

The Ins and Outs of Unemployment Shocks*

Carter Bryson[†]

April 12, 2024

Abstract

This paper quantifies the contribution of unemployment inflows and outflows to cyclical changes in the unemployment rate. I show that the time series behavior of worker flows implies that they exhibit a specific dynamic structure. I then implement a simple identification strategy motivated by this evidence in order to empirically separate changes in job separation and job finding. I find that both margins contribute significantly to unemployment volatility and that their interaction is important for understanding the business cycle dynamics of the unemployment rate. My results suggest that models of the labor market should aim to capture this interaction as well as cyclical variation in job loss.

Keywords: Worker Flows, Unemployment, Business Cycles, Sign Restrictions

JEL Classification: C32, E24, E32, J64

*I thank Theo Papageorgiou, Robert Ulbricht, and Hanno Foerster for their guidance and support. I thank Gary Cornwall and seminar participants at Salisbury University and the 2024 MEA Annual Meetings for helpful discussions and suggestions. The views expressed in this paper are solely those of the author and not necessarily those of the U.S. Bureau of Economic Analysis or the U.S. Department of Commerce.

[†]Office of the Chief Economist, U.S. Bureau of Economic Analysis. Email: wcarterbryson@gmail.com.

1 Introduction

What accounts for the rise in the unemployment rate during a recession? Economists have debated this seemingly straightforward question going back at least five decades.¹ Reaching a definitive conclusion has proved challenging as it requires quantifying the contributions of two different, but interlinked channels: in recessions, employed workers lose jobs, but unemployed workers also have trouble finding jobs. In this paper, I provide evidence that both inflows to and outflows from unemployment play an important role in the business cycle dynamics of the labor market.

I begin by examining the dynamic properties of unemployment inflows and outflows and present three pieces of motivating evidence. First, while the *rate* at which workers find jobs out of unemployment is procyclical, the *number* of workers that find jobs out of unemployment is countercyclical. Next, aggregate unemployment outflows (job finding) are, on average, predated by aggregate unemployment inflows (job separation). I find that the job finding rate displays the strongest cross-correlation with the lagged rather than the contemporaneous job separation rate; the same holds true for the relationship between the unemployment rate and the job separation rate.² Last, I test formally for Granger causality and find that unemployment inflows Granger-cause unemployment outflows. Hence, unemployment and worker flows display a specific dynamic structure. To my knowledge, the last of these three facts is new to the literature.

Motivated by these findings, I design a simple identification strategy in the context of a structural vector autoregression (SVAR) framework to isolate disturbances in labor market variables. The goal is to estimate the response of job finding, job separation, and unemployment to exogenous changes in these variables and to use the recovered shocks to decompose fluctuations in the unemployment rate. I use several techniques familiar to the SVAR literature to formally examine the role played by unemployment inflows and outflows in driving changes in unemployment.

A benefit of this approach is that it does not require that certain worker flow margins be held constant when examining the contribution of others. For instance, a spike in job separations may feed back on job creation incentives, in turn affecting the job finding rate. In this case, my methodology accounts for the endogenous reaction of the job finding rate

¹See early contributions by [Darby et al. \(1986\)](#) and [Blanchard and Diamond \(1990\)](#) as well as more recent studies by [Hall \(2005\)](#) and [Shimer \(2012\)](#).

²I use the term “job separation” to refer to movements of workers from employment to unemployment (*EU* flows). Of course, true job separations include instances where a worker moves from employment to not-in-the-labor-force (*EN* flows) or from one employer to another (employer-to-employer flows). However, labeling *EU* flows as “job separation” and *UE* flows as “job finding” provides a convenient shorthand. See [Section 2](#) for more details on terminology.

to an initial change in unemployment inflows. Although this procedure does not identify the underlying fundamental disturbances that drive business cycles, the properties of the unemployment inflow and outflow “shocks” that I recover shed light on the forces that contribute to the cyclical variance of unemployment.

The assumptions behind my identification strategy rest on the time series properties of worker flows I document as well as the relationship between worker flows and the distribution of unemployment duration. The evidence on the lead-lag structure of worker flows implies that job separations do not react immediately to changes in job finding or unemployment. I therefore restrict the impact response of the unemployment rate and the job finding rate to be zero following a shock to job separation. I supplement these zero restrictions with sign restrictions on the response of the unemployment rate and average unemployment duration to innovations in job finding and job separation. I assume that while shocks to job finding move the unemployment rate and average unemployment duration in the same direction, shocks to job separation move the unemployment rate and average unemployment duration in opposite directions.

Because my identification strategy combines both zero and sign restrictions, I use Bayesian techniques to estimate the model following recent advances in the literature ([Rubio-Ramírez et al., 2010](#); [Arias et al., 2018](#)). My baseline specification includes four variables: the job finding probability, the job separation probability, the unemployment rate, and average unemployment duration for the period 1948Q1–2019Q4. I recover four shocks in total, which I label as shocks to the Unemployment Inflow, Unemployment Outflow, Unemployment Level, and Unemployment Length.

I find that Unemployment Inflow shocks lead to large movements in all four series. In response to a (contractionary) Unemployment Inflow shock, the job separation probability increases significantly and remains elevated for almost two years. The job finding probability reacts little on impact, but then falls for several periods before slowly returning to trend. The unemployment rate rises on impact and displays a hump-shaped pattern: the initial increase in job separation increases unemployment and the subsequent decline in job finding causes it to rise further. Lastly, average unemployment duration falls on impact, but then follows an increasing, hump-shaped pattern as workers accumulate in the unemployment pool. In contrast, while (contractionary) Unemployment Outflow shocks lead to increases in the unemployment rate and average unemployment duration, they are of smaller magnitude and are not associated with increases in the job separation probability.

I also use local projections to assess the impact of unemployment inflow and outflow shocks on auxiliary labor market variables ([Jordà, 2005](#)). I find that while unemployment

inflow shocks are consistent with the business cycle properties of gross worker flows and vacancy posting, unemployment outflow shocks are not. In particular, a contractionary unemployment inflow shock leads to a significant and persistent increase in the total flow of workers between unemployment-and-employment and between employment-and-unemployment. Both of these series are countercyclical in the data, which points to changes in unemployment inflows as plausible drivers of gross worker flows. Additionally, unemployment inflow shocks lead to large changes in vacancy posting. I find that the level of vacancies falls by about 50% following a contractionary inflow shock. This pattern is consistent with Beveridge Curve dynamics, as unemployment and vacancies move in opposite directions. On the other hand, unemployment outflow shocks do not lead to significant movements in gross worker flows or vacancies.

Having documented the properties of unemployment inflow and outflow shocks, I turn to quantifying their contributions to fluctuations in unemployment. I find that on average, unemployment inflow shocks explain about 50% of the overall variance of the unemployment rate, while unemployment outflow shocks account for no more than 20%. However, I note several caveats to this result. First, their contributions are not constant across horizons; unemployment outflow shocks display larger contributions at shorter horizons. Second, the *interaction* between job finding and job separation in response to unemployment inflow shocks greatly affects this calculation. I find that shutting down their interaction such that the job separation probability does not endogenously affect the job finding probability (and vice versa) significantly decreases (increases) the contribution of unemployment inflow (outflow) shocks to unemployment variance. Lastly, the mix of unemployment inflow and outflow shocks has been different in different recessions. While unemployment inflow shocks primary drove the persistent rise in unemployment during the 2008 recession, unemployment outflow shocks played a greater role in labor market dynamics during the 2001 recession. This suggests that to understand the cyclical behavior of the unemployment rate, both margins should be considered in conjunction.

I conclude by discussing the implications of my findings for the development of macroeconomic models of the labor market. Since the influential work of [Shimer \(2005, 2012\)](#), much of the literature's focus shifted toward modeling forces that generate cyclical movements in the job finding rate, while treating the job separation rate as exogenous and acyclical. Though I am not the first to point this out, the dynamic relationship between unemployment inflows and outflows cautions against this approach.³ In addition to not being able to explain the increase in gross unemployment outflows during a recession,

³In particular, see discussions in [Elsby et al. \(2009\)](#), [Fujita and Ramey \(2009\)](#), [Fujita and Ramey \(2012\)](#), and [Mercan et al. \(2024\)](#), among others.

treating the job separation probability as fixed does not allow for a connection between changes in inflows and outflows by assumption. My results suggest that their interaction is a key feature of unemployment fluctuations that models of the labor market should aim to explicitly capture.

Related literature My paper contributes to both empirical and theoretical studies on the dynamics of labor market flows. In a seminar paper, [Darby et al. \(1986\)](#) find that the flow into unemployment is the main determinant of the unemployment rate and declare that “the ins win.” [Blanchard and Diamond \(1990\)](#) show that the amplitude of changes in the total inflow to unemployment is larger than the amplitude of changes in the total outflow, and conclude that reduced employment in a recession results from higher job destruction. However, both [Hall \(2005\)](#) and [Shimer \(2012\)](#) challenge the results of these earlier studies and argue that contrary to the conventional wisdom, unemployment instead rises in a recession because it becomes more difficult to find a job. In contrast to the two earlier studies that considered the *total number* of workers exiting or entering unemployment, [Hall \(2005\)](#) and [Shimer \(2012\)](#) emphasize the importance of examining the *probabilities* of unemployment inflows and outflows. In my study, I attempt to reconcile the evidence in both sets of papers by studying both gross worker flows and flow probabilities.

My results are in line with those of several papers that initially responded to the Hall-Shimer critique and warned that studying movements in the job finding rate alone paint an incomplete picture of labor market fluctuations. [Elsby et al. \(2009\)](#) decompose changes in the unemployment rate and argue that it is important to understand both procyclical movements in the job finding rate and countercyclical movements in the job separation rate. [Fujita and Ramey \(2009\)](#) re-examine gross flows data and find that allowing for a dynamic interaction between job finding and job separation leads to a larger contribution of the separation rate in unemployment dynamics. [Fujita \(2011\)](#) uses sign restrictions on the joint behavior of unemployment and vacancies and finds that the dynamics of job loss help explain the cyclical behavior of labor market variables. [Canova et al. \(2013\)](#) study the role of both investment-neutral and investment-specific technology shocks in labor market dynamics and find that the impact response of unemployment is primarily driven by the job separation rate, while the job finding rate accounts for most of its persistence. Their evidence is consistent with the interpretation that unemployment initially increases in a recession due to a wave of layoffs and remains high due to a slow recovery in the job finding rate. I take a different approach to identifying labor market disturbances, but the conclusions of my analysis are consistent with these studies.

More recently, several studies have outlined theoretical grounds for a link between

increased job separation rates and decreased job finding rates during a recession. For instance, [Ahn and Hamilton \(2020\)](#) and [Gregory et al. \(2022\)](#) argue that if workers with low individual job finding propensities disproportionately enter the unemployment pool in recessions, then the aggregate job finding rate will fall due to a composition effect. [Coles and Moghaddasi-Kelishomi \(2018\)](#), relax the free entry of vacancies assumption in the Diamond-Mortensen-Pissarides (DMP) framework and find a substantial role for job separation shocks in unemployment fluctuations. They show that if vacancy creation is less than infinitely elastic, as is assumed to be the case in the standard DMP model, then both labor market tightness and the job finding rate decline in response to an increase in job separations. [Engbom \(2021\)](#) hypothesizes that if unemployed workers apply to jobs less selectively than employed workers on average, then it becomes harder for firms to assess which workers are qualified in times of high unemployment, putting downward pressure on hiring incentives. [Mercan et al. \(2024\)](#) show that if new hires and incumbent workers are imperfect substitutes, then increases in the separation rate reduce job creation incentives because “congestion” in hiring decisions renders firms unable to fully absorb the resulting increase in unemployment. In my paper, I estimate a decline in the job finding probability after a bout of increased separations, providing empirical evidence for these mechanisms.

Lastly, my paper contributes to the literature on SVAR identification using a combination of zero and sign restrictions. In a seminal paper, [Rubio-Ramírez et al. \(2010\)](#) provide general conditions for identification in SVAR models and develop efficient algorithms for estimation and inference. [Arias et al. \(2018\)](#) build on the methodology in [Rubio-Ramírez et al. \(2010\)](#) and extend their results to allow for both sign and zero restrictions. [Arias et al. \(2019\)](#) apply this methodology in order to identify the effects of monetary policy shocks on output and prices. Though I make no theoretical contributions of my own, my study is the first to my knowledge to apply this methodology in the context of the literature on the ins and outs of unemployment.

Layout The rest of the paper is structured as follows. Section 2 describes how I measure worker flows and Section 3 presents motivating evidence for the identifying restrictions used in my empirical analysis. In Section 4, I present specific details on the empirical strategy and discuss the identification assumptions. In Section 5, I present estimates of the effect of unemployment shocks on different labor market variables and use the results to formally decompose overall fluctuations in unemployment. Section 6 discusses the implications of my findings for models of the labor market. Section 7 concludes.

2 Measuring Worker Flows

I measure worker transitions across labor market states using the Current Population Survey (CPS) as well as publicly available data from the Bureau of Labor Statistics (BLS). I employ two separate methodologies to measure worker flows, which both follow [Shimer \(2012\)](#). This section describes how I construct these measures and also serves to define key terms that I will use throughout the rest of the paper.

2.1 Two-State Model

First, I measure the inflow rate to unemployment and outflow rate from unemployment abstracting from movements into and out of the labor force. I follow [Shimer \(2012\)](#) and define the job finding probability F_t for some interval $[t, t + 1)$ as:

$$F_t = 1 - \frac{U_{t+1} - U_{t+1}^s}{U_t} \quad (1)$$

where U_t is total unemployment and U_t^s is total short-term unemployment in month t . Both series are published by the BLS starting in January 1948.⁴ Let $f_t \equiv -\ln(1 - F_t)$ be the associated job finding *rate*, which is the arrival rate of a Poisson process. The equation below implicitly defines the employment exit *rate* x_t

$$U_{t+1} = \frac{(1 - e^{-f_t - x_t})x_t}{f_t + x_t} L_t + e^{-f_t - x_t} U_t \quad (2)$$

where I assume that the size of the labor force L_t is constant. I then compute the associated employment exit probability $X_t \equiv 1 - e^{-x_t}$.

Though abstracting from changes in labor force participation certainly affects the level and cyclical nature of worker flows ([Elsby et al., 2015](#)), this approach allows me to use publicly available time series with a long time span. I use the series from the “two-state model” in my empirical analysis, as the larger sample size facilitates estimation.

2.2 Three-State Model

Second, I construct additional measures of transition rates between labor market states using CPS microdata. The CPS is published on a monthly basis and due to the rotating panel structure of the survey, roughly three quarters of individuals in a given month can

⁴Specifically, I use the series LNS13000000: Unemployment Level and LNS13008396: Number Unemployed for Less than 5 Weeks.

be linked to their survey responses in the previous month.⁵ This facilitates the constructions of measures of the *total number* of workers who transition between labor market states – which I will refer to as “gross flows” – as well as the *probability* that a worker transitions between labor market states – which I will refer to as “flow probabilities.”

I gather data from IPUMS CPS for the period January 1976 to December 2019 and link survey respondents across consecutive months using the unique identifier `CPSIDV` (Flood et al., 2022).⁶ This variable includes linking criteria that ensures individuals match on age, sex, and race characteristics. After linking individuals, information on their employment status in each month allows me to construct measures of both gross flows and flow probabilities. Specifically, I use the variable `EMPSTAT` to determine whether a given individual was employed (E), unemployed (U), or not-in-the-labor-force (N) in a particular month. I then compute weighted sums of the number of individuals who transition across labor market states using longitudinal weights provided by IPUMS CPS.

Defining Flows Let A and B denote two labor market states. Let $\mathbb{1}\{A_{i,t-1} \& B_{i,t}\}$ be an indicator for whether worker i was in state A at time $t - 1$ and in state B at time t . The aggregate gross flow of workers from labor market state A in month $t - 1$ to labor market state B in month t is defined as follows:

$$AB_t \equiv \sum_i \mathbb{1}\{A_{i,t-1} \& B_{i,t}\} \cdot w_{i,t-1}$$

where $w_{i,t-1}$ are weights corresponding to the probability of worker i being included in the sample. For instance, the aggregate gross flow UE_t is simply the (weighted) number of workers who were unemployed in time $t - 1$ and employed in time t .

The corresponding flow probability for workers between states A and B is given by:

$$P_t(AB) \equiv \frac{\sum_i \mathbb{1}\{A_{i,t-1} \& B_{i,t}\} \cdot w_{i,t-1}}{\sum_i \mathbb{1}\{A_{i,t-1}\} \cdot w_{i,t-1}}$$

In this equation, the expression $\sum_i \mathbb{1}\{A_{i,t-1}\} \cdot w_{i,t-1}$ is the weighted sum of workers in state A at time $t - 1$. For instance, the aggregate flow probability $P_t(UE)$ is simply the (weighted) number of workers who were unemployed in time $t - 1$ and employed in time t , divided by the (weighted) number of workers who were unemployed in time $t - 1$.

⁵In particular, households are initially interviewed for four consecutive months, then excluded from the sample for the following eight months, and then included in the CPS again for the subsequent four months.

⁶See the following link for more information on linking individuals across surveys in the IPUMS CPS data: https://cps.ipums.org/cps/cps_linking_documentation.shtml.

Lastly, I seasonally adjust all flow measures using the Census Bureau’s X-13-ARIMA-SEATS procedure. I also adjust flow probabilities for time aggregation bias following the methodology in [Shimer \(2012\)](#). I use the series produced by the “three state model” in the motivating evidence section of the paper.

Appendix [A](#) contains additional measurement details, including a small correction to the short-term unemployment series and the time aggregation adjustment procedure. Appendix Figures [A.1](#) and [A.2](#) show that I successfully replicate and extend the series from [Shimer \(2012\)](#)’s original analysis. Moreover, I show in Appendix Figure [A.3](#) that flow probabilities obtained using the two-state model and the three-state model have similar cyclical properties. Therefore, after documenting the dynamic properties of worker flows using series produced by the three-state model, I use the larger sample period of the flow series produced by the two-state model for my empirical analysis.

3 Motivating Evidence

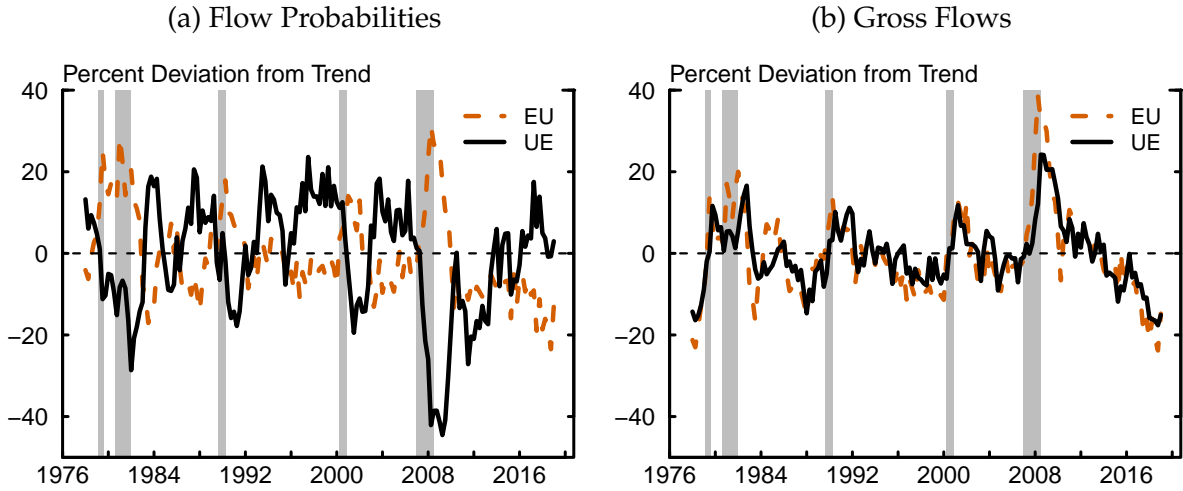
In this section, I present three stylized facts about labor market flows that motivate the assumptions behind my empirical analysis. First, I show that while the probability of transitioning from U -to- E is procyclical, the *number* of workers transitioning from U -to- E is countercyclical. Next, I find that both the unemployment rate and the U -to- E flow probability have the strongest correlation with *lagged* values of the E -to- U flow probability. Last, I formally test this relationship and find that unemployment inflows Granger cause unemployment outflows.

3.1 Countercyclicity of Gross Flows

Figure [1](#) displays the cyclical properties of of worker flows. I use the recently devised filtering method in [Hamilton \(2018\)](#) to extract the cyclical component of each series. Data are plotted in terms of percent deviations from the respective trend. Panel (a) shows flow probabilities while Panel (b) shows the corresponding gross flow series.

Panel (b) clearly indicates that both UE_t and EU_t gross flows are countercyclical, meaning that they rise around NBER recession dates and fall during business cycle expansions. Therefore, the *total number* of workers exiting unemployment, somewhat counter-intuitively, rises in recessions. An older literature that examined the cyclicity of worker flows also recognized this pattern ([Darby et al., 1986](#); [Davis, 1987](#); [Blanchard and Diamond, 1990](#); [Ritter, 1993](#)). Harkening back to the findings of this earlier literature, Figure [1](#) shows that the amplitude of fluctuations in EU_t is greater than that of UE_t .

Figure 1: Cyclical Properties of Worker Flows



Notes: Black, solid lines show U -to- E flows. Orange, dashed lines show E -to- U flows. Series Hamilton filtered with horizon h and lag p parameters set to 8 and 4 quarters, respectively. I take a quarterly average of monthly values before applying the filter. Monthly series adjusted for seasonal variation and time aggregation bias.

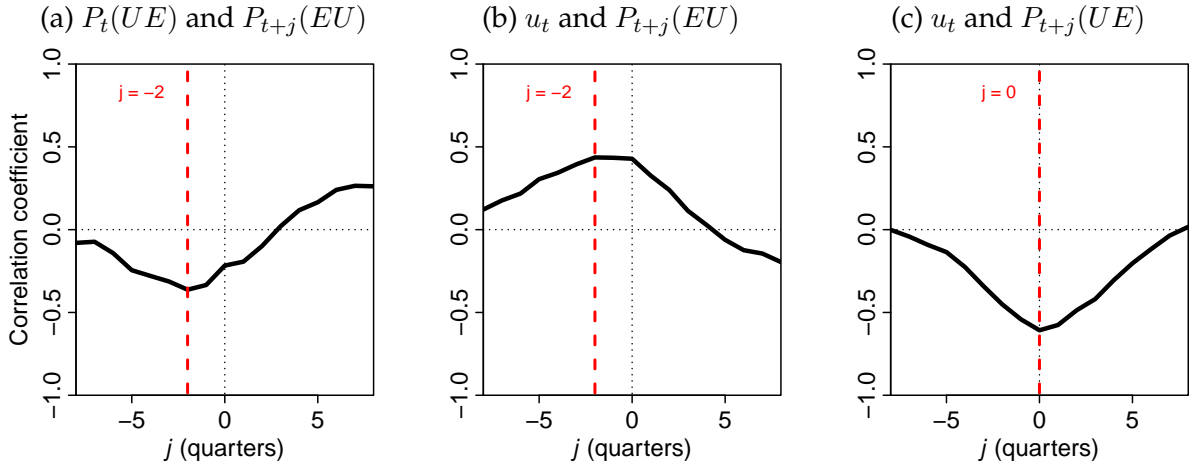
3.2 Cross Correlations of Flow Probabilities

I now investigate the lead-lag structure among flow probabilities and the unemployment rate. To do so, I compute cross-correlations of these variables at different horizons. In particular, I calculate the correlation coefficient between some variable X_t and some other variable Y_{t+j} , where I allow j to vary between -8 and 8 quarters and use a uniform sample size across different values of j . Figure 2 contains the results of this exercise.

Focusing first on the left panel of Figure 2a, the black, solid line shows that the unemployment-to-employment flow probability $P(UE)$ is negatively correlated with the employment-to-unemployment flow probability $P(EU)$, on average. Times when it is easier to find a job are times when a lower fraction of workers exit employment and vice versa. However, the correlation is not constant across different horizons. In particular, the red, dashed line shows that it reaches a low point at $j = -2$ quarters, indicating that $P(UE)$ is most strongly (negatively) correlated with previous values of $P(EU)$. Therefore, $P(EU)$ leads $P(UE)$ in time.

In the center panel, Figure 2b shows that a similar pattern exists in the relationship between the unemployment rate (u) and $P(EU)$. The cross-correlation between these two variable troughs at a value of $j = -2$ quarters. Therefore, movements in the E -to- U flow probability also lead movements in the unemployment rate. Lastly, turning to the right panel, Figure 2c shows that this is not the case with the U -to- E flow probability.

Figure 2: Cross-Correlations of Flow Probabilities and Unemployment



Notes: Series Hamilton-filtered with horizon h and lag p parameters set to 8 and 4 quarters, respectively. I take a quarterly average of monthly values before applying the filter. Monthly series adjusted for seasonal variation and time aggregation bias. Correlation coefficient is the Kendall rank correlation coefficient. Sample = 1980Q4–2017Q4.

Instead, the minimum correlation between u and $P(UE)$ occurs at $j = 0$. Therefore, the *contemporaneous* relationship between the U -to- E flow probability and unemployment is the strongest. Collectively, this evidence suggests that while changes in U -to- E flows occur in simultaneity with changes in the unemployment rate, changes in E -to- U flows predate changes in the unemployment rate.

3.3 Granger Causality in Gross Flows

Next, I further explore the dynamic structure of worker flows by conducting a formal test of whether certain labor market flows Granger-cause others. Following the approach outlined in [Hamilton \(1994\)](#), I conduct an econometric test for whether some time series variable x_t Granger-causes some other time series variable y_t .⁷ For a given lag length L , I estimate the following two equations by OLS.

$$y_t = \alpha_0 + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \dots + \alpha_L y_{t-L} + \beta_1 x_{t-1} + \beta_2 x_{t-2} + \dots + \beta_L x_{t-L} + u_t \quad (3)$$

$$y_t = \gamma_0 + \gamma_1 y_{t-1} + \gamma_2 y_{t-2} + \dots + \gamma_p y_{t-L} + e_t \quad (4)$$

Then, I conduct an F -test of the null hypothesis $H_0 : \beta_1 = \beta_2 = \dots = \beta_L = 0$. This tests whether the variable x_t is important in predicting values of y_t in a forecasting sense. If x_t Granger-causes y_t , then information contained in x_t will be useful in forecasting y_t

⁷I implement the in-sample F -test specified in [Hamilton \(1994\)](#) Chapter 11.2, pages 304–305.

above and beyond the information contained in past values of y_t . To implement the test, I construct the test statistic below

$$S_1 \equiv \frac{(RSS_0 - RSS_1)/L}{RSS_1/(T - 2L - 1)}$$

where $RSS_1 = \sum_{t=1}^T \hat{u}_t^2$ is the sum of squared residuals from the regression in Equation (3), $RSS_0 = \sum_{t=1}^T \hat{e}_t^2$ is the sum of squared residuals from the regression in Equation (4), and T is the number of time periods in the sample. I compare the value of the test statistic to the critical value from the F distribution with parameters L and $T - 2L - 1$.

Table 1: Does EU_t Granger Cause UE_t ?

	Raw Data				Hamilton-Filtered Data			
	$L = 1$	$L = 2$	$L = 3$	$L = 4$	$L = 1$	$L = 2$	$L = 3$	$L = 4$
S_1	72.14	37.24	22.71	16.42	73.86	37.85	23.46	17.49
Crit. Val.	6.80	4.74	3.91	3.44	6.80	4.74	3.91	3.44
p -value	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Reject?	Y	Y	Y	Y	Y	Y	Y	Y
BIC	3892.47	3900.34	3909.83	3918.69	857.01	865.19	873.88	883.61

Notes: Raw Data are quarterly averages of monthly values. Hamilton-Filtered Data are quarterly averages of monthly values, then Hamilton-filtered with quarterly parameters $h = 8$ and $p = 4$. Significance level $\alpha = 1\%$. Sample = 1979Q4–2019Q4.

Table 1 displays the results of this exercise using the gross flows data, where EU_t is the x variable (independent variable) and UE_t is the y variable (dependent variable) in the unrestricted regression, Equation (3). I choose a significance level of $\alpha = 1\%$ and report the test statistic, critical value, and p -value of the test, as well as whether the test rejects the null hypothesis. Note that under the null hypothesis, the independent variable *does not* Granger-cause the dependent variable. If the test rejects the null hypothesis, then EU_t is said to Granger-cause UE_t . I report the results of the test using several different lag lengths as well as using raw or filtered data. I also report the Bayesian information criteria (BIC) value for each lag length, which favors a value of $L = 1$. I find overwhelming support for the result that EU_t Granger-causes UE_t , with p -values indistinguishable from 0 in most instances.

On the other hand, UE_t gross flows do not appear to Granger cause EU_t gross flows. Table 2 reports the results of the test using UE_t as the independent variable and EU_t as the dependent variable. Across lag lengths and filtering methods considered, the null hypothesis of no Granger-causality is not rejected for the given significance level. Therefore, although the two series exhibit a very high contemporaneous correlation, a clear lead-lag

Table 2: Does UE_t Granger Cause EU_t ?

	Raw Data				Hamilton-Filtered Data			
	$L = 1$	$L = 2$	$L = 3$	$L = 4$	$L = 1$	$L = 2$	$L = 3$	$L = 4$
S_1	5.46	2.17	0.74	0.93	4.29	4.52	1.65	2.59
Crit. Val.	6.80	4.74	3.91	3.45	6.80	4.74	3.91	3.45
p -value	0.02	0.12	0.53	0.45	0.04	0.01	0.18	0.04
Reject?	N	N	N	N	N	N	N	N
BIC	3991.96	4001.65	4011.48	4010.38	990.52	995.93	1005.40	1005.18

Notes: Raw Data are quarterly averages of monthly values. Hamilton-Filtered Data are quarterly averages of monthly values, then Hamilton-filtered with quarterly parameters $h = 8$ and $p = 4$. Significance level $\alpha = 1\%$. Sample = 1979Q4–2019Q4.

structure exists between them. Namely, the number of workers exiting employment into unemployment (EU) Granger-causes the number of workers exiting unemployment into employment (UE).

4 Identifying Unemployment Shocks

The evidence presented in Section 3 motivates a particular identification strategy of the disturbances that affect unemployment over the business cycle. Namely, the lead-lag structure suggests that disturbances to unemployment and unemployment outflows affect unemployment inflows only with a lag. I identify innovations in worker flows using a mixture of zero and sign restrictions following the methodology of [Arias et al. \(2018\)](#). The subsections below detail the setup of the empirical model and the identification strategy.

4.1 Model Setup

Consider the empirical VAR model in reduced-form as follows

$$y_t = b_0 + \sum_{l=1}^L B_l y_{t-l} + \mu_t \quad (5)$$

where y_t is a $(k \times 1)$ vector of endogenous variables at time t , b_0 is a $(k \times 1)$ vector of intercept terms, B_l is a $(k \times k)$ matrix of coefficients for each lag length l , L is the total number of lags in the system, and μ_t is a $(k \times 1)$ vector of reduced-form residuals at time t . Assume that $\mu_t \sim \mathcal{N}(0, \Sigma)$, so that the residuals are normally distributed, with mean zero and $(k \times k)$ variance-covariance matrix Σ . Moreover, assume that the vector of reduced-form residuals μ_t is related to the $(k \times 1)$ vector of structural shocks ε_t in the

following manner:

$$\mu_t = A\varepsilon_t \tag{6}$$

Here, it is assumed that $\varepsilon_t \sim \mathcal{N}(0, I)$ so that the structural shocks are normally distributed, with mean zero, unit variance, and zero cross-correlation. A is a $(k \times k)$ matrix referred to as the “impact matrix” because it captures the immediate effect that structural shocks ε_t have on the endogenous variables in y_t .

Identifying the above system requires placing restrictions on the matrix A such that the structural shocks ε_t can be recovered from the reduced-form residuals μ_t . Note that due to the relationship between ε_t and μ_t , the following holds: $\Sigma = \text{Var}(\mu_t) = \text{Var}(A\varepsilon_t) = A\text{Var}(\varepsilon_t)A' = AIA' = AA'$. Therefore, since the matrix Σ is symmetric, it is sufficient to place $k(k - 1)/2$ restrictions on the matrix A in order to identify the structural shocks in Equation (6). I now turn to the method I use to set these restrictions and how they relate to the motivating evidence in Section 3.

4.2 Identification Strategy

I consider a four-variable system with the job separation probability X_t , the job finding probability F_t , the unemployment rate u_t , and average unemployment duration d_t , in that order. The ordering of the variables in the VAR is motivated by the evidence discussed in Section 3. In particular, the job separation probability leads both the job finding probability and the unemployment rate. This motivates the use of zero impact restrictions on innovations to the impact of these variables on X_t .

However, since there are 4 variables in the model, the system requires additional identifying restrictions.⁸ I appeal to the relationship between worker flows and the distribution of unemployment duration to impose these additional restrictions. Consider a decrease in the job finding probability. All else equal, the unemployment rate will rise because of the reduced flow of workers out of the unemployment pool. Additionally, since workers face a lower probability of leaving unemployment, they remain in the unemployment pool for longer. Therefore, innovations to the finding probability cause the unemployment rate and average unemployment duration to move in the same direction.

Consider an increase in the job separation probability. All else equal, the unemployment rate will rise because of the increased flow of workers into the unemployment pool. Additionally, since these workers have an unemployment duration of 0 by construction, the mean of the unemployment duration distribution must fall. Therefore, innovations to

⁸See the discussion in Section 2.2 of [Arias et al. \(2018\)](#).

Table 3: Identifying Restrictions

	Unemployment			
	Inflow	Outflow	Level	Length
Job separation probability	+	0	0	
Job finding probability		-		
Unemployment rate	+	+	+	
Average unemployment duration	-	+		+

Notes: As suggested by [Arias et al. \(2018\)](#), all restrictions are imposed on impact.

the job separation probability cause the unemployment rate and average unemployment duration to move in opposite directions.

Table 3 summarizes these restrictions. In total, I identify four shocks, which I label as shocks to the Unemployment Inflow, Unemployment Outflow, Unemployment Level, and Unemployment Length. Note that I also impose regularity conditions (along the diagonal) such that the shocks are contractionary. For instance, a shock to the unemployment outflow is associated with a fall in the job finding probability.

4.3 Estimation

I use quarterly data on X_t , F_t , u_t , and d_t from 1948Q1 to 2019Q4 and include 4 lags of the endogenous variables in my baseline specification.⁹ As discussed in Section 2, using the two-state model to measure the job finding and job separation probabilities allows for a larger sample, whereas using the three-state model produces worker flow series that begin in 1976Q1. I take quarterly averages of monthly values and filter the data using the [Hamilton \(2018\)](#) filter with the suggested quarterly parameters, $h = 8$ and $p = 4$. This ensures that the data are all expressed in the same units: percentage deviations from their respective trends.

Because my identification strategy uses a combination of zero and sign restrictions, I implement the methodology developed in [Arias et al. \(2018\)](#) to estimate the shocks.¹⁰ The idea is draw the reduced form coefficients B_i , the variance covariance matrix of the residuals Σ , and an orthogonal rotation matrix Q that maps the structural shocks into the reduced form residuals, and keep draws of these matrices if the resulting impulse response functions satisfy the sign and zero restrictions. In particular, conditional on a draw

⁹The BIC suggests using more than 12 lags, which would increase the computational load considerably. Results are robust to using 1, 4, 8, or 12 lags in the system.

¹⁰I make use of existing codes from the Empirical Macro Toolbox (EMT) developed by Filippo Ferroni and Fabio Canova. See [Ferroni and Canova \(2022\)](#) for a guide to the toolbox.

of (B_l, Σ) , the zero restrictions in the identification strategy impose linear restrictions on the columns of Q . Then, conditional on (B_l, Σ, Q) that satisfies the zero restrictions, one checks to see if the sign restrictions are satisfied and discards the draw if they are not. I use *conjugate* priors such that the prior and posterior densities come from the same family of distributions. The draws of (B_l, Σ) are from a normal-inverse-Wishart distribution and the draws of Q are from a uniform distribution, conditional on (B_l, Σ) . I implement the algorithm in Matlab using 10,000 draws.

5 The Effects of Unemployment Shocks

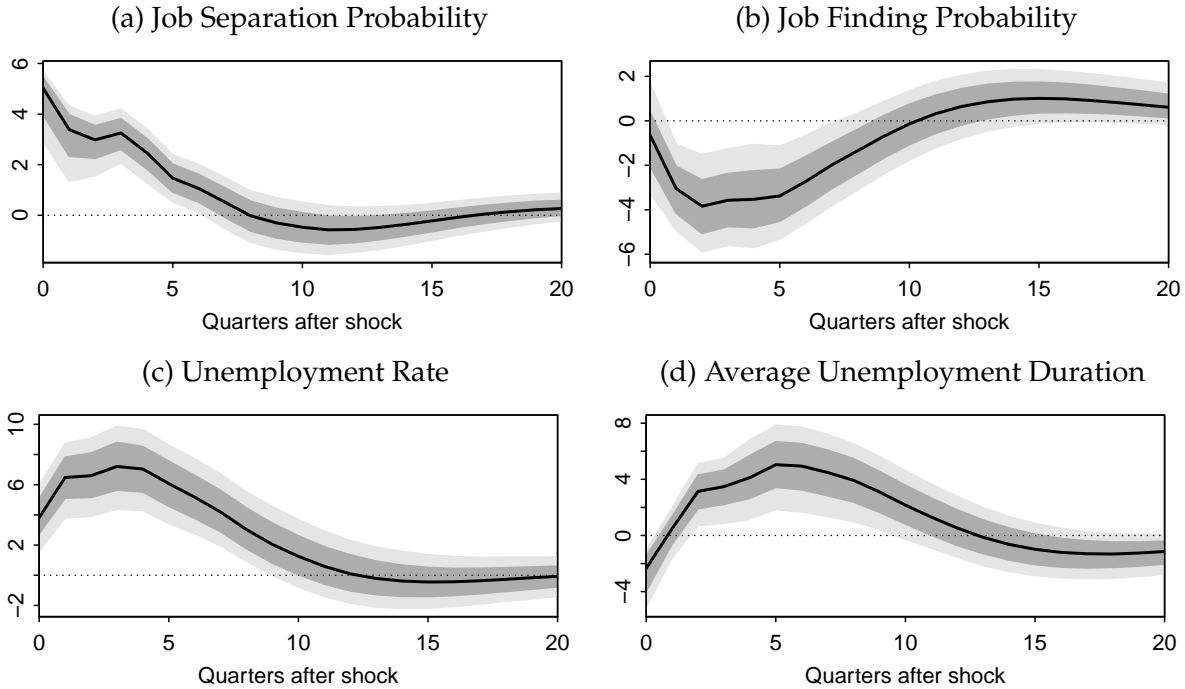
I now discuss the estimated effects of the unemployment shocks on the variables in the VAR system. Then, I use local projection methods to estimate the effects of these shocks on auxiliary labor market variables. I discuss how the estimated responses of the auxiliary variables shed light on the underlying forces at work. Next, I formally decompose the variance of each variable in the system into the components accounted for by each shock. Lastly, I compute the forecast error variance decomposition (FEVD) and historical decomposition (HD) for specific recessionary episodes and discuss their implications.

5.1 Response to Unemployment Inflow and Outflow Shocks

Figure 3 contains the estimated impulse response functions to a shock to the first variable in the VAR system. I label this shock an “unemployment inflow” shock, as it is associated with the job separation probability. The job separation probability rises by about 5% on impact and then remains elevated relative to trend for almost two years. The job finding probability falls slightly on impact, troughs at about -4% after 2-3 quarters, and then slowly returns to trend. The unemployment rate displays a hump-shaped pattern in that it initially rises, peaks about 4 quarters after the shock, and then subsequently returns to trend. Note that the response of the unemployment rate is driven by the joint dynamics of the job separation and job finding probabilities. An initial spike in job separations causes the unemployment rate to increase on impact; it then rises further because of the decrease in the job finding probability. Average unemployment duration initially falls due to spike in unemployment inflows, but subsequently rises as workers accumulate in the unemployment pool and the job finding probability falls.

Figure 4 shows the response of the variables in the system to a (contractionary) shock to the the second variable in the system, the job finding probability, which I label a “unemployment outflow” shock. Recall that under my identification assumptions, an unem-

Figure 3: Response to Unemployment Inflow Shock



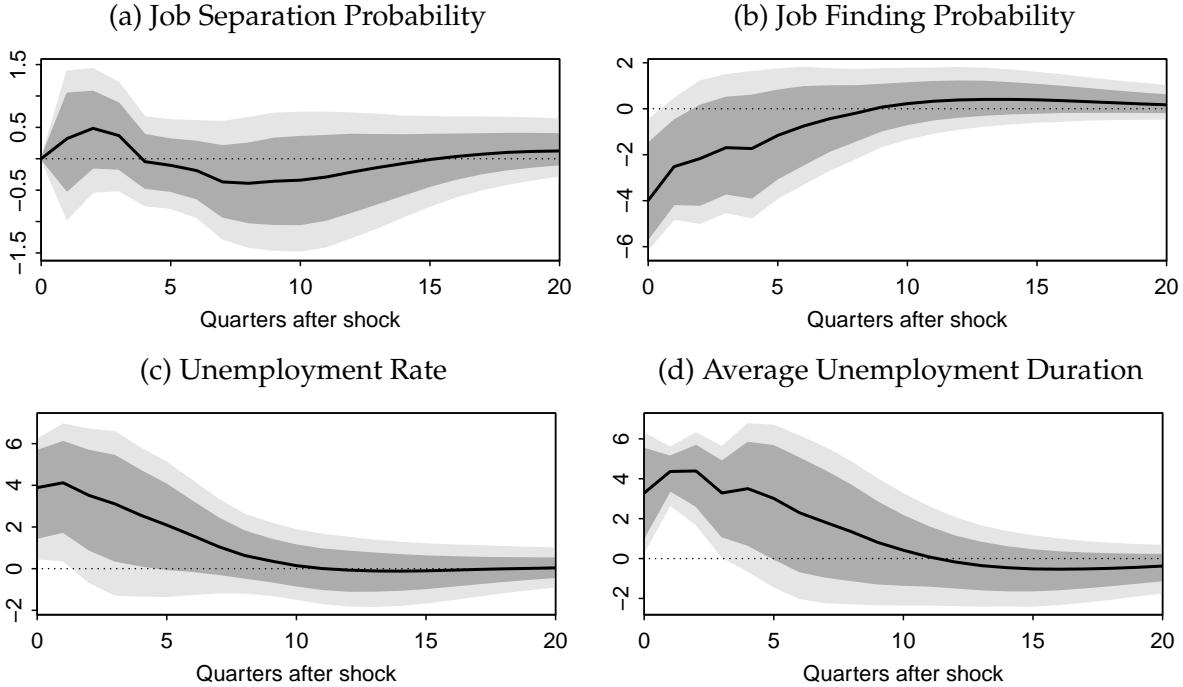
Notes: Response to a one standard deviation (contractionary) structural shock. Units are percentage deviations from trend. Black, solid lines display the median response. Dark gray bands represent the 68% credible set. Light gray bands represent the 90% credible set.

employment outflow shock may not affect the job separation probability on impact. In response to this shock, there is not a significant reaction of the job separation probability at any horizon, though separations do rise slightly in the first 1-4 quarters under the median impulse response function. The job finding probability falls on impact and slowly returns to trend, meaning that an initial fall in unemployment outflows has a persistent effect on the rate at which workers find jobs out of unemployment. Both the unemployment rate and the average unemployment duration rise significantly on impact and remain elevated for several quarters relative to trend.

5.2 Responses of Other Variables

Exploring the effects of the identified unemployment shocks on additional variables helps to unpack the underlying structural forces these shocks capture. I therefore use local projection methods (Jordà, 2005) to estimate the impact of unemployment inflow and outflow shocks on auxiliary labor market variables. Let $\hat{\varepsilon}_t^i$ represent an identified shock to variable i at time t . To recover the time series of each $\hat{\varepsilon}_t^i$, I use the median response implied by my sign restriction approach and invert the resulting impact matrix to back

Figure 4: Response to Unemployment Outflow Shock



Notes: Response to a one standard deviation (contractionary) structural shock. Units are percentage deviations from trend. Black, solid lines display the median response. Dark gray bands represent the 68% credible set. Light gray bands represent the 90% credible set.

out the structural shocks according to Equation 6. Then, I estimate the equation

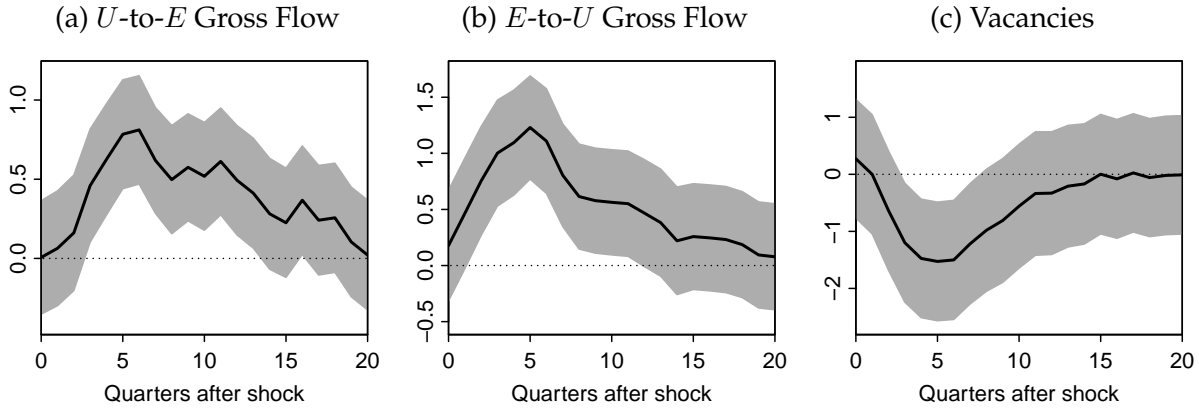
$$x_{t+h} = \alpha_h + \beta_h \hat{\varepsilon}_t^i + e_{t+h} \quad (7)$$

for different horizons h and for different outcome variables x_{t+h} . The coefficients β_h capture the response of the outcome variable to the shock at horizon h , providing an estimate the impulse response function (Plagborg-Møller and Wolf, 2021).

I examine the responses of U -to- E gross flows, E -to- U gross flows, and vacancies to unemployment inflow and outflow shocks. As shown in Figure 1, both gross flow series are countercyclical, so shocks that are associated with large recessionary increases in unemployment should also lead to increases in gross worker flows. The response of job openings is informative about the degree to which my estimated shocks generate Beveridge Curve dynamics, which has been a central focus of the search-and-matching literature. Collectively, examining these auxiliary impulse responses also serves as a check on the validity of the identification approach.

Figure 5 plots the response of gross flows and vacancies to a contractionary unemployment inflow shock. After an increase in unemployment inflows, the total number of

Figure 5: Response to Unemployment Inflow Shock




Notes: Gross flow series constructed as described in Section 2. Vacancy series is the help-wanted index from [Barnichon \(2010\)](#). Black solid lines show estimated β_h coefficients. Gray bands show 95% confidence intervals. Dependent variables are in log-levels so that units are percent changes.

workers moving from employment to unemployment (center panel) initially rises. It then increases to a peak change of almost 1.5% around 5 quarters before returning to its initial level. The gross flow of workers between unemployment and employment (left panel) does not change on impact, but eventually rises by just under 1%, also peaking around 5 quarters after the shock. Note that these patterns are consistent with the evidence presented in [Figure 1](#); both EU and UE gross flows rise in contractionary periods and EU gross flows rise by more. Lastly, vacancies decline considerably in response to a contractionary unemployment inflow shock, falling by almost 2% a year after the shock. This pattern is consistent with not only a decline in labor market tightness during a recession (as the unemployment rate also rises), but also the decrease in the job finding probability in response to an inflow shock shown in [Figure 3](#). In sum, unemployment outflow shocks seem to be associated with realistic behavior of labor market variables during a recession.

[Figure 6](#) plots the responses of the same three variables to an unemployment outflow shock. In contrast, unemployment outflow shocks do not seem to replicate the empirical patterns referenced above. Though the point estimates show that EU gross flows change by a greater magnitude than UE gross flows, neither response is meaningfully large at the given level of statistical significance. The same is true of the response of vacancies, as the estimated impulse response function is not significant for any horizon. It therefore seems that both unemployment inflow and unemployment outflow shocks capture features of the recessionary labor market, but unemployment inflow shocks are more promising candidates for generating realistic unemployment dynamics, as they are consistent with additional business cycle patterns of key variables.

Figure 6: Response to Unemployment Outflow Shock



Notes: Gross flow series constructed as described in Section 2. Vacancy series is the help-wanted index from [Barnichon \(2010\)](#). Black solid lines show estimated β_h coefficients. Gray bands show 95% confidence intervals. Dependent variables are in log-levels so that units are percent changes.

5.3 Contribution to Unemployment Volatility

I now assess the contribution of each structural shock to the volatility of unemployment. To do so, I compute the variance of each of the endogenous variables under specific shock combinations. I derive a closed form expression for the variance of any particular variable in a $VAR(p)$ model in Appendix B and use the formulas in this section.

Let $\text{Var}(y_t|\varepsilon_t^i)$ denote the variance of the variables in the vector y_t under a particular structural shock ε_t^i . The total variance of y_t is equal to the sum across i of $\text{Var}(y_t|\varepsilon_t^i)$ because the structural shocks are independent. Therefore, we can decompose the variance of each endogenous variable into the components accounted for by each structural shock. This procedure is similar in spirit to the forecast error variance decomposition (FEVD), which I discuss below, but it does not fix a particular horizon.

Table 4: Shock Contribution to Endogenous Variable Variance

	Inflow	Outflow	Level	Length	Total
Job separation probability	57.44	6.23	11.05	25.28	100%
Job finding probability	31.98	20.9	28.59	18.54	100%
Unemployment rate	50.33	16.32	19.05	14.3	100%
Average unemployment duration	32.23	24.07	25.39	18.32	100%

Notes: Each element of the table contains the percentage of the overall variance of a given variable explained by a particular shock under the median of the estimated coefficients. Total column shows the sum across the Inflow, Outflow, Level, and Length columns.

Table 4 shows the results of this decomposition for each variable in my baseline specification. As noted above, summing across the structural shocks returns the overall vari-

ance of the endogenous variables, so the columns sum to 100% by construction. Focusing first on the the unemployment inflow shock, we can see that it explains a large fraction of each of the variables in the system. It naturally explains a large portion of the variance of the job separation probability, but also plays a large role in the variance of the job finding probability and average unemployment duration. This is because, as discussed above, job finding falls and average unemployment duration eventually rises substantially in response to an increase in unemployment inflows (Figure 3). Unemployment outflow shocks also play a considerable role in changes in the job finding probability and average unemployment duration.

Next, the row corresponding to the unemployment rate shows that the unemployment inflow shock explains the majority of the variance in the unemployment rate. In particular, unemployment inflow shocks account for over 50% of unemployment rate fluctuations, while unemployment outflow shocks account for only about 16%, in line with the contributions of the other two shocks. It would therefore be tempting to conclude from this decomposition alone that job separation (the “ins” of unemployment) plays a larger role in unemployment fluctuations than job finding (the “outs” of unemployment). However, this decomposition also relies on the *endogenous* responses of the variables in the VAR system, which confound the roles played individually by the job finding and job separation probabilities. I therefore turn to the next decomposition to separate their contributions.

Shutting Down Interaction Table 5 shows the results from the same decomposition, shutting down interaction effects between the job finding and job separation probabilities. In particular, I set the VAR coefficients on lags of F_t in the X_t equation and on lags of X_t in the F_t equation to 0. If part of the large role of unemployment inflow shocks in the variance of the unemployment rate is due to the endogenous, negative response of the job finding probability, we should see a smaller role for these shocks in unemployment fluctuations. The third row of the table shows that this is indeed the case, where the contribution of unemployment inflow (outflow) shocks decreases (increases) from 50% (16%) to 40% (19%). Note that in percentage terms, both of these changes are similar in magnitude, about 20% in either case. Moreover, the contribution of the unemployment inflow shock to the variance of the job finding probability declines as well, from over 31% to just under 29%. Clearly, the interaction between job finding and job separation affects any conclusions drawn the contribution of both margins to unemployment rate movements.

The results of these decompositions caution against the interpretation that changes

Table 5: Shock Contribution to Endogenous Variable Variance, No Interaction

	Inflow	Outflow	Level	Length	Total
Job separation probability	50.66	11.42	12.86	25.06	100%
Job finding probability	28.64	20.86	25.69	24.81	100%
Unemployment rate	37.53	19.49	22.3	20.68	100%
Average unemployment duration	29.35	22.52	24.68	23.44	100%

Notes: Each element of the table contains the percentage of the overall variance of a given variable explained by a particular shock under the median of the estimated coefficients, with the interaction between job finding and job separation turned off. Total column shows the sum across the Inflow, Outflow, Level, and Length columns.

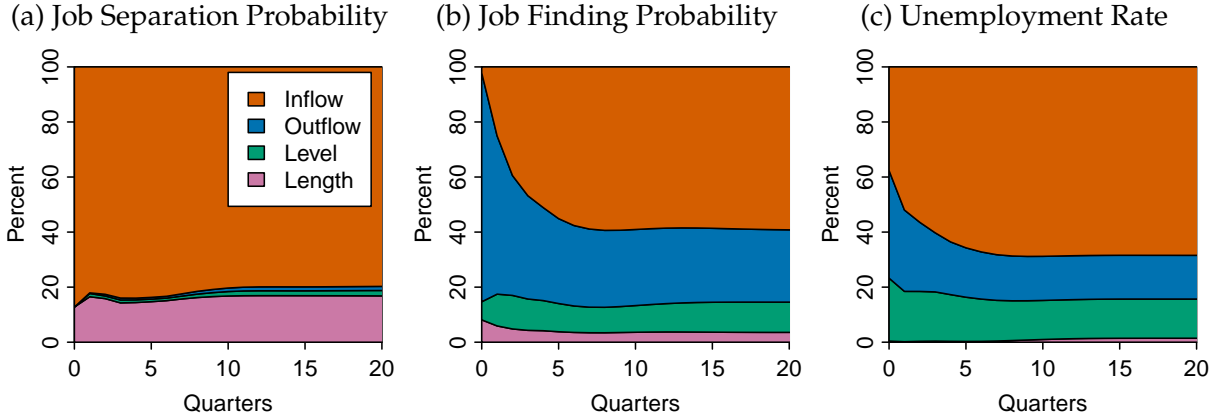
in the job finding probability are solely or primarily responsible for business cycle fluctuations in the unemployment rate. Whereas [Shimer \(2012\)](#) finds that the contribution of job finding is 2-3 times larger than the contribution of job separation to unemployment volatility, my results are more in line with those of [Fujita and Ramey \(2009\)](#), who find that fluctuations in the separation probability explain between 40 and 50% of fluctuations in unemployment. I find that while the job separation probability has a quantitatively larger contribution, both margins play a substantial role in driving unemployment dynamics. Moreover, allowing for an innovation in job separations to endogenously cause job finding to respond changes the conclusions regarding the contribution of each margin to overall unemployment volatility.

5.4 Forecast Error Variance Decomposition

I also compute the forecast error variance decomposition (FEVD) to further explore the role of the shocks at different horizons. [Figure 7](#) plots this decomposition for the main variables of interest: the job separation probability, the job finding probability, and the unemployment rate. As is also apparent from [Table 4](#), unemployment inflow shocks play a large role in driving unemployment dynamics. The left panel shows that they contribute to about 80% of fluctuations in the job separation probability in a 5-year window. The remaining 20% are attributed to unemployment length shocks, which are associated with innovations to average unemployment duration. This result likely stems from the inverse relationship between job separations and the unemployment duration distribution. Times when workers are flowing into unemployment are times when average unemployment duration are falling and vice versa.

The center panel of [Figure 7](#) shows that unemployment inflow shocks also drive the bulk of both job finding and unemployment fluctuations, but their contribution changes

Figure 7: Forecast Error Variance Decomposition



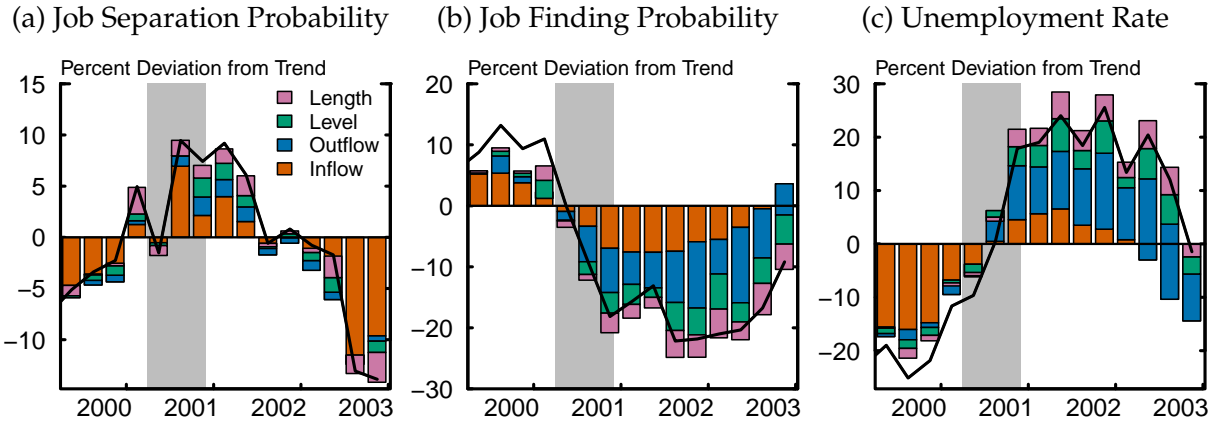
Notes: Figure shows the percentage of the unpredictable component of each variable accounted for by different shocks across horizons, calculated using pointwise median impulse response functions.

significantly across horizons. Indeed, unemployment *outflow* shocks play a significant role in both job finding and unemployment changes at shorter horizons. They comprise the majority share job finding probability fluctuations at shorter horizons (1-4 quarters), a pattern masked by the results presented in Table 4. At longer horizons, however, unemployment inflow shocks dominate. The large contribution of these shocks at longer horizons results from the endogenous response of the job finding probability to an unemployment inflow shock (Figure 3), whereby the job finding probability falls for several quarters following a spike in *E-to-U* flows. The FEVD also makes clear that while both job finding and job separation play an important role in unemployment fluctuations in an *accounting* sense, it is *shocks* to unemployment inflows that are quantitatively more relevant. Identifying innovations in worker flows therefore serves to distinguish the two margins, providing a more complete picture of the forces responsible for driving the business cycle dynamics of unemployment.

5.5 Historical Decomposition

Lastly, I assess the contribution of unemployment shocks to labor market dynamics during different historical episodes. While the variance decomposition and forecast error variance decomposition show that unemployment inflow shocks play a larger role than unemployment outflow shocks on average, their contributions may not be constant across recessions. Each U.S. recession has been driven by a different combination of macroeconomic disturbances. Because my identified unemployment shocks likely capture a confluence of underlying structural forces, they should also display different roles in different

Figure 8: Historical Decomposition, 2001 Recession



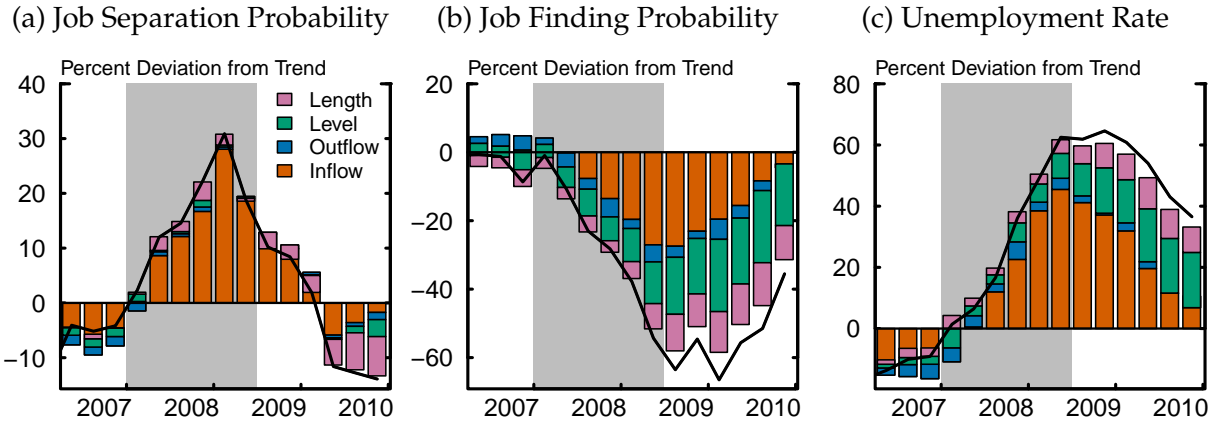
Notes: Figure shows shock contribution to the evolution of each variable, calculated by taking the median HD across draws. Black, solid line shows the actual level of the given variable, which may not line up exactly due to the contribution of the initial condition.

economic downturns. I therefore compute the historical decomposition (HD) and compare the role of unemployment inflow and outflow shocks in the 2001 and 2008 recessions.

Figure 8 plots the HD for the 2001 recession. From the left panel, we can see that unemployment inflow shocks primarily drive changes in the job separation probability, consistent with the results above. However, the center panel reveals that both unemployment inflow and unemployment outflow shocks played a large role in driving job finding dynamics in this recession. Early in the recession, unemployment outflow shocks accounted for the majority share of the sharp drop in the job finding probability. After the end of the recession according to the NBER recession dates (shown in dark, grey bars in the figure), both unemployment inflow and outflow shocks accounted for a depressed job finding probability into the early stages of recovery. As the job finding probability recovered after 2002, the role of unemployment inflow shocks diminished, disappearing almost completely by the end of 2003.

Turning to the right panel, we can see that unemployment outflow shocks accounted for the largest portion of the increase in the unemployment rate during this period. While unemployment inflow shocks also played a role, their contribution becomes negligible by the end of 2002. This result is a direct implication of the findings above: unemployment inflow shocks lead to *endogenous* movements in the job finding probability (Figure 3); because they are not as important in driving the job finding probability during this period (Figure 8b), unemployment outflow shocks must dominate. Therefore, this recession was characterized more so by a depressed job finding rate than by an increased separation rate. The same may not apply to *all* recessions, however.

Figure 9: Historical Decomposition, 2008 Recession



Notes: Figure shows shock contribution to the evolution of each variable, calculated by taking the median HD across draws. Black, solid line shows the actual level of the given variable, which may not line up exactly due to the contribution of the initial condition.

For instance, Figure 9 shows that the opposite conclusion holds for the 2008 recession. Unemployment inflow shocks accounted for the majority share of fluctuations in the job separation probability, job finding probability, and unemployment rate during this period. As discussed, this is because they endogenously lead to drops in the job finding probability, putting further upward pressure on the unemployment rate. This recession was characterized by a large and persistent increase in the unemployment rate and a slow recovery. The historical decomposition of this period reveals that this was due to elevated unemployment inflows and the associated decline in the speed with which unemployed workers found new jobs. Interestingly, unemployment level shocks also played a large role in unemployment dynamics during this period. I interpret these shocks as capturing any forces that result in changes in the unemployment rate holding both E -to- U flows and U -to- E flows constant. For instance, they may capture movements into and out of the labor force. Their large role in unemployment fluctuations during the 2008 recession and subsequent recovery makes sense in the context of the large decline in labor force participation during this period. Indeed, the right panel shows that they had become the dominant force in driving the unemployment rate by the end of 2010.

Collectively, this analysis reveals that studying the role of unemployment inflows and unemployment outflows in conjunction paints a more accurate picture of labor market dynamics. Though the variance decomposition exercises above revealed that unemployment inflows shocks are the quantitatively dominant force on average, they played drastically different roles in different recessions. Moreover, assessing the endogenous response of labor market variables to unemployment shocks is crucial for understanding

their effects. Therefore, focusing only on a particular margin of worker flows may lead to misleading conclusions about unemployment fluctuations.

6 Implications for Models of the Labor Market

I now discuss the implications of my empirical results for theoretical models of search and unemployment in labor markets. First, I briefly discuss the various underlying forces that my crude structural shock estimates likely capture.

6.1 Discussion of Structural Shocks

Though I refer to the job separation shock and the job finding shock as “structural shocks,” I do not think of them as fundamental, underlying macroeconomic disturbances. Rather, they are attempts to isolate changes in job separation and job finding probabilities that are plausibly exogenous to prevailing labor market conditions. They likely capture the responses of worker flows to deeper changes in the aggregate economy, but in this sense, studying the response of labor market variables to these shocks is informative about what underlying structural forces are at play. In particular, my empirical approach does not impose how job finding reacts to an innovation in job separations. Instead, I estimate across several specifications that the job finding probability falls after an increase in job separations. Moreover, I do not restrict the job finding probability to be constant when assessing the contribution of job separations to unemployment volatility. I find that this alters the degree to which job finding versus job separation accounts for movements in the unemployment rate. Therefore, though my identification strategy is simplistic and likely does not capture the true structural origins of labor market fluctuations, it nonetheless yields empirical moments that serve to discipline macroeconomic models of the labor market.

6.2 Theories of Unemployment

The evidence I present in this paper shows that changes in the inflow to unemployment should not be ignored in analyzing the cyclical behavior of the labor market. The strong pattern of Granger causality that I establish in Section 3 suggests that countercyclical E -to- U flows help give rise to unemployment fluctuations. My empirical methodology formalizes the role of job separations in driving unemployment, and I find that job separation shocks have a large and significant contribution to unemployment volatility. Therefore,

treating job separations as exogenous and acyclical in models of the labor market misses one of the key factors responsible for increased unemployment in recessions.

Moreover, allowing for a dynamic interaction between job separation and job finding affects the degree to which changes in either margin contribute to the overall variance of unemployment. I estimate that the job finding reacts significantly to job separation shocks and vice versa. My results point to a role for increased separations in driving down the job finding probability in recessions, potentially through one of the mechanisms referenced in the literature review. In this sense, my results are in line with the conclusions of other studies that economists should jointly consider the behavior of unemployment inflows and outflows. For instance, [Elsby et al. \(2009\)](#) declare that “everyone’s winner” and rejects attempts to analyze the two margins separately.

In terms of economic modeling, my paper suggests a return to the considerations that motivated the original [Mortensen and Pissarides \(1994\)](#) study. These authors included an endogenous separation margin in their model exactly to explain the rise in unemployment inflows at the beginning of an economic downturn. However, in subsequent analyses, the endogenous separation margin dropped out of these models and focus turned instead to explaining what factors determine fluctuations in the job finding rate. For instance, [Shimer \(2005\)](#), which concluded that productivity shocks in the DMP framework cannot quantitatively generate unemployment fluctuations in line with those in the data, notably does not feature endogenous separations. [Gavazza et al. \(2018\)](#) motivate their paper by appealing to the [Shimer \(2012\)](#) result that “swings in the job-finding rate account for the bulk of unemployment fluctuations.”

Studies that reintroduce a role for endogenous job separation (e.g. [den Haan et al., 2000](#); [Fujita and Ramey, 2012](#); [Coles and Moghaddasi-Kelishomi, 2018](#), among others) find that the picture is more complicated. These papers show that by making small modifications to the DMP framework, the model can capture labor market fluctuations both qualitatively and quantitatively. A recent study by [Hall and Kudlyak \(2022\)](#) draws attention to a persistently elevated separation rate as a potential feedback mechanism to explain why unemployment falls slowly after a recession. My results support the notion that the job separation margin should be modeled explicitly in theories of unemployment.

7 Conclusion

In this paper, I develop a methodology to assess how inflows to versus outflows from the pool of unemployed workers drive the dynamics of the unemployment rate. In a reces-

sion, unemployment may rise if it is harder for workers to find jobs or if more workers lose jobs. This question has been widely studied by macro and labor economists, but to my knowledge, no consensus exists regarding which channel is more important.

My analysis uses two different types of worker flows series that measure how often workers move between employment and unemployment. While gross flows measure the total number of workers who transition between labor market states, flow probabilities measure the likelihood these transitions occur. I show that regardless of the series used, job separation precedes job finding in time. Importantly, however, the probability a worker transitions from unemployment to employment is procyclical while the gross flow of workers from unemployment to employment is countercyclical.

I estimate a VAR model under an identification strategy motivated by this evidence and use my estimates to decompose changes in the unemployment rate into the contributions of job finding and separation. I find that both margins account for a significant fraction of unemployment volatility. Unlike models of unemployment fluctuations that abstract from endogenous movements in job separation, my results suggest that economists should endeavor to understand the reasons why more workers lose jobs during recessions in conjunction with the reasons why jobs become harder to find.

References

- Ahn, Hie Joo and James D. Hamilton**, "Heterogeneity and Unemployment Dynamics," *Journal of Business & Economic Statistics*, 2020, 38 (3), 554–569.
- Arias, Jonas E., Dario Caldara, and Juan F. Rubio-Ramírez**, "The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure," *Journal of Monetary Economics*, 2019, 101, 1–13.
- , **Juan F. Rubio-Ramírez, and Daniel F. Waggoner**, "Inference Based on Structural Vector Autoregressions Identified With Sign and Zero Restrictions: Theory and Applications," *Econometrica*, 2018, 86 (2), 685–720.
- Barnichon, Regis**, "Building a Composite Help-Wanted Index," *Economics Letters*, 2010, 109, 175–178.
- Blanchard, Olivier Jean and Peter Diamond**, "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity*, 1990, 1990 (2), 85–155.
- Canova, Fabio, David Lopez-Salido, and Claudio Michelacci**, "The Ins and Outs of Unemployment: An Analysis Conditional on Technology Shocks," *The Economic Journal*, 2013, 123 (569), 515–539.
- Coles, Melvyn G. and Ali Moghaddasi-Kelishomi**, "Do Job Destruction Shocks Matter in the Theory of Unemployment?," *American Economic Journal: Macroeconomics*, 2018, 10 (3), 118–136.
- Darby, Michael R., John C. Haltiwanger, and Mark W. Plant**, "The Ins and Outs of Unemployment: The Ins Win," *NBER Working Paper No. 1997*, 1986.
- Davis, Steven J.**, "Fluctuations in the Pace of Labor Reallocation," *Carnegie-Rochester Conference Series on Public Policy*, 1987, 27, 335–402.
- den Haan, Wouter J., Garey Ramey, and Joel Watson**, "Job Destruction and Propagation of Shocks," *American Economic Review*, 2000, 90 (3), 482–498.
- Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin**, "On the Importance of the Participation Margin for Labor Market Fluctuations," *Journal of Monetary Economics*, 2015, 72, 64–82.
- , **Ryan Michaels, and Gary Solon**, "The Ins and Outs of Cyclical Unemployment," *American Economic Journal: Macroeconomics*, 2009, 1 (1), 85–110.

- Engbom, Niklas**, “Contagious Unemployment,” *Working Paper*, 2021.
- Ferroni, Filippo and Fabio Canova**, “A Hitchhiker’s Guide to Empirical Macro Models,” *Working Paper*, 2022.
- Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, and Michael Westberry**, “Integrated Public Use Microdata Series, Current Population Survey: Version 10.0 [dataset],” Technical Report, Minneapolis, MN: IPUMS 2022.
- Fujita, Shigeru**, “Dynamics of Worker Flows and Vacancies: Evidence from the Sign Restriction Approach,” *Journal of Applied Econometrics*, 2011, 26, 89–121.
- **and Garey Ramey**, “The Cyclicity of Separation and Job Finding Rates,” *International Economic Review*, 2009, 50 (2), 415–430.
- **and –**, “Exogenous versus Endogenous Separation,” *American Economic Journal: Macroeconomics*, 2012, 4 (4), 68–93.
- Gavazza, Alessandro, Simon Mongey, and Giovanni L. Violante**, “Aggregate Recruiting Intensity,” *American Economic Review*, 2018, 108 (8), 2088–2127.
- Gregory, Victoria, Guido Menzio, and David Wiczer**, “The Alpha Beta Gamma of the Labor Market,” *Working Paper*, 2022.
- Hall, Robert E.**, “Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market,” *The Review of Economics and Statistics*, 2005, 87 (3), 397–307.
- **and Marianna Kudlyak**, “Why Has the US Economy Recovered So Consistently from Every Recession in the Past 70 Years?,” *NBER Macroeconomics Annual*, 2022, 36, 1–55.
- Hamilton, James D.**, *Time Series Analysis*, Princeton, NJ: Princeton University Press, 1994.
- , “Why You Should Never Use the Hodrick-Prescott Filter,” *The Review of Economics and Statistics*, 2018, 100 (5), 831–843.
- Jordà, Òscar**, “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 2005, 95 (1), 161–182.
- Mercan, Yusuf, Benjamin Schoefer, and Petr Sedláček**, “A Congestion Theory of Unemployment Fluctuations,” *American Economic Journal: Macroeconomics*, 2024, 16 (1), 238–285.

- Mortensen, Dale T. and Christopher A. Pissarides**, "Job Creation and Job Destruction in the Theory of Unemployment," *The Review of Economic Studies*, 1994, 61 (3), 397–415.
- Plagborg-Møller, Mikkel and Christian K. Wolf**, "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 2021, 89 (2), 955–980.
- Ritter, Joseph A.**, "Measuring Labor Market Dynamics: Gross Flows of Workers and Jobs," *Federal Reserve Bank of St. Louis Review*, 1993, 75 (6), 39–57.
- Rubio-Ramírez, Juan F., Daniel F. Waggoner, and Tao Zha**, "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *The Review of Economic Studies*, 2010, 77 (2), 665–696.
- Shimer, Robert**, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 2005, 95 (1), 25–49.
- , "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 2012, 15 (2), 127–148.

A Details on Measurement

A.1 Time Aggregation Bias

Let P_t denote the monthly, discrete time Markov transition matrix between labor market states. In my application, this is a 3×3 matrix with non-negative entries and columns that sum to 1.¹¹ The continuous time counterpart to P_t , which I denote Λ_t , can be recovered from the formula below.

$$\Lambda_t = O_t N_t O_t^{-1} \quad (\text{A.1})$$

In this formula, O_t is the matrix of eigenvectors of P_t and N_t is a diagonal matrix containing the *natural logarithm* of the corresponding eigenvalues. Such a decomposition is possible as long as P_t has distinct, real, and positive eigenvalues, which is the case in my sample and also holds in the [Shimer \(2012\)](#) data. Lastly, I construct the time aggregation adjusted, monthly, discrete time Markov transition matrix by converting the continuous time transition rates back to transition probabilities. Let π_t^{AB} denote an element in Π_t and let λ_t^{AB} denote an element in Λ_t for two different labor market states A and B . Then,

$$\pi_t^{AB} = 1 - \exp(-\lambda_t^{AB}) \quad (\text{A.2})$$

A.2 Short-Term Unemployment

The 1994 CPS redesign introduced a discontinuity in the short-term unemployment rate series U_t^s . Beginning in 1994, the CPS introduced dependent interviewing techniques such that only individuals in rotation groups 1 and 5, the “incoming rotation groups,” are asked explicitly about their unemployment duration, while unemployment duration for workers in other rotation groups is imputed. The discontinuity produces a structural break in both the job finding probability F_t and the job separation probability X_t in 1994 that is visible in the level of each series. Following [Shimer \(2012\)](#), I use CPS microdata to correct for this discontinuity.¹² I download data on unemployment duration from IPUMS CPS and keep only individuals in the incoming rotation groups. I then compute the total number of workers with unemployment duration less than 5 weeks using the appropriate sample weights. Therefore, my short-term unemployment series U_t^s contains the published short-term unemployment series from the BLS before 1994, and the estimated series using CPS microdata after 1994. Appendix Figure [A.1](#) shows both the uncorrected and corrected flow probabilities, which match those used in [Shimer \(2012\)](#).

¹¹Note that this is because I exclude persons with missing labor force status is either month $t - 1$ or month t .

¹²See [Shimer \(2012\)](#) Appendix A for additional detail on how to implement the correction.

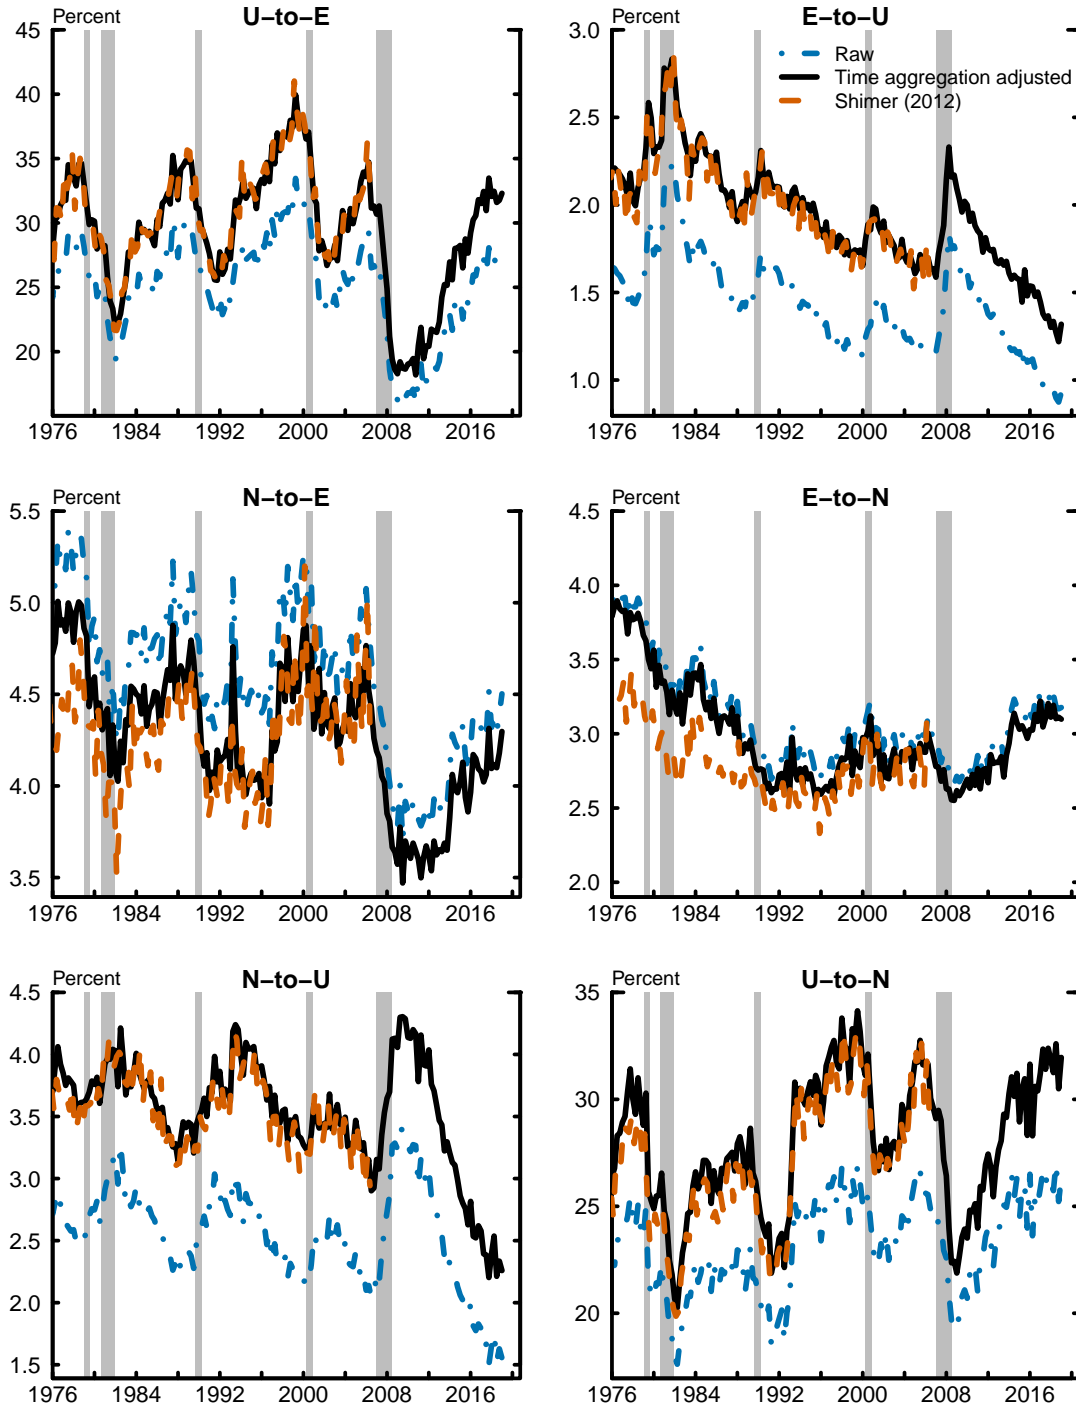
A.3 Replication of Shimer (2012)

Figure A.1: Transition Probabilities in the 2-State Model



Notes: Black, solid lines show author's calculations based on the methodology described in Section 2. Orange, dashed lines show data downloaded from Robert Shimer's website (<https://sites.google.com/site/robertshimer/research/flows>).

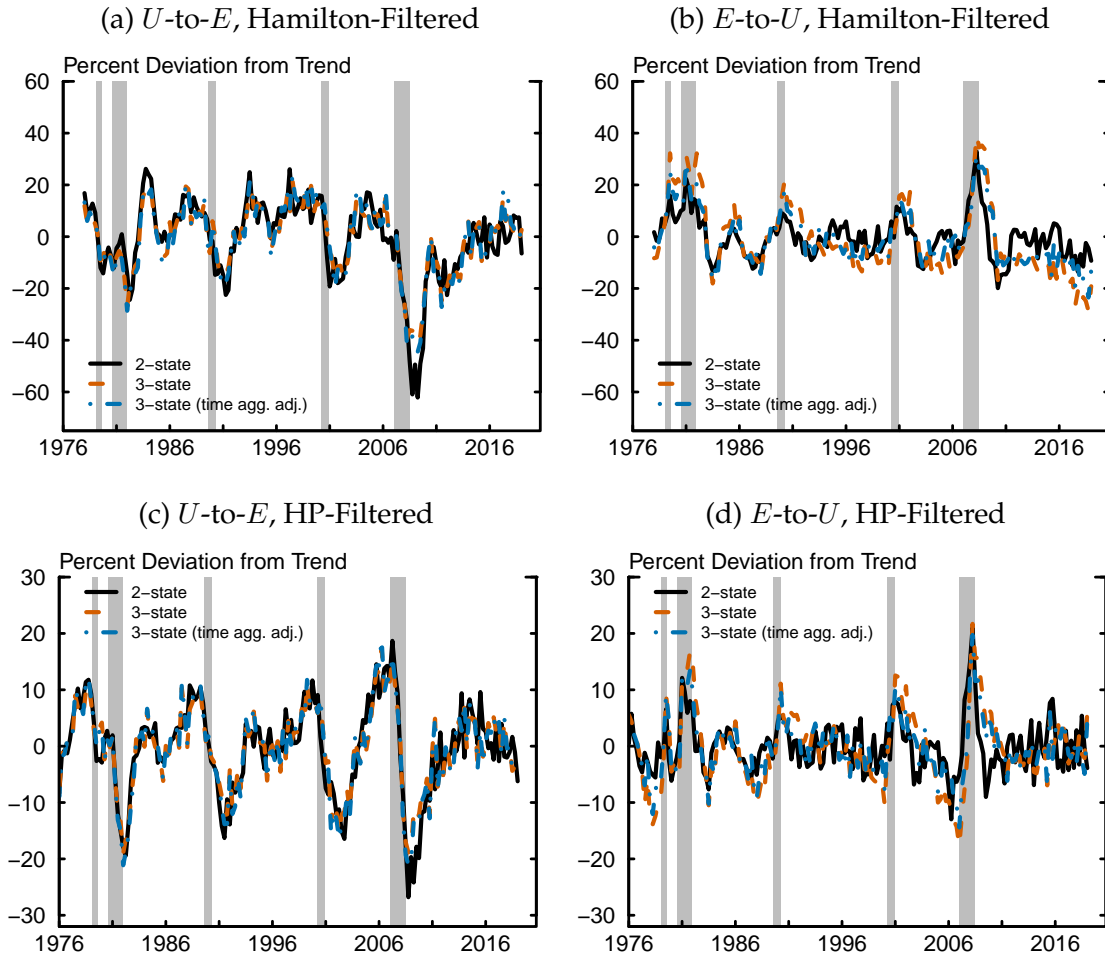
Figure A.2: Transition Probabilities in the 3-State Model



Notes: Black, solid lines show author's calculations based on the methodology described in Section 2. Orange, dashed lines show data downloaded from Robert Shimer's website (<https://sites.google.com/site/robertshimer/research/flows>). Blue, dash-dot lines show author's calculations, unadjusted for time aggregation bias.

A.4 Cyclical Properties of Flows Probabilities, Different Samples

Figure A.3: Detrended Flow Probabilities



Notes: Author's calculations based on the methodology described in Section 2. Black, solid lines show flow probabilities for the two-state model. Orange, dashed lines show flow probabilities for the three-state model, unadjusted for time aggregation bias. Blue, dash-dot lines show flow probabilities for the three-state model, adjusted for time aggregation bias. I take a quarterly average of monthly values before applying the respective filter. I set the horizon h and lag p parameters of the Hamilton filter to 8 and 4 quarters, respectively. I set the smoothing parameter of the HP filter to 1,600.

B Derivations and Proofs

This appendix contains additional detail on the decompositions I perform in Section 5 of the paper. Below, I show how to compute the shock contribution to the variance of the endogenous variables in closed form. I also provide the formulas I use to construct the forecast error variance decomposition and historical decomposition.

B.1 Closed-Form Solution for Variance of VAR(p)

Let the vector autoregression (VAR) model of order p be given by the equation below.¹³

$$Y_t = B_0 + B_1 Y_{t-1} + B_2 Y_{t-2} + \cdots + B_p Y_{t-p} + U_t \quad (\text{B.3})$$

Let n be the number of variables in the model. In the equation above, Y_t is a $n \times 1$ column vector of endogenous variables, U_t is a $n \times 1$ column vector of (reduced-form) residuals, B_0 is a $n \times 1$ column vector of intercept terms, and B_1, \dots, B_p are $n \times n$ matrices of coefficients.

Companion Form First, we transform the VAR(p) model above into a VAR(1) model by constructing its companion form. Let the variance-covariance matrix of the residuals U_t be given by $\text{Var}(U_t) = \Sigma$. Let \mathbf{Y}_t denote the $np \times 1$ *stacked* column vector of endogenous variables and let \mathbf{U}_t denote the $np \times 1$ *stacked* column vector of residuals. Let \mathbf{B}_0 denote the $np \times 1$ *stacked* column vector of intercept terms and \mathbf{B}_1 denote the $np \times np$ matrix of coefficients. These objects and their relationship to Equation (B.3) are given below.

$$\underbrace{\begin{bmatrix} Y_t \\ Y_{t-1} \\ \vdots \\ Y_{t-p} \end{bmatrix}}_{\mathbf{Y}_t} = \underbrace{\begin{bmatrix} B_0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_{\mathbf{B}_0} + \underbrace{\begin{bmatrix} B_1 & B_2 & \cdots & B_p \\ I_n & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & I_n & 0 \end{bmatrix}}_{\mathbf{B}_1} \underbrace{\begin{bmatrix} Y_{t-1} \\ Y_{t-2} \\ \vdots \\ Y_{t-p-1} \end{bmatrix}}_{\mathbf{Y}_{t-1}} + \underbrace{\begin{bmatrix} U_t \\ 0 \\ \vdots \\ 0 \end{bmatrix}}_{\mathbf{U}_t}$$

Note that the first row of this relation simply restates Equation (B.3). Each subsequent row states that $\mathbf{Y}_{t-j} = \mathbf{Y}_{t-j}$ for $j = 1, \dots, p$. Therefore, the companion form is given by

$$\mathbf{Y}_t = \mathbf{B}_0 + \mathbf{B}_1 \mathbf{Y}_{t-1} + \mathbf{U}_t \quad (\text{B.4})$$

Notice that this is a VAR model of order 1.

¹³Note that the notation used to define the VAR is slightly different than in the main text.

Variance Matrix of Endogenous Variables Define the following expression for the variance of the vector of endogenous variables.

$$\mathbf{\Gamma} \equiv \text{Var}(\mathbf{Y}_t) = \begin{bmatrix} \text{Var}(Y_t) & \text{Cov}(Y_t, Y_{t-1}) & \dots & \text{Cov}(Y_t, Y_{t-p}) \\ \text{Cov}(Y_{t-1}, Y_t) & \text{Var}(Y_{t-1}) & \dots & \text{Cov}(Y_{t-1}, Y_{t-p}) \\ \vdots & \ddots & \ddots & \vdots \\ \text{Cov}(Y_{t-p}, Y_t) & \dots & \dots & \text{Var}(Y_{t-p}) \end{bmatrix}$$

The matrix $\mathbf{\Gamma}$ is a $np \times np$ variance-covariance matrix whose entries $\text{Cov}(Y_{t-i}, Y_{t-j})$ are themselves $n \times n$ variance-covariance matrices. Notice that the diagonal elements of $\mathbf{\Gamma}$ (i.e. where $i = j$) are simply the variance of the vector Y_t . When working with a (covariance-) stationary model, $\text{Var}(Y_t) = \text{Var}(Y_{t-1}) = \dots = \text{Var}(Y_{t-p})$ in population. The off-diagonal elements are simply equal to the covariance of Y_t with lags of itself. Hence, $\mathbf{\Gamma}$ has the familiar symmetry property of variance-covariance matrices.

Variance Matrix of Residuals Define the following expression for the variance of the vector of residuals.

$$\mathbf{\Omega} \equiv \text{Var}(\mathbf{U}_t) = \begin{bmatrix} \text{Var}(U_t) & \dots & 0 \\ 0 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 0 \end{bmatrix} = \begin{bmatrix} \Sigma & \dots & 0 \\ 0 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & 0 \end{bmatrix}$$

The matrix $\mathbf{\Omega}$ is a $np \times np$ variance-covariance matrix where the first element is simply the variance of Y_t . All other elements are zero because of the definition of \mathbf{U}_t .

Definition B.1. *The variance of the endogenous variables in the VAR(p) model is related to (i) the VAR coefficients and (ii) the variance of the residuals through the formula:*

$$\text{vec}(\mathbf{\Gamma}) = (\mathbb{I}_M - \mathbf{B}_1 \otimes \mathbf{B}_1)^{-1} \text{vec}(\mathbf{\Omega}) \quad (\text{B.5})$$

where $M = (np)^2 \times (np)^2$, \otimes is the Kronecker product, and $(\cdot)^{-1}$ represents the matrix inverse.¹⁴

Proof. Start with the definition of $\mathbf{\Gamma} = \text{Var}(\mathbf{Y}_t)$. Then, plug in the expression for the

¹⁴See Hamilton (1994) Chapter 10.2, pages 265-266 for a similar derivation.

companion form in Equation (B.4) and reduce.

$$\begin{aligned}
\Gamma &= \text{Var}(\mathbf{Y}_t) \\
&= \text{Var}(\mathbf{B}_0 + \mathbf{B}_1 \mathbf{Y}_{t-1} + \mathbf{U}_t) \\
&= \text{Var}(\mathbf{B}_0) + \text{Var}(\mathbf{B}_1 \mathbf{Y}_{t-1}) + \text{Var}(\mathbf{U}_t) \\
&= 0 + \mathbf{B}_1 \underbrace{\text{Var}(\mathbf{Y}_{t-1})}_{\Gamma} \mathbf{B}_1' + \mathbf{\Omega}
\end{aligned}$$

We have that $\Gamma = \mathbf{B}_1 \Gamma \mathbf{B}_1' + \mathbf{\Omega}$. To solve for Γ , we make use of the “vec” operator $\text{vec}(\cdot)$.

$$\begin{aligned}
\text{vec}(\Gamma) &= \text{vec}(\mathbf{B}_1 \Gamma \mathbf{B}_1' + \mathbf{\Omega}) \\
&= \text{vec}(\mathbf{B}_1 \Gamma \mathbf{B}_1') + \text{vec}(\mathbf{\Omega}) \\
&= (\mathbf{B}_1 \otimes \mathbf{B}_1) \text{vec}(\Gamma) + \text{vec}(\mathbf{\Omega})
\end{aligned}$$

Solving for $\text{vec}(\Gamma)$ in the final line yields the desired result in Equation (B.5). \square

Estimation of Closed-Form Variance We can now make use of the above result contained in Equation (B.5) to compute the sample variance of a VAR(p) model in closed form. The algorithm for doing so is as follows:

1. Estimate the VAR(p) model given by Equation (B.3).
2. Obtain the estimated coefficients $\hat{B}_1, \dots, \hat{B}_p$ and estimated variance-covariance matrix of the residuals $\hat{\Sigma}$.
3. Construct the companion form coefficient matrix $\hat{\mathbf{B}}_1$ and companion form variance-covariance matrix of the residuals $\hat{\mathbf{\Omega}}$.
4. Compute $\text{vec}(\hat{\Gamma})$ using $\hat{\mathbf{B}}_1, \hat{\mathbf{\Omega}}$, and Equation (B.5). Then compute $\hat{\Gamma} = \text{vec}^{-1}(\text{vec}(\hat{\Gamma}))$.
5. The variance-covariance matrix of Y_t is the first n rows and n columns of $\hat{\Gamma}$.

Therefore, we can find a closed-form estimate of the variance of the VAR(p) model by first constructing $\hat{\Gamma}$, and then taking the appropriate rows and columns.

$$\hat{\Gamma}_0 \equiv \hat{\text{Var}}(Y_t) = \hat{\Gamma}_{(1:n, 1:n)} \tag{B.6}$$

Computing Conditional Variance We can use the relationship above to compute the variance of the endogenous variables *conditional* on a given structural shock. Notice that because $U_t = AE_t$,

$$\Sigma \equiv \text{Var}(U_t) = \text{Var}(AE_t) = A \underbrace{\text{Var}(E_t)}_{\mathbb{I}_n} A' = AA'$$

Hence, once we have such a matrix A , we can replace Σ with AA' in the formulas above to compute the variance of Y_t . To compute the variance of Y_t conditional on some structural shock, we can follow the following procedure:

1. Obtain A using one of the identification procedures discussed below.
2. To compute the variance conditional on the j th structural shock, construct $A_j = Ae_j$, where $e_j = [0, \dots, 0, 1, 0, \dots, 0]'$ is the column-wise basis vector, with 1 in the j th row and 0 elsewhere.
3. Replace Σ with $A_j A_j'$ in the formulas above.

Let $\text{Var}(Y_t|e_j)$ denote the variance of the endogenous variables conditional on the j th structural shock (i.e. where we have constructed Ω using $A_j A_j'$ instead of Σ). Because variance is additive and the structural shocks are independent, we have that

$$\text{Var}(Y_t) = \sum_j \text{Var}(Y_t|e_j)$$

For VAR models estimated with sign restrictions, we set $A = CQ$, where Q is an orthogonal matrix that satisfies the sign restriction and C is the Cholesky decomposition of Σ .

B.2 Formulas for Decompositions

Forecast Error Variance Decomposition (FEVD) Let $\bar{\phi}_{i,j,h}$ be the median impulse response function of variable i to shock j at horizon h . $\bar{\phi}_{i,j,h}$ is obtained by computing the impulse response function for each (i, j, h) and each draw d , and taking the pointwise median across draws. Then, the FEVD is given by the formula below.

$$FEVD_{i,j,h} = \frac{\sum_{a=1}^h (\bar{\phi}_{i,j,a})^2}{\sum_j \sum_{a=1}^h (\bar{\phi}_{i,j,a})^2} \quad (\text{B.7})$$

Historical Decomposition (HD) The historical decomposition allows one to analyze the counterfactual evolution of each endogenous variable if only certain shocks had hit the economy during the sample period. Let $(B_l, \Sigma, Q)_d$ denote a draw of the reduced form coefficients B_l , the variance covariance matrix of the residuals Σ , and an orthogonal rotation matrix Q that maps the structural shocks into the reduced form residuals. Let $\hat{\mu}_{t,d}$ denote the vector of reduced form residuals associated with this draw and let $\hat{\varepsilon}_{t,d} = A_d^{-1}\hat{\mu}_{t,d}$ denote the vector of structural shocks associated with this draw, where $A_d = C_d Q_d$ is the impact matrix and C_d is the Cholesky decomposition of Σ_d . Now, iterate to form

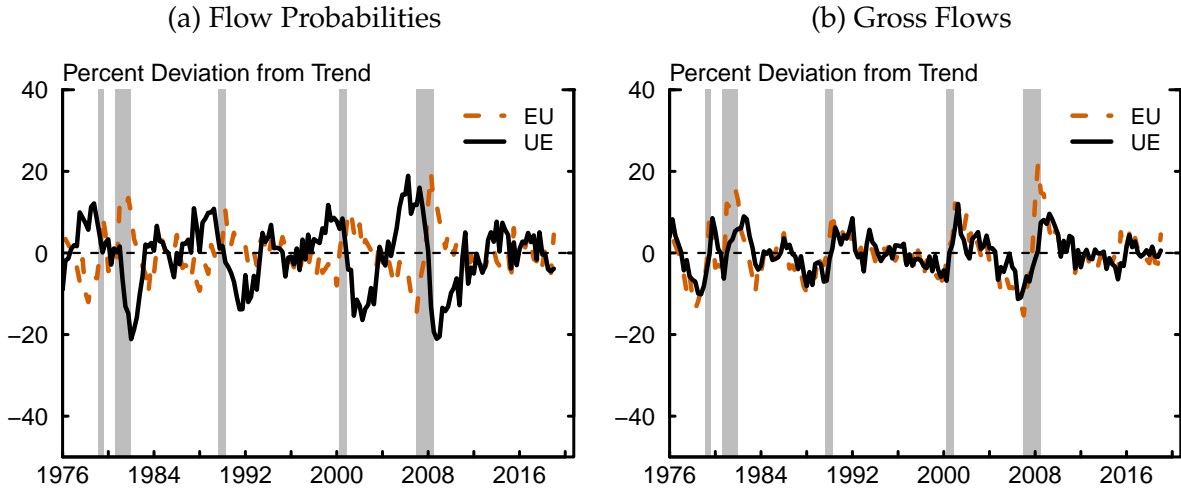
$$y_{t,j,d} = b_{0,d} + \sum_{l=1}^L B_{l,d} y_{t-l,j,d} + \hat{\varepsilon}_{t,d} A_d \mathbf{e}_j$$

for each time period t in the sample period, where \mathbf{e}_j is the j th basis vector, starting from an initial condition of $y_{0,j,d} = [0, \dots, 0]'$. The pointwise median of $y_{t,j,d}$ across draws d is the evolution of y_t under only shock j . Plotting $y_{t,j}$ across time for all shocks $j = 1, \dots, n$ shows the HD.

C Additional Figures and Tables


C.1 Motivating Evidence Robustness

Figure C.4: Cyclical Properties of Worker Flows (HP-Filtered)



Notes: Black, solid lines show U -to- E flows. Orange, dashed lines show E -to- U flows. Series HP-filtered with smoothing parameter 1,600. I take a quarterly average of monthly values before applying the filter. Monthly series adjusted for seasonal variation and time aggregation bias.

Figure C.5: Cross-Correlations of Flow Probabilities (HP-Filtered)



Notes: Series HP-filtered with smoothing parameter 1,600. I take a quarterly average of monthly values before applying the filter. Monthly series adjusted for seasonal variation and time aggregation bias. Correlation coefficient is the Kendall rank correlation coefficient. Sample = 1980Q4–2017Q4.