

NBER WORKING PAPER SERIES

THE LABOR DEMAND AND LABOR SUPPLY CHANNELS OF MONETARY POLICY

Sebastian Graves
Christopher K. Huckfeldt
Eric T. Swanson

Working Paper 31770
<http://www.nber.org/papers/w31770>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2023, Revised October 2024

We thank Hie Joo Ahn, Felipe Alves, Ryan Chahrour, Bart Hobijn, André Kurmann, Kris Nimark, Andrea Prestipino, Ayşegül Şahin, and Gianluca Violante, as well as participants at various seminars and conferences, for many helpful comments. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, any other person associated with the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Sebastian Graves, Christopher K. Huckfeldt, and Eric T. Swanson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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 October 2024
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 non-employment, 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
International Finance Division
Federal Reserve Board
20th & Constitution Ave., NW
Washington, DC 20551
chris.huckfeldt@gmail.com

Eric T. Swanson
Department of Economics
University of California at Irvine
3151 Social Science Plaza
Irvine, CA 92697-5100
and NBER
eric.swanson@uci.edu

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 non-employment, 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); in statements by other monetary policymakers around the world; and in the New Keynesian (NK) literature, where the standard assumption of sticky wages in a neoclassical

We thank Hie Joo Ahn, Felipe Alves, Ryan Chahrour, Bart Hobijn, André Kurmann, Kris Nimark, Andrea Prestipino, Ayşegül Şahin, and Gianluca Violante, as well as participants at various seminars and conferences, for many helpful comments. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System. *Graves*: University of Cambridge, sg566@cam.ac.uk. *Huckfeldt*: Board of Governors of the Federal Reserve System, chris.huckfeldt@frb.gov. *Swanson*: University of California, Irvine, and NBER, eric.swanson@uci.edu. First version: February 2023. This version: October 9, 2024.

labor market precludes any significant quantitative role for labor supply considerations to affect the response of employment to a monetary policy shock.¹

In contrast to the consensus 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 driven by labor supply considerations. While the response of all labor market flows can be thought of as being determined in general equilibrium, certain labor market flows are more directly reflective of labor supply insofar as they are initiated by the worker. 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 non-employment 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 shows that a large and procyclical component of flows from E to N is due to quits.

We then estimate the response of labor market flows to exogenous changes in monetary policy. There are three main reasons to look at the conditional response of labor market flows to an exogenous demand shock (like a monetary policy shock) as opposed to their unconditional behavior. First, knowing how labor market flows respond to a monetary policy shock improves our understanding of how monetary policy affects the economy. Second, the conditional response of labor market flows to a cleanly identified demand shock can provide insights that would be obscured in the unconditional setting, which combines the responses to many types of shocks (technology shocks, monetary policy shocks, preference shocks, etcetera). For example, it would not be surprising if supply-driven flows varied in response to a preference shock, while it is more surprising if they respond to a monetary policy shock. Third, the conditional responses of labor market flows to different shocks can similarly provide more stringent tests of a given model. Indeed, in the final part of the paper, we use our impulse responses as targets to estimate a heterogeneous agent model with frictional labor markets and an active labor supply margin.

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

¹Christiano (2011), Broer et al. (2020) and Wolf (2023) offer detailed discussions of this property of the sticky-wage NK model. See also the discussion below.

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 consensus 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 the standard interpretation of lower labor demand amidst a weakening economy.² However, in contrast to the consensus view, 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 non-employment. We further identify a significant reduction in quits from employment to nonparticipation. Intuitively, this response of supply-driven flows is consistent with an income effect, where households increase their labor supply in a weakening economy to maintain their consumption, as in the classic “added worker effect” literature of Lundberg (1985) and others; as well as a precautionary response to increased income risk.

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 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

²We use the terminology “flows” and “transition probabilities” interchangeably throughout.

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 70 percent larger if we simulate the model holding labor supply policy functions fixed at steady-state.

Thus, our empirical estimates and model indicate an important role for labor supply in explaining the response of employment to a monetary policy shock. As such, our findings sharply contrast with the typical approach of the sticky-wage NK literature, where “demand-determined” labor precludes any quantitatively meaningful role for labor supply (as discussed in detail by Christiano (2011), Broer et al. (2020), Auclert, Bardóczy and Rognlie (2021), and Wolf (2023)). Insofar as our estimates and model indicate an important role for labor supply in explaining the responses of labor market flows (and thus labor market stocks) to a monetary policy shock, our paper highlights a potentially important shortcoming of the transmission mechanism in such models.

We believe our evidence yields insights beyond improving our general understanding of the monetary transmission mechanism: for example, labor supply appeared to 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 EU and EN 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. 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). The responses of labor market flows are typically imprecisely estimated for the subset of papers in this literature analyzing U.S. data, likely due to the interaction of the well-known difficulty of identifying suitable measures of monetary policy shocks and the additional sampling noise introduced from the longitudinal linking of individual-level data

necessary to construct measures of labor market flows. To mitigate these issues, we follow the approach of Bauer and Swanson (2023a,b) in adopting an instrument that incorporates additional interest rate variation around Fed Chair speeches and is orthogonalized with respect to recent macroeconomic and financial market news. In doing so, not only do we avoid standard critiques of monetary policy SVARs (e.g., Ramey, 2016), we also obtain an instrument that is both more relevant and more exogenous than those used by other authors, giving us more precise and less biased estimates. Thus, our paper offers new benchmark estimates for the response of U.S. labor market flows to a monetary policy shock.

The greater accuracy of our estimated impulse responses also allows us to extend the above literature in several ways, by: 1) studying the quit and layoff components of flows from employment to non-employment; 2) considering the role of composition effects in the responses of labor market flows; 3) studying heterogeneity in impulse responses of disaggregated labor market flows; and 4) analyzing the quantitative importance of the response of supply-driven labor market flows.

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 and unemployment (e.g., Davis, Faberman and Haltiwanger, 2006; Shimer, 2012). In studying flows between unemployment and nonparticipation, we follow Elsby et al. (2015), who assess the importance of these flows for cyclical variation in unemployment and labor force participation. We additionally study the effect of these flows on the employment-population ratio, showing that flows between U and N are quantitatively more important for employment dynamics than unemployment dynamics in response to a monetary policy shock.

A distinctive contribution of our paper to the empirical literature is to introduce a novel measure of quits from employment to nonparticipation in the CPS, which we show constitutes a large component of the total flow of workers from employment to non-employment. Beyond using this new data to document the unconditional cyclical behavior of quits to non-employment, we also estimate 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. Although macroeconomic models à la Krusell et al. (2017) (such as ours) offer direct predictions for quits to nonparticipation, these implications typically go unexplored, likely due to a lack of data.

While conditional responses of labor market flows to monetary policy shocks are of independent interest, our estimates highlight some differences from the unconditional cyclical behavior of labor market flows: 1) our estimates show a more important role for EU flows compared to the unconditional literature (Appendix C.7); 2) we find a less prominent role

for composition in explaining the response of labor market flows (Section 4.2); and 3) we find no response of job-to-job flows to a monetary policy shock (Appendix C.3.2).

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 our estimated increase in labor supply significantly dampens a decline in labor force participation that would have been *much larger* in the absence of that supply 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 model in Section 6 also exhibits these features.

Our paper also relates to the sticky-wage New Keynesian literature, which typically assumes that workers are assigned hours by a representative union (e.g., Christiano, Eichenbaum and Evans, 2005; Smets and Wouters, 2007; Auclert, Rognlie and Straub, 2020) or includes search and matching frictions but assumes labor to be supplied inelastically (e.g., Gertler, Sala and Trigari, 2008; Christiano, Eichenbaum and Trabandt, 2016; Ravn and Sterk, 2017; Graves, 2023). For standard parameterizations of models in the former case, the combination of sticky wages and a representative union limits the labor supply response of households to such an extent that employment is effectively demand-determined, as discussed by Christiano (2011), Broer et al. (2020), and Wolf (2023); whereas in the latter case, labor supply considerations are ruled out altogether.³

The literature typically offers two justifications for the assumptions of a limited household labor supply response: First, the intensive margin of labor (i.e., hours) shows very little conditional or unconditional business cycle variation, which is taken as evidence for the limited importance of labor supply.⁴ Second, estimates of the marginal propensity to earn (MPE) from idiosyncratic transfers are small, suggesting that labor supply is relatively inelastic (Auclert et al., 2021). Our analysis tackles both of these justifications. First, even if hours variation is limited, we still document considerable variation along the extensive margin of labor supply (i.e., flows between U and N and quits from E to N). Second, we show that a low MPE from an idiosyncratic transfer does not necessarily imply a small labor supply response to an aggregate shock: in particular, our model matches the substantial response of supply-driven labor market flows to a monetary contraction, but also generates a small MPE consistent with typical estimates in the literature.

³Note that analytic sticky-price NK models with fixed real rates, such as Woodford (2011) and Angeletos, Lian and Wolf (2023), also exhibit demand-determined labor.

⁴For an alternative view, see Bilbiie, Primiceri and Tambalotti (2023) and Cantore et al. (2023), who consider TANK models incorporating cross-sectional variation in hours within a representative labor union setting.

Finally, our paper is complementary to the contemporaneous work of Alves and Violante (2023), who extend a framework similar to that in Krusell et al. (2017) into a rich HANK model to study how monetary policy rules influence long-run inequality. 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 an increase in layoffs and a decrease in quits in response to a contractionary monetary policy shock offers validation for their results.

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 are plausibly driven by 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 then related by the Markov process

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

TABLE 1. Cyclicalities of Labor Market Stocks

	Employment- Population Ratio	Unemployment Rate	Participation Rate
$\text{mean}(x)$	61.14	6.19	65.16
$\text{std}(x)/\text{std}(Y)$	0.72	8.25	0.23
$\text{corr}(x, Y)$	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 quarterly averages. The sample is 1978-2019.

TABLE 2. Cyclicalities of Labor Market Flows

	EU	EN	UE	UN	NE	NU
$\text{mean}(x)$	0.014	0.030	0.255	0.226	0.046	0.025
$\text{std}(x)/\text{std}(Y)$	5.20	2.46	5.69	4.14	3.00	5.22
$\text{corr}(x, Y)$	-0.83	0.49	0.78	0.71	0.65	-0.68

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

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

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

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 The properties of these transition probabilities have been well documented in the literature (e.g., Shimer, 2012; Elsby et al.,

⁵We seasonally adjust each flow using the X-13ARIMA-SEATS software provided by the Census Bureau. Given our subsequent focus on quits and layoffs from non-employment, we do not adjust for time aggregation bias. Our results are robust to corrections for time aggregation, where such corrections are possible (e.g. see Appendix Figure C.1). See Figure A.1 for plots of each transition probability for our sample.

TABLE 3. Components of EU and EN Flows

	EU Flows				EN Flows			
	Total	Quits	Layoffs	Other	Total	Quits	Layoffs	Other
$\text{mean}(x)$	0.014	0.002	0.008	0.004	0.030	0.012	0.003	0.015
$\text{std}(x)/\text{std}(Y)$	5.20	8.11	8.03	5.43	2.46	5.88	14.42	4.80
$\text{corr}(x, Y)$	-0.83	0.60	-0.83	-0.54	0.49	0.53	-0.44	0.25

Note: The process for decomposing EU and EN 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 quarterly averages. The sample is 1978-2019.

2015; Krusell et al., 2017). Here we simply note that we consider flows between nonparticipation and unemployment as being driven by supply considerations, given that such flows are initiated by workers. The procyclicality of UN flows and countercyclicality of NU flows can be interpreted as evidence of greater job-seeking behavior among the non-employed during downturns and account for around one-third of cyclical variation in the unemployment rate (Elsby et al., 2015). Finally, note that the average UE probability is more than five times greater than the average NE 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 EU and EN transitions are driven by labor supply choices, we decompose EU and EN 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 non-employment as supply-driven.

Although many authors have studied the cyclical and composition of EU flows, far less attention has been paid to EN 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 EN flows into quits and layoffs. This decomposition is more complicated to construct than that for EU 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 EN transitions into quits and layoffs, including evidence showing that the subsequent labor market transition probabilities of individuals who quit to non-employment differ significantly from those of individuals who are laid off.

⁶For 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.

The left panel of Table 3 summarizes the size and cyclical properties of the subcomponents of EU flows. About 60% of EU 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 similarly volatile, they are procyclical. The remaining 25–30% of EU 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 EN flows. EN layoffs are countercyclical and EN quits are procyclical, as is the case for EU flows. However, in contrast to EU flows, quits represent a much larger share of EN flows than layoffs, implying a much more important role for both the magnitude and cyclical property of quits to non-employment than has been previously recognized. Indeed, the portion of EN flows that can be identified as quits is of similar magnitude to the entirety of EU flows. 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.⁷

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 six monthly macroeconomic variables: the log of industrial production, the unemployment rate, the labor force participation rate, the log of the consumer price index, the Gilchrist and Zakrajšek (2012) excess bond premium, and the two-year Treasury yield.⁸ This specification is very similar to Bauer and Swanson (2023b), except that we include labor force participation as an additional variable, given our focus on the labor market (and we will also extend this core VAR to include labor market flow variables,

⁷Faberman and Justiniano (2015) explain their use of the JOLTS quit rate as a proxy for the job-to-job transition rate from the finding of Elsby, Hobijn and Sahin (2010) that only 16% of quits lead to unemployment. Our findings suggest that a non-trivial fraction of JOLTS quits may reflect quits to nonparticipation rather than job-to-job transitions.

⁸Industrial 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. We include the GZ excess bond premium for consistency with Bauer and Swanson (2023b) and because Caldara and Herbst (2019) found including a credit spread is important for the estimation of monetary policy VARs. 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.

below). We stack these six core variables into a vector Y_t and estimate the reduced-form VAR,

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

where α is a constant, $B(L)$ a matrix polynomial in the lag operator, and u_t is a 6×1 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

$$u_t = S\varepsilon_t, \quad (4)$$

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} .⁹ 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 and Fed Chair speeches, converted to a monthly series by summing over all the high-frequency surprises within each month.¹⁰ In order for z_t to be a valid instrument for ε_t^{mp} , it must satisfy an instrument *relevance* condition,

$$E[z_t \varepsilon_t^{mp}] \neq 0, \quad (5)$$

and an instrument *exogeneity* condition,

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

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

⁹If 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.

¹⁰High-frequency interest rate changes around FOMC announcements and Fed Chair speeches are from Swanson and Jayawickrema (2023) and include all 323 FOMC announcements from 1988–2019 and all 404 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 somewhat 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.

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 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).¹¹

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

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

via two-stage least squares, using z_t as the instrument for Y_t^{2y} .¹² 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).¹³

Finally, we follow the prescriptions of Bauer and Swanson (2023a,b) and adjust our high-frequency instrument z_t 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

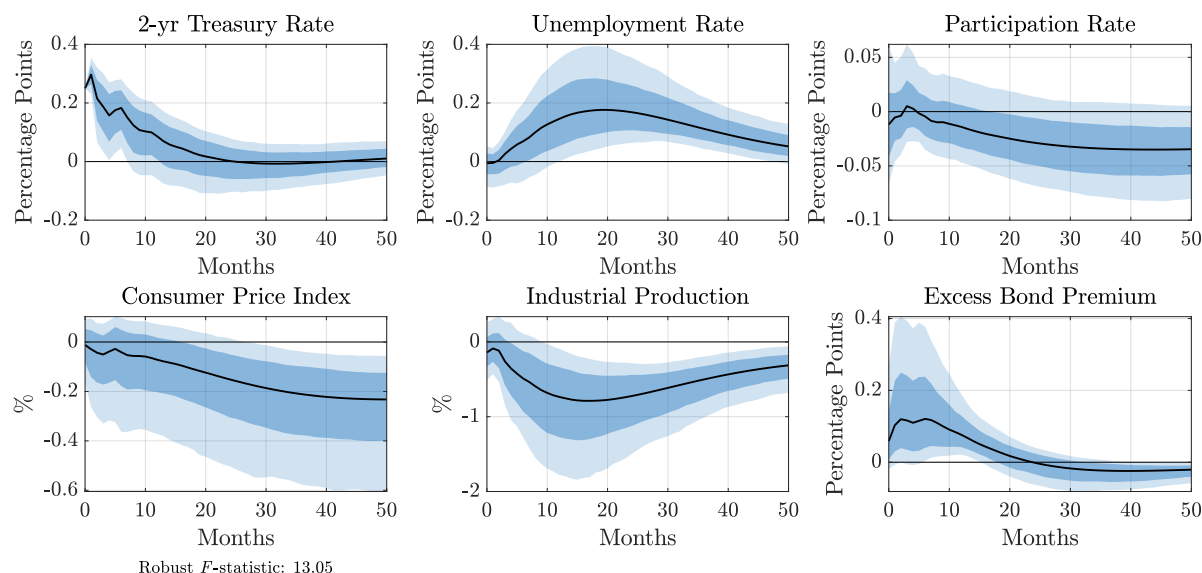
We present several sets of results. First, we report baseline impulse response functions (IRFs) for the core six-variable VAR described above. Second, we extend this core VAR to include labor market flow variables and report IRFs for labor market flows. Third, we augment the core VAR to include the quit and layoff subcomponents of EU and EN flows to provide

¹¹Swanson and Jayawickrema (2023) use narrow intradaily windows around these announcements and are careful to avoid overlapping with any other macroeconomic data releases.

¹²One 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.

¹³Note 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 1988:2–2019:12, while we estimate $B(L)$ over the longer sample 1978:1–2019:12.

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



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.

additional evidence of the response of supply-driven flows. Finally, we augment our core VAR with additional variables to further understand the response of the participation rate.

3.1. Baseline VAR Impulse Responses to a Monetary Policy Shock. Estimated IRFs from the core six-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 13.05, 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. The Gilchrist and Zakrajšek (2012) excess bond premium, in the bottom right panel, increases by 5bp on impact and rises 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 begins to decline 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, -1 percent for industrial production, and -0.2 percent for the CPI. These responses

are similar to those from monetary policy VARs estimated by other 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 the labor force participation rate. Although speeches by monetary policymakers increasingly include references to labor force participation to convey the economy’s proximity to “maximum employment,” the response of labor force participation to a monetary policy shock has received less study than that of other labor market variables. Our estimates suggest that a contractionary monetary policy shock generates a slow-moving decline in labor force participation, which begins to fall 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 higher labor supply, but an even larger decrease due to higher unemployment interacting with the high average level of UN relative to EN transitions (22.6 percent vs. 3 percent).


3.2. Responses of Labor Market Flows to a Monetary Policy Shock. We next extend our core six-variable VAR to include labor market flows. Extending the VAR to include all six labor market flows (EN, EU, NE, NU, UE, and UN) 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 seven-variable VAR—the six variables in the baseline VAR, above, plus the labor market flow variable listed at the top of the panel.¹⁴ 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 25.5 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 UE and EU 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 (UE) in the top middle panel of Figure 2 falls significantly in response to the monetary tightening, consistent with a fall in hiring, while the transition rate from employment to unemployment (EU) in the top left panel

¹⁴IRFs for the six 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.

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.

increases significantly, consistent with an increase in layoffs. This latter increase may seem small at first glance—about 0.025 percentage points at its peak—but it is sizeable relative to the steady-state flow of about 1.4 percent each month.¹⁵ Moreover, the increase in EU flows is highly persistent, especially compared to the more transitory increase in EU 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 (NU) 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 non-employment to look for a job *increases* significantly. Simultaneously, the symmetric flow from unemployment to nonparticipation (UN) in the top right panel *declines*. The increase in NU and decrease in UN flows tilts the composition of non-employment (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 model, in Section 6). In Appendix C.3.1, we present additional

¹⁵Because 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.

evidence from two “intensive margins” of job search—the number of search methods used by the unemployed and the fraction of nonparticipants that report wanting a job—that are also consistent with an increase in labor supply.

The flow from employment to nonparticipation (EN) in the bottom left panel of Figure 2 declines modestly. We show in the next section that a labor supply response is crucial for explaining why the EN rate declines in response to a contractionary shock, while the EU rate rises significantly. Finally, the flow from nonparticipation to employment (NE) in the bottom middle panel responds similarly to the UE flow, but by a smaller amount.

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


Note that the few other papers that study the responses of labor market flows to monetary policy shocks in the U.S. show far less conclusive estimates. For example, White’s (2018) estimated responses of UN and NU flows using Romer and Romer (2004) shocks are not significant from zero at the 68% level for most horizons; and Faia, Shabalina and Wiczer (2023) find ambiguous evidence for the response of UE flows to high-frequency monetary surprises around FOMC announcements, contrasting not only with our estimates of a significant decline following a monetary contraction, but also with the literature documenting substantially procyclical unconditional variation in UE flows (e.g., Shimer, 2012).

We speculate that the problem of imprecise (and possibly biased) estimates in the other papers stems from well-documented issues associated with identifying exogenous and relevant measures of monetary policy surprises, compounded by the presence of additional sampling noise from the longitudinal linking of individual-level data necessary to construct measures of labor market flows. By following the approach of Bauer and Swanson (2023a,b), our instrument is both more relevant and more exogenous than those used by these other authors, giving us more precise and less biased estimates of the response of labor market flows.¹⁷

¹⁶In 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.

¹⁷We offer an example of how our estimates are more relevant and exogenous in Appendix C.2. In particular, when we estimate the response of UE flows to a contractionary monetary policy shock using non-orthogonalized monetary policy surprises around FOMC announcements à la Gertler and Karadi (2015), we get point estimates close to zero and large standard errors, consistent with Faia, Shabalina and Wiczer (2023). As in Bauer and Swanson (2023b), we interpret these results as reflecting problems with exogeneity (due to non-orthogonalization) and relevance (due to looking only at FOMC announcements).

FIGURE 3. Decomposition of EU and EN Responses



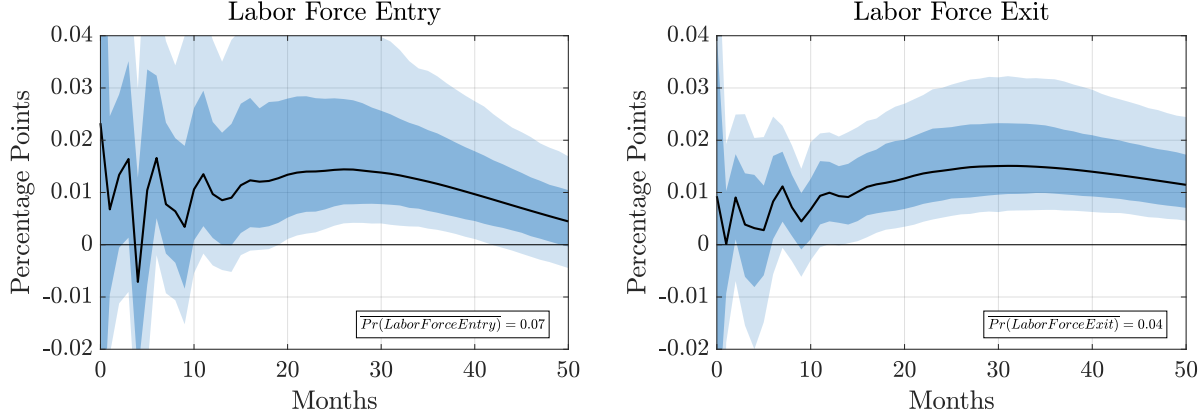
Note: Estimated impulse responses to a 25bp monetary policy tightening shock, computed by appending the given labor market flow variable to the baseline VAR from Figure 1, where “E” denotes employment, “U” denotes unemployment, and “N” denotes nonparticipation. Solid black lines report impulse response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

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 EU and EN flows (defined in Section 2.2) to a 25bp monetary policy tightening. Each of these variables is appended to our core six-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, reinforcing the evidence of an increase in labor supply found in the response of UN and NU flows (as we formalize in the discussion of our model in Section 6). The portion of EU flows that cannot be definitively attributed to layoffs or quits increases modestly, while the unattributed EN flow rate declines slightly.¹⁸ As layoffs represent a much larger fraction of EU flows than quits, the overall response of EU flows tracks that of the layoffs component. The opposite is true for EN flows: the modest decline in the overall EN rate in response to a contractionary monetary policy shock

¹⁸While 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).

FIGURE 4. Response of Labor Force Entry and Exit 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. 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.

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 EU transition) or enter nonparticipation (as an EN transition). The impulse responses in Figure 3 show a proportionally larger increase in EU layoffs than EN layoffs, 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.

3.4. Labor Force Entry and Exit. Sections 3.2 and 3.3 document that supply-driven labor market flows respond in a manner consistent with an increase in labor supply. However, we also showed in Section 3.1 that a contractionary monetary policy shock leads to a sluggish and modest decline in labor force participation, a response that is sometimes interpreted as consistent with a decline in labor supply (e.g., Galí et al., 2012). Here, we study how labor force entry and exit (and their constituent flows) generate the estimated decline in participation from a monetary contraction. Note first that labor force entry and exit rates satisfy:

$$(\text{Labor Force Entry Rate})_t = NU_t + NE_t, \quad (8)$$

$$(\text{Labor Force Exit Rate})_t = u_{t-1} \cdot UN_t + (1 - u_{t-1}) \cdot EN_t, \quad (9)$$

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. The labor force exit rate is the weighted average of UN and EN flows, with weights u_{t-1} and $(1 - u_{t-1})$ corresponding to the fractions of the labor force that are unemployed and employed, respectively.

Figure 4 plots the estimated impulse response functions of the labor force entry and exit rates, again each computed by appending the given variable to the baseline VAR from Section 3.1. Labor force entry and exit both increase in response to a contractionary monetary policy shock, with the increase in entry putting upward pressure on participation and the increase in exit doing the opposite. Thus, 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.

Equation (8) shows that the dynamics of labor force entry reflect those of NU and NE flows. Recall from Section 3.2 that NU flows increase in response to a contractionary monetary policy shock, while NE flows decrease. Thus, from equation (8), the increase in workers moving from nonparticipation to unemployment (NU) more than offsets the decline in job-finding from nonparticipation (NE), generating an overall increase in labor force entry. We have argued that NU flows are driven primarily by labor supply considerations, so the increase in labor force entry in Figure 4 is consistent with an increase in labor supply.

Equation (9) shows that the dynamics of labor force exit reflect not just those of UN and EN flows, but also those of the unemployment rate u_t itself. Intuitively, an increase in unemployment after a contractionary monetary policy shock tilts the distribution of the labor force towards unemployment, and UN transitions are much more common than EN transitions—about 22.6 percent per month vs. 3 percent. This places substantial upward pressure on labor force exits. In fact, recall from Section 3.2 that both UN and EN flows decrease in response to a monetary contraction, which puts downward pressure on labor force exit. Thus, the increase in labor force exit in Figure 4 in response to a monetary contraction—and therefore, the decline in participation—is entirely driven by the increase in unemployment.

Thus, our findings accord with Hobijn and Şahin (2021), who show that the unconditional dynamics of participation are driven by unemployment. 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.

3.5. Additional Results. Appendix C.3 presents additional results on the response of other labor market variables. We show that a contractionary monetary policy shock leads to increases in two “intensive margins” of job search: the number of search methods used by the unemployed and the fraction of nonparticipants that report wanting a job. We also show that the job-to-job transition rate does not respond significantly. Finally, we show that a monetary contraction leads to a decline in vacancies and a modest rise in real wages.

4. COMPOSITION AND HETEROGENEITY

The impulse response functions for supply-driven labor market flows in Figures 2 and 3 are consistent with a quantitatively important 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 response of supply-driven labor market flows 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


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

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 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

¹⁹We 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.

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.

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 six-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), the impulse responses in Figure 5 are largely unchanged. An exception is the IRF for UN flows, which decreases by roughly half as much when holding composition fixed. This suggests that part of the decline in UN 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 UN flows from the end of an expansion through a recession are due to changes in the composition of the unemployed.²⁰ In Section 5, below, we discuss further how controlling for composition has little effect on our finding that supply-driven flows are quantitatively important for the response of employment to a monetary policy shock.

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

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. Inset boxes report average transition rates.

4.2. Heterogeneity. While the above section shows that our results on the quantitative importance of supply-driven flows are robust to controlling for composition effects, it does not preclude heterogeneous labor supply responses across different types of workers.

In particular, we study here the labor supply responses of higher- and lower-educated workers.²¹ Lower-educated workers typically have fewer financial assets with which to smooth consumption. We also establish that lower-educated workers face more severe reductions in employment in response to a monetary policy contraction, due in large part to a greater increase in the probability of being laid off. Under the interpretation that the aggregate response of supply-driven flows to a monetary contraction can be understood through precautionary behavior, we should expect a greater response of such flows from lower-educated workers. We show that this is indeed the case: lower-educated workers exhibit a greater response of supply-driven flows in the face of a monetary policy contraction, most evidently through a decrease in quits to nonparticipation.

Figure 6 reports impulse responses of the employment-population ratios and EU and EN flows for higher- and lower-educated workers, computed by extending our baseline six-variable VAR one variable at a time, as in previous sections. The left column of Figure 6

²¹We 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.

reports IRFs for the employment-population ratio of both groups. 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 and remains significant even 50 months after the shock.

The middle and right columns of Figure 6 report the responses of EU and EN flow rates for each education group. The increase in EU flows following a monetary contraction is again about twice as large for lower-educated workers. Splitting by education also shows that the decline in EN flows—which we have shown is driven by a decline in quits to non-employment—is concentrated among lower-educated workers. There is little discernible drop for higher-educated workers.²² The larger decrease in quits to non-employment among lower-educated workers is consistent with a greater response of household labor supply.

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, labor supply responses show important differences across groups: lower-educated workers seem to adjust their labor supply more aggressively to offset the employment impact of a monetary policy shock. Third, to the extent that this supply response is driven not only by a larger increase in layoffs but also through lower asset holdings, our findings suggest that the wealth distribution helps shape the aggregate labor supply response to a monetary policy shock.

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 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 responses of UN and NU flows. Third, we shut down the response of quits to non-employment. Fourth, we shut down the responses of UN and NU flows and quits to non-employment both.

²²This is not to say that there is no labor supply response of more educated individuals: Figure C.11 of the Appendix shows a labor supply response among the higher-educated in the form of higher NU flows and lower UN flows in response to a contractionary monetary policy shock.

²³Appendix C reports analogous results for the responses of the unemployment rate and labor force participation rate—see Figures C.14 and C.15.

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 both UN and NU rates are held constant. The red dot-dashed line shows the response if quits to U or N are held constant. The blue dotted line shows the response if all supply-driven flows are held constant.

Figure 7 plots the results. In the baseline scenario (the solid black line), employment falls a little more than 0.1 percent after about 20 months. In the second scenario, holding UN and NU flows fixed at their average values (the dashed green line), the fall in employment is almost 60% larger than in the baseline. As discussed in Section 3.3, the increase in NU flows and decrease in UN flows after a monetary contraction tilts the distribution of the non-employed from N towards U, and the UE transition rate is much higher than the NE rate (about 25.5 percent per month vs. 4.6 percent). Thus, fixing $U \leftrightarrow N$ flows at their average levels reduces the rate at which workers move from non-employment to employment after the shock, generating the more pronounced employment decline in Figure 7.

In our third scenario, we shut down the response of quits to non-employment (the dashed red line). In this case, employment falls about 40% more than in the baseline, as we are now turning off the significant decline in quits to nonparticipation identified in Section 3.3.

Fourth, we shut down both the response of $U \leftrightarrow N$ flows and quits to non-employment (the dotted blue line) and find that employment declines roughly twice as much relative to baseline.²⁴ Thus, we conclude that the response of supply-driven flows is quantitatively

²⁴Note 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.

very important for the overall response of employment to a monetary policy shock: if we counterfactually hold those labor-supply-driven flows constant at their average values, the response of employment would be about twice as large.

Finally, in Appendix A we show the results of the same exercise for the unemployment and labor force participation rates. We find that the labor force participation rate would decline significantly more without the response of supply-driven labor market flows, consistent with our discussion in Section 3.4. Additionally, in Appendix C.4, we repeat the exercise using our composition-adjusted flows from Section 4. We find that, when controlling for composition, holding supply-driven labor market flows constant still amplifies the decline in employment after a contractionary monetary policy shock by about 75%. This reflects the fact that, while the response of UN flows is partly muted when holding composition constant, composition adjustment has little effect on the other important supply-driven flows: NU flows and quits to non-employment.

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

In our empirical analysis, above, we documented a response of labor market flows consistent with an increase in labor supply in response to a contractionary monetary policy shock. In Section 4, we presented evidence that our main empirical results in Figures 2 and 3 are robust to controlling for cyclical changes in labor force composition and thus reflect changes in individual labor supply. In this section, we develop a simple, structural heterogeneous agent model that is also consistent with this interpretation.

More specifically, 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 (2023).²⁵

We consider the effect of a monetary policy shock in the model by feeding in the response of 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 response of agents in the model and their implications for aggregate labor market flows. 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.

²⁵Two 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).

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. This close fit allows us to then construct a counterfactual path of employment, holding labor supply policy functions at their steady-state value. In this scenario, employment declines about 70% more than in the baseline model. This exercise provides further evidence that our empirical results are driven by true changes in labor supply at the individual level and highlights the features that general-equilibrium New Keynesian models are likely to require in order to be consistent with these results.

Importantly, we also establish that our model is quantitatively consistent with micro-level evidence on marginal propensities to consume (MPCs) and marginal propensities to earn (MPEs). In particular, an existing literature has estimated small MPEs from the labor earnings response to idiosyncratic windfalls. These estimates have been interpreted as justification for treating employment as demand-determined in the NK framework (e.g., Auclert et al., 2021), thereby implicitly ruling out the type of labor supply response that we document earlier in the paper. By simultaneously matching the estimates of labor market flows presented in Figure 2 alongside existing estimates of MPCs and MPEs, our model establishes that a modest labor supply response to an idiosyncratic transfer does not rule out a large labor supply response to an aggregate shock.

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 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, non-employed but eligible for unemployment insurance (UI), or non-employed and not eligible for 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 non-employment.

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 $V_E(a, z)$, $V_U(a, z, \kappa)$, and $V_N(a, z, \kappa)$ 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.²⁶

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:

$$V_E(a, z) = \max_{c, a'} \left\{ u(c) + \beta \max \left\{ \mathbb{E} V_N(a', z', \kappa'), \mathbb{E} [\delta_L V_U(a', z', \kappa') + (1 - \delta_L) V_E(a', z')] \right\} \right\} \quad (11)$$

subject to

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

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 non-employment. 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 non-employment.

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 $V_U(a, z, \kappa)$ satisfies

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

subject to

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

where ψ denotes the flow utility value of non-employment, 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 $(1 - \kappa)\psi$

²⁶We refer to UI-eligible and UI-ineligible workers with subscripts U and N , reflecting that UI-eligible non-employed have greater incentive to search (and thus are more likely to be classified as unemployed, U) compared to the UI-ineligible, who are thus are more likely to be nonparticipants, N .

and continuation value $\mathcal{V}_U^s(a', z)$ —and not searching—in which case the worker receives flow utility ψ and continuation value $\mathcal{V}_U^{ns}(a', z)$.

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

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

$$\mathcal{V}_U^{ns}(a', z) = f_{ns} \cdot \max\{\mathbb{E} V_E(a', z'), \mathbb{E} V_N(a', z', \kappa')\} + (1 - f_{ns}) \mathbb{E} V_N(a', z', \kappa'), \quad (16)$$

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:

$$\tilde{V}_U(a, z, \kappa) = \delta_{UI} V_N(a, z, \kappa) + (1 - \delta_{UI}) V_U(a, z, \kappa). \quad (17)$$

The presence of \tilde{V}_U 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

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

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

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

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

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}_N^s and \mathcal{V}_N^{ns} (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 In Section 6.4 we will discuss the recent literature that supports this target. We set \bar{R} to imply a steady-state 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 job searchers receive offers, consistent with the average transition rate among nonparticipants who want a job relative to the average UE 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 (2023).²⁹

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 job-finding rate f_s in the model is distinct from the UE 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 effects in the model of an unanticipated monetary policy shock, which leads to changes in the real interest rate, job-finding rate, layoff rate and wage that match those shown in Figures 1–2 (i.e., the responses to a 25bp monetary tightening). 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: $J(\theta_{EST}) = \{EU_t, EN_t, UE_t, UN_t, NE_t, NU_t\}_{t=0}^{50}$.³⁰

Denoting by \hat{J} the impulse response functions estimated in the data, our estimator is

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

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 4.³¹

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

²⁸We report NE|Want in Table B.2 of the Appendix.

²⁹In 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).

³⁰We 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 C.4 describes this approach and provides further computational details.

³¹Our estimates for the labor productivity process imply a less persistent and more volatile process for labor income than seen in the data. In light of our reduced-form interpretation of this process as capturing all shocks that affect a worker's overall willingness to work, our estimates suggest a separate and potentially important 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	Steady-State Wage	1	Normalization
α	Efficiency of Passive Search	0.6	Job-finding rate from N
ϕ	UI Replacement Rate	0.50	Graves (2023)
$\bar{\phi}$	Maximum UI Payments	1.85	Graves (2023)
τ	Labor Income Tax Rate	0.33	Auclert et al. (2021)
T	Lump-sum Transfer	0.24	Auclert et al. (2021)
Estimated			
Parameter	Description	Value	Standard Error
ρ_z	Persistence of Labor Productivity	0.960	(0.004)
σ_z	Standard Deviation of Labor Productivity	0.362	(0.023)
μ_κ	Mean Value of Search Cost	0.783	(0.105)
σ_κ	Dispersion of Search Cost	0.167	(0.022)
ψ	Value of Leisure	0.421	(0.107)
δ	Steady-State Layoff Rate	0.019	(0.002)
f_s	Steady-State Job-Finding Rate	0.273	(0.028)


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

Our treatment of the response of layoff and job-finding rates as inputs determined outside of the model is consistent with our focus on labor supply and affords our analysis greater clarity.³² Moreover, while there has been important research incorporating job creation and job destruction into HANK search and matching (HANK-SAM) models, it is unclear what precise features would be required in order to match the hump-shaped and highly persistent responses of layoffs and job-finding rates implied by the response of EU and UE rates shown in Figure 2, particularly for a heterogeneous agent model such as we consider here. Thus, our approach highlights the extent to which the labor supply forces present in our model allow us to obtain a close fit to the data, without concern that the model's failure (or ability) to fit the data may reflect a simple mis-specification of the model's labor demand block.

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 and verify that the model matches the steady-state labor market flows from the data. We then calculate the MPC and MPE from an idiosyncratic transfer in the model and compare it to empirical estimates in the literature.

³²In 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.

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.

6.4.1. *Labor Supply Policy Functions and Steady-State Labor Market Flows.* Here, we document the substantial heterogeneity in steady-state optimal labor supply policies, consistent with the notion that different workers maintain different levels of attachment to the labor force as a function of their underlying characteristics. Then, we show that, after aggregating labor supply decisions within labor market states, the implied labor market flows in the stationary distribution of the model match the data.

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 non-employment, 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-most panel of the figure, individuals are more likely to quit from employment to nonparticipation (thereby making an EN transition) the lower their productivity and the higher their wealth. Comparing the left and right panels confirms that states where employed workers quit their jobs closely correspond to those where workers will not accept a job offer if non-employed. The middle panel of Figure 8 then shows that individuals are more likely to search when they have low wealth or high productivity. Note that a UI-eligible worker is a worker who was either employed and laid off, or unemployed and unsuccessful at finding a job. Thus, the middle panel illustrates which workers will move to unemployment (versus nonparticipation) after a layoff, and which workers will continue

TABLE 5. Average Labor Market Transition Rates

Transition Rate	Model	Data
EU	0.0143	0.0143
EN	0.0297	0.0296
UE	0.2547	0.2547
UN	0.2260	0.2262
NE	0.0462	0.0461
NU	0.0253	0.0252

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

searching from unemployment (versus stopping search and moving to nonparticipation) after a period of unsuccessful job search.

The middle and right-most panels of Figure 8 make evident the existence of many intermediate combinations of wealth and productivity where individuals do not search but will accept a job if they receive an offer. When such individuals find and accept a job, it is recorded as an NE transition in the model. The model’s ability to match the lower NE rate in the data (compared to the UE rate) obtains both through the lower job-finding probability of non-searchers and the distribution of non-searching workers who will not accept a job when one is offered.


Figure 8 makes it apparent that the ability of our model to match aggregate labor market flows is mediated not only through the distribution of workers across labor market states, but also through the distribution of workers over labor productivity and wealth. Table 5 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.

6.4.2. The Marginal Propensities to Consume and Earn. We now describe how our model matches estimates of the marginal propensity to consume (MPC) and the marginal propensity to earn (MPE), quantities describing how agents adjust their consumption and labor earnings in response to an unanticipated transfer. Our model generates a quarterly MPC of 0.073 and an annual MPE of 0.027.³³

While lower than estimates from the earlier literature (e.g., Parker et al., 2013), the MPC implied by our quantitative model 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

³³We calculate these from an unexpected transfer equivalent to approximately \$500 to agents in the model.

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

nondurable consumption—should lie in a range between about 0.07 and 0.114.³⁴ Our MPC falls in this range.

Auclert, Bardóczy and Rognlie (2021) define the MPE as “the negative of the level of the response of earned income to a one-time, unexpected unit payment, in the period of the payment” and document a tight connection between MPCs and MPEs in models with frictionless labor markets and an intensive margin of labor supply. Drawing from an empirical literature estimating labor supply responses to lottery winnings, Auclert et al. (2021) conclude that the annual MPE lies between 0 and 0.04. The annual MPE from our model falls comfortably within this range.³⁵


Having established that the model (i) generates realistic variation in labor force attachment within labor market states, and (ii) implies MPCs and MPEs consistent with those estimated in the existing literature, we proceed to study the model’s ability to fit the response of labor market flows to a contractionary monetary policy shock.

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 reproduces Figure 2 with the

³⁴Laibson, 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.

³⁵The presence of labor market frictions allows our model to break the tight link between MPCs and MPEs described by Auclert et al. (2021).

FIGURE 10. Decomposition of EU and EN Responses: Model and 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.

model-generated impulse response functions overlaying the empirical estimates as dashed magenta lines. The labor market flows implied by the estimated model are remarkably close to those of the data, with the models' impulse responses lying within the 68% confidence bands for almost all labor market flows and horizons.

Note that the close fit of the model to the data is not guaranteed, but is instead 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. The contribution of labor supply decisions to the fit of the impulse responses is most evident in the responses of UN and NU flows (the right column of the figure), showing a decrease in UN transitions and increase in NU transitions in the model and the data. Figure D.2 shows that the response of NU and UN flows can be understood as largely driven by the response of job-finding and layoff rates (rather than by the response of the real interest rate or wages).

To understand how the model generates an increase in NU flows (and a decrease in UN flows) after a monetary contraction, consider a non-employed worker in the model who finances their consumption through previously accumulated savings. The reduction in job-finding probabilities from unemployment and nonparticipation (i.e., NE and UE) increases

the expected duration of non-employment, forcing such non-employed workers to stretch their savings over a longer period, thereby reducing expected consumption. By engaging in active job search, the worker can shorten their expected spell of joblessness, mitigating the necessary reduction in consumption over their non-employment spell and ameliorating other forms of consumption risk.

Although the model offers implications for the response of quits to nonparticipation, these are not easily seen in Figure 9, which only shows the response of aggregate EN flows. Figure 10 compares the 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 of the model here is also good, with the model doing particularly well in matching the the prolonged decline in EN quits after a monetary contraction. As in the data, the model implies that the share of laid-off workers who move to unemployment (rather than nonparticipation) increases. Indeed, both in the model and in the data the layoff rate to nonparticipation declines on impact. Figures D.2 and D.3 of Appendix D.3 shows that the response of EN flows is largely driven by the the response of job-finding and layoff rates to a contractionary monetary policy shock.


The fall in EN quits in the model can be understood as individuals prolonging their employment spells to maintain consumption. Given the fall in job-finding probabilities from non-employment (i.e., NE and UE), the expected duration of non-employment increases: thus, a worker who may have otherwise been able to temporarily quit to non-employment and maintain their consumption through accumulated savings may no longer able to do so, given the necessity of stretching the same amount of savings over a longer anticipated spell of nonemployment.

6.6. The Role of Labor Supply. The previous section shows that the model offers an excellent fit to our estimates of the response of labor market flows to a contractionary monetary policy shock. We also describe how the response of supply-driven labor market flows can be understood to reflect an increase in labor supply among workers possessing limited assets by which to self-insure and smooth consumption. Here, we establish that the model's ability to match the data occurs through such a labor supply response, rather than a shift in the composition of workers across labor market states.

More specifically, it is in theory possible that the response of these flows could be occurring due to changes in the distribution of employed and non-employed individuals, interacting with the heterogeneity in labor supply policy functions shown in Figure 8. For example, an increase in layoffs could shift the the composition of workers in non-employment towards workers who are more attached to the labor force.

Here, we show that the ability of our model to match the response of labor market flows reflects a broad-based increase in labor supply (as opposed to a shift in the composition of

FIGURE 11. The Employment Response with Fixed Labor Supply



Note: The black solid line shows the overall response of the employment-population ratio to a contractionary monetary policy shock in the estimated model. The blue dotted line shows the response if labor supply policy functions are held at their steady-state values.

non-employed workers towards those with greater labor market attachment). We solve for a counterfactual path of aggregate employment where labor supply policy functions (as a function of worker characteristics) do not respond to the shock, while consumption/savings policy functions evolve as in the baseline model. To the extent that the ability of the model to match the data is generated through composition (rather than a shift in labor supply behavior), the counterfactual fixed-labor-supply path of employment should closely resemble the response of employment implied in the baseline model.

Figure 11 shows the response of employment under the full model and under the restricted fixed-labor-supply counterfactual.³⁶ The two paths diverge considerably, with employment declining by about 70% more under the fixed-labor-supply counterfactual. The steeper decline in the counterfactual employment path indicates that the baseline model fits the data through a broad-based increase in labor supply, via reductions in quits to non-employment and increases in search activity.³⁷

³⁶Figure D.1 in the Appendix confirms that the path of employment in the baseline model is very close to that estimated in the data.

³⁷Note 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.

7. CONCLUSION

This paper offers new empirical evidence of a sizeable 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 non-employment and stimulates job-seeking behavior among the non-employed. 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 EN 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. These results contrast with much of the New Keynesian literature, where labor supply plays almost no role in the response of employment to monetary policy. Thus, our paper highlights a potentially important shortcoming of such models.

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 points to an important role for labor supply in moderating the response of employment to a monetary policy shock. Additionally, we show that our model is consistent with existing empirical evidence on the responses of consumption and labor earnings to idiosyncratic transfers. 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.

Indeed, incorporating an active labor supply margin in New Keynesian models may prove to be 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”).³⁸ Existing models are well-placed to consider the effects of such stimulus on consumption, but less suited to considering how such policies may affect labor supply, or how the labor supply response to such policies might have contributed to the rise in inflation.

We thus view endogenizing the layoff and job-finding rates in a fully-fledged New Keynesian model with an active labor supply margin as a promising topic for future research.

³⁸Adopting our methodology for decomposing flows from employment to nonparticipation, Michaels (2024) documents that quits to non-employment account for a disproportionate share of the increase in quits during this period.

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.
- Alves, Felipe and Giovanni L. Violante**, “Some Like it Hot: Inclusive Monetary Policy Under Okun’s Hypothesis,” 2023.
- Angeletos, George-Marios, Chen Lian, and Christian K Wolf**, “Can Deficits Finance Themselves?,” Working Paper 31185, National Bureau of Economic Research April 2023.
- 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.
- 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.
- Bilbiie, Florin O, Giorgio Primiceri, and Andrea Tambalotti**, “Inequality and Business Cycles,” Working Paper 31729, National Bureau of Economic Research September 2023.
- 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.
- , **Niels-Jakob Harbo Hansen, Per Krusell, and Erik Öberg**, “The New Keynesian transmission mechanism: A heterogeneous-agent perspective,” *The Review of Economic Studies*, 2020, 87 (1), 77–101.
- Cairó, Isabel, Shigeru Fujita, and Camilo Morales-Jiménez**, “The cyclicity of labor force participation flows: The role of labor supply elasticities and wage rigidity,” *Review of Economic Dynamics*, 2022, 43, 197–216.
- Caldara, Dario and Edward Herbst**, “Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs,” *American Economic Journal: Macroeconomics*, January 2019, 11 (1), 157–192.
- 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.
- , **Martin S. Eichenbaum**, and **Mathias Trabandt**, “Unemployment and Business Cycles,” *Econometrica*, 2016, 84 (4), 1523–1569.
- , **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.
- Elsby, Michael W. L., Bart Hobijn, and Aysegül Sahin**, “The Labor Market in the Great Recession,” *Brookings Papers on Economic Activity*, 2010, 41 (1 (Spring)), 1–69.
- , **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 Aysegül 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.
- 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.
- 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.
- , **Luca Sala, and Antonella Trigari**, “An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining,” *Journal of Money, Credit and Banking*, December 2008, 40 (8), 1713–1764.
- Gilchrist, Simon and Egon Zakrajšek**, “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, June 2012, 102 (4), 1692–1720.
- Graves, Sebastian**, “Does Unemployment Risk Affect Business Cycle Dynamics?,” 2023.
- Hobijn, Bart and Aysegül 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 Aysegül 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.
- Lundberg, Shelly**, “The Added Worker Effect,” *Journal of Labor Economics*, January 1985, *3* (1), 11–37.
- Michaels, Ryan**, “What Explains the Great Resignation?,” *Economic Inquiry*, Q2 2024, *9* (2), 10–18.
- 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 McClelland**, “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, October 2013, *103* (6), 2530–53.
- 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.
- Ravn, Morten O. and Vincent Sterk**, “Job uncertainty and deep recessions,” *Journal of Monetary Economics*, 2017, *90*, 125–141.
- 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.
- Smets, Frank and Rafael Wouters**, “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, June 2007, *97* (3), 586–606.
- 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,” 2023. Working Paper, University of California, Irvine.
- White, Neil**, “Gross Worker Flows, Job Loss, and Monetary Policy,” 2018.

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

Woodford, Michael, “Simple Analytics of the Government Expenditure Multiplier,” *American Economic Journal: Macroeconomics*, January 2011, 3 (1), 1–35.