labor market status one year prior (e.g., employment, unemployment, or nonparticipation). Such further classification requires studying CPS respondents in rotation groups five through eight and, as shown by Ahn and Hamilton (2022), workers in later rotation groups are a non-representative sample, displaying lower unemployment rates. Thus, we cannot compare the response of flows from such a sample with those in Figure 2. In Appendix C.4, we show that our conclusions regarding the importance of composition are unchanged when considering the full set of compositional characteristics from Elsby et al. (2015), but that the IRFs of labor market flows are slightly different, consistent with Ahn and Hamilton's findings.

the impulse responses in Figure 5 are largely unchanged. An exception is the IRF for U-to-N flows, which decreases by roughly half as much when holding composition fixed. This suggests that part of the decline in U-to-N flows in response to a monetary contraction does reflect a change in the composition of the unemployed towards workers with greater labor force attachment. While our estimates of the role of composition are somewhat smaller, our findings here echo those of Elsby et al. (2015), who calculate that roughly 75% of the change in U-to-N flows from the end of an expansion through a recession are due to changes in the composition of the unemployed.[20] In Section 6.6, we discuss a similar role for composition in our model that is important for matching the response of U-to-N flows. In Section 5, below, we discuss further how controlling for composition has little effect on our finding that the response of supply-driven flows is quantitatively important for the response of employment to a monetary policy shock.

4.2. **Heterogeneity.** While the results above show that our findings on the response of supply-driven flows are largely robust to controlling for composition effects, they do not preclude heterogeneous responses across different types of workers. Here we study the responses of higher- and lower-educated workers.[21] We study heterogeneity in labor supply responses across other dimensions in Appendix C.5.

   Figure 6 reports impulse responses of employment, the E-to-U flow, the U-to-E flow, and the quit rate from E-to-N for higher- and lower-educated workers, computed by extending our baseline seven-variable VAR one variable at a time, as in previous sections. The left column of Figure 6 shows that employment of higher-educated workers responds modestly to the contraction, reaching a maximum decline of about 0.15 percent at 20 months, while the fall in employment for lower-educated workers is roughly twice as large.
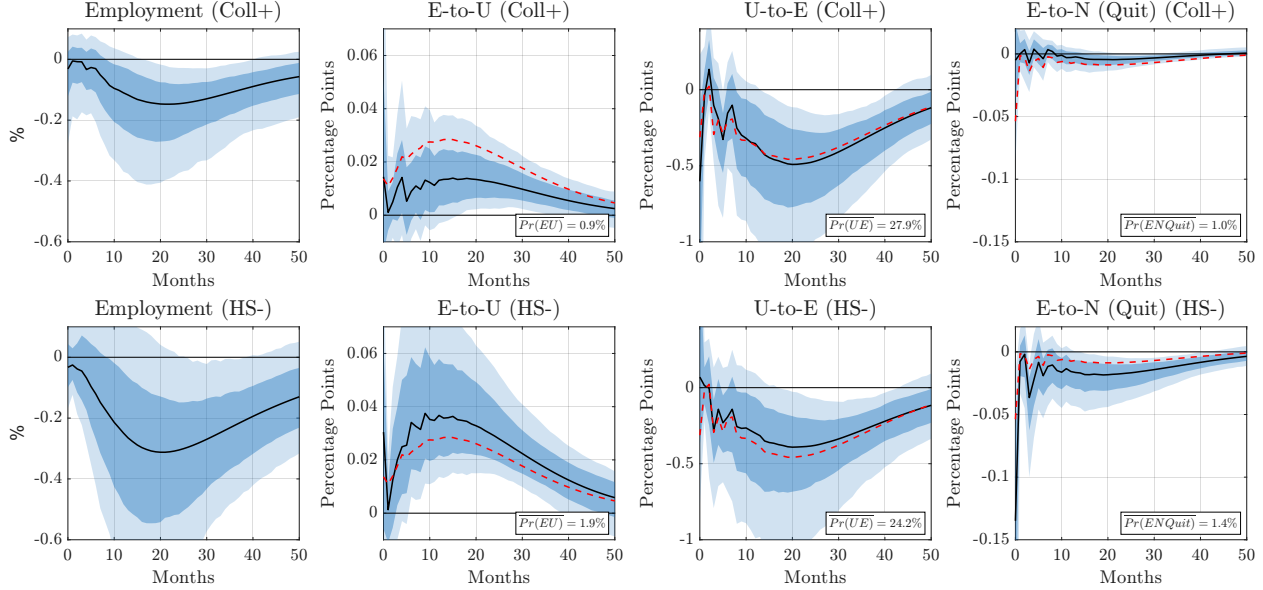
   The second column of Figure 6 reports the responses of E-to-U rates for each education group. The increase in E-to-U flows following a monetary contraction is again about twice as large for lower-educated workers. Meanwhile, the third column shows very similar responses of the U-to-E rate, suggesting that the greater response of layoffs drives the greater employment losses among lower-educated workers. The fourth column plots the response of the quit rate from E-to-N for each group, showing that the decline in quits is concentrated among lower-educated workers, with no discernible drop for higher-educated workers.[22] As we document in Section 6.7, our structural model rationalizes this heterogeneous response in

[20]We conjecture that the greater role for composition found by Elsby et al. (2015) partly reflects their focus on the evolution of U-to-N flows from the *end of an expansion* over the course of a recession, whereas we calculate the impulse response of U-to-N flows starting from steady state (similar to Shimer, 2012).

[21]We classify an individual as higher-educated if they have attended at least some college, and lower-educated if their maximum educational attainment is a high-school diploma or less.

[22]This is not to say that there is no response of supply-driven flows for more educated individuals: In Appendix C we show that the N-to-U rate increases and U-to-N rate decreases for both education groups.

FIGURE 6. Responses by Education

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given variable to the baseline VAR from Figure 1. The top row reports results for individuals with at least some college education (Coll+); the bottom row reports results for individuals with at most a high-school diploma (HS-). Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Red dashed lines report aggregate responses. Inset boxes report average transition rates. Columns share a common scale to facilitate quantitative comparison.

quits as reflecting a greater sensitivity of less-productive workers (who have correspondingly fewer assets) to falls in the job-finding rate.

We see three important takeaways from these estimates. First, monetary policy shocks do not hit all workers equally: lower-educated workers see greater employment declines from a monetary policy contraction, in part from a more responsive layoff margin. Second, the response of supply-driven flows shows important differences across groups, particularly the likelihood of quitting from employment to nonparticipation. Third, to the extent that this heterogeneous response in quits is driven in part by lower asset holdings among the less-educated, our findings suggest that the wealth distribution helps shape the aggregate labor supply response to a monetary policy shock, an explanation that we confirm using the structural model in Section 6.

5. FLOW-BASED ACCOUNTING FOR THE DYNAMICS OF EMPLOYMENT

We now quantify the importance of supply-driven labor market flows for the overall response of employment to a contractionary monetary policy shock.[23] Following Shimer (2012) and Elsby et al. (2015), we account for the contribution of a particular flow by computing the

---

[23]Appendix C reports analogous results for the responses of the unemployment rate and labor force participation rate—see Figures C.15b and C.15c.

FIGURE 7. Flow-Based Accounting for Employment

Note: The black solid line shows the overall response of the employment-population ratio to a contractionary monetary policy shock. The green dashed line shows the response if quits to U or N are held constant. The red dot-dashed line shows the response if both U-to-N and N-to-U rates are held constant. The blue dotted line shows the response if all supply-driven flows are held constant.

hypothetical response of employment when the given flow is held fixed at its average value, using equations (1)–(2). The difference between the implied hypothetical response of employment and the actual employment response provides us with a measure of the quantitative importance of the flow for the employment response.

We perform this analysis using four scenarios: First, a baseline scenario that reports the employment response when all flows respond as estimated in our VAR in Section 3.2. Second, we shut down the response of quits to nonemployment. Third, we shut down the responses of U-to-N and N-to-U flows. Fourth, we shut both the response of quits to nonemployment and the responses of U-to-N and N-to-U flows.

Figure 7 plots the results. In the baseline scenario (the solid black line), employment falls about 0.15 percent after about 20 months. In our second scenario, we shut down the response of quits to nonemployment (the dashed green line). Employment falls about 40% more than in the baseline. As shown in Section 3.3, quits to nonparticipation decline sharply after a monetary contraction: workers become less willing to voluntarily leave employment when labor market conditions deteriorate. Holding this flow fixed removes this margin of adjustment, amplifying the employment decline.

In the third scenario, holding U-to-N and N-to-U flows fixed at their average values (the red dot-dashed line), the fall in employment is almost 60% larger than in the baseline. As discussed in Section 3.3, the increase in N-to-U flows and decrease in U-to-N flows after a monetary contraction tilts the distribution of the non-employed from N towards U, and the

U-to-E transition rate is much higher than the N-to-E rate (25.4 percent per month vs. 4.5 percent). Thus, fixing U↔N flows at their average levels reduces the rate at which workers move from nonemployment to employment after the shock, generating the more pronounced employment decline in Figure 7.

Fourth, we shut down both the response of quits to nonemployment and U↔N flows (the dotted blue line). Employment now declines roughly twice as much as in the baseline, indicating that the response of supply-driven flows is quantitatively important for the overall employment response to a monetary policy shock.[24]

Appendix C.6 shows the results of the same exercise for the unemployment and labor force participation rates. Consistent with our discussion in Section 3.4, the labor force participation rate would decline significantly more without the response of supply-driven labor market flows. Additionally, in Appendix C, we repeat the exercise using our composition-adjusted flows from Section 4. Even after controlling for composition, the decline in employment is about 75% larger when supply-driven labor market flows are held constant. Controlling for composition partly mutes the response of U-to-N flows, but has little effect on the other important supply-driven flows: N-to-U flows and quits to nonparticipation.

6. A STRUCTURAL MODEL OF THE LABOR SUPPLY RESPONSE TO MONETARY POLICY

In the empirical analysis above, we documented a response of labor market flows suggestive of an increase in labor supply following a contractionary monetary policy shock. In Section 4, we showed that this response is not driven by cyclical changes in labor force composition and thus plausibly reflects changes in individual labor supply behavior. In this section, we develop a heterogeneous agent model consistent with this interpretation.

We study the household block of a heterogeneous agent New Keynesian model. We consider an incomplete-markets setting with labor market frictions where individuals make decisions over their consumption, saving, and labor supply: whether to quit, search for, or accept a job. Thus, our model builds upon the framework developed in Krusell et al. (2017) and subsequently adapted elsewhere in the literature, including Alves and Violante (2025).[25]

We consider the effect of a monetary policy shock in the model by feeding in the exact response of the job-finding rate, the layoff rate, the real interest rate, and wages to a contractionary monetary policy shock, studying the labor supply and consumption/savings

---

[24]Note here we are not including the decline in "other separations" to nonparticipation in the labor supply response. This is a conservative assumption, given that such separations, which include retirements as well as individuals that are "tired of working", have similar cyclical properties to quits to nonparticipation and are of a similar magnitude.

[25]Two other recent papers incorporating labor supply decisions into search models include Cairó, Fujita and Morales-Jiménez (2022) and Ferraro and Fiori (2023). These papers consider models somewhat different from our own: the former considers a representative-agent framework abstracting from quits, whereas the latter considers a framework with risk-neutral households whose flow value of leisure is assumed to be procyclical. Thus neither model admits a precautionary motive for labor supply (as we consider here).

response of agents in the model and their implications for aggregate labor market flows. Our approach of treating these particular transition rates and prices as inputs determined outside the model permits a sharper focus on labor supply and allows us to interpret our results without concern about issues of fit relating to the demand or monetary blocks that would naturally arise from estimating a full general equilibrium New Keynesian model.[26] Moreover, the decomposition exercises we perform to assess the role of labor supply in the model (described below) are conditional on these paths, so any model generating them as equilibrium outcomes would yield the same conclusions.

We estimate the model's key parameters to minimize the distance between the impulse response functions for the six labor market transition rates in the model and those reported in Figure 3. Our model is able to closely match both the average level of each of these six transition rates and their dynamic response to a monetary policy shock. Importantly, we also establish that our model does this while being quantitatively consistent with micro-level evidence on marginal propensities to consume (MPCs) and marginal propensities to earn (MPEs).

Having established that the model matches the data, we use it as an accounting framework to decompose the response of employment following a contractionary monetary policy shock. Specifically, we consider the response of employment to each component of the shock individually, and also construct counterfactual paths of employment, holding labor supply policy functions at their steady-state value. This exercise shows that such policy functions change significantly: when they are held fixed, employment declines about 80% more than in the baseline model.

We identify the decline in the job-finding rate as the key driver of the change in labor supply policy functions. By increasing the expected duration of nonemployment, a fall in the job-finding rate increases the drop in consumption associated with moving from employment to nonemployment, and thus lowers the relative value of nonemployment, reducing the probability of quits from employment and increasing the probability that the non-employed search for and accept job offers. The model shows that this effect is more pronounced for low-productivity workers, who are naturally closer to the borrowing constraint, rationalizing the heterogeneous response of quits among low- and high-education workers in the data.

6.1. **Setting.** Time is discrete with an infinite horizon. There is a unit measure of individuals who make decisions over consumption and labor supply subject to a no-borrowing

---

[26]In this sense, our analysis is analogous to Krusell et al. (2017), whose quantitative analysis treats business cycle variation in job-finding rates and layoffs as exogenously determined by demand-side considerations. One limitation of this approach is that we are unable to quantify two-way feedback loops between labor supply (e.g., the decision to search from nonemployment) and labor demand (e.g., vacancy posting), which would require a fully-specified general equilibrium model. In Appendix F we take one step in this direction by estimating a matching function and instead treating vacancies as the primitive. Our experiments there suggest that the feedback from labor supply to the job-finding rate may be relatively modest.

constraint and a number of exogenous shocks: First, individual labor productivity z follows an AR(1) process in logs,

$$\log z' = \rho \log z + \epsilon'_z, \quad \epsilon'_z \sim N(0, \sigma_z^2)$$

with $\rho \in [0, 1)$, where employed workers receive total labor income that is proportional to their productivity. We interpret this process as capturing not only shocks to earnings, but also any other shock that affects an individual's willingness to work.

Second, for non-employed individuals, the cost of active job search, κ, is i.i.d. each period following a logistic distribution with mean μ_κ and scale s_κ. Finally, both employed and non-employed individuals make labor supply decisions in the presence of shocks to their labor market status. An individual's labor market state is either employed (E), non-employed but eligible for unemployment insurance (UI), or non-employed and not eligible for unemployment insurance (No UI).

Employed individuals may choose to quit their job and move to UI-ineligible non-employment in the following period. If they don't quit, they may still be laid off with probability δ_L, in which case they will move to UI-eligible nonemployment.

Non-employed individuals can choose to actively search for a job. Active search raises an individual's job-finding rate, but is subject to the stochastic cost described above. For the UI-eligible non-employed, active search is also required for there to be any possibility of UI-eligibility in the next period. If non-employed individuals receive a job offer, they face a job acceptance decision. For the UI-eligible non-employed who search and either do not receive a job offer, or who reject one, there is an exogenous probability of UI expiry, δ_{UI}.

6.2. **Value Functions.** Let $V_E(a, z)$, $V_{UI}(a, z, \kappa)$, and $V_{NoUI}(a, z, \kappa)$ represent the values of being employed, UI-eligible non-employed, and UI-ineligible non-employed, defined over assets a, productivity z, and, in the case of non-employed agents, their cost of active job search, κ. For clarity, we describe here the model equations in steady-state and thus suppress time subscripts.

An employed worker chooses consumption c, asset holdings a', and whether or not to quit. Accordingly, $V_E(a, z)$ can be expressed as follows:

$$V_E(a, z) = \max_{c, a'} \left\{ u(c) + \beta \max \left\{ \mathbb{E} \, V_{NoUI}(a', z', \kappa'), \mathbb{E} \left[\delta_L V_{UI}(a', z', \kappa') + (1 - \delta_L) V_E(a', z') \right] \right\} \right\}$$

$$\tag{11}$$

subject to

$$c + a' = \bar{R}a + (1 - \tau) wz + T, \quad a' \geq 0, \tag{12}$$

where the mathematical expectation operator $\mathbb{E}$ is conditional on the state (a, z), the max operator is over the values of quitting and not quitting, β denotes the worker's discount

factor, $\bar{R}$ the gross real interest rate, τ the tax rate on labor income, w the real wage, and T a real lump-sum transfer. If a worker quits, she moves to UI-ineligible nonemployment. If she chooses not to quit, she still faces a risk of losing her job exogenously with probability δ_L, in which case she moves to UI-eligible nonemployment.

A UI-eligible non-employed worker chooses consumption c, asset holdings a', whether or not to search, and whether or not to accept a job should she receive an offer. Thus, the value $V_{UI}(a, z, \kappa)$ satisfies

$$V_{UI}(a, z, \kappa) = \max_{c,a'}\left\{u(c) + \max\left\{(1 - \kappa)\psi + \beta\mathcal{V}_{UI}^s(a', z), \psi + \beta\mathcal{V}_{UI}^{ns}(a', z)\right\}\right\} \tag{13}$$

subject to

$$c + a' = \bar{R}a + (1 - \tau)\min\{\phi w z, \bar{\phi}\} + T, \quad a' \geq 0, \tag{14}$$

where ψ denotes the flow utility value of nonemployment, ϕ denotes the replacement rate from unemployment insurance, with a maximum possible UI benefit of $\bar{\phi}$. The max operator is taken over the values of searching—in which case the worker receives flow utility $(1 - \kappa)\psi$ and continuation value $\mathcal{V}_{UI}^s(a', z)$—and not searching—in which case the worker receives flow utility ψ and continuation value $\mathcal{V}_{UI}^{ns}(a', z)$.

The terms $\mathcal{V}_{UI}^s(a', z)$ and $\mathcal{V}_{UI}^{ns}(a', z)$ appearing in (13) reflect the expected continuation values associated with searching and not searching:

$$\mathcal{V}_{UI}^s(a', z) = f_s \cdot \max\{\mathbb{E}\,V_E(a', z'), \mathbb{E}\,\tilde{V}_{UI}(a', z', \kappa')\} + (1 - f_s)\,\mathbb{E}\,\tilde{V}_{UI}(a', z', \kappa') \tag{15}$$

$$\mathcal{V}_{UI}^{ns}(a', z) = f_{ns} \cdot \max\{\mathbb{E}\,V_E(a', z'), \mathbb{E}\,V_{NoUI}(a', z', \kappa')\} + (1 - f_{ns})\,\mathbb{E}\,V_{NoUI}(a', z', \kappa'), \tag{16}$$

where searchers find jobs at a higher probability, $f_s > f_{ns} > 0$, and $\tilde{V}_U(a, z, \kappa)$ expresses the expected value of unemployment taking into account the realization of exogenous benefit exhaustion:

$$\tilde{V}_{UI}(a, z, \kappa) = \delta_{UI}V_{NoUI}(a, z, \kappa) + (1 - \delta_{UI})V_{UI}(a, z, \kappa). \tag{17}$$

The presence of $\tilde{V}_{UI}$ in (15) but not (16) reflects that the worker is additionally incentivized to search to retain access to UI benefits. Note that the values defined by (15) and (16) encode the worker's optimal decision of whether or not to accept a job offer (if received) through the max operator.

Finally, a UI-ineligible non-employed worker faces the same menu of decisions as a UI-eligible non-employed individual, with

$$V_{NoUI}(a, z, \kappa) = \max_{c, a'} \left\{ u(c) + \max \left\{ (1 - \kappa)\psi + \beta \mathcal{V}^s_{NoUI}(a', z), \psi + \beta \mathcal{V}^{ns}_{NoUI}(a', z) \right\} \right\} \quad (18)$$

subject to

$$c + a' = \bar{R}a + T, \quad a' \geq 0, \quad (19)$$

and the law of motion for z and the distribution of κ, where

$$\mathcal{V}^s_{NoUI}(a', z) = f_s \cdot \max\{\mathbb{E}\, V_E(a', z'), \mathbb{E}\, V_{NoUI}(a', z', \kappa')\} + (1 - f_s)\, \mathbb{E}\, V_{NoUI}(a', z', \kappa') \quad (20)$$

$$\mathcal{V}^{ns}_{NoUI}(a', z) = f_{ns} \cdot \max\{\mathbb{E}\, V_E(a', z'), \mathbb{E}\, V_{NoUI}(a', z', \kappa')\} + (1 - f_{ns})\, \mathbb{E}\, V_{NoUI}(a', z', \kappa'). \quad (21)$$

As in the case of the UI-eligible non-employed, a UI-ineligible non-employed individual faces a tradeoff between the potential utility benefit of not searching with higher job-finding probabilities, summarized by the terms $\mathcal{V}^s_{NoUI}$ and $\mathcal{V}^{ns}_{NoUI}$ (similar to (15) and (16), but without implications for future UI benefits).

6.3. **Calibration and Estimation.** Our estimation procedure broadly follows Christiano, Eichenbaum and Evans (2005) and Auclert, Rognlie and Straub (2020): we set a number of our model parameters based on clear external evidence or consensus in the literature and estimate the remainder to minimize the distance between the model and our empirical impulse response functions. The model period is one month. We assume $u(c) = c^{1-\gamma}/(1-\gamma)$ and $f_{ns} = \alpha f_s$.

The parameters that we calibrate are $\theta_{EXT} \equiv \{\gamma, \beta, \bar{R}, w, \alpha, \delta_{UI}, \phi, \bar{\phi}, \tau, T\}$. We set $\gamma = 2$, a standard value, and $\beta = 0.988$, in order to generate a quarterly MPC in the model in the range of 7-8%.[27] We set $\bar{R}$ to imply a steady-state annual real interest rate of 1% and normalize the real wage w to 1. We calibrate the rate at which nonparticipants receive job offers to 60% of that at which active searchers receive offers, consistent with the average transition rate among nonparticipants who want a job relative to the average U-to-E transition rate.[28] We set $\delta_{UI} = \frac{1}{6}$, implying that UI lasts six months on average, as in most states in the US in normal times. The UI replacement rate and upper bound, income tax rate, and lump-sum transfer are set to match US evidence, as described in Auclert, Bardóczy and Rognlie (2021) and Graves (2025).[29]

---

[27]See Kaplan and Violante (2022) for a discussion of methods of calibrating discount factors and effects on implied MPCs.

[28]We report N-to-E | Want in Table B.2 of the Appendix.

[29]In the stationary distribution of our model, unemployment insurance is equal to 1% of labor compensation. This is marginally above the figure of 0.75% reported in Krusell et al. (2017).

We estimate the remaining parameters, $\theta_{EST} \equiv \{\rho_z, \sigma_z, \mu_\kappa, \sigma_\kappa, \psi, \delta_L, f_s\}$, which govern the idiosyncratic productivity and search cost processes, the value of leisure, and the steady-state layoff and job-finding rates. In most cases, these parameters do not have a clear mapping to a single moment, and thus must be jointly estimated. For example, the average job-finding rate f_s in the model is distinct from the U-to-E rate in the data, as an endogenous fraction of job offers will be rejected by unemployed agents in the model.

Starting from steady state, we consider the effect in the model of an unanticipated monetary policy shock that changes the real interest rate, real wage, layoff rate and job-finding rate to match the estimates in Section 3. The paths of the real interest rate and real wage are taken from the response of the equivalent nominal series in Figures 1 and 4, deflated by the response of CPI in Figure 1. The response of the layoff rate is the sum of the response of the layoff components of E-to-U and E-to-N flows in Figure 3. Finally, the path of the job-finding rate is chosen to match the response of the U-to-E rate in Figure 2.[30]

We feed these paths $\{R_t, f_{s,t}, \delta_{L,t}, w_t\}_{t=0}^T$ into the model and calculate the responses of the six labor market transition rates: $J(\theta_{EST}) = \{EU_t, EN_t, UE_t, UN_t, NE_t, NU_t\}_{t=0}^{50}$. These series are in levels, and thus take into account both the steady-state value of each rate and its response to the aggregate shock.[31] Denoting by $\hat{J}$ the impulse response functions estimated in the data, our estimator is

$$\min_{\theta_{EST}} (J(\theta_{EST}) - \hat{J})' \Sigma^{-1} (J(\theta_{EST}) - \hat{J}), \tag{22}$$

where Σ is a diagonal matrix containing the estimated variances of the empirical impulse responses. As in Christiano et al. (2005) and Auclert et al. (2020), standard errors for our estimated parameters are calculated using the delta method. The externally calibrated and internally estimated parameters are reported in Table 4.[32]

6.4. **Results: Steady State.** Before studying the model's dynamic responses to a monetary policy shock, we briefly discuss its steady-state properties. First, we consider the model's labor supply policy functions. We then calculate the MPC and MPE from an idiosyncratic transfer in the model and compare it to empirical estimates in the literature. In Appendix E we show that the model closely matches the steady-state values of all six labor market transition rates.

---

[30]We use an iterative procedure to find this path for the job-finding rate, given that it does not map exactly to the U-to-E rate due to a small fraction of the unemployed that reject job offers in the model.

[31]We truncate the responses after 50 months, as in Figure 2. In order to generate smooth responses of the six transition rates, we introduce very small taste shocks for the discrete choices that individuals face over quitting or accepting jobs. Appendix E describes this approach and provides further computational details.

[32]Our estimates for the labor productivity process imply less persistence and more volatility than typically found in labor income data. Under our reduced-form interpretation of this process as capturing all shocks that affect a worker's willingness to work, this suggests a role for shocks to the value of leisure or job-related amenities.

TABLE 4. Model Parameters

	Calibrated		
Parameter	Description	Value	Source/Target
β	Discount factor	0.988	Quarterly MPC of 7-8%
R	Steady state real interest rate	1.001	1% Annual
γ	Risk Aversion Coefficient	2	Standard value
δ^{UI}	Benefit Exhaustion Probability	0.167	Expected duration of UI
w	Wage	1	Normalization
α	Efficiency of Passive Search	0.6	N-to-E | want a job
ϕ	UI Replacement Rate	0.50	Graves (2025)
$\bar{\phi}$	Maximum UI payments	2.15	Graves (2025)
τ	Labor income tax rate	0.33	Auclert et al. (2021)
T	Lump-sum Transfer	0.28	Auclert et al. (2021)

	Estimated		
Parameter	Description	Value	Standard Error
ρ_z	Persistence of Labor Productivity	0.961	(0.013)
σ_z	Standard Deviation of Labor Productivity	0.392	(0.025)
μ_κ	Mean value of Search Cost	0.878	(0.181)
σ_κ	Dispersion of Search cost	0.188	(0.041)
ψ	Value of Leisure	0.318	(0.215)
δ	Steady-state Layoff Rate	0.019	(0.003)
f_s	Steady-state Job-Finding Rate	0.272	(0.029)

Note: Standard errors for estimated parameters are calculated using the delta method. See text for details.

FIGURE 8. Labor Supply Policy Functions

Note: The left plot shows the probability that an employed individual quits their job, for different levels of assets on the x-axis and labor productivity on the y-axis. The middle plot shows the probability that a UI-eligible individual searches for a job, before the realization of their search cost. The right plot shows the probability that a UI-eligible individual accepts a job. For exposition, we truncate the asset grid at twice average wealth.

6.4.1. *Labor Supply Policy Functions.* Figure 8 plots the probability of quitting a job (for employed workers), and searching for and accepting a job (for UI-eligible workers) at different levels of idiosyncratic productivity and assets. The policy functions show considerable heterogeneity in the propensity of workers to quit to nonemployment, search for a job, or accept a job as a function of their wealth and labor productivity. Thus, the model displays substantial variation in labor supply policies indicative of a worker's degree of "labor force attachment" within labor market states.

As indicated in the left panel, individuals are more likely to quit from employment to nonparticipation (thereby making an E-to-N transition) the lower their productivity and the higher their wealth. The middle panel shows that individuals are more likely to search when they have low wealth or high productivity. A worker is UI-eligible if she was laid off in the prior period, or if she maintained active search since being laid off while also maintaining exogenous UI eligibility. Thus, the middle panel illustrates which workers will move to unemployment (versus nonparticipation) after a layoff, and which workers will continue searching from unemployment (versus stopping search and moving to nonparticipation) after a period of unsuccessful job search.

The right panel shows that workers are more likely to accept a job offer when they have low wealth or high productivity—mirroring the quit decision in the left panel. The middle and right panels together show intermediate combinations of wealth and productivity where individuals do not search but will accept a job if offered. The model matches the lower N-to-E rate of such workers through two channels: the lower job-finding probability of non-searchers and the share of non-searchers who reject offers. More broadly, matching all aggregate flows depends on the distribution over productivity and wealth, not just across labor market states.

6.4.2. *The Marginal Propensities to Consume and Earn.* We calculate the marginal propensity to consume (MPC) and marginal propensity to earn (MPE) from an unexpected transfer equivalent to approximately $500. Our model generates a quarterly MPC of 7.3%. While lower than estimates from the earlier literature (e.g., Parker et al., 2013), this is consistent with more recent studies that have identified potential biases in the earlier estimates (e.g., Borusyak, Jaravel and Spiess, 2024; Orchard, Ramey and Wieland, 2023; Boehm, Fize and Jaravel, 2024). These more recent studies suggest that the "notional MPC"—the appropriate MPC to target for a model of nondurable consumption—should lie in a range between about 7-11%.[33] Our MPC falls in this range.

For MPEs, we rely on Golosov et al. (2023), who find that households reduce annual earnings by $2.3 per $100 of lottery winnings on average. The MPE is larger for smaller

---

[33]Laibson, Maxted and Moll (2022) develop the "notional MPC" to address the fact that empirical estimates often include spending on durables, while most models do not. See the discussion of Table III of Boehm et al. (2024) for more details of recent estimates.

lotteries, with earnings falling by almost $6 for the smallest lotteries they consider. This is the most relevant comparison for our transfer size.[34] In our model, earnings fall by $4.7 per $100 of transfer, consistent with this evidence.[35]

6.5. **Results: Model Dynamics.** We now turn to the dynamic properties of the model in response to a contractionary monetary policy shock. Figure 9 combines the empirical estimates from Section 3 with the model-generated impulse response functions overlayed as dashed magenta lines.

Figure 9(A) shows that the model achieves a good fit to the targeted moments from the data: the labor market flows implied by the estimated model are remarkably close to those estimated from data, with the models' impulse responses lying within the 68% confidence bands for almost all labor market flows and horizons. The close fit of the model to the data is not guaranteed but achieved through the optimal labor supply response of workers to the change in the path of job-finding probabilities, layoff rates, interest rates, and wages. For example, the model matches the shallower drop in N-to-E flows versus U-to-E flows in part as nonparticipants become more likely to accept jobs. The model also matches the responses of U-to-N and N-to-U flows: the probability that an unemployed individual stops searching falls, while the probability that a nonparticipant starts searching rises.

Figure 9(B) compares the model's responses of quits and layoffs to unemployment and nonparticipation to our empirical estimates from Figure 3. While these impulse responses are not targeted in the estimation, the fit is good. The model matches the prolonged decline in quits to nonparticipation after a monetary contraction, as well as the much larger rise in the E-to-U layoff rate than the E-to-N layoff rate—implying that the share of laid-off workers who move to unemployment (rather than nonparticipation) increases, itself evidence of an increased willingness to work.

Finally, Figure 9(C) shows the model-implied response of the unemployment rate, labor force participation rate, and employment rate, which show a close fit against those from the data. As described in the discussion of equation (2), given an initial condition, the dynamics of stocks are determined entirely by those of flows. Hence, the goodness-of-fit here reflects that of labor market flows in Figure 9(A).[36]
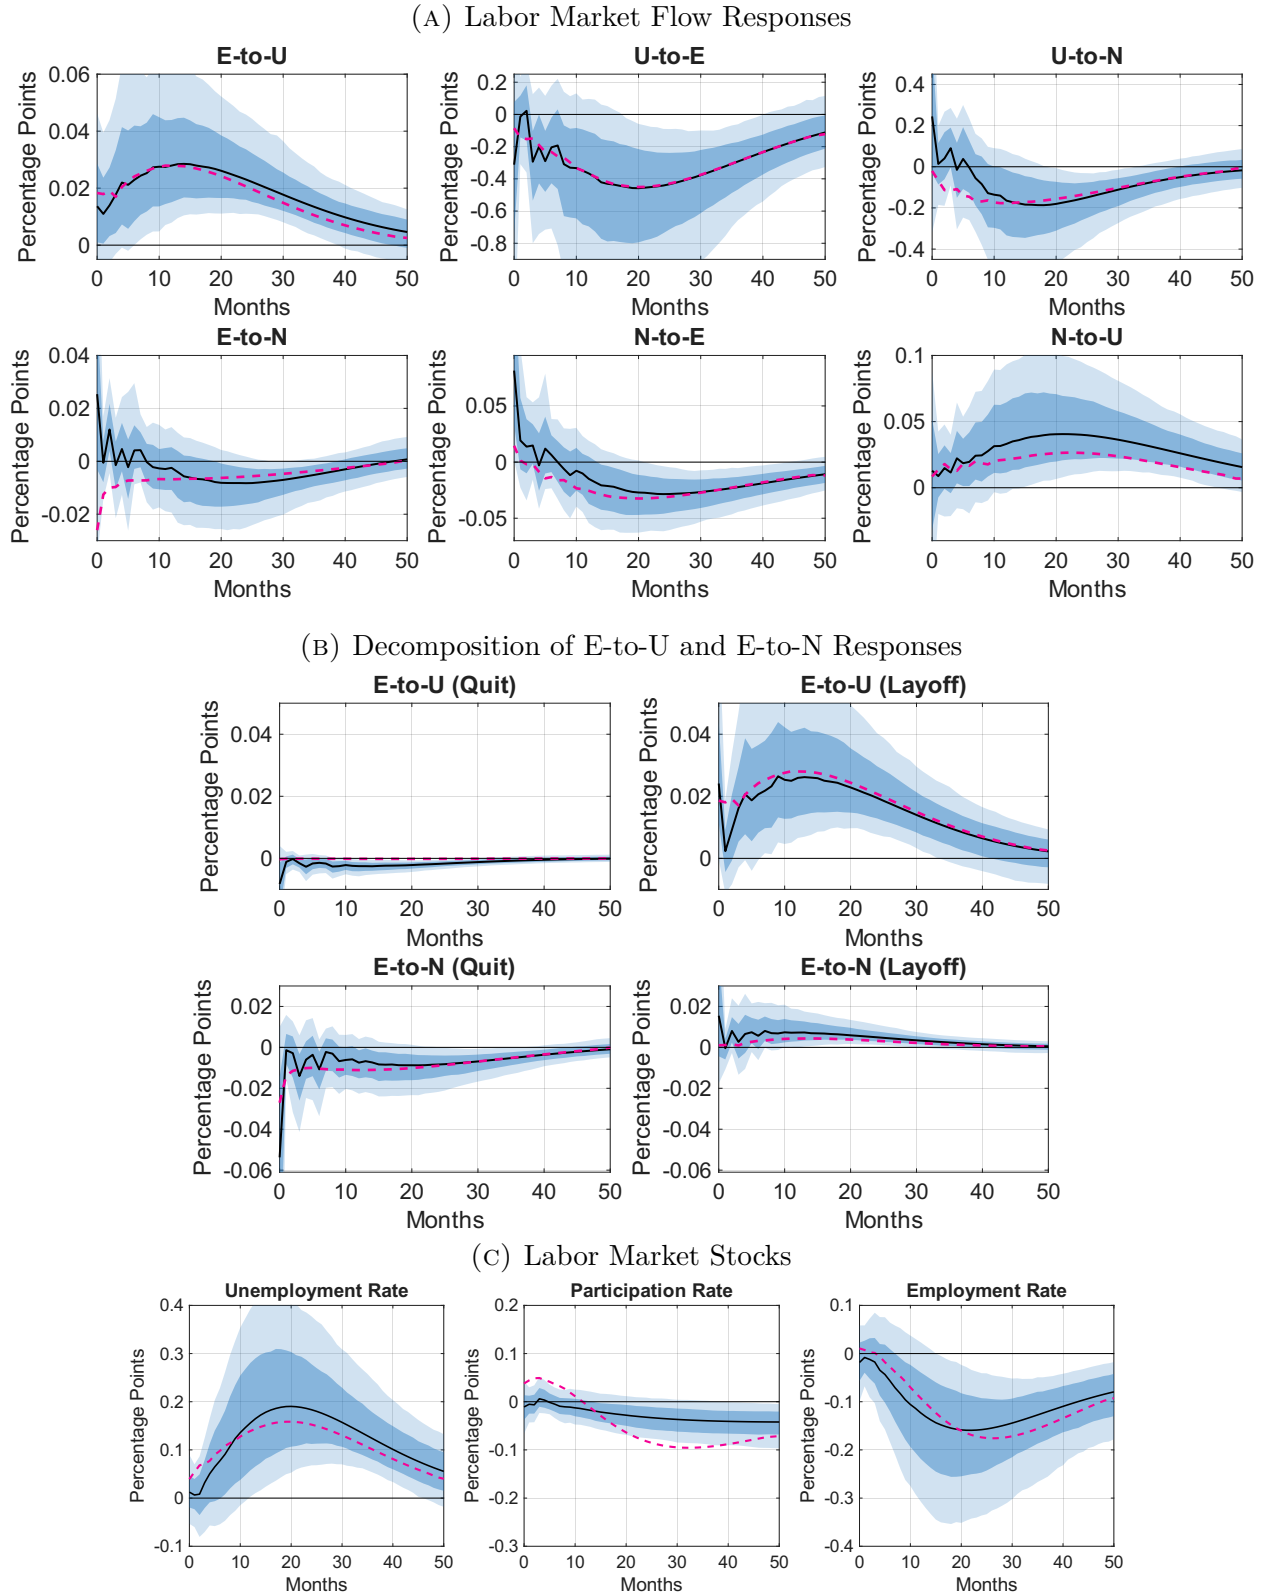
6.6. **Decomposing A Monetary Policy Shock: The Role of Labor Supply.** Having shown that the model offers an excellent fit to our estimates of the response of labor market

---

[34]These calculations refer to the average response of earnings in each of the five years following the lottery win. See Figure B.6 of the Online Appendix of Golosov et al. (2023) for the comparison across lottery sizes.
[35]Auclert et al. (2021) document a tight connection between MPCs and MPEs in models with frictionless labor markets and an intensive margin of labor supply. The presence of labor market frictions allows our model to break this link and produce an MPE which is significantly smaller than the MPC, in line with the data.
[36]We replace the participation rate in the main VAR with the employment rate to estimate the response of the latter.

FIGURE 9. Response of Labor Market Flows and Stocks: Model and Data

(A) Labor Market Flow Responses

(B) Decomposition of E-to-U and E-to-N Responses

(C) Labor Market Stocks

Note: Estimated impulse responses to a 25bp monetary policy tightening shock computed by appending the given labor market variable to the baseline VAR from Figure 1 where necessary. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Dashed magenta lines report impulse response functions from the estimated model. See text for details.

flows and stocks to a contractionary monetary policy shock, we now use the model to understand the key drivers of the response of employment. First, we show the separate roles of the four components of a monetary policy shock in our model: the response of the job-finding rate, layoff rate, real interest rate and real wage. We also show the extent to which these components, and the shock overall, induces a shift in labor supply policy functions. Our main finding is that a decline in the job-finding rate leads to a significant shift in labor supply policy functions and an increase in the aggregate propensity to work, as we explain further in Section 6.7.

The top panel of Figure 10 shows the separate contribution of each of the four components of the monetary policy shock for the response of employment. The black lines show the baseline version of the model, while the dotted and dashed lines show the response of employment under counterfactuals where various labor supply policy functions (as a function of worker characteristics) do not respond to the shock. These counterfactuals allow us to identify the particular role of shifts in labor supply policy functions in generating the overall response of employment, as opposed to changes in exogenous transition probabilities or shifts in the distribution of households across idiosyncratic states.
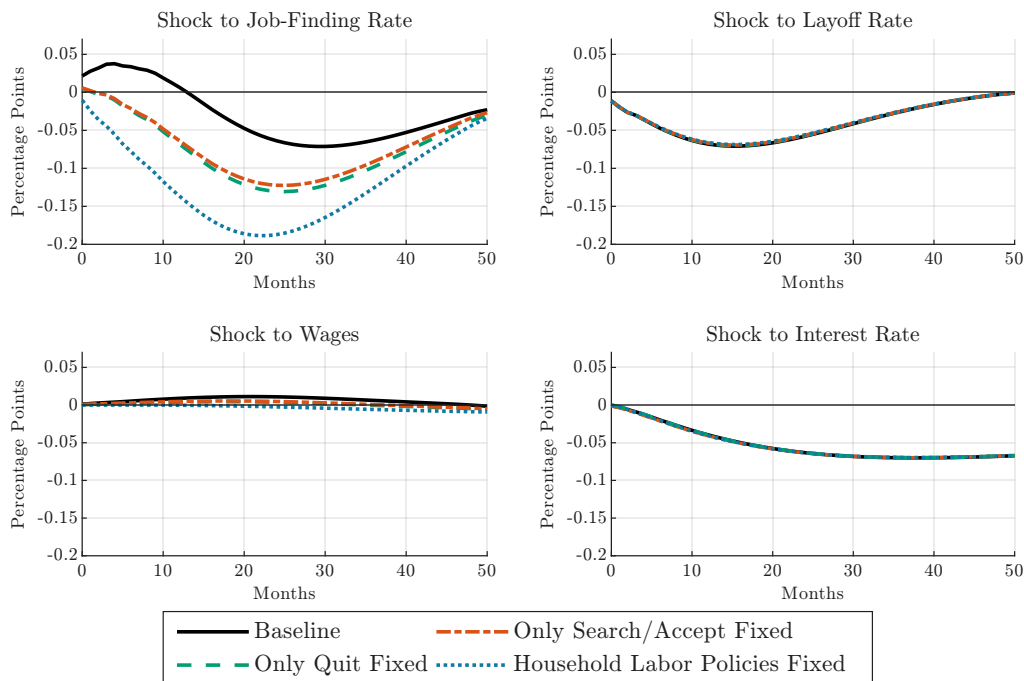
The top row of Figure 10(A) shows that the decline in the job-finding rate and the increase in the layoff rate contribute roughly equally to the overall decline in employment, albeit at different horizons (solid black lines). However, their effects on labor supply differ markedly. When we hold labor supply policies fixed, the employment decline from the job-finding rate shock alone is significantly larger—workers respond to the lower job-finding rate by reducing quits and increasing search effort, moderating the overall decline in employment. In contrast, the layoff rate shock has almost no effect on labor supply policy functions: the solid and dashed lines nearly coincide. We explain in Section 6.7 why the fall in the job-finding rate drives such a shift in labor supply policy functions.

There are also a number of interesting features of the bottom row of Figure 10(A). First, we find that the modest response of real wages that we estimate in the data has a correspondingly modest role in determining the path of employment. Despite this, it is also clear that shocks to real wages do shift labor supply policy functions: the small increase in employment from the shock to wages is nullified when labor supply policies are held fixed. Finally, we find that the increase in the real interest rate, both due to higher nominal rates and also the decline in inflation, leads to a small but persistent decline in employment in the model. As the dashed and solid lines lie on top of each other in this panel, we see that this occurs not because of changes in labor supply policy functions but because this shock moves the distribution of households towards states in which they hold sufficient wealth that they are more likely to quit and less likely to search for or accept jobs.[37]

---

[37]We abstract from a separate but related channel, whereby a monetary policy shock affects asset prices as in, for example, Melcangi and Sterk (2025).

FIGURE 10. Decomposing the Employment Response to a Monetary Policy Shock

(A) Separate Components of A Monetary Policy Shock

(B) Response to All Four Components

Note: The black solid line shows the overall response of the employment-population ratio to (the components of) a contractionary monetary policy shock in the estimated model. The blue dotted line shows the response if all labor supply policy functions are held at their steady-state values. The green dashed line shows the response if only the quit policy function is held constant. The red dot-dashed line shows the response if only the search/accept policy functions are held constant.

While here we focus on the employment response, Figure E.1 of Appendix E.3 shows the effect of each component of the shock on each labor market flow. Consistent with our findings here, the job-finding rate plays the primary role in shaping the response of supply-driven flows. Other components of the shock matter as well, but through different channels. For example, the increase in layoffs is important for matching the response of flows from unemployment to nonparticipation. Layoffs shift the composition of the unemployed towards workers with greater labor force attachment—higher productivity and/or lower wealth—which lowers the U-to-N rate. The results above and in the next section confirm that this compositional effect, rather than a change in labor supply policy functions, explains the decline in the U-to-N rate.

In Figure 10(B) we repeat the counterfactual exercise of panel (A) but allowing all four shocks to occur simultaneously. We find that employment declines by around 80% more when labor supply policy functions are held fixed, with a roughly equal role of quits and search/accept decisions. Thus, through the lens of the model we see that an increase in aggregate propensity to work significantly attenuates the decline in employment following a contractionary monetary policy shock. In the final section we will unpack this change in labor supply further.[38]
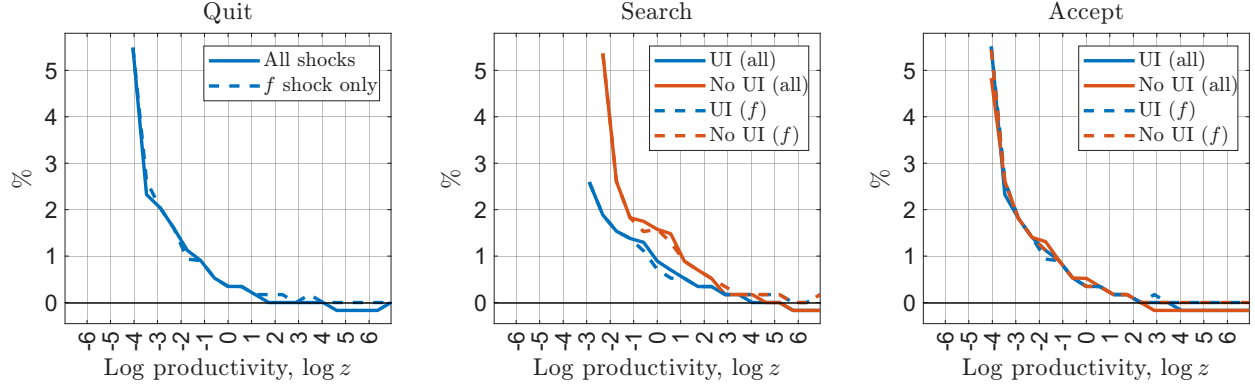
6.7. **Explaining the response of labor supply.** The results in the previous section show that the decline in the job-finding rate following a contractionary monetary policy shock generates a significant shift in labor supply policy functions that is not seen in response to any of the other components of the monetary policy shock. Here, we explain why changes in the job-finding rate shape labor supply decisions. We also show how the response of the job-finding rate generates heterogeneous responses of labor supply decisions by idiosyncratic productivity and wealth, consistent with our estimates from the data.

We begin by showing how labor supply policy functions shift across the distribution. Note from Figure 8 that, for each level of idiosyncratic labor productivity z, the relative value of quitting is increasing in wealth, a; whereas the relative value of searching and accepting is decreasing in a.[39] Thus, for each level of idiosyncratic productivity z, we define the *marginal quitter* by the lowest level of wealth at which an employed worker is willing to quit; the *marginal searcher* by the maximal level of wealth at which a non-employed worker is willing

---

[38]Note that this analysis is distinct from the flow-based accounting exercise in Section 5, where we held supply-driven labor market flows fixed at their steady-state levels. Here, we instead hold individual labor supply policies fixed, potentially allowing aggregate labor market flows to vary as the composition of the labor force changes. While distinct, both exercises are complementary in illustrating the importance of labor supply considerations for the employment response to monetary policy.

[39]Even for the highest productivity level there are (very high) wealth levels where such individuals quit, do not search and would not accept a job offer. This is not shown in Figure 8 as we truncate the asset grid for exposition.

FIGURE 11. Change in Wealth of Marginal Quitter, Searcher and Accepter After Shock

Note: The figure plots the percent change in wealth identifying the marginal quitter, searcher, and accepter between $t = 0$ and $t = 1$ (when the aggregate shock occurs). The dashed lines shows the change in wealth for a shock only to the path of the job-finding rate. The absence of a value associated to a level of productivity indicates either that (a) all workers quit, (b) no workers search, or (c) no workers accept at that level of productivity.

TABLE 5. Consumption and Wealth of the Median Marginal Quitter: SS and Shock

|	Wealth	Cons. Drop (%, SS)	Cons. Drop (%, Shock)
Marginal Quitter (SS)	62.03	1.50	1.65
Marginal Quitter (Shock)	62.25	1.40	1.57
% Change	0.35		

Note: This table shows the level of wealth and the consumption drop upon moving to nonemployment for the "marginal quitter" at the median level of productivity. The first row corresponds to the marginal quitter in steady-state (SS). The second corresponds to the marginal quitter after the shock. As the median productivity and the steady-state wage in the model are both equal to one, wealth is scaled relative to monthly pre-tax income.

to actively search for a job; and the *marginal accepter* by the maximal level of wealth at which a non-employed worker is willing to accept a job offer.

Figure 11 shows the percent change in the level of wealth for the marginal quitter, searcher, and accepter on impact following the monetary policy shock in the model. The solid line plots the change induced by all of the shocks (job-finding, layoffs, real interest rate, and real wages), whereas the dashed line plots the change induced by the shock to the job-finding rate alone. The absence of a value associated to a level of productivity in a subplot indicates the absence of a marginal worker at that productivity.

In general, the contractionary monetary policy shock leads to an increase in the wealth of the marginal quitter, searcher, and accepter. A higher wealth threshold for quitting means fewer employed workers lie above it, reducing quits; higher thresholds for searching and accepting mean more non-employed workers lie below them, increasing search and acceptance

FIGURE 12. Heterogeneity in Quit Response: Model vs Data

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given labor market flow variable to the baseline VAR from Figure 1. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Dashed magenta lines report impulse response functions from the estimated model. See text for details.

probabilities. These shifts are consistent with a decrease in the relative value of nonemployment across the state space. The fact that the solid lines lie very close to the dashed lines is further evidence that changes in the job-finding rate are the key driver of shifts in labor supply policy functions.

Intuitively, a reduction in the job-finding rate increases the expected duration of nonemployment. A worker's assets must therefore stretch over a longer period without labor income, reducing per-period consumption during nonemployment. Some workers in employment who previously found it optimal to quit to nonemployment will no longer do so; and additional workers in nonemployment will find it optimal to either search or to accept a job offer. The wealth threshold defining the marginal quitter therefore rises.

Table 5 explores this logic for the marginal quitter at the median level of productivity. In steady state, this individual has wealth of 62.03 and incurs a 1.50% drop in consumption upon quitting to nonemployment (compensated by an increase in value from additional leisure). Were that same worker to counterfactually still quit after the realization of the shock, she would incur a larger 1.65% decline in consumption, consistent with the need to reduce consumption further to maintain a greater buffer stock of savings over a longer expected spell of nonemployment. After the shock, however, the wealth level defining the marginal quitter at this level of productivity rises by 0.35% to 62.25. This wealthier worker would have incurred only a 1.40% consumption drop from quitting in steady state, and incurs a 1.57% drop post-shock.

Figure 11 also shows a proportionally larger change in the wealth of the marginal quitter (or accepter) for low levels of productivity. Low-productivity workers are in general closer to the borrowing constraint, and thus their consumption is particularly responsive to changes in their expected duration of nonemployment.[40] This greater sensitivity of the quit threshold for lower-productivity workers is reminiscent of findings from Section 4.2, where we showed that quits from E-to-N fall significantly more for lower-educated workers than for higher-educated workers in response to a contractionary monetary policy shock.

Treating education in the data as an imperfect proxy for productivity, Figure 12 compares the model-implied impulse responses of quits from E-to-N for below- and above-median productivity workers to those of lower- and higher-educated workers in the data. The heterogeneous response of quits is very similar in the model and the data: quits decline sharply for low-education workers in the data and low-productivity workers in the model, whereas we see virtually no response for their high-education and high-productivity counterparts.

7. Conclusion

This paper offers new empirical evidence of a sizable response of supply-driven labor market flows to a contractionary monetary policy shock. Using high-frequency identified monetary policy shocks from FOMC announcements and Fed Chair speeches, we show that a contractionary monetary policy shock decreases the rate at which workers quit jobs to nonemployment and stimulates job-seeking behavior among the nonemployed. In doing so, we develop a novel decomposition of transitions from employment to nonparticipation into quits and layoffs, and we offer new evidence that a large and procyclical component of E-to-N flows reflects quits.

Our estimates imply a quantitatively important role for labor supply considerations in shaping the employment response to a monetary policy shock: Holding the response of such supply-driven labor market flows fixed, the decline in employment from a monetary contraction would be about twice as large. Thus, our paper highlights a potentially important role for labor supply in the monetary transmission mechanism.

To better understand our new empirical findings, we estimate a heterogeneous agent model with frictional labor markets and an active labor supply margin. The estimated model provides an excellent fit to our new empirical evidence and shows that labor supply significantly dampens the response of employment to a monetary policy shock. Importantly, the model matches not only aggregate labor market dynamics but also micro-level evidence on marginal propensities to consume and earn. We show that the change in labor supply policy functions within the model is driven primarily by the change in the job-finding rate, whereas

---

[40]We see similar evidence of heterogeneous responses in search and acceptance behavior of the non-employed; however, these responses do not map as neatly into particular labor market flows: the first is partially obscured by compositional issues (as discussed in Section 4), whereas the second is unobserved.

other components of the monetary policy shock affect aggregate labor supply primarily through compositional considerations. Our finding that the Main Business Cycle shock generates similar co-movements suggests that these insights may apply more generally.

Given its ability to match both micro and macro facts, we view our modeling framework as a promising foundation for general equilibrium analysis of monetary policy and other policy interventions. For example, incorporating an active labor supply margin in New Keynesian models may prove helpful for understanding the recent U.S. labor market experience since the pandemic: a sequence of unprecedentedly large stimulus payments in 2020 and 2021 was soon followed by a period of weak labor force participation and an unexpectedly high quit rate (the "Great Resignation"). Existing models are well-placed to consider the effects of such stimulus on consumption, but might be less suited to considering how such policies affect labor supply, or how the labor supply response to such policies might have contributed to the rise in inflation.

We thus view extending our framework to a fully-fledged New Keynesian model—endogenizing the job-finding rate through vacancy posting and allowing for endogenous layoffs—as a promising direction for future research.

<h2 style="text-align:center">REFERENCES</h2>

Ahn, Hie Joo and James D. Hamilton, "Measuring labor-force participation and the incidence and duration of unemployment," *Review of Economic Dynamics*, 2022, *44*, 1–32.

Alves, Felipe, "The Job Ladder in HANK: Inspecting the Mechanism," January 2025.

_____ **and Giovanni L. Violante**, "Monetary Policy Under Okun's Hypothesis," 2025.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas, "Business-cycle anatomy," *American Economic Review*, 2020, *110* (10), 3030–3070.

Aruoba, S. Borağan and Thomas Drechsel, "Identifying Monetary Policy Shocks: A Natural Language Approach," *American Economic Journal: Macroeconomics*, 2026, p. forthcoming.

Auclert, Adrien, Bence Bardóczy, and Matthew Rognlie, "MPCs, MPEs, and Multipliers: A Trilemma for New Keynesian Models," *The Review of Economics and Statistics*, 07 2021, pp. 1–41.

_____ **, Matthew Rognlie, and Ludwig Straub**, "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Technical Report, National Bureau of Economic Research 2020.

Barnichon, Regis, "Building a composite help-wanted index," *Economics Letters*, 2010, *109* (3), 175–178.

Bauer, Michael D. and Eric T. Swanson, "An Alternative Explanation for the 'Fed Information Effect'," *American Economic Review*, 2023, *113*, 664–700.

_____ **and** _____ , "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual*, 2023, *37* (1), 87–155.

Birinci, Serdar, Fatih Karahan, Yusuf Mercan, and Kurt See, "Labor Market Shocks and Monetary Policy," Technical Report RWP 24-04, Federal Reserve Bank of St. Louis Working Paper 2024.

Blanco, Andrés, Andres Drenik, Christian Moser, and Emilio Zaratiegui, "A Theory of Labor Markets with Inefficient Turnover," Working Paper 32409, National Bureau of Economic Research May 2024.

Boehm, Johannes, Etienne Fize, and Xavier Jaravel, "Five facts about MPCs: Evidence from a randomized experiment," CEP Discussion Papers dp1998, Centre for Economic Performance, LSE May 2024.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *The Review of Economic Studies*, 02 2024.

Broer, Tobias, John Kramer, and Kurt Mitman, "The curious incidence of monetary policy shocks across the income distribution," Technical Report 2021.

Cairó, Isabel, Shigeru Fujita, and Camilo Morales-Jiménez, "The cyclicality of labor force participation flows: The role of labor supply elasticities and wage rigidity," *Review of Economic Dynamics*, 2022, *43*, 197–216.

Cantore, Cristiano, Filippo Ferroni, Haroon Mumtaz, and Angeliki Theophilopoulou, "A tail of labor supply and a tale of monetary policy," 2023.

Christiano, Lawrence J., "Comment on "Unemployment in an Estimated New Keynesian Model"'," in "NBER Macroeconomics Annual 2011, Volume 26" NBER Chapters, National Bureau of Economic Research, Inc, January-J 2011, pp. 361–380.

—, **Martin Eichenbaum, and Charles L. Evans**, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy, University of Chicago Press*, February 2005, *113* (1), 1–45.

—, **Mathias Trabandt, and Karl Walentin**, "Involuntary unemployment and the business cycle," *Review of Economic Dynamics*, 2021, *39*, 26–54.

Cochrane, John and Monika Piazzesi, "The Fed and Interest Rates—A High-Frequency Identification," *American Economic Review*, 2002, *92* (2), 90–95.

Coglianese, John M, Maria Olsson, and Christina Patterson, "Monetary Policy and the Labor Market: a Quasi-Experiment in Sweden," 2023.

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*, Summer 2006, *20* (3), 3–26.

Eeckhout, Jan and Alireza Sepahsalari, "The Effect of Wealth on Worker Productivity," *The Review of Economic Studies*, 07 2023, *91* (3), 1584–1633.

Ellieroth, Kathrin and Amanda Michaud, "Quits, Layoffs, and Labor Supply," Technical Report 94, Federal Reserve Bank of Minneapolis Opportunity & Inclusive Growth Institute Working Paper January 2026.

Elsby, Michael W. L., Ryan Michaels, and Gary Solon, "The Ins and Outs of Cyclical Unemployment," *American Economic Journal: Macroeconomics*, January 2009, *1* (1), 84–110.

Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin, "On the importance of the participation margin for labor market fluctuations," *Journal of Monetary Economics*, 2015, *72*, 64–82.

Faberman, R. Jason and Alejandro Justiniano, "Job Switching and Wage Growth," *Chicago Fed Letter*, 2015.

Faccini, Renato and Leonardo Melosi, "Job-to-Job Mobility and Inflation," *The Review of Economics and Statistics*, 07 2025, *107* (4), 1027–1041.

Faia, Ester, Ekaterina Shabalina, and David Wiczer, "Heterogeneous Effects of Monetary Policy Across Income and Race: The Labour Mobility Channel," CEPR Discussion Papers 17741, C.E.P.R. Discussion Papers October 2023.

Faust, Jon, Eric T. Swanson, and Jonathan Wright, "Identifying VARs Based on High Frequency Futures Data," *Journal of Monetary Economics*, 2004, *51* (6), 1107–1131.

———— , **John Rogers, Eric T. Swanson, and Jonathan Wright**, "Identifying the Effects of Monetary Policy Shocks on Exchange Rates Using High Frequency Data," *Journal of the European Economic Association*, 2003, *1* (5), 1031–1057.

Ferraro, Domenico and Giuseppe Fiori, "Search Frictions, Labor Supply, and the Asymmetric Business Cycle," *Journal of Money, Credit and Banking*, 2023, *55* (1), 5–42.

Fujita, Shigeru, Giuseppe Moscarini, and Fabien Postel-Vinay, "Measuring Employer-to-Employer Reallocation," Working Paper 27525, National Bureau of Economic Research July 2020.

———— , ———— , **and** ———— , "Measuring Employer-to-Employer Reallocation," *American Economic Journal: Macroeconomics*, July 2024, *16* (3), 1–51.

Galí, Jordi, "Notes for a New Guide to Keynes (I): Wages, Aggregate Demand, and Employment," *Journal of the European Economic Association*, 10 2013, *11* (5), 973–1003.

Galí, Jordi, Frank Smets, and Rafael Wouters, "Unemployment in an Estimated New Keynesian Model," *NBER Macroeconomics Annual*, 2012, *26*, 329–360.

Gertler, Mark and Peter Karadi, "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 44–76.

Gilchrist, Simon and Egon Zakrajšek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, June 2012, *102* (4), 1692–1720.

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky, "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income*," *The Quarterly Journal of Economics*, 10 2023, p. qjad053.

Graves, Sebastian, "Does Unemployment Risk Affect Business Cycle Dynamics?," *American Economic Journal: Macroeconomics*, April 2025, *17* (2), 65–100.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante, "The rise of US earnings inequality: Does the cycle drive the trend?," *Review of Economic Dynamics*, 2020, *37*, S181–S204.

Hobijn, Bart and Ayşegül Şahin, "Maximum Employment and the Participation Cycle," Working Paper 29222, National Bureau of Economic Research September 2021.

Jentsch, Carsten and Kurt G. Lunsford, "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States: Comment," *American Economic Review*, July 2019, *109* (7), 2655–2678.

Kaplan, Greg and Giovanni L. Violante, "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Review of Economics*, 2022, *14* (Volume 14, 2022), 747–775.

Krusell, Per, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin, "Gross Worker Flows over the Business Cycle.," *The American Economic Review*, 2017, *107* (11), 3447 – 3476.

Laibson, David, Peter Maxted, and Benjamin Moll, "A Simple Mapping from MPCs to MPXs," Working Paper 29664, National Bureau of Economic Research January 2022.

Melcangi, Davide and Vincent Sterk, "Stock market participation, inequality, and

monetary policy," *Review of Economic Studies*, 2025, *92* (4), 2656–2690.

Michaels, Ryan, "What Explains the Great Resignation?," *Economic Inquiry*, Q2 2024, *9* (2), 10–18.

Moscarini, Giuseppe and Fabien Postel-Vinay, "The Job Ladder: Inflation vs. Reallocation," Working Paper 31466, National Bureau of Economic Research July 2023.

Mukoyama, Toshihiko, Christina Patterson, and Ayşegül Şahin, "Job search behavior over the business cycle," *American Economic Journal: Macroeconomics*, 2018, *10* (1), 190–215.

Orchard, Jacob, Valerie A Ramey, and Johannes F Wieland, "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates," Working Paper 31584, National Bureau of Economic Research August 2023.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert Mc-Clelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, October 2013, *103* (6), 2530–53.

Parker, Simon C., Yacine Belghitar, and Tim Barmby, "Wage Uncertainty and the Labour Supply of Self-Employed Workers," *The Economic Journal*, 2005, *115* (502), C190–C207.

Pijoan-Mas, Josep, "Precautionary savings or working longer hours?," *Review of Economic Dynamics*, 2006, *9* (2), 326–352.

Ramey, Valerie, "Macroeconomic Shocks and Their Propagation," in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics, Vol. 2*, North-Holland, 2016, chapter 2, pp. 71–162.

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

Shimer, Robert, "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, April 2012, *15* (2), 127–148.

Stock, James and Motohiro Yogo, *Testing for Weak Instruments in Linear IV Regression*, New York: Cambridge University Press,

Stock, James H. and Mark W. Watson, "Disentangling the Channels of the 2007–09 Recession," *Brookings Papers on Economic Activity*, 2012, *Spring*, 81–130.

‗‗‗‗‗‗ **and** ‗‗‗‗‗‗ , "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *Economic Journal*, 2018, *128* (610), 917–948.

Swanson, Eric T. and John C. Williams, "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," *American Economic Review*, October 2014, *104* (10), 3154–3185.

‗‗‗‗‗‗ **and Vishuddhi Jayawickrema**, "Speeches by the Fed Chair Are More Important than FOMC Announcements: An Improved High-Frequency Measure of U.S. Monetary Policy Shocks," 2024. Working Paper, University of California, Irvine.

White, Neil, "Gross Worker Flows, Job Loss, and Monetary Policy," 2018.

Wolf, Christian K, "Appendix to "The missing intercept: A demand equivalence approach"," *American Economic Review*, 2023, *113* (8), 2232–2269.

‗‗‗‗‗‗ , "The missing intercept: A demand equivalence approach," *American Economic Review*, 2023, *113* (8), 2232–2269.

Online Appendix to "The Labor Demand and Labor Supply Channels of Monetary Policy"

Sebastian Graves, Christopher Huckfeldt, and Eric T. Swanson

Contents

Online Appendix to "The Labor Demand and Labor Supply Channels of Monetary Policy"

Appendix A. Time series: Labor Market Flows and Intensive Margins of Job Search

Figure A.1 shows the time series of labor market flows, decomposed E-to-U and E-to-N flows, and intensive margin measures. Sections 2.1 and 2.2 discuss our measures of labor market flows; Appendix B provides details on the E-to-U and E-to-N decomposition.

We measure the intensive margin of job search using the number of distinct search methods reported by unemployed workers, following Mukoyama et al. (2018). To construct a consistent series across the 1994 CPS redesign (which increased possible methods from 6 to 12), we group responses into 5 categories: public employment agency, private employment agency, friends/relatives, employer contact/interview, and other active methods.[41] Relative to Mukoyama et al. (2018), we construct a consistent measure of the number of search methods starting from 1978, rather than 1994, shown in the left panel of Figure A.1c.

We measure the intensive margin for nonparticipants using the fraction of such workers who report wanting a job. Before 1994, nonparticipants were only asked if they wanted a job in the outgoing rotation group. The answers were "Yes", "Maybe, it depends", "No", or "Don't know". From 1994 this question was asked to all nonparticipants and the answers changed to "Yes, or maybe, it depends", "No", "Retired", "Disabled", or "Unable to work". Given this change, we group "Yes" and "Maybe, it depends" as "Yes" and all other answers as "No". This produces a consistent series across the 1994 redesign (right panel of Figure A.1c).

Appendix B. Measurement of Quits and Layoffs from E-to-U and E-to-N

We develop a novel decomposition of E-to-N flows measured from the CPS into quits and layoffs. We denote worker-initiated separations as "quits", firm-initiated separations as "layoffs", and use "other separations" for ambiguous cases (e.g., fixed-term jobs ending). After describing our decomposition methodology, we validate that quits and layoffs capture economically distinct phenomena and discuss robustness to measurement issues. The time series for our decomposition of E-to-U and E-to-N transition rates are shown in Figure A.1b.

---

[41]In principle, "placed or answered ads" is a sixth method that is included both before and after 1994. However, we have found that the number of individuals reporting this method dropped sharply after 1994. This is likely explained by the introduction of "Sent out resumes/filled out applications" as a possible search method at this time.

FIGURE A.1. Time Series of Labor Market Flows and Intensive Margins

(A) Labor Market Flows

(B) Decomposition of E-to-U and E-to-N Flows

(C) Intensive Margins of Labor Supply

Note: All series are smoothed using a centered 5-month moving average. See text for details on variable construction.

B.1. **Data Construction.**

B.1.1. *Decomposition of E-to-U Flows: Quits versus Layoffs.* The E-to-U decomposition is straightforward. We label an E-to-U transition as a quit if the reason for unemployment is "job leaver" and as a layoff if the reason for unemployment is "job loser/on layoff" or "other job loser". We label as "other separations" transitions where the reason for unemployment is "temporary job ended", "re-entrant" or "new entrant".[42]

We label the end of temporary or seasonal jobs as "other separations." Compared to the ending of an open-term job, there is no clear economic rationale for labeling the ending of fixed-term job as a quit or a layoff. However, while it is simple to separately categorize such E-to-U transitions for the majority of our data, "temporary job ended" was removed as a possible response from the survey from 1989 to 1993. An inspection of the data shows that during this period such transitions were labeled as layoffs. Thus, we estimate the share of E-to-U transitions due to temporary jobs ending for each month between 1989 and 1993, and then remove this share from that which is initially defined as layoffs during this period.

To implement this procedure, we run a regression of the share of E-to-U transitions due to temporary jobs ending on all six labor market transition rates, month dummies and a time trend for the period from January 1978 to December 1988. The R^2 of this regression is 0.58, implying that the share of E-to-U transitions due to temporary jobs ending is largely predictable. We use this regression to predict the share of E-to-U separations due to temporary jobs ending from January 1989 to December 1993. Finally, we adjust down the share of E-to-U separations due to layoffs in this period accordingly. This adjustment is minor: "temporary job ended" accounts for only 13% of E-to-U transitions when available.

B.1.2. *Decomposition of E-to-N Flows: Quits versus Layoffs.* The decomposition of E-to-N flows is more involved: to our knowledge, our paper is the first to use the CPS to develop a harmonized measure of E-to-N quits and layoffs suitable for time series analysis.

A subset of CPS respondents in an Outgoing Rotation Groups (ORG) identified to be nonparticipants are asked the reason that they left their last job. However, the particular subset has changed over time. Since 1994, this question is asked to individuals in the outgoing rotation group who are: (1) not in the labor force, (2) neither retired nor disabled and (3) who report working in the past 12 months. Prior to 1994 this question was asked to individuals in the outgoing rotation group who are: (1) not in the labor force and (2) who reported working in the past five years. Moreover, the possible answers to the question also changed slightly starting in 1994, as discussed below.[43]

---

[42]An individual moving from E-to-U should be neither a "re-entrant" nor a "new entrant". Thus, these reasons appear to be measurement error. They account for around 15% of E-to-U transitions in our sample.
[43]For technical background on these changes, see U.S. Census Bureau (2019), pg. 111.

TABLE B.1. Labor Market Transition Observations

	E-to-U	Quits	Layoffs	Other	E-to-N	ORG	Quits	Layoffs	Other	U-to-E	U-to-N	N-to-E	N-to-U
Mean	632.8	82.7	366.9	183.4	1372.3	410.8	102.4	21.4	165.6	711.4	621.2	1209.0	629.5
Std. Dev.	146.4	27.5	92.7	59.5	160.3	45.0	34.4	7.8	20.5	150.6	135.8	157.4	137.9

Note: This table reports the unweighted mean and standard deviation of the number of observations in the CPS making each transition each month. ORG denotes individuals making an E-to-N transition in an outgoing rotation group.

To create a harmonized series, we restrict our attention to individuals who report having worked in the past 12 months.[44] We label an E-to-N transition as a quit if the reason for leaving the job is "personal, family or school" or "unsatisfactory work arrangements".[45] We label an E-to-N transition as a layoff if the reason for leaving the job is "slack work or business conditions". We label all remaining E-to-N transitions as other separations.[46] After 1994 we assume that individuals who make an E-to-N transition and either report being retired or disabled would have given this as their reason for leaving their job had they been asked the question. Consequently, such transitions are defined as neither quits nor layoffs. Finally, as our sample is only ever a fraction of all E-to-N transitions, in all periods we calculate the share of E-to-N transitions in each classification and then multiply this by the overall E-to-N transition rate to complete our decomposition.

Table B.1 reports the mean and standard deviation of the number of individuals making each labor market transition per month. Whereas all E-to-U transitions can be decomposed into quits and layoffs, only individuals completing E-to-N transitions while in the ORG are asked their reason for leaving, resulting in fewer observations for the E-to-N decomposition.

B.2. **Further Evidence for Economic Relevance of Quit/Layoff Distinction.** Here, we provide additional evidence that the distinction between quits and layoffs is economically meaningful at the individual level, by documenting that the subsequent labor market transition probabilities for individuals who quit to either unemployment or nonparticipation are notably different from those of individuals who are laid off.

The top panel of Table B.2 shows transition probabilities of workers who entered unemployment from employment in the previous month either due to a quit (i.e., E−U(Quit)) or a layoff (i.e., E−U(Layoff)). Workers making E−U(Quit) transitions have slightly higher reemployment probabilities *and* significantly higher probabilities of entering nonparticipation than workers making E−U(Layoff) transitions.[47] This suggests that individuals quitting to

---

[44]In principle, all individuals that make E-to-N transitions should report having worked in the past 12 months. In practice, a minority do not, as we discuss later.

[45]These are the possible answers from before 1994. After 1994 we define such transitions analogously.

[46]Other E-to-N separations include retirements, individuals reporting disability, and the end of temporary seasonal or non-seasonal jobs.

[47]We can reject the null hypothesis that the two rows of transition probabilities given in Table B.2 are equal using a chi-squared goodness-of-fit test with a p-value that is less than 0.01%.

TABLE B.2. Post-Separation Transition Rates and Labor Force Attachment: Quits vs Layoffs

Panel A: Post-E-to-U Transition Rates

| | To | | |
From	E	U	N
E − U(Quit)	0.448	0.399	0.153
E − U(Layoff)	0.426	0.468	0.106

Panel B: Post-E-to-N Labor Force Attachment

	Average Probability
Want Job | E-N(Quit)	0.210
Want Job | E-N(Layoff)	0.515
NE | Want Job	0.145
NE | Do Not Want Job	0.037
NU | Want Job	0.172
NU | Do Not Want Job	0.012

Note: Panel A shows transition rates for individuals in their first month of unemployment following an employment spell. Panel B shows the probability that individuals want a job after an E-to-N transition (first two rows) and transition probabilities for all nonparticipants conditional on wanting a job (last four rows). Reasons for job separation are defined in Appendix B.1.1.

unemployment likely fall into two groups: The first are individuals who appear to have strong employment prospects when they quit to unemployment, and thus move back to employment at a high rate. The second are individuals with low attachment to the labor market, who thus move to nonparticipation at a higher rate than individuals laid off to unemployment.

The same exercise is not possible for E-to-N quits and layoffs, as nonparticipants are only asked their reason for leaving their last job if they are in the outgoing rotation group, and thus we do not see their employment status the following month. However, we are able to provide evidence that such individuals likely have very different subsequent labor market transition probabilities. The bottom panel of Table B.2 shows that those who are laid off to nonparticipation are more than twice as likely to report that they want a job as those who quit to nonparticipation, and that nonparticipants who want a job are around four (fourteen) times more likely to move to employment (unemployment) in the next month than nonparticipants who report that they do not want work. This suggests that people who quit to nonparticipation are less attached to the labor market than individuals laid off to nonparticipation, and thus are more likely to stay there.

B.3. **Robustness: E-to-U Flows.** Shimer (2012) points out potential inconsistencies in the measurement of quits and layoffs to unemployment in the CPS, noting that, prior to the 1994 survey redesign, a portion of E-to-U quitters who are newly unemployed in month t

TABLE B.3. Measurement of Quit and Layoff Status in E−U−U Sequences

Panel A: Sequences of Reasons for Unemployment

Sample period	P(Quit | Layoff)	P(Layoff | Quit)
pre-Redesign	0.039	0.208
post-Redesign	0.007	0.026

Panel B: Transition Rates (Pre-1994 Redesign)

		To		
	From	E	U	N
(*a*)	E − U(Quit) − U(Layoff)	0.339	0.553	0.108
(*b*)	E − U(Quit) − U(Quit)	0.343	0.536	0.121
(*c*)	E − U(Layoff) − U(Quit)	0.352	0.557	0.091
(*d*)	E − U(Layoff) − U(Layoff)	0.264	0.667	0.068

Note: Panel A shows the probability of individuals switching their reported reason for unemployment from layoff to quit (first column) or quit to layoff (second column), before and after the 1994 CPS redesign. Panel B shows transition rates for individuals in their second month of unemployment following an employment spell, split by reason for unemployment (pre-1994 only). Reasons for unemployment are defined in Appendix B.1.1.

and remain unemployed in month $t+1$ then report having been laid off; and a much smaller portion of those laid off to unemployment in month t that remain unemployed in month $t+1$ then report having quit. In this section, we investigate these issues and show that they present only minor concerns for our measures of E-to-U quits and layoffs.

We reproduce evidence akin to Shimer (2012) in the top panel of Table B.3. For individuals with an E-U-U labor market sequence, around 4% of those who initially report having been laid off subsequently report having quit their job, before the 1994 redesign of the CPS. Switching is higher among those who initially report having quit: around 20% of such individuals subsequently report having been laid off. After the redesign of the survey the likelihood of switching in either direction drops dramatically. Note, the patterns from Table B.3 have two possible interpretations: First, that quits and layoffs are measured inaccurately in the CPS, as suggested by Shimer (2012). Second, the patterns could be explained by the existence of short-term jobs that are not picked up by the monthly CPS survey. Although we cannot easily distinguish between these two explanations, we next provide evidence that such switching is not quantitatively relevant for our measures of E-to-U quits and layoffs.

The bottom panel of Table B.3 reports subsequent transition rates for workers having previously made an E−U−U transition during the period prior to the 1994 CPS redesign, with four separate rows for each sequence of reasons for unemployment across the two months, with rows as follows: (*a*) E−U(Quit)−U(Layoff), (*b*) E−U(Quit)−U(Quit), (*c*) E−U(Layoff)−U(Quit), and (*d*) E−U(Layoff)−U(Layoff).

In rows(a) and (b), we compare labor market transitions of workers who quit to unemployment and then report a layoff (i.e., E−U(Quit)−U(Layoff)) to those of workers who quit to unemployment and continue to report a quit (i.e., E−U(Quit)−U(Quit)). Table B.3 shows that around 20% of individuals who initially report having quit and remain in unemployment then report a layoff. However, the subsequent labor market transitions of workers reporting quit-layoff are very similar to those of individuals who continue to report quit-quit, as seen by comparing rows (a) and (b). Indeed, using a chi-squared goodness-of-fit test, we cannot reject the null hypothesis that the two rows are the same, with a p-value of 0.72. Hence, for such individuals, we find that only the reason for unemployment reported in the first month is relevant for predicting future employment transitions, offering validation for our measure of E-to-U quits.

In rows (c) and (d), we compare labor market transitions of workers who are laid-off to unemployment and then report a quit to those of workers who are laid-off to unemployment and continue to report a layoff. We find that transition patterns of E−U(Layoff)−U(Quit) workers are notably different from those of E−U(Layoff)−U(Layoff) workers. However, even if such layoff-quit transitions represent measurement error, they are relatively uncommon: only around 4% of E−U−U workers who initially report being laid off then report having quit in the following month (as shown in Table B.3). Thus, this group accounts for a small enough proportion of total E−U layoffs as to be considered quantitatively insignificant.

B.4. **Robustness: E-to-N Flows.** Recall, our measurement of quits and layoffs for E-to-N transitions is based on a variable specific to respondents in the ORG that reports the reason why the individual left their previous job. For approximately 30 percent of such transitions the value of this variable is missing. The red line in Figure B.1 shows the time-series for this fraction. The proportion of E-to-N transitions with a missing reason rose from about 20 percent in the early 1980s to around a third by the early 2000s and has been relatively stable since. Here, we offer evidence that this is not a concern for our measure of E-to-N transitions.

Since 1994, nonparticipants are only asked their reason for leaving their last job if they report that this job occurred during the past 12 months. For individuals who are coded as working in this required time period, there is no missing data on the reason for leaving their job. Thus, data appears to be missing because some fraction of workers recorded making transitions from E in month t to N in month $t+1$ are coded in month $t+1$ as not having worked in the past year. While this could reflect spurious E-to-N transitions—where employment status was mismeasured in month t, and the individuals truly never were employed in the past 12 months—we argue below that such spurious transitions could only reflect a small minority of the missing data; and instead, that workers are erroneously recorded as not having worked in the prior year.

FIGURE B.1. Fraction of E-to-N Transitions With Missing Reason

Note: The red line shows the proportion of individuals making an E-to-N transition for which there is missing data on the reason for leaving the last job. The blue line shows the same calculation for individuals that were employed in each of the first three months before moving to nonparticipation. Series are smoothed using a centered 5-month moving average.

First, the share of E-to-N transitions with missing data does not vary significantly across worker subgroups where one might expect variation (e.g., non-self-employed, self-respondents, full-time workers). Moreover, there is no discernible discontinuity in missing data around 1994, despite the change from a five-year to one-year lookback window.

Next, we compare the incidence of missing data for all E-to-N transitions to the subset of individuals who report three months of employment prior to their transition to nonparticipation (i.e., EEEN workers). The latter is plotted in the blue line in Figure B.1. EEEN workers are presumably more likely to have truly been employed before their transition to nonparticipation. While the incidence of missing data is slightly smaller for these individuals, still around 25% of observations are missing.

Finally, we examine the subsample of individuals included in the Job Tenure Supplement in the month before they moved to nonparticipation. If we restrict the sample to such individuals who report having worked at their current job for at least one year when answering the Job Tenure Supplement, we still find that, one month later, around 30% of such individuals are classified as having not worked in the past 12 months.

Thus, while it is possible that some individuals are misclassified as employed in the month before they are nonparticipants, the evidence suggests that the more plausible source of measurement error stems from workers being incorrectly coded as not having worked in the previous 12 months after 1994 (and previous five years prior to 1994).

APPENDIX C. ADDITIONAL EMPIRICAL RESULTS

In this section, we first discuss empirical estimates based on alternative measures of monetary policy shocks. We then discuss results using our baseline monetary policy shock series, but estimate impulse responses either using local projections or in a Bayesian VAR

framework. Finally, we provide additional results regarding composition and heterogeneity, as well as additional accounting exercises.

C.1. **Alternative Monetary Policy Shocks.** We are aware of only two other papers estimating the response of labor market flows to monetary policy shocks in the United States: White (2018), whose estimates using Romer and Romer (2004) shocks are largely insignificant, and Faia et al. (2023), who estimate impact effects but not impulse responses at longer horizons.

Our baseline shocks from Swanson and Jayawickrema (2024)—incorporated into our estimation following Bauer and Swanson (2023a,b)—are constructed to address econometric problems associated with prior HFI surprises. To investigate robustness, we also estimate our specification using shocks from Aruoba and Drechsel (2026), who develop a complementary approach for policy-rule residual shocks, expanding on Romer and Romer (2004) by incorporating additional information from the text of the *Greenbook*. The two shock series—each designed independently to confront separate critiques of instrument exogeneity and relevance—yield very similar point estimates for labor market flows.

We also document the importance of these methodological advances by showing estimates produced by earlier approaches. Using HFI surprises from FOMC announcements alone, we find the attenuation bias documented by Bauer and Swanson (2023b); orthogonalizing these shocks removes the bias but results in a weak-instrument problem, with a first-stage F-statistic below 3. Using Romer and Romer (2004) shocks, our estimates are imprecise with wide confidence intervals. These results highlight the value of the additional variation from Chair speeches in our baseline HFI approach and the additional Greenbook controls in AD.

C.1.1. *FOMC Announcement Shocks.* Figure C.1 shows impulse responses using high-frequency shocks from FOMC announcements that are not orthogonalized. The results are attenuated relative to our baseline, consistent with the bias documented by Bauer and Swanson (2023b). Figure C.2 orthogonalizes the same shocks with respect to recent macroeconomic news. The attenuation bias is removed, but the confidence intervals widen significantly and the first-stage F-statistic falls below 3.

C.1.2. *Policy Rule Based Shocks and Combinations of Shocks.* Figure C.3 shows impulse responses using Romer and Romer (2004) shocks, updated by Wieland and Yang (2020). We begin the shock sample in October 1982 for consistency with Aruoba and Drechsel (2026) and to restrict attention to the period in which the Federal Funds rate was the main policy target. The confidence intervals are wide and the F-statistic is low, though the point estimates are broadly consistent with our baseline.

Figure C.4 uses the shocks from Aruoba and Drechsel (2026), who extend Romer and Romer (2004) by controlling for additional information in the text of the Greenbook. The

estimates are very similar to our baseline. Figure C.5 uses both our baseline shocks and those of Aruoba and Drechsel (2026) as multiple external instruments. The point estimates remain similar, with no gain in precision, likely because the non-overlapping time periods reduce the first-stage sample size.

C.2. **Alternative Specifications.** In this section we show robustness of our results to using a local projection approach, or estimating responses in a Bayesian VAR.

C.2.1. *Local Projections.* First, we estimate responses using local projections rather than the VAR. Specifically, for each dependent variable y_i and horizon h we estimate the following regression:

$$y_{i,t+h} = \alpha_i^h + \beta_i^h \varepsilon_t^{MP} + \sum_{j=1}^{p} \gamma_{i,j}^h \varepsilon_{t-j}^{MP} + \sum_{j=1}^{p} \mathbf{\Gamma}_{i,j}^{h\prime} \mathbf{X}_{t-j} + u_{i,t+h}, \qquad h = 0, 1, \ldots, H \qquad \text{(C.1)}$$

where $y_{i,t+h}$ is outcome variable i at horizon h, ε_t^{MP} is the monetary policy shock and $\mathbf{X}_{t-j}$ is a vector containing all the endogenous variables in the equivalent VAR. As in the VAR, we set the lag length $p = 6$. The impulse response function is given by the sequence $\{\beta_i^h\}_{h=0}^{H}$.

Results for the main variables and labor market flows are shown in Figure C.6. We also provide estimates using the smooth local projection procedure of Barnichon and Brownlees (2019) in Figure C.7. In Panel A of each Figure, we see that the local projection estimates for unemployment, industrial production and vacancies are similar to those from the VAR until around 18 months. After this point the local projection estimates continue to expand and then stabilise, while the responses in the VAR begin to revert. Panel B of each Figure shows a similar pattern for the response of labor market flows. In Figure C.8 we compare the baseline local projection estimates with those from the VAR by repeating the accounting exercise of Figure 7. Overall, the pattern is very similar: holding supply-driven flows constant roughly doubles the response of employment.

C.2.2. *Bayesian VAR.* Next, we use a Bayesian approach to estimate the VAR, following Miranda-Agrippino and Ricco (2021). This approach allows us to include a large set of relevant variables. Thus, we include all labor market flows, as well as the additional variables from Figure 4 that are available for the full sample. Finally, we also add capacity utilization, consumer sentiment, and the S&P stock index. We also increase the lag length in the VAR from 6 to 12. Figure C.9 reports the results and compares them to those from baseline estimates (where available). Overall, we find very similar results.

C.3. **Response of Fiscal Variables.** In Figure C.10 we consider the response of various fiscal variables. The first four (spending, transfers, receipts and debt) are measured at the federal level. We find little evidence of a significant response of federal spending or transfer

payments. However, there is a significant decline in federal tax receipts, which is accompanied by a rise in federal debt. These results are consistent with those in Bouscasse and Hong (2023) using monetary policy shocks following the Romer and Romer (2004) methodology. The limited response of transfers may seem surprising. However, the vast majority of such transfers are due to Social Security and Medicare, while the majority of unemployment insurance is financed at the state level. In the final panel we plot the response of total UI payments (both state and federal) and find that such payments significantly increase, consistent with the increase in layoffs to unemployment estimated in Figure 3 and with the response of UI payments in our model.

C.4. **Time-Aggregation and Composition.** Figure C.11 shows impulse responses for labor market flows corrected for time aggregation, as in Shimer (2012) and Elsby et al. (2015). There are no notable differences compared to Figure 2.

Figure C.12 shows compositionally-adjusted flows using the full set of controls in Elsby, Hobijn and Şahin (2015), which adds labor market status one year prior to the controls for age, gender, educational attainment, and reason for unemployment. Conditioning on status one year prior restricts the sample to individuals in the fifth to eighth CPS interviews, who are not representative of the overall CPS sample Ahn and Hamilton (2022).
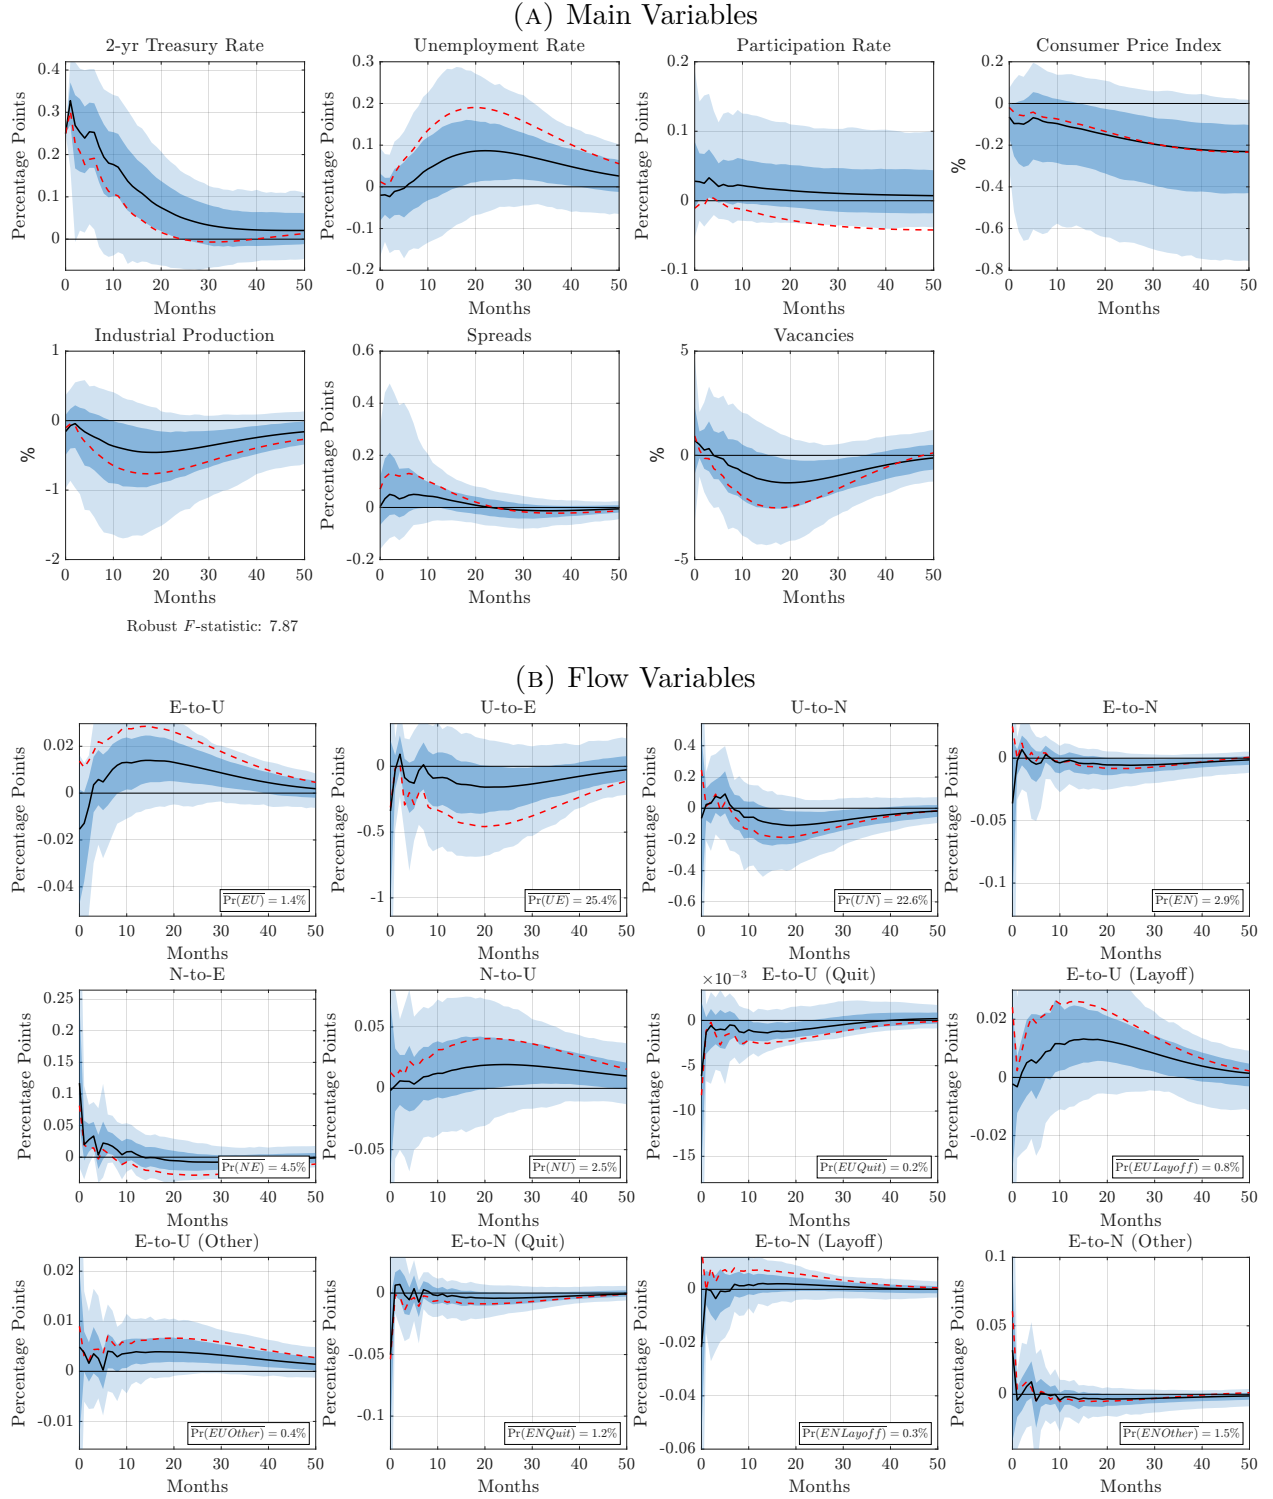
Despite this sample difference, the responses are qualitatively similar to those in Figure 5: composition adjustment dampens the response of the U-to-N rate by around half. Quantitative comparisons are more difficult. For example, employed individuals in the Figure C.12 sample have lower baseline transition rates to both unemployment and nonparticipation than those in the full sample (compare inset boxes across Figures 2 and C.12).

C.5. **Heterogeneity.** Figures C.13 and C.14 show impulse responses of labor market flows for higher- and lower-educated workers. Less-educated workers see larger movements in both E-to-U and E-to-N rates, and the decline in the N-to-E rate is entirely accounted for by this group.

Table C.1 reports the response of employment two years after the shock across demographic groups.[48] Each column replicates the accounting exercise from Figure 7, holding various supply-driven flows constant. The response of supply-driven flows dampens the employment decline for all groups, but there is notable heterogeneity: the decline in quits to non-employment (and thus the effect of holding quits constant) is largest for individuals who are young, renters, less-educated, or Black.

---

[48]As wage data is only available in the CPS outgoing rotation groups for employed individuals, we use a machine learning method to predict wage groups for all individuals, based on a large number of demographic and geographic variables.

FIGURE C.1. FOMC Announcement Shocks, Unadjusted

(A) Main Variables

Robust F-statistic: 7.87

(B) Flow Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using unadjusted FOMC announcement shocks. Panel (a) report the responses of the variables in the main VAR and reports the robust F-statistic for the first stage. Panel (b) reports results computed by appending the given flow to the main VAR. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Red dashed lines report the results from Figures 1 to 3.

FIGURE C.2. FOMC Announcement Shocks, Orthogonalized

(A) Main Variables

(B) Flow Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using orthogonalized FOMC announcement shocks. Panel (a) report the responses of the variables in the main VAR and reports the robust F-statistic for the first stage. Panel (b) reports results computed by appending the given flow to the main VAR. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Red dashed lines report the results from Figures 1 to 3.
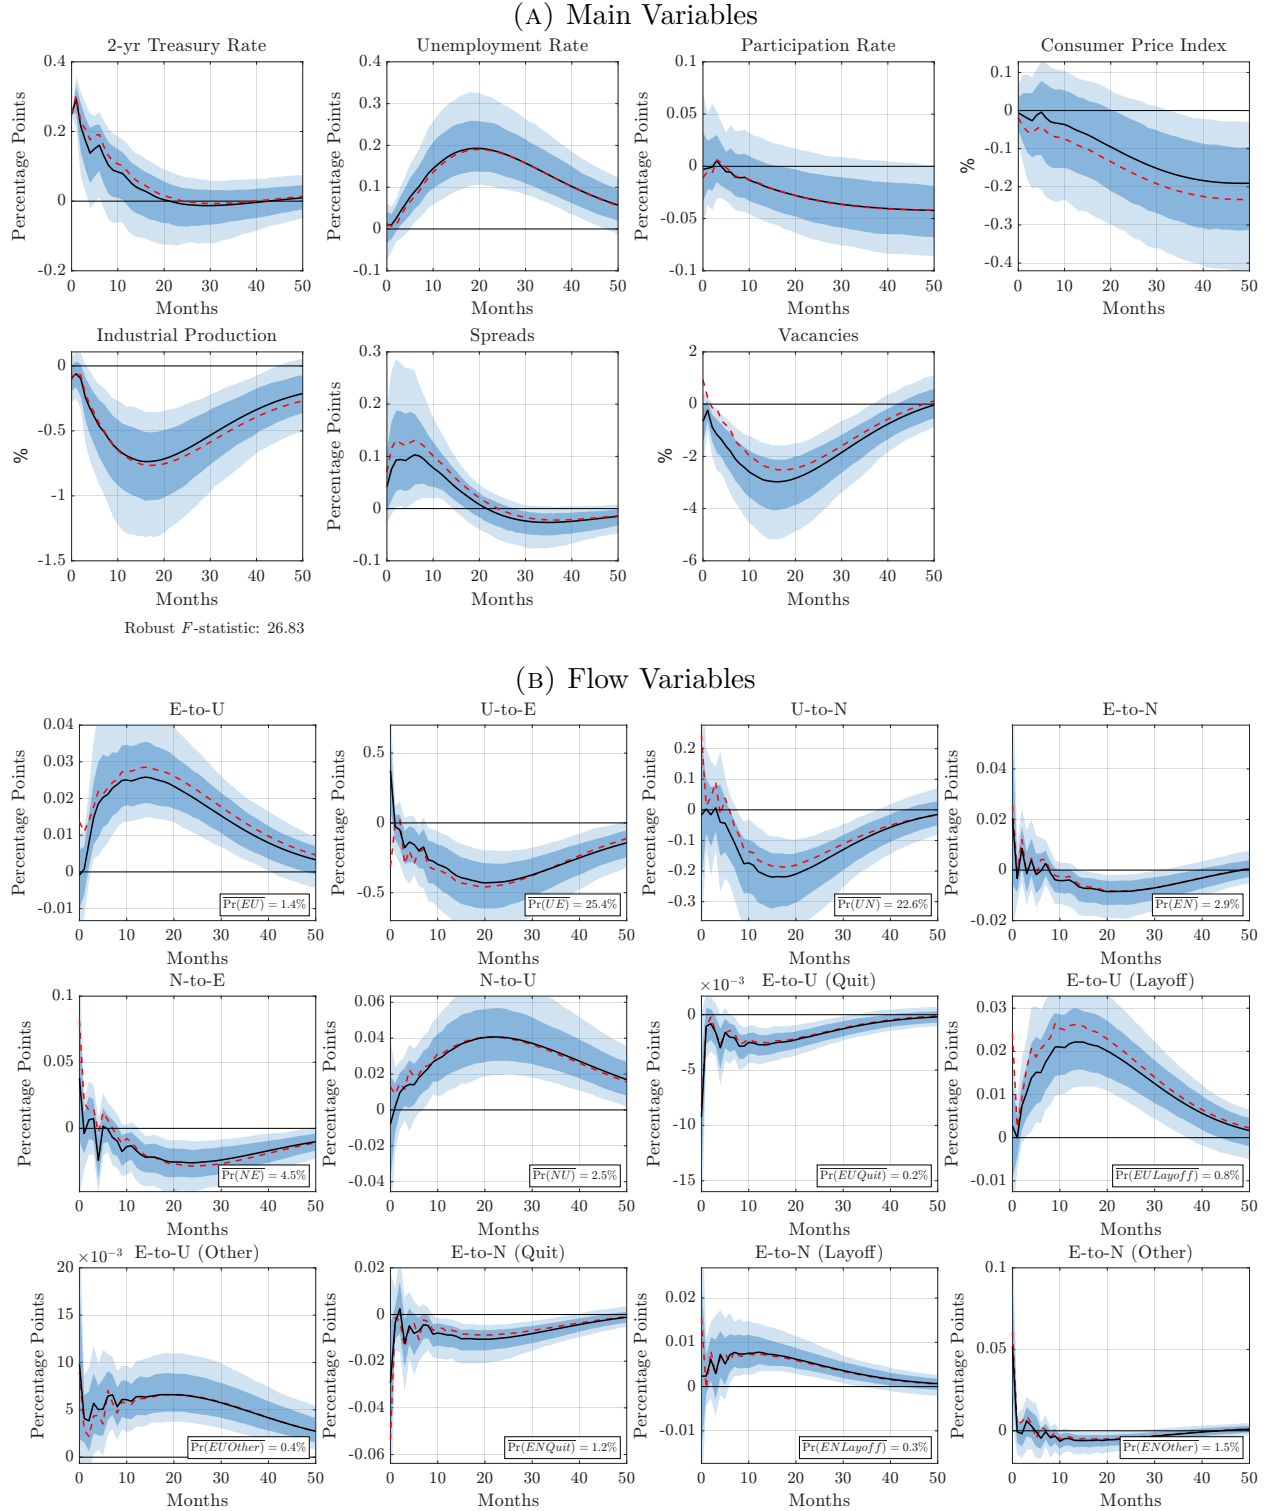
FIGURE C.3. Romer-Romer (2004) Shocks

(A) Main Variables

Robust *F*-statistic: 4.23

(B) Flow Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using Romer and Romer (2004) shocks. Panel (a) report the responses of the variables in the main VAR and reports the robust F-statistic for the first stage. Panel (b) reports results computed by appending the given flow to the main VAR. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Red dashed lines report the results from Figures 1 to 3.

FIGURE C.4. Aruoba-Drechsel (2026) Shocks

(A) Main Variables

Robust *F*-statistic: 26.83

(B) Flow Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using Aruoba and Drechsel (2026) shocks. Panel (a) report the responses of the variables in the main VAR and reports the robust F-statistic for the first stage. Panel (b) reports results computed by appending the given flow to the main VAR. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Red dashed lines report the results from Figures 1 to 3.

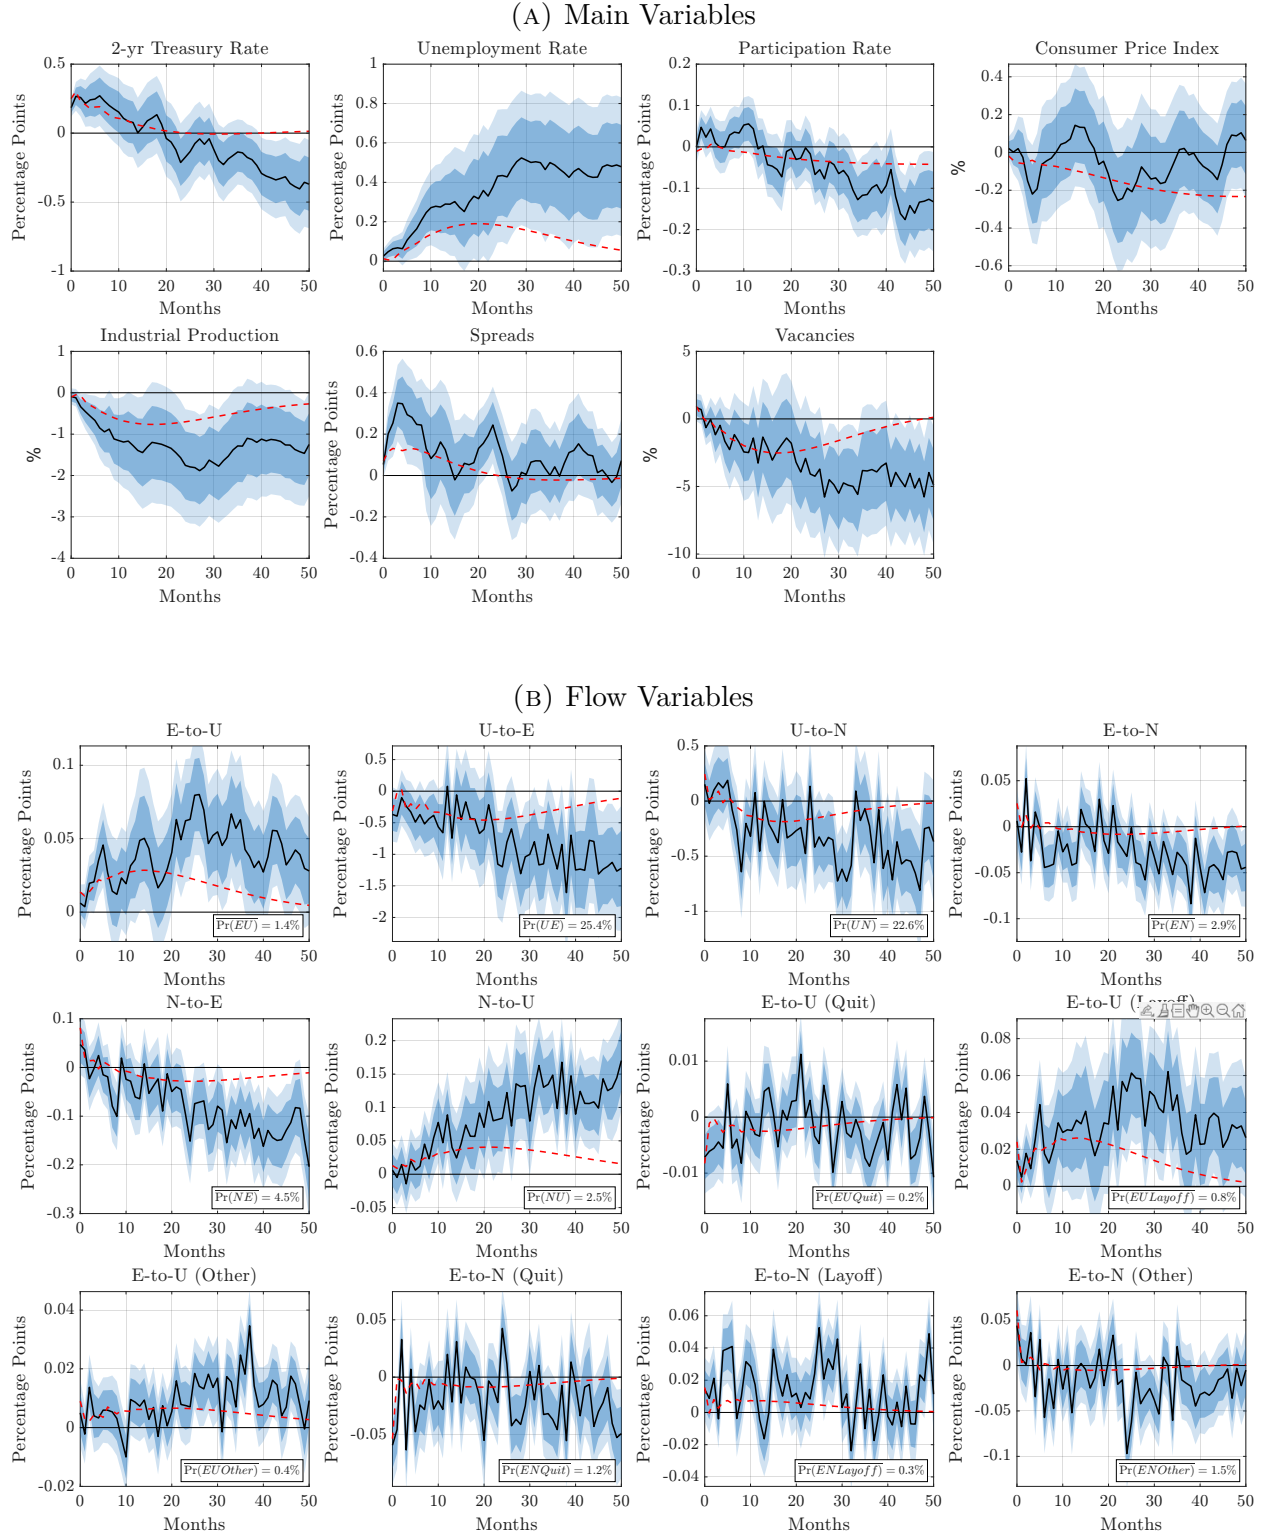FIGURE C.5. Both Baseline and Aruoba-Drechsel (2026) Shocks

(A) Main Variables

Robust F-statistic: 18.37

(B) Flow Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using both our baseline shocks and those of Aruoba and Drechsel (2026) simultaneously. Panel (a) report the responses of the variables in the main VAR and reports the robust F-statistic for the first stage. Panel (b) reports results computed by appending the given flow to the main VAR. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Red dashed lines report the results from Figures 1 to 3.

FIGURE C.6. Results using Baseline Shocks and Local Projections

(A) Main Variables

(B) Flow Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using baseline shocks in a local projection framework. Panel (a) report the responses of the variables from the main VAR. Panel (b) reports results for labor market flows. Solid black lines report impulse response functions while dark and light shaded regions report 68% and 90% confidence intervals constructed using robust standard errors. Red dashed lines report the results from Figures 1 to 3.

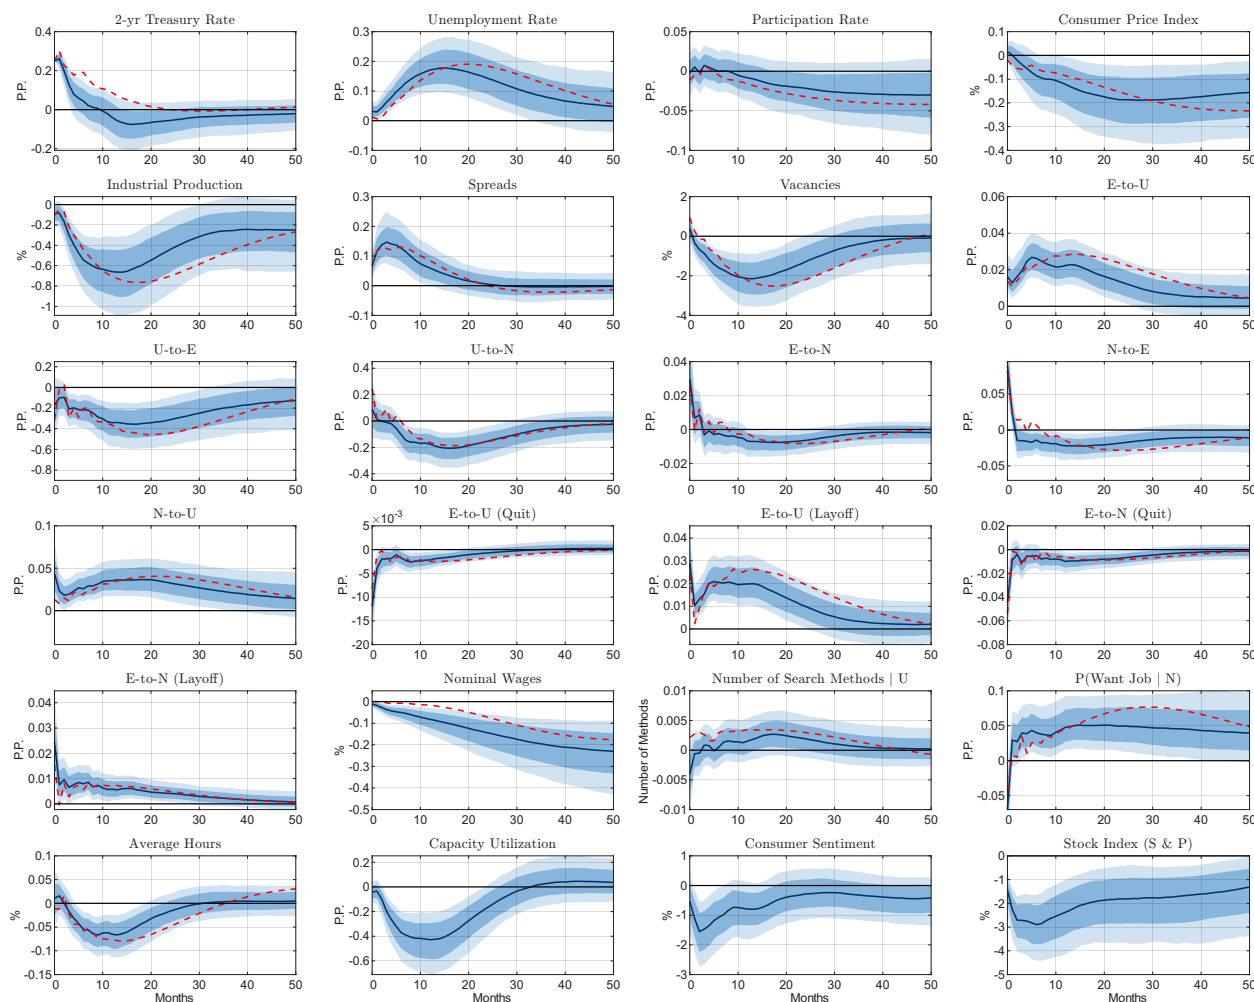FIGURE C.7. Results using Baseline Shocks and Smooth Local Projections

(A) Main Variables

(B) Flow Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using baseline shocks in a smooth local projection framework. Panel (a) report the responses of the variables from the main VAR. Panel (b) reports results for labor market flows. Solid black lines report impulse response functions while dark and light shaded regions report 68% and 90% confidence intervals constructed using robust standard errors. Red dashed lines report the results from Figures 1 to 3.

FIGURE C.8. Accounting for Employment: Local Projection Estimates

Note: The black solid line shows the overall response of the employment-population ratio to a contractionary monetary policy shock, estimated using local projections. The green dashed line shows the response if quits to U or N are held constant. The red dot-dashed line shows the response if both U-to-N and N-to-U rates are held constant. The blue dotted line shows the response if all supply-driven flows are held constant.

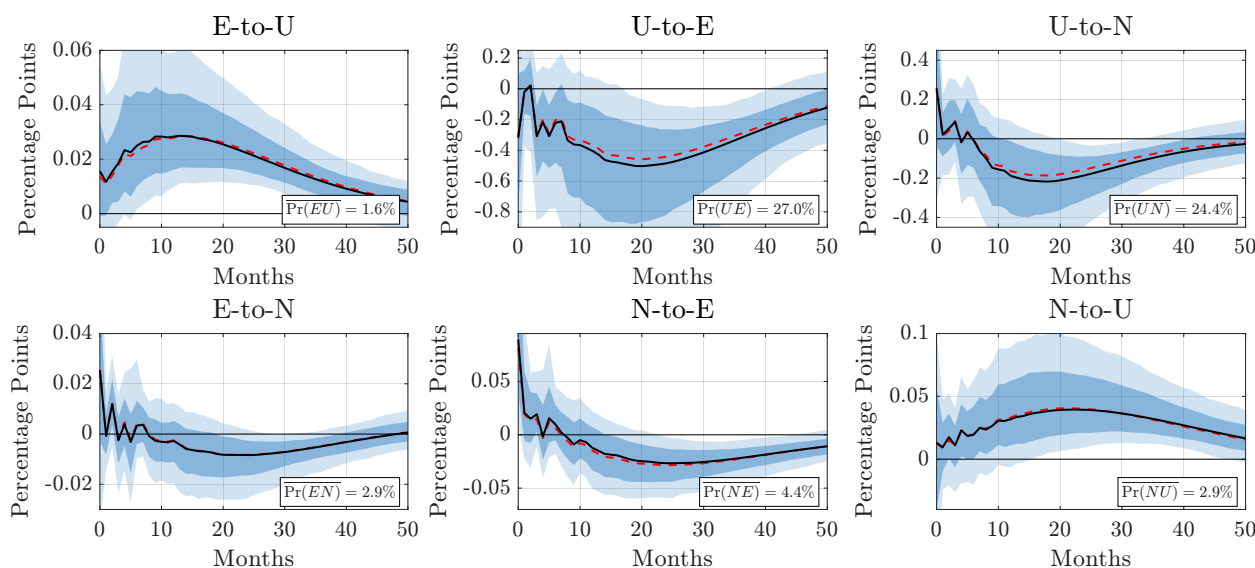FIGURE C.9. Estimates Using a Large-Scale Bayesian VAR

Note: Estimated impulse responses to a 25bp monetary policy tightening shock using baseline shocks and Bayesian methods with 12 lags. Solid black lines report posterior median responses while dark and light shaded regions report 68% and 90% posterior coverage bands. Red dashed lines report the results using frequentist methods from Figures 1 to 4 where available.
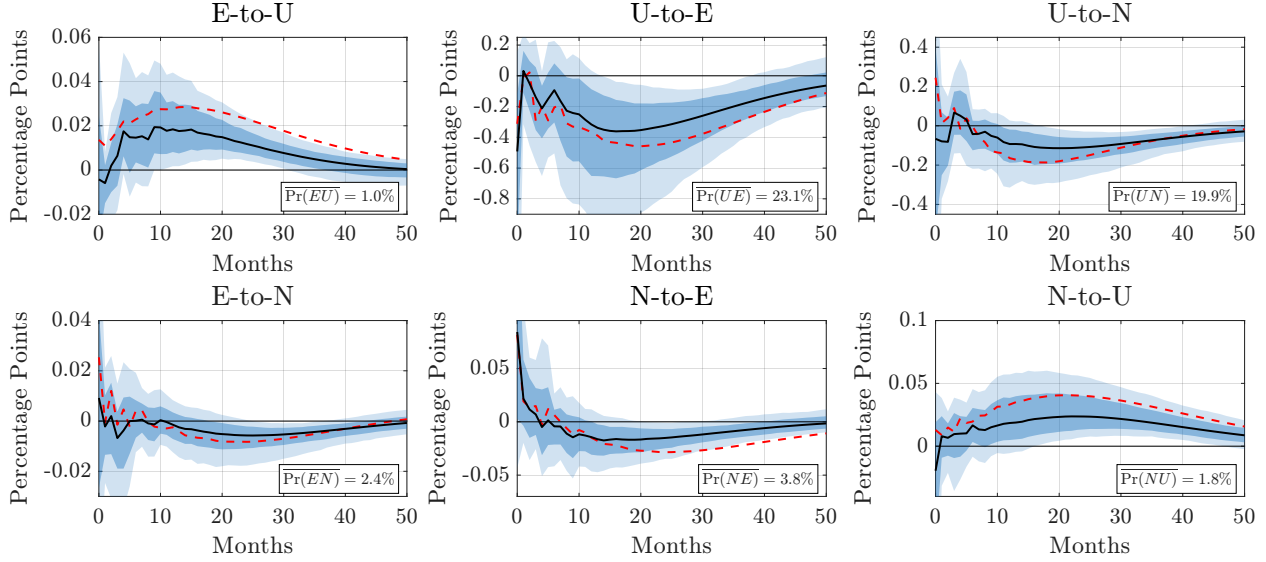
FIGURE C.10. Response of Fiscal Variables

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given variable to the baseline VAR from Figure 1. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. For UI payments, the dashed magenta line reports the response from the estimated model.

FIGURE C.11. Response of Time-Aggregation Corrected Labor Market Flows


Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given labor market flow variable, corrected for time-aggregation, to the baseline VAR from Figure 1. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals.

FIGURE C.12. Response of Composition-Adjusted Flows: Full EHS Controls

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given labor market flow variable to the baseline VAR from Figure 1. Solid black lines report impulse response functions for composition-adjusted flows, while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals for composition-adjusted flows. Dashed red lines report impulse responses for unadjusted flows with the same sample of individuals.

FIGURE C.13. Labor Market Flows: Higher-Educated

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given variable to the baseline VAR from Figure 1. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Inset boxes report average transition rates.

FIGURE C.14. Labor Market Flows: Lower-Educated

Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given variable to the baseline VAR from Figure 1. Solid black lines report impulse response functions while dark and light shaded regions report bootstrapped 68% and 90% confidence intervals. Inset boxes report average transition rates.

TABLE C.1. Response of Employment After 2 Years

| Sample	Baseline	Quits Constant	N↔U Constant	Quits and N↔U Constant
Aggregate	−0.140	−0.196	−0.205	−0.261
Male	−0.148	−0.200	−0.225	−0.277
Male, Single	−0.237	−0.299	−0.287	−0.350
Male, Married	−0.108	−0.138	−0.183	−0.214
Female	−0.108	−0.163	−0.156	−0.211
Female, Single	−0.165	−0.219	−0.198	−0.252
Female, Married	−0.095	−0.131	−0.140	−0.176
Age 16-24	−0.227	−0.310	−0.264	−0.347
Age 25-54	−0.111	−0.151	−0.166	−0.207
Age 55+	−0.064	−0.071	−0.101	−0.108
White	−0.118	−0.165	−0.177	−0.224
Black	−0.265	−0.351	−0.339	−0.425
High School or Less	−0.137	−0.228	−0.215	−0.306
Some College or More	−0.031	−0.057	−0.077	−0.103
Low Wage	−0.157	−0.226	−0.224	−0.293
High Wage	−0.091	−0.118	−0.138	−0.166
Home Owner	−0.103	−0.145	−0.186	−0.228
Renter	−0.217	−0.295	−0.302	−0.380
No Children	−0.163	−0.213	−0.224	−0.274
Children	−0.121	−0.152	−0.191	−0.222

Note: All variables use the same sample as the baseline VAR apart from those relating to housing tenure, for which the sample is January 1989 to June 2017.

FIGURE C.15. Flow-Based Accounting for Labor Market Aggregates

(A) Employment (Fixed Comp.)   (B) Unemployment Rate   (C) Participation Rate

Note: The black solid line shows the overall response to a contractionary monetary policy shock. The green dashed line shows the response if quits to U or N are held constant. The red dot-dashed line shows the response if both U-to-N and N-to-U rates are held constant. The blue dotted line shows the response if all supply-driven flows are held constant.

FIGURE C.16. Flow-Based Accounting: One Flow at a Time

(A) Unemployment Rate   (B) Employment Rate   (C) Participation Rate

Note: The black solid line shows the overall response of each labor market stock to a contractionary monetary policy shock. The alternative lines show the response if the specified flow rate is held constant at its average level.

C.6. **Additional Results from Flow-Based Accounting.** Panel A of Figure C.15 repeats our flow-based accounting exercise for employment using compositionally-adjusted flows. The results are very similar to Figure 7. Panels B and C apply the same methodology to unemployment and participation. Quits are more important for shaping the response of employment than of unemployment or participation, while flows between U and N matter roughly equally for all three.

C.7. **The Ins and Outs of Unemployment, Employment and Participation.** Figure C.16 applies the accounting methodology on a flow-by-flow basis. The solid black line shows the baseline response of each labor market stock; the alternative lines show hypothetical responses when we hold each transition probability constant. The distance between a hypothetical and the baseline indicates the importance of that flow.

For unemployment (panel A), the counterfactuals holding E-to-U and U-to-E constant reach roughly similar peak levels. This contrasts with exercises examining unconditional

variation: Shimer (2012) concludes that U-to-E flows account for three quarters of unconditional variation in unemployment. Panels B and C show the analogous decompositions for employment and participation.

Appendix D. Response of Flows to Main Business Cycle Shock

An important question is whether our findings on the role of labor supply are specific to monetary policy shocks. To address this, we consider the "Main Business-Cycle Shock" of Angeletos et al. (2020), identified by maximizing its contribution to the volatility of unemployment over business cycle frequencies.

In Figure D.1, we append our labor market flows to the ten variables in Figure 1 of Angeletos et al. (2020), replacing GDP with Industrial Production for comparability with our baseline results. We rescale the shock so that the maximum increase in the unemployment rate is 0.2 percentage points, matching Figure 1. The labor market flows—including quits and layoffs—respond very similarly to Figures 2 and 3.

This suggests that the labor supply response we document is not specific to monetary policy shocks, but rather operates in response to shocks that drive business cycle fluctuations in unemployment more broadly. Our model offers an explanation: the key driver of the labor

FIGURE D.1. Response of Variables to Main Business Cycle Shock

Note: We take quarterly averages of the labor market flows and append them to the VAR used by Angeletos et al. (2020). See text for details.

supply response is the change in the job-finding rate (Figure 10), so we would expect similar responses to any aggregate shock that significantly moves job-finding rates.

Appendix E. Model Appendix

E.1. **Timing.** The timing of the model within each period is summarized as follows:

(1) All individuals draw a new value of productivity, z. Non-employed individuals draw an i.i.d. search cost, κ.

(2) All individuals make consumption/saving decisions. The employed choose whether or not to quit their job. The non-employed choose whether or not to search.

(3) Employed individuals who do not quit are exogenously laid off with probability δ. Non-employed individuals receive job offers with probability f_s or f_{ns}.

(4) Non-employed individuals who receive job offers decide whether or not to accept.

(5) UI-eligible non-employed individuals who search and either do not receive a job offer or do not accept an offer are subject to UI expiry with probability δ_{UI}.

E.2. **Additional Computational Details.** Our solution method is as follows:

- We discretize the productivity process using the method of Rouwenhorst (1995) using 25 gridpoints. We discretize the asset grid using 200 gridpoints.

- We solve the consumption/saving problem at each gridpoint using golden-section search, linearly interpolating value functions where required.

- Given the distribution of the search cost, we can calculate the probability that an individual at a given (a, z) point will search.

- With the policy functions in hand, we simulate the model on the same discrete grid using non-stochastic simulation and iterate to the stationary distribution.

- When solving for the response to an aggregate shock, we apply standard methods for dealing with an unanticipated"MIT" shock.

In order to generate smooth responses of labor market transition rates, while simulating the model on a discrete grid, we introduce very small "taste shocks" which perturb the quit and job acceptance decisions that agents face. In particular, each period employed agents make their quit decision after drawing a taste shock, ϵ_Q from a logistic distribution. They then make their quit decision taking this into account:

$$V_E(a, z, \epsilon_Q) = \max_{c, a'}\left\{u(c) + \beta \max\left\{\mathbb{E}\, V_{NoUI}(a', z', \kappa') + \epsilon_Q, \mathbb{E}\left[\delta_L V_{UI}(a', z', \kappa') + (1 - \delta_L)V_E(a', z')\right]\right\}\right\}$$
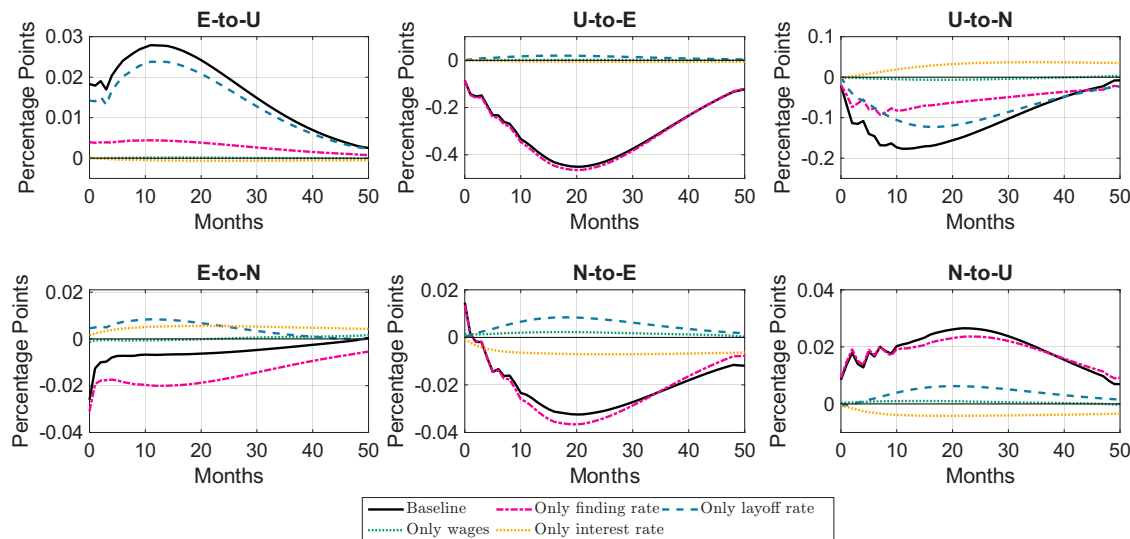
$$(E.2)$$

The scale of these shocks is chosen to be very small: The taste shocks are drawn from a logistic distribution with mean 0 and scale 0.005. (This scale is nearly two orders of magnitude smaller than that which we estimate for the stochastic search cost, κ.) The

TABLE E.1. Average Labor Market Transition Rates

	EU	EN	UE	UN	NE	NU
Model	0.0142	0.0292	0.2540	0.2259	0.0451	0.0252
Data	0.0142	0.0292	0.2539	0.2258	0.0451	0.0251

Note: Transition rates are calculated in the stationary distribution of the model.

FIGURE E.1. Response of Labor Market Flows to Each Shock Component

Note: The black solid line shows the overall response of each flow in the model. The alternative lines show the response if the given component of the shock is fed into the model on its own.

economic significance of these shocks is minimal: In the stationary distribution only 1.4% of employed individuals have a quit probability (before the realization of their quit taste shock) that is between 0.01% and 99.99%. We introduce equivalent taste shocks when non-employed workers have a decision on whether to accept a job.

E.3. **Additional Results From the Model.**

E.3.1. *Steady-State Flows.* Table E.1 compares the steady-state labor market transition rates in the model with their average values in the data. The model closely matches the data for all six transition rates.

E.3.2. *Each Component of the Monetary Policy Shock.* Figure E.1 shows the response of labor market flows when we feed in the four components of the monetary policy shock one at a time. The fall in the job-finding rate is the key driver of the decline in quits from employment to nonparticipation and the increase in search effort by the non-employed (evident in the rise in the N-to-U rate). The change in the layoff rate has little effect on flows other than E-to-U, which it directly affects. The exception is U-to-N, which falls

notably when only the layoff rate changes. Since the layoff rate has almost no effect on labor supply policy functions (Figure 10), this decline reflects a composition effect: layoffs shift the pool of unemployed toward workers with greater labor force attachment. This mirrors the composition effect estimated in Figure 5.

Appendix F. Matching Function

One limitation of our approach in Section 6 is the assumption that the job-finding rate is exogenous. In theory, the job-finding rate is determined by vacancies and "effective search effort", and thus changes in labor supply may affect the job-finding rate. To investigate this possibility, we use our impulse responses to estimate a standard matching function, allowing us to quantify the relative importance of the fall in vacancies and the rise in effective search effort in explaining the decline in the job-finding rate. We find that feedback from labor supply to the job-finding rate is limited.

Consider a standard Cobb-Douglas matching function,

$$M(S,V) = \mu S^{1-\eta}V^{\eta}, \qquad\qquad (F.3)$$

where S is "effective search" and V is the number of vacancies. In our model, all non-employed receive job offers, and nonparticipants receive job offers at a rate that is α times that of the unemployed, so that effective search takes the form $S = U + \alpha N$.

In the left panel of Figure F.1, we plot the response of effective search in the data following a contractionary monetary policy shock. The rise in effective search is driven by the rise in unemployment. In the middle panel, we use the responses of effective search, vacancies and the U-to-E rate to estimate the parameters of the matching function.[49] Our estimated matching function ($\eta = 0.66$) implies a path of the job-finding rate (the blue line) that closely fits the U-to-E rate in the data. The two dotted lines consider alternative paths of the job-finding rate, holding either vacancies or effective search constant at their average levels. This exercise implies that the main driver of the decline in the job-finding rate is the decline in vacancies, rather than the increase in effective search.

In the right panel of Figure F.1, we use our estimated model to assess the scope for feedback from labor supply to the job-finding rate. The line labeled "Model Search" feeds the path of effective search from the model into the estimated matching function; its close fit to the solid blue line confirms that the model replicates the path of effective search in the data. The line labeled "Labor Policies Fixed" instead uses the path of effective search from

---

[49]This model-free approach assumes that the job-finding rate from unemployment is equal to the U-to-E rate, and thus that all unemployed individuals accept job offers. In our model, a small fraction of the unemployed choose to reject offers. We could take this into account when estimating the matching function, with very similar results.

FIGURE F.1. Effective Search and Estimated Matching Function

Note: The left panel plots the response of effective search. The middle and right panels plot the path of the U-to-E rate under various scenarios. See text for details.

the version of the model where labor supply policy functions are held at steady state (as in Figure 10).

The similarity between "Model Search" and "Labor Policies Fixed" suggests that feedback from labor supply to the job-finding rate is minimal. If anything, our findings imply that our counterfactual in Figure 10 *understates* the employment decline when labor supply policies are held fixed. With fixed policies, quits to non-employment do not decline, so effective search rises by more, which would further lower the job-finding rate and amplify the employment decline.

REFERENCES

Ahn, Hie Joo and James D. Hamilton, "Measuring labor-force participation and the incidence and duration of unemployment," *Review of Economic Dynamics*, 2022, *44*, 1–32.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas, "Business-cycle anatomy," *American Economic Review*, 2020, *110* (10), 3030–3070.

Aruoba, S. Borağan and Thomas Drechsel, "Identifying Monetary Policy Shocks: A Natural Language Approach," *American Economic Journal: Macroeconomics*, 2026, p. forthcoming.

Barnichon, Regis and Christian Brownlees, "Impulse response estimation by smooth local projections," *Review of Economics and Statistics*, 2019, *101* (3), 522–530.

Bauer, Michael D. and Eric T. Swanson, "An Alternative Explanation for the 'Fed Information Effect'," *American Economic Review*, 2023, *113*, 664–700.

⸻ **and** ⸻ , "A Reassessment of Monetary Policy Surprises and High-Frequency Identification," *NBER Macroeconomics Annual*, 2023, *37* (1), 87–155.

Bouscasse, Paul and Seungki Hong, "Monetary–Fiscal Interactions in the United States," Working Paper, Unpublished working paper Sep 2023. First version: June 2023; current draft: September 2023.

Elsby, Michael W.L., Bart Hobijn, and Ayşegül Şahin, "On the importance of the participation margin for labor market fluctuations," *Journal of Monetary Economics*, 2015,

72, 64–82.

Faia, Ester, Ekaterina Shabalina, and David Wiczer, "Heterogeneous Effects of Monetary Policy Across Income and Race: The Labour Mobility Channel," CEPR Discussion Papers 17741, C.E.P.R. Discussion Papers October 2023.

Miranda-Agrippino, Silvia and Giovanni Ricco, "The transmission of monetary policy shocks," *American Economic Journal: Macroeconomics*, 2021, *13* (3), 74–107.

Mukoyama, Toshihiko, Christina Patterson, and Ayşegül Şahin, "Job search behavior over the business cycle," *American Economic Journal: Macroeconomics*, 2018, *10* (1), 190–215.

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

Rouwenhorst, K Geert, "Asset pricing implications of equilibrium business cycle models," *Frontiers of business cycle research*, 1995, *1*, 294–330.

Shimer, Robert, "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, April 2012, *15* (2), 127–148.

Swanson, Eric T. and Vishuddhi Jayawickrema, "Speeches by the Fed Chair Are More Important than FOMC Announcements: An Improved High-Frequency Measure of U.S. Monetary Policy Shocks," 2024. Working Paper, University of California, Irvine.

U.S. Census Bureau, "Design and Methodology," Technical Report CPS TP77, U.S. Bureau of Labor Statistics 2019.

White, Neil, "Gross Worker Flows, Job Loss, and Monetary Policy," 2018.

Wieland, Johannes F and Mu-Jeung Yang, "Financial dampening," *Journal of Money, Credit and Banking*, 2020, *52* (1), 79–113.