Finance and Economics Discussion Series

Federal Reserve Board, Washington, D.C.
ISSN 1936-2854 (Print)
ISSN 2767-3898 (Online)

(Re-)Connecting Inflation and the Labor Market: A Tale of Two Curves

Hie Joo Ahn and Jeremy B. Rudd

2024-050

Please cite this paper as:
Ahn, Hie Joo, and Jeremy B. Rudd (2024). "(Re-)Connecting Inflation and the Labor Market: A Tale of Two Curves," Finance and Economics Discussion Series 2024-050. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2024.050.

(Re-)Connecting Inflation and the Labor Market: A Tale of Two Curves[*]

Hie Joo Ahn[†] Jeremy B. Rudd[‡]

June 28, 2024

Abstract

We propose an empirical framework in which shocks to worker reallocation, aggregate activity, and labor supply drive the joint dynamics of labor market outcomes and inflation, and where reallocation shocks take two forms depending on whether they result from quits or from job loss. In order to link our approach with previous theoretical and empirical work, we extend the procedure for estimating a Bayesian sign-restricted VAR so that priors can be directly imposed on the VAR's impact matrix. We find that structural shocks that shift the Beveridge curve have different effects on inflation. Our model allows us to fully decompose movements of or along the empirical Beveridge curve in terms of the contribution of each shock and also allows us to estimate the Phillips correlation associated with each shock; our results imply that observed Beveridge and Phillips correlations can change over time depending on what types of structural shocks predominate in a given period. Applying our model to the pandemic-related recession and recovery, we find that reallocation shocks were a key source of labor market dynamics during this period and explain how a post-pandemic "soft landing," in which inflation declined without a significant rise in unemployment, was possible.

JEL Codes: C11, C32, E24, E31, E32.

Key words: Sign-restricted VAR, Bayesian methods, Labor market dynamics, Reallocation, Inflation, Beveridge curve, Phillips curve, Covid-19

[*]The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. We gratefully acknowledge invaluable guidance from James Hamilton and useful comments and suggestions from Gianni Amisano, Travis Berge, Mark Bognanni, Thorsten Drautzburg, Andrew Figura, and Lam Nguyen. Lucas Moyon provided excellent research assistance. All errors are our own.

[†]Ahn: Federal Reserve Board of Governors, hiejoo.ahn@frb.gov

[‡]Rudd: Federal Reserve Board of Governors, jeremy.b.rudd@frb.gov

"Over the past thirty years, macroeconomists thinking about aggregate labor market dynamics have organized their thoughts around two relations, the Phillips curve and the Beveridge curve. The Beveridge curve, the relation between unemployment and vacancies, has very much played second fiddle. We think that emphasis is wrong. The Beveridge relation comes conceptually first and contains essential information about the functioning of the labor market and the shocks that affect it." – Blanchard and Diamond (1989)

Introduction

The Covid-19 pandemic saw a sudden surge in the unemployment rate that was accompanied by a drop in labor force participation. While the unemployment rate declined gradually and the participation rate remained low in the wake of the pandemic recession, job vacancy postings rose to a record-high level and both wage and price inflation rose sharply and persistently. The unprecedented shocks to the labor market created equally unprecedented movements in the Beveridge curve, leading one observer to describe it as inebriated.[1] More precisely, empirical Beveridge curves shifted out dramatically—first with the rapid rise in unemployment, then with the prolonged increase in job vacancies—suggesting to some that the pandemic induced a large and persistent reallocation shock in the labor market.[2] Hidden in the outward shift, however, was a reduction in labor supply that shifted the Beveridge curve inward (Elsby, Hobijn, and Şahin", 2015); to make matters even more complicated, shocks to aggregate activity resulted in cyclical movements *along* the Beveridge curve.

Emerging evidence also suggests that the pandemic might have strengthened the link between real activity and inflation.[3] Plausibly, decreased labor supply combined with increased labor reallocation can be inflationary, with the former raising wage inflation and the latter raising the structural unemployment rate (and so narrowing the unemployment gap, all else equal). As a result, the increased importance of these structural shocks might have changed the inflation–unemployment tradeoff that would be captured by a reduced-form Phillips curve. This possibility shows that common structural shocks in the labor market can link the Beveridge curve and the Phillips curve, implying a closer relation between these two pillars of macroeconomic dynamics than is commonly realized. Furthermore, joint consideration of the Beveridge curve and Phillips curve should help us to better identify the cyclical state of the economy than

[1] Catherine Rampell commented on Twitter (now X) that the "Beveridge curve is drunk" (see Gallant, Kroft, Lange, and Notowidigdo, 2020).

[2] See, for example, Barrero, Bloom, and Davis (2020) and Barrero, Bloom, Davis, and Meyer (2021).

[3] See, for example, Hobijn, Miles, Royal, and Zhang (2023), Montag and Villar (2022), Smith, Timmermann, and Wright (2023), and Ahn and Smith (2024).

simply using the Phillips curve alone, especially given how flat this latter relation appears to have become in recent decades (Constâncio, 2015).

A number of studies of labor market dynamics attempt to measure the reallocation of workers and jobs, the effects of aggregate cyclical dynamics, or both (e.g., Lilien, 1982; Abraham and Katz, 1986; Davis, 1987; Murphy and Topel, 1987; Hamilton, 1988; Blanchard and Diamond, 1989; Davis and Haltiwanger, 1999a).[4] In a well-known paper, Blanchard and Diamond (1989) described labor market dynamics and the Beveridge curve in terms of shocks to aggregate activity, job reallocation, and labor supply. They proposed an identification strategy where a negative aggregate activity shock lowers job vacancies but raises unemployment (a movement along the Beveridge curve); a positive labor reallocation shock raises both unemployment and vacancies (an outward shift in the Beveridge curve); and a positive labor supply shock raises labor force participation and unemployment (also an outward shift in the Beveridge curve).

Blanchard and Diamond's empirical model essentially amounts to a type of sign-restricted VAR, but the methodology for estimating sign-restricted VARs was not fully understood at the time their paper was written.[5] More-recent work has developed a suitable statistical approach for estimating sign-restricted VARs; we exploit it in our paper to explore the empirical link between the Beveridge curve and the Phillips curve, with a particular focus on the effects of structural labor-market shocks on inflation. The structural shocks are *identified* in a manner similar to that of Blanchard and Diamond (1989). However, we *estimate* these shocks using a Bayesian sign-restricted VAR along the lines of Baumeister and Hamilton (2015), implementing a modified version of their algorithm that allows any restrictions on the structural parameters to be directly imposed on the impact matrix (as opposed to the matrix's inverse). Our paper is the first to combine information from the Beveridge and Phillips curves to investigate the joint dynamics of labor markets and inflation in the context of a Bayesian sign-restricted VAR; to demonstrate its usefulness, we employ this framework to examine labor market and inflation developments during the pandemic period and subsequent recovery.

Though the identification of structural shocks draws on Blanchard and Diamond's characterization of labor market dynamics, our approach differs from theirs in several important ways. First, we introduce

[4]The appendix includes a detailed literature review. Separately, Ferrante, Graves, and Iacoviello (2023) use a multi-sector model with sticky prices, labor reallocation costs, and input-output linkages to examine the effect of the shift in consumer demand from services to goods that occurred during the Covid-19 pandemic, and Garcia-Cabo, Lipinska, and Navarro (2023) use a multi-sector search-and-matching model with on-the-job human capital accumulation to study the effects of labor market policies in the presence of reallocation shocks .

[5]More precisely, Blanchard and Diamond (1989) transformed their model into an over-identified system by imposing further restrictions on the structural parameters such that the implied impulse responses had the desired sign patterns.

wage and price inflation into the model, which provides additional useful information about the state of the labor market. Second, we distinguish between two different types of reallocation shocks—one that raises involuntary separations and another that raises quits—by using the reported reason for unemployment from the Current Population Survey (CPS).[6]

Allowing for different types of reallocation shocks is necessary if we want to fully understand movements in the Beveridge curve and what these movements imply for wage and price inflation. Consider a case where a reallocation shock creates a number of workers who lose their jobs but whose skills are ill-matched to what employers currently demand. More job vacancies will not lower unemployment for these workers, so the observed Beveridge curve will shift outward. Since this rise in unemployment is attributable to a reduction in matching efficiency or a deterioration in labor market functioning, we associate a change in the unemployment rate driven by job-loss reallocation with a change in structural unemployment. To the extent that these workers lower their desired wage in the face of unsuccessful job searches, increased unemployment among these workers is likely to be associated with downward pressure on wage growth. Through the lens of the Beveridge curve, we would correctly perceive that an increase in structural unemployment has taken place.[7] But if we only think in terms of a standard Phillips curve, we are more likely to attribute the rise in actual unemployment and reduction in wage growth to a negative cyclical shock. In this case, therefore, the Beveridge curve correctly signals an increase in structural unemployment while the Phillips curve does not.[8]

Relatedly, if there is an increase in the number of job quitters who seek out and then find better jobs with higher wages (a possibility explored by Moscarini and Postel-Vinay, 2023), then the short-term unemployment rate will rise and the Beveridge curve will shift out. Typically, job quitters will experience a short spell of unemployment because of job-search frictions, and so we associate a change in the unemployment rate caused by quits reallocation with a change in frictional unemployment. At the same time, this development will exert *upward* pressure on wage growth that will also be hard to square with the Phillips curve. Because each type of reallocation shock shifts the Beveridge curve out, conventional

[6]In a set of robustness exercises, we draw on Davis and Haltiwanger (1999a) and also consider whether measures of job creation and job destruction allow us to better pin down these structural shocks.

[7]Michaillat and Saez (2022) use a theoretical argument based on the Beveridge curve to compute estimates of the *efficient* unemployment rate at full employment.

[8]In principle, the rise in structural unemployment could eventually be captured by a Phillips curve with a time-varying intercept after the reduction in workers' desired or reservation wages stopped putting downward pressure on actual wage growth. In practice, though, attendant changes in labor productivity and price inflation (together with the usual issues that surround statistical filtering exercises) will make this inference extremely difficult, especially in real time.

procedures will not be able to identify which shock has hit, and will also fail to provide a clean read on the implications for wage and price dynamics. In other words, we can still identify a combined reallocation shock from observed movements in vacancies and unemployment, but it will have importantly different effects on inflation depending on which of its components predominates. However, these shocks *can* be disentangled—and their inflation consequences separately explored—if we bring additional information to bear.

We find evidence for two types of reallocation shocks that are associated with changes in the noncyclical rate of unemployment (NCRU) and that do in fact have different effects on inflation: Increased reallocation that is associated with job loss (job-loss reallocation) lowers wage growth, while reallocation that involves higher quits (quits reallocation) raises it. Note that only quits reallocation shocks create the positive association between inflation and the unemployment rate that a user of the Phillips curve would interpret as a change in the NCRU or alternatively with the unemployment rate that is consistent with the absence of wage and price pressures.[9] Our results carry important implications for standard approaches used to measure the NCRU: Inasmuch as frictional unemployment and structural unemployment have different effects on inflation, we won't be able to identify a change in the NCRU using the Phillips curve alone. And this problem will be exacerbated if the noncyclical structural shocks induce similar Phillips correlations.

Overall, we find that shocks to aggregate activity are the main source of labor market fluctuations over the business cycle, and that these shocks result in a persistent (and statistically significant) increase in both wage and price inflation.[10] (We associate aggregate activity shocks with changes in aggregate demand; hence, shocks like these are informative about the slope of the Phillips curve.) Our results indicate that the Phillips curve is alive and well, with the data revealing an economically and statistically significant cyclical tradeoff between real activity and inflation. Notably, however, job-loss reallocation shocks and labor supply shocks each imply an unemployment–inflation tradeoff that is comparable in size to the one observed for cyclical shocks. Meanwhile, quits reallocation shocks, as previously noted, create a positive correlation between unemployment and inflation. Put differently, the observed relation

[9]Conceptually, the noncyclical unemployment rate combines the frictional and structural unemployment rates. (For more background on the concepts involved, see `https://research.stlouisfed.org/publications/page1-econ/2016/02/01/making-sense-of-unemployment-data/`.) We do not take a stance on whether the NCRU should coincide with Friedman's original definition of the natural rate of unemployment as the unemployment rate that would obtain at the Walrasian equilibrium of the labor market (Friedman, 1968), or even whether this rate should matter for wage and price determination. Rather, we are interested in how an empirically estimated NCRU is linked to observed inflation dynamics.

[10]The importance of the aggregate activity shock as a source of labor market dynamics is in line with Abraham and Katz (1986).

between the unemployment rate and inflation will depend on what sort of structural shock is hitting the labor market, so any change in the incidence or size of specific structural shocks will cause an estimated reduced-form Phillips correlation to vary.

Applying our framework to the pandemic period, we find that an unprecedented combination of three structural shocks—to aggregate activity, to job-loss reallocation, and to labor supply—drove observed labor market dynamics. Of note, job-loss reallocation shocks were important drivers of both the unemployment rate and the labor force participation rate over this period.[11] These shocks, then, were mainly responsible for the rapid outward shift in the Beveridge curve during the pandemic recession and the apparent inward shift during the subsequent recovery. The portion of the unemployment rate driven by job-loss reallocation shocks rose rapidly at the onset of the pandemic and then fell gradually; in 2022 and 2023, it was similar to its pre-pandemic level. This portion, combined with small contributions from shocks to wage growth that themselves align with the changes in reservation wages observed over this period, implies a path for the NCRU that gradually declines following the pandemic recession.[12] In addition, the job-loss reallocation shocks contributed to the rapid decline and subsequent recovery in the labor force participation rate and put downside pressure on wage and price inflation. Relative to the job-loss reallocation shocks, the quits reallocation shocks played only a small role.

Finally, aggregate activity shocks drove the robust recovery of the participation rate, the unemployment rate, and job postings after the pandemic recession; by the end of 2023, the estimated contribution of the aggregate activity shocks to unemployment implied that the cyclical state of the labor market had returned to where it was on the eve of the pandemic.[13] For the labor force participation rate, negative labor supply developments (such as early retirements) put downward pressure on the participation rate through the end of 2023, masking the boost from the aggregate activity shocks.

Our analysis speaks to the 2022 debate between Blanchard, Domash, and Summers (BDS) and Figura and Waller (FW) over the risk that a reduction in labor demand would cause an outsized rise in the unemployment rate—a so-called hard landing. According to BDS, the relatively flat Beveridge relation (in

[11] This parsing is consistent with recent studies that find that the pandemic resulted in significant and persistent job reallocation (e.g., Barrero et al., 2020; Barrero et al., 2021).

[12] The New York Fed's Survey of Consumer Expectations collects data on reservation wages. See `https://www.newyorkfed.org/microeconomics/sce/labor#` for details.

[13] Our estimate of the contribution of aggregate activity shocks to unemployment is conceptually similar to the unemployment-rate gap that Crump, Eusepi, Giannoni, and Şahin (2024) estimate using a very different methodology. Both our study and theirs find that in 2022 and 2023 the level of aggregate activity was holding down the unemployment rate by more than it did prior to the pandemic. However, our estimate indicates less cyclical tightness in the labor market than what the values from Crump et al. (2024) would suggest.

vacancy–unemployment space) that emerged during the pandemic signalled that even a modest reduction in labor demand would be associated with a large increase in the unemployment rate. By contrast, FW claimed that the observed pattern of vacancies and unemployment resulted from a surge in inflows to unemployment, and that the slope of the Beveridge curve itself was steeper than what simple plots of the data suggested. As a result, FW concluded that a drop in labor demand didn't have to yield a dramatic rise in the unemployment rate, making a *soft* landing more likely. Resolving this debate hinges on identifying the slope of the Beveridge curve—the portion of vacancy and unemployment rate movements attributable to cyclical shocks—along with the contributions of the various structural shocks that would cause the Beveridge curve to shift.[14]

Our methodology supports the conclusion of FW, but for a reason that neither they nor BDS considered; namely, that job-loss reallocation shocks were quantitatively important in the recent episode. During the pandemic, the bulk of the increase in vacancy postings was driven by reallocation shocks, and hence the portion of the rise in the vacancy rate that can be attributed to the economy's strong cyclical position is smaller than what a simple read of the data might suggest. Our historical decomposition shows that in 2022 and 2023 the unwinding of the job-loss reallocation shocks together with a softening in labor demand lowered the vacancy rate but had offsetting effects on the unemployment rate; in particular, reduced job-loss reallocation lowered the unemployment rate at the same time that a slackening of labor demand raised it. These results demonstrate that a full description of the labor market over this period can only be obtained by carefully identifying the underlying shocks that drove its dynamics.

Our analysis also provides a way to empirically link the Beveridge curve with the Phillips curve. Most descriptions of inflation relate it to the amount of "slack" in labor and product markets, where slack is measured as the difference between the actual unemployment rate and a benchmark unemployment rate that is determined by structural features of these markets (an "unemployment gap"). But there is no clear theoretical connection between the equilibrium unemployment rate that falls out of a search-and-matching model and the benchmark rate that is relevant for price and wage determination. And while we might think that changes in the unemployment rate that arise from noncyclical sources should have some connection with this benchmark rate, in practice empirical estimates of these changes obtained

[14] Given the importance of the topic, numerous studies have examined the key drivers of Beveridge curve dynamics over business-cycle episodes (e.g., Elsby, Michaels, and Ratner, 2015, Barnichon and Figura, 2015, Ahn and Crane, 2020, Barlevy, Faberman, Hobijn, and Şahin, 2023). Our own approach uses a novel statistical framework to identify these drivers and to interpret the observed behavior of the labor market over history (including during and after the pandemic period).

from conventional approaches appear not to matter for inflation. As we show, a more careful accounting of Beveridge curve dynamics does in fact yield noncyclical shocks that influence wage and price inflation; by implication, the observed slope of a reduced-form Phillips curve over a given period will depend on the incidence and magnitude of these labor market shocks.

The remainder of the paper includes the following material. Section 1 introduces a modified algorithm for estimating a Bayesian sign-restricted VAR; Section 2 describes our empirical model; and Section 3 discusses the data that we use and what we specify as our priors. Section 4 presents estimation results for the pre-pandemic period, and Section 5 gives a detailed description of the sources of labor market dynamics during the pandemic recession and subsequent recovery. Finally, Section 6 discusses various implications of our findings.

1 A modified estimation algorithm for Bayesian sign-restricted VARs

In a 2015 paper, Baumeister and Hamilton (BH) propose an analytical characterization for Bayesian inference in a structural VAR model, together with an estimation algorithm. Here we briefly review their method and introduce the modified approach and estimation procedure that we use.[15]

1.1 The Baumeister–Hamilton method for estimating sign-restricted VARs

Consider a structural model of the form:

$$\mathbf{A}\mathbf{y}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u}_t \tag{1.1}$$

where $\mathbf{y}_t$ is an $(n \times 1)$ vector of observed variables at time t, $\mathbf{A}$ is an $(n \times n)$ matrix characterizing contemporaneous structural relationships among the elements in $\mathbf{y}_t$, $\mathbf{x}_{t-1}$ is a $((mn+1) \times 1)$ vector containing a constant and m lags of $\mathbf{y}$, and $\mathbf{u}_t$ is an $(n \times 1)$ vector of structural shocks with variance matrix $\mathbf{D}$. Note that $\mathbf{A}$ and $\mathbf{B}$ will characterize the causal relationships among variables.

Suppose that $\mathbf{A}$ is invertible, and denote this inverse as $\mathbf{H}$. The structural model (1.1) can then be

[15] Relatedly, work by Nguyen (2022, 2024) extends the methodology of Baumeister and Hamilton (2015) to allow for external instruments and regime changes.

rewritten as the following reduced form:

$$\mathbf{y}_t = \mathbf{\Phi}\mathbf{x}_{t-1} + \boldsymbol{\varepsilon}_t \tag{1.2}$$

where

$$\mathbf{\Phi} = \mathbf{A}^{-1}\mathbf{B} = \mathbf{H}\mathbf{B} \tag{1.3}$$

$$\boldsymbol{\varepsilon}_t = \mathbf{A}^{-1}\mathbf{u}_t = \mathbf{H}\mathbf{u}_t \tag{1.4}$$

$$\mathbf{\Omega} = E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t') = (\mathbf{A}^{-1})\mathbf{D}(\mathbf{A}^{-1})' = \mathbf{H}\mathbf{D}\mathbf{H}'. \tag{1.5}$$

The reduced-form parameters are estimated from:

$$\begin{aligned}
\hat{\mathbf{\Phi}}_T &= \left(\sum_{t=1}^{T} \mathbf{y}_t\mathbf{x}_{t-1}'\right)\left(\sum_{t=1}^{T} \mathbf{x}_{t-1}\mathbf{x}_{t-1}'\right)^{-1} \\
\hat{\mathbf{\Omega}}_T &= T^{-1}\sum_{t=1}^{T} \hat{\boldsymbol{\varepsilon}}_t\hat{\boldsymbol{\varepsilon}}_t' \\
\hat{\boldsymbol{\varepsilon}}_t &= \mathbf{y}_t - \hat{\mathbf{\Phi}}_T\mathbf{x}_{t-1}.
\end{aligned}$$

BH show how a prior density that a researcher has about the structural parameters, denoted as $p(\mathbf{A}, \mathbf{D}, \mathbf{B})$, can be updated once the researcher observes the data $\mathbf{Y}_T = (\mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_T)$. First, a prior on $\mathbf{A}$ is specified; this captures prior information on the contemporaneous structural coefficients, such as each parameter's sign. Next, a prior on $p(\mathbf{D}, \mathbf{B}|\mathbf{A}) = p(\mathbf{D}|\mathbf{A})p(\mathbf{B}|\mathbf{A}, \mathbf{D})$ is formed using natural conjugate priors for $\mathbf{B}$ and $\mathbf{D}$. Specifically, BH employ gamma priors for the reciprocals of the diagonal elements of $\mathbf{D}$ conditional on $\mathbf{A}$, and normal priors for the coefficients in $\mathbf{B}$. They derive the likelihood function of $\mathbf{Y}_T$ conditional on $\mathbf{A}$,$\mathbf{D}$, and $\mathbf{B}$ under the assumption of Gaussian residuals; their Proposition 1 demonstrates that the posterior distribution $p(\mathbf{A}, \mathbf{D}, \mathbf{B}|\mathbf{Y}_T)$ can be written as $p(\mathbf{A}|\mathbf{Y}_T)p(\mathbf{D}|\mathbf{A}, \mathbf{Y}_T)p(\mathbf{B}|\mathbf{A}, \mathbf{D}, \mathbf{Y}_T)$.

The BH estimation algorithm involves three steps. In the first step, a researcher specifies their priors for $\mathbf{A}$,$\mathbf{D}$, and $\mathbf{B}$ (BH use truncated or full Student t distributions to characterize the prior distribution of $\mathbf{A}$). Priors for the reciprocals of the diagonal elements of $\mathbf{D}$ conditional on $\mathbf{A}$ are specified as inverse-gamma distributions, with prior means chosen to be equal to the reciprocals of the diagonal elements of $\mathbf{A}\hat{\mathbf{S}}\mathbf{A}'$ (where $\hat{\mathbf{S}}$ is the sample variance matrix of estimated residuals that are obtained from univariate AR models). A normal distribution is used as the prior for the structural coefficients $\mathbf{B}$ conditional on

$\mathbf{A}$ and $\mathbf{D}$, $p(\mathbf{B}|\mathbf{A},\mathbf{D})$; following Doan, Litterman, and Sims (1984), the priors on the lag coefficients in $\mathbf{B}$ shrink toward zero as the lag length increases. Then, in the second step, the researcher computes the target function that characterizes the joint distribution of $\mathbf{A}$,$\mathbf{B}$, and $\mathbf{D}$.[16]

In the third step, the posterior distributions are computed as follows. First, the parameters constituting the $\mathbf{A}$ matrix, denoted by $\boldsymbol{\alpha}$, are generated from the posterior distribution of $\boldsymbol{\alpha}$ given the data using a random-walk Metropolis–Hastings algorithm. Given the posterior draws of $\boldsymbol{\alpha}$ and the data, draws of the inverse of the diagonal elements in $\mathbf{D}$ are generated from a gamma distribution. Next, conditional on the posterior draws of $\boldsymbol{\alpha}$, $\mathbf{D}$, and the data, draws of $\mathbf{B}$ are generated from a normal distribution. With the posterior draws of $\mathbf{A}$,$\mathbf{B}$, and $\mathbf{D}$, one can also make inferences on impulse responses and compute historical decompositions and variance decompositions in a way that is analogous to a frequentist approach (see Baumeister and Hamilton, 2018).

Note that BH's algorithm forms priors on the elements of interest in $\mathbf{A}$.[17] In spite of this flexible feature, the algorithm is not always applicable if a researcher has specific priors on parameters in the impact matrix $\mathbf{A}^{-1}(\equiv \mathbf{H})$. One can still form priors on the parameters in $\mathbf{A}$ after inverting the impact matrix: Baumeister and Hamilton (2018) and Baumeister and Hamilton (2019) illustrate how a researcher can find restrictions on $\mathbf{H}$ that are compatible with those on $\mathbf{A}$, thereby allowing them to apply the original Baumeister–Hamilton method. However, the priors on $\mathbf{A}$ often become intractable after the matrix is inverted; this is particularly likely to occur when the impact matrix is large, or when the priors on the parameters in the impact matrix have complicated functional relationships. So it is more straightforward and convenient to form prior beliefs on the parameters in the impact matrix itself. We show how this can be done in the next section.

1.2 The modified Baumeister–Hamilton method

This section describes the modified algorithm that we use. The analytical characterization is closely based on the Baumeister–Hamilton method, but is adapted to our purpose.[18]

As long as $\mathbf{H}$ is invertible, there exists a unique $\mathbf{A}$; in other words, there is a one-to-one mapping

[16] From the target function, one can estimate the initial guess for the posterior mean of $\boldsymbol{\alpha}$ (the parameters constituting the $\mathbf{A}$ matrix) and the scale of the posterior distribution of $\boldsymbol{\alpha}$.

[17] This approach is somewhat different from previous empirical studies that use sign-restricted VARs (e.g., Baumeister and Peersman, 2013).

[18] Since the algorithm is an extension of the one presented in Baumeister and Hamilton (2015), our exposition is mostly borrowed from that paper.

between the distribution of $\mathbf{H}$ and that of $\mathbf{A}$. Therefore, given a candidate draw of $\mathbf{H}$ we can find the corresponding $\mathbf{A}$, which in turn allows us to modify the BH algorithm in a way that lets us characterize our prior beliefs on $\mathbf{H}$.

Specifically, Proposition 1 in Baumeister and Hamilton (2015) is modified as follows. Let $p(\mathbf{H},\mathbf{D},\mathbf{B})$ summarize a researcher's prior beliefs about the structural parameters. The prior is characterized by

$$p(\mathbf{H},\mathbf{D},\mathbf{B}) = p(\mathbf{H})\prod_{i=1}^{n}[p(d_{ii}|\mathbf{H})p(b_i|\mathbf{D},\mathbf{H})], \tag{1.6}$$

where $p(\mathbf{H})$ is a prior distribution of parameters in $\mathbf{H}$ and d_{ii} is the ith diagonal element of $\mathbf{D}$. Natural conjugate distributions are employed for the priors of $\mathbf{D}$ and $\mathbf{B}$. For the reciprocal of d_{ii}, a $\Gamma(\kappa_i,\tau_i)$ prior is specified and each diagonal element is taken to be independent across equations. The distribution $p(\mathbf{D}|\mathbf{H})$ is described by:

$$p(\mathbf{D}|\mathbf{H}) = \prod_{i=1}^{n} p(d_{ii}|\mathbf{H})$$

$$\prod_{i=1}^{n} p(d_{ii}^{-1}|\mathbf{H}) = \begin{cases} \frac{\tau_i^{\kappa_i}}{\Gamma(\kappa_i)}(d_{ii}^{-1})^{\kappa_i-1}\exp(-\tau_i d_{ii}^{-1}) & \text{for} \quad d_{ii}^{-1}\geq 0 \\ 0 \qquad \text{otherwise.} \end{cases}$$

The prior mean of d_{ii}^{-1} is κ_i/τ_i and its variance is $\kappa_i/(\tau_i)^2$.[19]

Normal priors are adopted for the coefficients in $\mathbf{B}$. Let $\mathbf{b}_i'$ denote the ith row of $\mathbf{B}$. If the coefficients are taken independently from $N(\mathbf{m}_i, d_{ii}\mathbf{M}_i)$, we have

$$p(\mathbf{B}|\mathbf{D},\mathbf{H}) = \prod_{i=1}^{n} p(\mathbf{b}_i|\mathbf{D},\mathbf{H}),$$

$$p(\mathbf{b}_i|\mathbf{D},\mathbf{H}) = \frac{1}{(2\pi)^{k/2}|d_{ii}\mathbf{M}_i|^{1/2}} \times \exp[-(1/2)(\mathbf{b}_i-\mathbf{m}_i)'(d_{ii}\mathbf{M}_{ii})^{-1}(\mathbf{b}_i-\mathbf{m}_i)],$$

where the prior mean for $\mathbf{b}_i$ is $\mathbf{m}_i$ and the variance associated with the prior is given by $d_{ii}\mathbf{M}_{ii}$.

The likelihood value $p(\mathbf{Y}_T|\mathbf{H},\mathbf{D},\mathbf{B})$ is written as

$$\begin{aligned} p(\mathbf{Y}_T|\mathbf{H},\mathbf{D},\mathbf{B}) \quad &= \quad (2\pi)^{-Tn/2}|\det(\mathbf{H}^{-1})|^{T}|\mathbf{D}|^{-T/2} \\ &\quad \times[-(1/2)\sum_{t=1}^{T}(\mathbf{H}^{-1}\mathbf{y}_t-\mathbf{B}\mathbf{x}_{t-1})'\mathbf{D}^{-1}(\mathbf{H}^{-1}\mathbf{y}_t-\mathbf{B}\mathbf{x}_{t-1})], \end{aligned}$$

[19] As in Baumeister and Hamilton (2015), κ_i and τ_i can be a function of $\mathbf{H}$. In the implementation algorithm, we consider a case where τ_i is a function of $\mathbf{H}$.

and the posterior distribution of the structural parameters $p(\mathbf{H},\mathbf{D},\mathbf{B}|\mathbf{Y}_T)$ is given by

$$p(\mathbf{H},\mathbf{D},\mathbf{B}|\mathbf{Y}_T) = p(\mathbf{H}|\mathbf{Y}_T)\,p(\mathbf{D}|\mathbf{H},\mathbf{Y}_T)\,p(\mathbf{B}|\mathbf{H},\mathbf{D},\mathbf{Y}_T). \tag{1.7}$$

Let $\tilde{\mathbf{Y}}_i$ and $\tilde{\mathbf{X}}_i$ be defined as follows:

$$\tilde{\mathbf{Y}}_i = [\mathbf{y}_1'\mathbf{a}_i \quad \cdots \quad \mathbf{y}_T'\mathbf{a}_i \quad \mathbf{m}_i'\mathbf{P}_i]' \tag{1.8}$$

$$\tilde{\mathbf{X}}_i = [\mathbf{x}_0' \quad \cdots \quad \mathbf{x}_{T-1}' \quad \mathbf{P}_i]' \tag{1.9}$$

for $\mathbf{a}_i$ the ith row of $\mathbf{H}^{-1}$ and for $\mathbf{P}_i$ the Cholesky factor of $\mathbf{M}_i^{-1}$. The analytical characterizations of the posterior distributions on the right-hand side of equation (1.7) are:

$$p(\mathbf{H}|\mathbf{Y}_T) = \frac{k_T\,p(\mathbf{H})[\det(\mathbf{H}^{-1}\hat{\boldsymbol{\Omega}}_T\mathbf{H}^{-1\prime})]^{T/2}}{\prod_{i=1}^{n}[(2\tau_i^*/T)]^{\kappa_i^*}}\prod_{i=1}^{n}\left\{\frac{|\mathbf{M_i^*}|^{1/2}}{|\mathbf{M_i}|^{1/2}}\frac{\tau_i^{\kappa_i}}{\Gamma(\kappa_i)}\Gamma(\kappa_i^*)\right\}, \tag{1.10}$$

$$p(\mathbf{D}|\mathbf{H},\mathbf{Y}_T) = \prod_{i=1}^{n}\gamma(d_{ii}^{-1};\kappa_i^*,\tau_i^*),$$

$$p(\mathbf{B}|\mathbf{H},\mathbf{D},\mathbf{Y}_T) = \prod_{i=1}^{n}\phi(\mathbf{b}_i;\mathbf{m}_i^*,d_{ii}\mathbf{M}_i^*),$$

where

$$\begin{aligned}
\mathbf{m}_i^* &= (\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i')^{-1}(\tilde{\mathbf{X}}_i'\tilde{\mathbf{Y}}_i),\\
\mathbf{M}_i^* &= (\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i')^{-1},\\
\kappa_i^* &= \kappa_i + (T/2)\\
\tau_i^* &= \tau_i + (\zeta_i^*/2)\\
\zeta_i^* &= (\tilde{\mathbf{Y}}_i'\tilde{\mathbf{Y}}_i) - (\tilde{\mathbf{Y}}_i'\tilde{\mathbf{X}}_i)(\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i')^{-1}(\tilde{\mathbf{X}}_i'\tilde{\mathbf{Y}}_i),
\end{aligned}$$

and k_T is the constant for which equation (1.10) integrates to one.

The estimation algorithm follows BH's algorithm closely, except for the alterations needed to incorporate our modification to their method. The algorithm involves setting parameters for the prior distribution;

computing the target function; and computing the posterior distribution. (Section A.2 of the appendix provides the technical details.)

1.3 Computing the impulse responses

Here we describe how impulse response functions are computed under the modified algorithm.

Rewrite the reduced-form model (1.2) as

$$\mathbf{y}_t = \boldsymbol{c} + \boldsymbol{\Phi}_1 \boldsymbol{y}_{t-1} + \boldsymbol{\Phi}_2 \boldsymbol{y}_{t-2} + \cdots + \boldsymbol{\Phi}_p \boldsymbol{y}_{t-p} + \boldsymbol{\varepsilon}_t, \tag{1.11}$$

where $\boldsymbol{c}$ denotes an $(n \times 1)$ vector of constants and $\boldsymbol{\Phi}_j$ an $(n \times n)$ matrix of autoregressive coefficients for $j = 1, 2, \cdots, p$. Collecting $\boldsymbol{\Phi}_j$'s in $\boldsymbol{\Phi}$ such that $\boldsymbol{\Phi} = [\boldsymbol{\Phi}_{\mathbf{1}}, \boldsymbol{\Phi}_{\mathbf{2}}, \cdots]$, we have

$$\boldsymbol{HB} = \boldsymbol{\Phi}.$$

Equation (1.11) can be expressed as a sum of the history of $\boldsymbol{\varepsilon}$ in a vector $MA(\infty)$ representation,

$$\mathbf{y}_t = \boldsymbol{\mu} + \boldsymbol{\varepsilon}_t + \boldsymbol{\Psi}_1 \boldsymbol{\varepsilon}_{t-1} + \boldsymbol{\Psi}_2 \boldsymbol{\varepsilon}_{t-2} + \cdots.$$

The nonorthogonalized impulse-response function at horizon s is

$$\boldsymbol{\Psi}_s = \frac{\partial \mathbf{y}_{t+s}}{\partial \boldsymbol{\varepsilon}'_t}.$$

Note that

$$\begin{aligned}
\boldsymbol{\Psi}_0 &= \boldsymbol{I}_n \\
\boldsymbol{\Psi}_1 &= \boldsymbol{\Phi}_1 \\
\boldsymbol{\Psi}_2 &= \boldsymbol{\Phi}_1 \boldsymbol{\Psi}_1 + \boldsymbol{\Phi}_2 \\
&\cdots \\
\boldsymbol{\Psi}_s &= \boldsymbol{\Phi}_1 \boldsymbol{\Psi}_{s-1} + \boldsymbol{\Phi}_2 \boldsymbol{\Psi}_{s-2} + \cdots + \boldsymbol{\Phi}_p \boldsymbol{\Psi}_{s-p} \qquad (1.12)
\end{aligned}$$

for $s = 1, 2, \cdots$. (For $s < 0$, $\boldsymbol{\Psi}_s = 0$.) The *structural* impulse-response functions are given by

$$\boldsymbol{J}_s = \boldsymbol{\Psi}_s \boldsymbol{H}, \tag{1.13}$$

with the jth column of $\boldsymbol{J}_s$ capturing the effect of the jth structural shock s-periods ahead.

In the modified Bayesian SVAR, these structural impulse responses are computed as follows. With the triple of posterior draws $\{\boldsymbol{H}(\boldsymbol{h}^{(l)}), \boldsymbol{D}^{(l)}, \boldsymbol{B}^{(l)}\}_{l=L+1}^{2L}$ of a sample size L drawn from the posterior distribution $p(\boldsymbol{H}, \boldsymbol{D}, \boldsymbol{B}|\boldsymbol{Y}_T)$ (after discarding the first L draws), compute

$$\boldsymbol{\Phi}^{(l)} = \boldsymbol{H}(\boldsymbol{h}^{(l)})\boldsymbol{B}^{(l)}.$$

For each l, the impulse responses with respect to structural shocks at horizon $s = 0, 1, \cdots, S$ are then computed from

$$\boldsymbol{J}_s^{(l)} = \boldsymbol{\Psi}_s^{(l)} \boldsymbol{H}^{(l)}.$$

With the collection of $\{\boldsymbol{J}_s^{(l)}\}_{l=L+1}^{2L}$, Baumeister and Hamilton (2018) show how one can form inferences about impulse responses in a similar setting.

2 Linking the Beveridge curve and the Phillips curve

This section describes a Bayesian sign-restricted VAR model that allows us to empirically link Beveridge curve dynamics with the dynamics of wage and price inflation. We start by reviewing Blanchard and Diamond's empirical model of the Beveridge curve, and then propose a new model in which their single reallocation shock takes two different forms depending on whether the shock raises the unemployment of job losers or that of job quitters. The connection with inflation dynamics is directly explored by including wage and price inflation in the VAR system.

2.1 The Blanchard and Diamond (1989) model

Blanchard and Diamond's model identifies three structural shocks—a reallocation shock, an aggregate activity shock, and a labor supply shock—and analyzes the dynamic effects of each shock on unemployment, vacancies, and the labor force. The structural shocks are identified from the sign of the effect that each one has on the three labor market variables in the period it hits. Table 1 summarizes the identification scheme.[20]

Table 1: IDENTIFICATION OF STRUCTURAL SHOCKS IN THE BLANCHARD–DIAMOND MODEL

	To a positive shock to …		
	Reallocation	Aggregate activity	Labor supply
Response of the following variables …			
Labor force participation rate	-	+	+
Vacancy rate	+	+	**0**
Unemployment rate	+	-	+

A positive reallocation shock destroys jobs in some sectors but creates jobs in other sectors, thereby raising the unemployment and vacancy rates but lowering the labor force participation rate. A positive aggregate activity shock raises vacancy postings and labor force participation, and lowers the unemployment rate. Finally, a positive labor supply innovation raises the participation and unemployment rates, with no effect on the vacancy rate at impact.[21]

These identifying assumptions are summarized as follows:

$$\begin{bmatrix} l_t \\ v_t \\ u_t \end{bmatrix} = \underbrace{\begin{bmatrix} -\theta & \delta & 1 \\ \beta & 1 & 0 \\ \gamma & -\epsilon & \omega \end{bmatrix}}_{\mathbf{H}} \begin{bmatrix} u_t^r \\ u_t^c \\ u_t^s \end{bmatrix} \tag{2.1}$$

where l_t, v_t, and u_t are residuals from a reduced-form VAR that includes the labor force participation rate (L_t), the vacancy rate (V_t), and the unemployment rate (U_t), and where u_t^r, u_t^c, and u_t^s denote the reallocation, aggregate activity ("cyclical"), and labor supply shocks, respectively. The parameters θ, β, and γ capture the contemporaneous effect of a reallocation shock on the labor force, vacancies, and unemployment; the parameters δ and ϵ capture the contemporaneous effect of an aggregate activity shock on the labor force and unemployment; and the parameter ω captures the effect of a labor supply

[20] Note that we use the unemployment, vacancy, and labor force participation *rates*, rather than the transformed *levels* that Blanchard and Diamond used.

[21] For more details on the identifying assumptions, see Blanchard and Diamond (1989).

shock on unemployment. The parameters themselves are all defined to be positive.[22]

2.2 Two unemployment states with wage and price inflation

We extend the Blanchard–Diamond approach along two dimensions. First, we separately consider the unemployment rate of job quitters (U_t^q) and that of other unemployed individuals (U_t^l), as well as the labor force participation rate (L_t) and the vacancy rate (V_t).[23] The second unemployed group, U_t^l, is composed of job losers and entrants to the labor force.[24] We interpret U_t^l as largely representing job losers, because job losers account for most of this category while those who return to the labor force tend to be individuals who previously dropped out of the labor force after losing a job (Ahn, 2023). As defined, the sum of the two unemployed groups yields aggregate unemployment. These two categories of unemployed workers let us identify two types of reallocation shocks, each of which likely results in different labor market outcomes and each of which has different implications for the NCRU. In particular, the quits reallocation shock can result in job seekers who experience frictional unemployment when switching jobs, while the job-loss reallocation shock can result in job seekers who experience a spell of structural unemployment owing to inefficient job searches, worker–job mismatch, or a transition into long-term unemployment.

Our second extension is to include measures of wage growth (W_t) and price inflation (Π_t) in the model, which lets us examine the effect of each structural shock on inflation within a coherent Bayesian SVAR framework. The identifying assumptions for this extended model are summarized in Table 2.[25]

With these extensions, the dependent variables ($\mathbf{y}_t$), structural shocks ($\mathbf{u}_t$), and impact matrix ($\mathbf{H}$) can be written as:

$$\begin{aligned}
\mathbf{y}_t &= [L_t, V_t, U_t^l, U_t^q, W_t, \Pi_t]' \\
\mathbf{u}_t &= [u_t^l, u_t^q, u_t^c, u_t^s, u_t^w, u_t^\pi]'
\end{aligned}$$

[22] Section A.3 of the appendix reports estimation results for this three-variable model using the modified Baumeister–Hamilton procedure.

[23] The labor force participation rate is defined to be the civilian labor force as a share of the population; the remaining variables are shares of the civilian labor force.

[24] Here, we classify people whose temporary job ended as job losers.

[25] The identification scheme in Table 2 explicitly imposes a number of zero restrictions. We view these zero restrictions to be reasonable in a monthly model like ours on the grounds that the presence of nominal rigidities should dampen most of the within-month effects that the own shocks to wages and prices would have on the labor market. However, we do permit the own shocks to wages and prices to have within-month effects on the unemployment rate for job losers; these effects could reflect such things as changes in the minimum wage or wage expectations, and possibly the emergence of supply bottlenecks.

Table 2: IDENTIFICATION OF STRUCTURAL SHOCKS IN THE EXTENDED MODEL

	To a positive shock to …					
	Job-loss reallocation	Quits reallocation	Aggregate activity	Labor supply	Wage -specific	Price -specific
Response of the following variables …						
Labor force participation rate	-	-	+	+	**0**	**0**
Vacancy rate	+	+	+	**0**	**0**	**0**
Unemployment rate (job losers)	+	**0**	-	+	+	**?**
Unemployment rate (job quitters)	**0**	+	**?**	**0**	**0**	**0**
Wage inflation	**?**	+	+	-	+	**0**
Price inflation	**?**	+	+	-	**0**	+

$$
\mathbf{H} = \begin{bmatrix} -\theta_l & -\theta_q & \delta & 1 & 0 & 0 \\ \beta_l & \beta_q & 1 & 0 & 0 & 0 \\ 1 & 0 & -\epsilon & \omega & \chi_w & \chi_p \\ 0 & 1 & \zeta_c & 0 & 0 & 0 \\ \phi_l & \phi_q & \phi_c & -\phi_s & 1 & 0 \\ \psi_l & \psi_q & \psi_c & -\psi_s & 0 & 1 \end{bmatrix}. \tag{2.2}
$$

As before, the parameters in **H** are themselves all defined to be positive so the signs of the parameters represent the sign restrictions that we impose on the model.

The first two columns characterize the effects on $\mathbf{y}_t$ of the two different reallocation shocks in the period that they hit. The first column captures the effect of u_t^l, the "job-loss reallocation shock." We normalize the effect of this shock on the unemployment rate of job losers to be one (third row, first column), with the parameters θ_l and β_l capturing the effects on the labor force participation rate and the vacancy rate.[26]

The parameters in the second column give the effects on $\mathbf{y}_t$ of a reallocation shock that is associated with job quits (u_t^q)—the "quits reallocation shock"—in the period when it hits. Note that the effect of u_t^q on the unemployment rate of job quitters is normalized to be one at impact but u_t^q does not affect U_t^l at

[26] Normalizing the impact matrix in this way does not restrict the variances of the innovations to the endogenous variables, as the variances of the structural shocks are adjusted to match the variances of the residuals in the reduced-form VAR.

impact; the parameters θ_q and β_q capture the effects of u_t^q on the labor force participation rate and the vacancy rate.

The third column captures the effects of the aggregate activity shock (u_t^c). The shock's effect on the vacancy rate is normalized to be one at impact; the parameter δ gives the effect of u_t^c on L_t. We assume that at impact the magnitude of the activity shock's effect on U_t^l is ϵ, and that its effect on U_t^q is ζ_c. Note that the cyclicality of the unemployment rate for job quitters is likely muted because of two offsetting factors: Quits are themselves procyclical, with the number of newly unemployed job quitters rising during an economic expansion, while the job-finding rate of unemployed job quitters also rises during an expansion, thus lowering their unemployment rate. Because the sign of the combined effect on unemployment is ambiguous, we allow for this uncertainty when we specify our prior for ζ_c.

The fourth column captures the effects of a labor supply shock (u_t^s). The shock's effect on the participation rate is normalized to be one. The parameter $\omega > 0$ captures the effect on U_t^l as U_t^l includes the unemployment rate of entrants to the labor force, including those who enter the labor force for the first time.

We also estimate the impact effects on wage and price inflation of the two reallocation shocks, the aggregate activity shock, and the labor supply shocks. In this case, we do not impose particular sign restrictions on the impact effects of the job-loss reallocation shocks (ϕ_l and ψ_l). However, we assume that the impact effects of the quits reallocation (ϕ_q and ψ_q) and aggregate activity shocks (ϕ_c and ψ_c) are positive, but that the impact effects of the labor supply shocks (ϕ_s and ψ_s) are negative.[27] We let the data tell us about the effects of job-loss reallocation shocks on wage and price inflation because there is no well-established theory to guide us; however, Moscarini and Postel-Vinay (2023) argue that reallocation driven by quits can raise price inflation, which is reflected in our sign restrictions on ϕ_q and ψ_q.[28]

The last two columns capture the effects on $\mathbf{y}_t$ of shocks specific to wage and price inflation (u_t^w and u_t^π) in the period that they hit. The effect of u_t^w (u_t^π) on W_t (Π_t) at impact is normalized to be one. In addition, we allow these shocks to affect the unemployment rate of job losers in the period that they hit. The impact effect of the wage-specific shock (χ_w) is assumed to be positive on the grounds that an innovation to wages that results from factors such as a rise in reservation wages or in the minimum wage

[27] As a robustness check, we considered a case where the labor supply shock does not affect wage and price inflation at impact (this alternative assumption might be justified on the grounds that nominal rigidities are sufficiently widespread in the economy that they would dampen a within-month effect). The resulting estimates, which are reported in the appendix, are very similar to what we obtain under our baseline assumptions.

[28] We examined whether the estimation results were robust to assuming no effects on impact (that is, ϕ_q and $\psi_q = 0$); they are.

are likely to raise the unemployment rate of job losers as well as the nominal wage. However, we do not restrict the sign of the price-specific shock's effect on the unemployment rate at impact (this is because supply bottlenecks—which were an important factor raising inflation during the pandemic—can result in job losses and so raise the unemployment rate of job losers, while monetary policy actions that *reduce* inflation are also likely to raise the unemployment rate of job losers).

The baseline model is a monthly VAR in the six variables $\mathbf{y}_t$. We use eight lags to ensure that the dynamics of $\mathbf{y}_t$ are sufficiently well captured; our full sample period runs from January 1967 to December 2023. Owing to large transitory swings in many of these series during the pandemic, we first estimate the model for a sample period that ends in December 2019; in Section 5 we provide more detail about how we handle the period that includes the Covid-19 pandemic and subsequent economic recovery.[29]

3 Estimation

The data used for the empirical analyses are detailed in Section 3.1; Section 3.2 documents our choice of priors.

3.1 Data

We use monthly data on the labor force participation rate and the unemployment rate by reason for unemployment from the Current Population Survey. We consider two reasons for unemployment: quits (job leavers) and any other reason including involuntary separation and entrance into the labor force.[30] The unemployment data by reason for unemployment are available starting in January 1967. For data on vacancy postings, we use the composite measure of job vacancies compiled by Barnichon (2010) because this series covers a long sample period and extends prior to the period in which the Job Openings and Labor Turnover Survey (JOLTS) becomes available (December 2000). This series ends in September 2021; we extend it through December 2023 with the growth rate of total vacancy postings from the JOLTS data.

We use the 12-month percent change in average hourly earnings of production and nonsupervisory workers (AHE) as the measure of wage inflation. The AHE data are monthly and available from Jan-

[29] Our model is a stationary VAR model, so any trends in the variables are captured in the residuals. We do not first difference the variables in order to keep the analysis as close as possible to Blanchard and Diamond (1989), and given Christopher Sims's dictum that the I(0) assumption is not necessary for VAR models while a first difference can remove important identifying information (Cochrane, 2005).

[30] Entrance into the labor force includes both new entrance (e.g., college graduates) and re-entrance (e.g., job losers who leave and then return to the labor force).

uary 1965-forward; among the available monthly wage indicators, it is the only series that extends back into the 1960s. AHE is likely affected by shifts in the composition of the workforce; however, to the extent that this sort of bias is unrelated to underlying labor market strength, it can be captured with the model's wage-specific disturbances.[31]

For price inflation, we use the change in the core market-based PCE price index from the Bureau of Economic Analysis (BEA). We use the market-based core as our principal price measure because several nonmarket components of consumption are priced using input-cost indexes that are in turn derived from wage or compensation measures. However, official data for the core market-based PCE price index are only published starting in 1987. To compute a market-based series prior to 1987, we strip out the prices of core nonmarket PCE components from the published overall core PCE price index, where the definition of "nonmarket" and the Fisher aggregation procedure that we use mimic the ones used by the BEA. The market-based PCE inflation series that we use for the extended-sample VARs subtracts out Blinder and Rudd's (2013) estimates of the effects of the Nixon-era price controls; in our robustness checks, we also consider whether the results change importantly when we use a market-based core price index that makes no adjustment for the 1970s controls.[32]

3.2 Prior specification

We characterize the prior for each structural parameter with a Student t distribution whose mode is close to the corresponding calibrated value in Blanchard and Diamond (1989) or to a value that is updated to reflect more-recent observations.[33] For the majority of parameters, the scale parameter is set to 0.3 and three degrees of freedom are used so as to allow for sizable uncertainty around the modal value.[34] For β_l, we use five degrees of freedom and set the scale parameter to be 0.2. This prior moves the model away from overstating the role that job-loss reallocation plays in driving movements in the participation rate, reflecting our belief that these shocks are less likely to be responsible for much of the variation in the participation rate. For χ_p, we employ a prior scale parameter of 0.2 with two degrees of freedom, reflecting our prior belief, albeit with a fair amount of uncertainty, that the impact effect of an own-shock

[31] For this reason, we view the potential composition effects in AHE as less of a concern in estimating the consequences of structural labor market shocks.

[32] See Section A.4.1.

[33] Doing so reduces the number of generated draws that the estimation algorithm rejects.

[34] A Student t distribution with scale parameter 0.3 and three degrees of freedom represents a fairly weak prior belief (Baumeister and Hamilton, 2021).

Table 3: PRIORS ON PARAMETERS IN THE IMPACT MATRIX

Parameter	Distribution	Prior mode	Prior scale	Sign of parameter	Sign in **H**
θ_l	Student t (3 d.f.)	0.5	0.3	+	−
β_l	$\cdots$ (5 d.f.)	0.1	0.2	+	+
ϕ_l	$\cdots$	- 0.2	0.3	NA	NA
ψ_l	$\cdots$	- 0.2	0.3	NA	NA
θ_q	$\cdots$	0.9	0.3	+	−
β_q	$\cdots$	0.2	0.3	+	+
ϕ_q	$\cdots$	-0.1	0.3	+	+
ψ_q	$\cdots$	0.0	0.3	+	+
δ	$\cdots$	0.4	0.3	+	+
ϵ	$\cdots$	0.8	0.3	+	−
ζ_c	$\cdots$	0.01	0.3	NA	NA
ϕ_c	$\cdots$	0.5	0.3	+	+
ψ_c	$\cdots$	0.3	0.3	+	+
ω	$\cdots$	0.5	0.3	+	+
ϕ_s	$\cdots$	0.0	0.3	+	−
ψ_s	$\cdots$	-0.1	0.3	+	−
χ_w	$\cdots$	0.1	0.3	+	+
χ_p	$\cdots$ (2 d.f.)	0.05	0.2	NA	NA

Notes: "d.f." stands for degrees of freedom. "NA" stands for not available.

to price inflation on the unemployment rate of job losers is likely to be small. Last, for ϕ_q and ψ_q we impose prior modes of negative 0.1 despite the sign restrictions in order to give the prior distribution a higher density at zero than it would have with a modal value of zero.[35]

The modal value for each structural parameter's prior is summarized in Table 3. If a parameter enters **H** with a particular sign we impose that sign restriction; as previously mentioned, we do not restrict the signs of ϕ_l and ψ_l.[36]

4 Estimation results

Section 4.1 reports the estimation results for the baseline model, Section 4.2 describes the impulse response functions, and Section 4.3 reports historical decompositions for the period prior to the Covid-19 pandemic.

[35]These priors for β_l, χ_p, ϕ_q and ψ_q help to facilitate the draws to converge to the target distribution.

[36]Note that the determinant of **H** is smaller than negative one, which sidesteps a problem considered by Baumeister and Hamilton (2018) in which the determinant of the impact matrix changes signs (this would lead it to become infinite in allowable regions of the parameter space).

4.1 Posterior distributions of structural parameters

Figure 1 displays the prior and posterior distributions of the model's structural parameters. The posterior distributions capture the (absolute) magnitudes of the structural shocks' contemporaneous effects.

First, the job-loss reallocation shock results in a sizable decrease in the participation rate ($-\theta_l$) and an outward shift of the Beveridge curve (β_l). Second, the quits reallocation shock has effects that are smaller ($-\theta_q$ and β_q) than those of job-loss reallocation shocks. The posterior median of θ_q is smaller than that for θ_l, which suggests that workers who voluntarily leave their jobs are less likely to exit the labor force than workers who lose their jobs. Third, the positive aggregate activity shock has a non-negligible positive effect on the labor force participation rate (δ) and a sizable negative effect on the unemployment rate of job losers ($-\epsilon$).[37] However, the aggregate activity shock has a nearly zero effect on the unemployment rate of job quitters at impact (ζ_c), though ζ_c is not sign-restricted.

The parameter ω captures the contemporaneous effect of a labor supply shock on the unemployment rate of job losers. The shock's impact effect on the unemployment rate is smaller than its impact effect on the participation rate (1). Our interpretation is that a positive labor supply shock likely also involves out-of-the-labor-force workers moving directly into employment.

Next, the impact effects of a job-loss reallocation shock on wage inflation (ϕ_l) and price inflation (ψ_l) are both negative, though these parameters are not sign-restricted. The negative impact effects on wage inflation (ϕ_l) are sizable, and the negative impact effects on price inflation (ψ_l) are smaller but not negligible. Meanwhile, the positive impact effects of a quits reallocation shock on wage inflation (ϕ_q) and price inflation (ψ_q) are both small. The positive impact effects of an aggregate activity shock on wage inflation (ϕ_c) are sizable, though those on price inflation (ψ_c) are small. The differential effects of the activity shock on wage and price inflation are consistent with the empirical fact that wage Phillips curves are typically much steeper than price Phillips curves. The negative impact effects of a labor supply shock on wage inflation (-ϕ_s) and price inflation (-ψ_s) are overall small, but not negligible.

Last, the positive impact effects of a wage-specific shock on the unemployment rate of job losers (χ_w) is minor. Notably, the price-specific shock results in positive and non-negligible impact effects on the unemployment rate of job losers (χ_p), despite the absence of sign restrictions on this parameter.

[37] The posterior distributions of β_l and ϵ are not sensitive to changes in the corresponding priors.

4.2 Impulse responses

Figure 2 reports the impulse responses for each variable following each structural shock. (In the panel titles, "$x < y$" denotes the effect of shock y on variable x.) The first column reports the impulse responses following a job-loss reallocation shock. The job-loss reallocation shock lowers the participation rate persistently (first row), while the shock raises the vacancy rate (second row) and has persistent positive effects on the unemployment rate of job losers (third row). This shock lowers the unemployment rate of job quitters for a few months after the impact, but the statistical significance of the response declines rapidly (fourth row). The job-loss reallocation shock lowers AHE growth persistently (fifth row); the shock's effects on market-based core inflation are negative for about two years and are only statistically significant for a little over a year (sixth row).

The second column reports the impulse responses following a quits reallocation shock. The quits reallocation shock has a small negative effect on the participation rate at impact; further out, the shock leads to a rise in participation (first row). This shock also persistently raises the vacancy rate (second row) and the unemployment rate of job quitters (fourth row). The shock lowers the unemployment rate of job losers, but the effects are not statistically significant (third row). The quits reallocation shock results in a statistically significant increase in AHE growth and price inflation (fifth and sixth rows), which is noticeably different from the effects of a job-loss reallocation shock.

The third column reports the impulse responses that result from an aggregate activity shock. This shock has persistent effects on all of the variables: The aggregate activity shock raises the participation rate, vacancy rate, AHE growth, and price inflation, and lowers both unemployment rates. These effects are statistically significant for quite some time; in particular, the responses for wage and price inflation and the unemployment rate of job losers remain significant 3 to 4 years after the activity shock occurs. The response of wage inflation is a little faster and larger than that of price inflation; this result is also consistent with the wage Phillips curve's being steeper than the price Phillips curve.

The fourth column reports the impulse responses following a labor supply shock. This shock raises the participation rate and both unemployment rates (first, third, and fourth rows), reflecting the increased number of job seekers drawn to the labor force. A labor supply shock lowers wage inflation (fifth row); it also lowers market-based core inflation, but the effects are not statistically significant (sixth row). Effects of labor supply shocks on the vacancy rate are not statistically significant either (second row).

The final two columns report the impulse responses following own shocks to wage growth and price inflation, respectively. Starting with the last column, the own shock to price inflation behaves like a negative shock to the labor market: It lowers the participation rate and the vacancy rate and raises both unemployment rates. In addition, we find a positive pass-through of price inflation to wage growth (bottom row in the sixth column).[38]

The fifth column reports the impulse responses following an own shock to wage inflation. A positive shock to wage growth raises the unemployment rate of job losers (third row) and the vacancy rate (second row). Elsewhere, this shock lowers the participation rate (first row) and the unemployment rate of job quitters (fourth row). The own shock to wage inflation has a positive effect on price inflation, though the magnitude is relatively small (sixth row).[39] Most of these responses have credible sets that do not exclude zero for any appreciable length of time, with the responses of the vacancy rate and the unemployment rate of job quitters serving as exceptions. Qualitatively, much of the effect of the wage shock on the other labor market variables reflects its eventually leading to an increase in price inflation.

[38] Given its negative effect on the labor market, we might interpret the own shock to price inflation as a cost-push shock (such as an oil shock). Alternatively, the effects of the shock might reflect the monetary authority's response to inflation (with higher inflation yielding tighter policy and hence weaker real activity).

[39] Both sets of pass-through results are consistent with those reported by Peneva and Rudd (2017) for an ECI-based measure of unit labor costs.

Figure 1: PRIOR AND POSTERIOR DISTRIBUTIONS FOR STRUCTURAL PARAMETERS

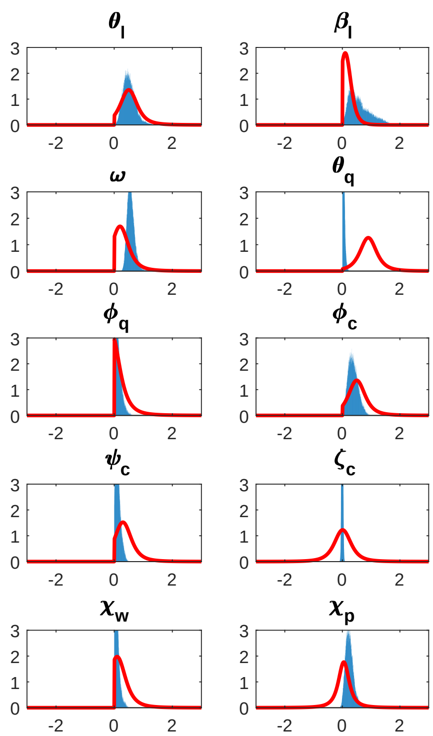

Note: This figure shows the prior and posterior distributions of the structural parameters in **H**. Red lines are prior distributions; blue regions are posterior distributions.
Source: Authors' calculations.

Figure 2: IMPULSE RESPONSES

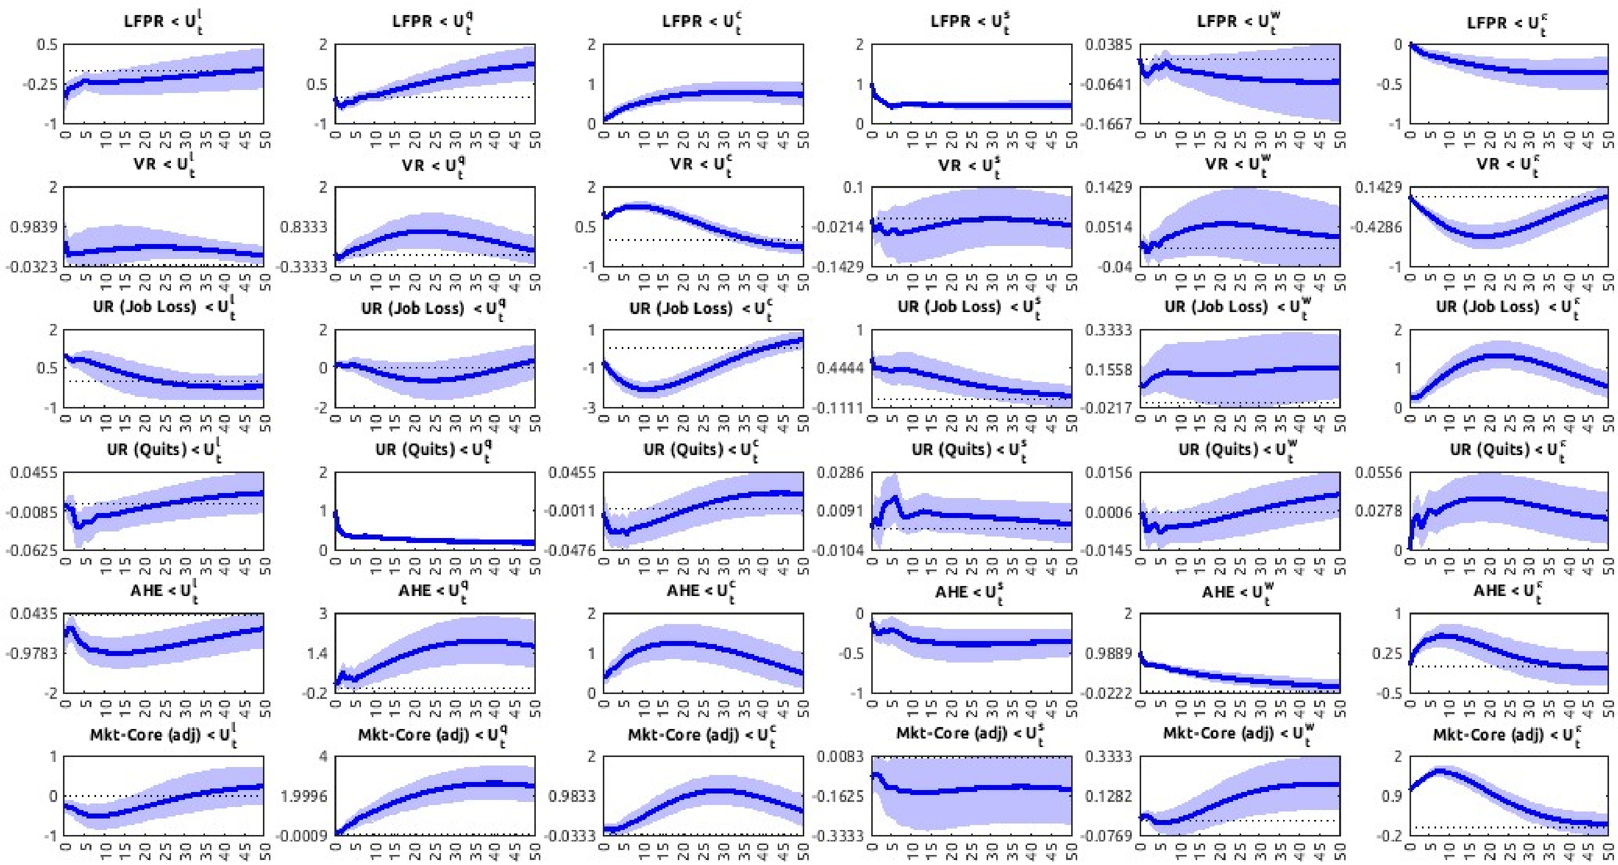

Notes: This figure displays the estimated impulse responses from our baseline system (X-axis: number of months after the shock hits; Y-axis: the magnitude of the response). The posterior median is shown as a solid line; the dotted lines indicate zero; the shaded area gives the 68 percent credible set; and $x < y$ indicates the effect of shock y on variable x. *LFPR* : the labor force participation rate; *VR* : the vacancy rate; *UR (Job Loss)* : the unemployment rate of job losers and entrants to the labor force; *UR (Quits)* : the unemployment rate of job quitters; *AHE* : Average hourly earnings; *Mkt-Core(adj)* : Adjusted market core inflation. U_t^l: job-loss reallocation shock; U_t^q: job-quits reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock; U_t^w: wage-specific shock; U_t^π: inflation-specific shock.
Source: Authors' calculations.

4.3 Historical decompositions

Figure 3 reports the historical decomposition of each variable from August 1967 to December 2019.[40] Overall, the aggregate activity shock (u_t^c) accounts for most of the variation in the labor force participation rate, the vacancy rate, and the job-loser unemployment rate. This result is consistent with much earlier work by Abraham and Katz (1986), who find that aggregate cyclical shocks, not reallocation shocks, are the main driver of labor market fluctuations in the period they consider.

The labor supply shock (u_t^s) is responsible for an important portion of the movements in the participation rate; again, though, the aggregate activity shock is the dominant influence on participation over the business cycle (first column in Figure 3). This suggests that both demand and supply make important contributions to the labor force participation rate, with demand being the main source of cyclical dynamics.

Job-loss reallocation shocks raise the unemployment rate of job losers by about three-quarters of a percentage point during the Great Recession; the labor supply shock raises the unemployment rate of job losers by half of a percentage point during this period. The sum of these magnitudes is broadly in line with the increase in the short-term natural rate that the Congressional Budget Office (CBO) and Federal Reserve Board staff identified during this episode.[41] In this context, the rise in the unemployment rate from the job-loss reallocation shocks likely represents increased reallocational frictions, including mismatch between workers and jobs, while the portion driven by labor supply shocks might reflect the effect of the extended unemployment insurance (UI) benefits that were put in place during this period.[42] If so, the portion of the unemployment rate driven by job-loss reallocation and labor supply shocks would represent a change in structural unemployment over this period.

The unemployment rate of job quitters is mainly driven by the quits reallocation shock; the other structural shocks have only a limited influence (first column in Figure 4). Apparently, then, movements in the unemployment rate of job quitters are driven by factors that are largely independent of those that drive the unemployment rate of job losers.

For AHE growth and market-based core inflation (second and third columns in Figure 4), each series'

[40]Note that the historical decompositions capture the contribution of each structural shock to the dynamics of the endogenous variables, but do not capture the trend of each variable.

[41]The rise in the CBO natural rate is computed from their (now discontinued) short-term natural rate series. The Federal Reserve Board staff estimates come from publicly available material posted on the Federal Reserve Bank of Philadelphia's website.

[42]The extended UI benefits might have motivated long-term unemployed individuals to continue their job searches, raising the unemployment rate.

own shock accounts for most of its variability, with the aggregate activity shock coming in as a distant second.[43] The importance of the aggregate activity shock for the realized inflation and unemployment rates is, of course, consistent with the view that shocks to aggregate activity are responsible for generating a discernable (and statistically significant) Phillips correlation. Finally, it is noteworthy that quits reallocation shocks are important contributors to wage and price inflation from the 1970s to the 1990s. This finding is consistent with a model of the labor market in which wage changes are realized by finding a new employer rather than by negotiating with one's current employer.

4.4 Robustness checks

We examine the robustness of our estimates along the following three dimensions. First, we use market-based core inflation without an adjustment for the effects of the 1970s price controls in order to see whether this adjustment influences our results. Second, we relax the sign restrictions for the effects of a quits reallocation shock on wage and price inflation. The motivation is that some evidence suggests that workers who are prone to involuntary separations or who have unstable jobs are likely to quit their jobs (Ahn, Hobijn, and Şahin, 2023), and these workers may not necessarily experience wage increases after they quit. Overall, the estimation results from these alternative specifications are consistent with the results from our baseline specification. These results are documented in Section A.4 of the appendix.[44]

[43] Results like these are typically found in the empirical Phillips curve literature, even when additional variables are included to explicitly control for supply shocks.

[44] We further consider how the estimates change if we impose additional zero restrictions on the impact effects that own shocks to wages and prices have on the unemployment rate of job losers and zero restrictions on the impact effects that labor supply shocks have on wage and price inflation. This set of results is available upon request.

Figure 3: HISTORICAL DECOMPOSITION (1)

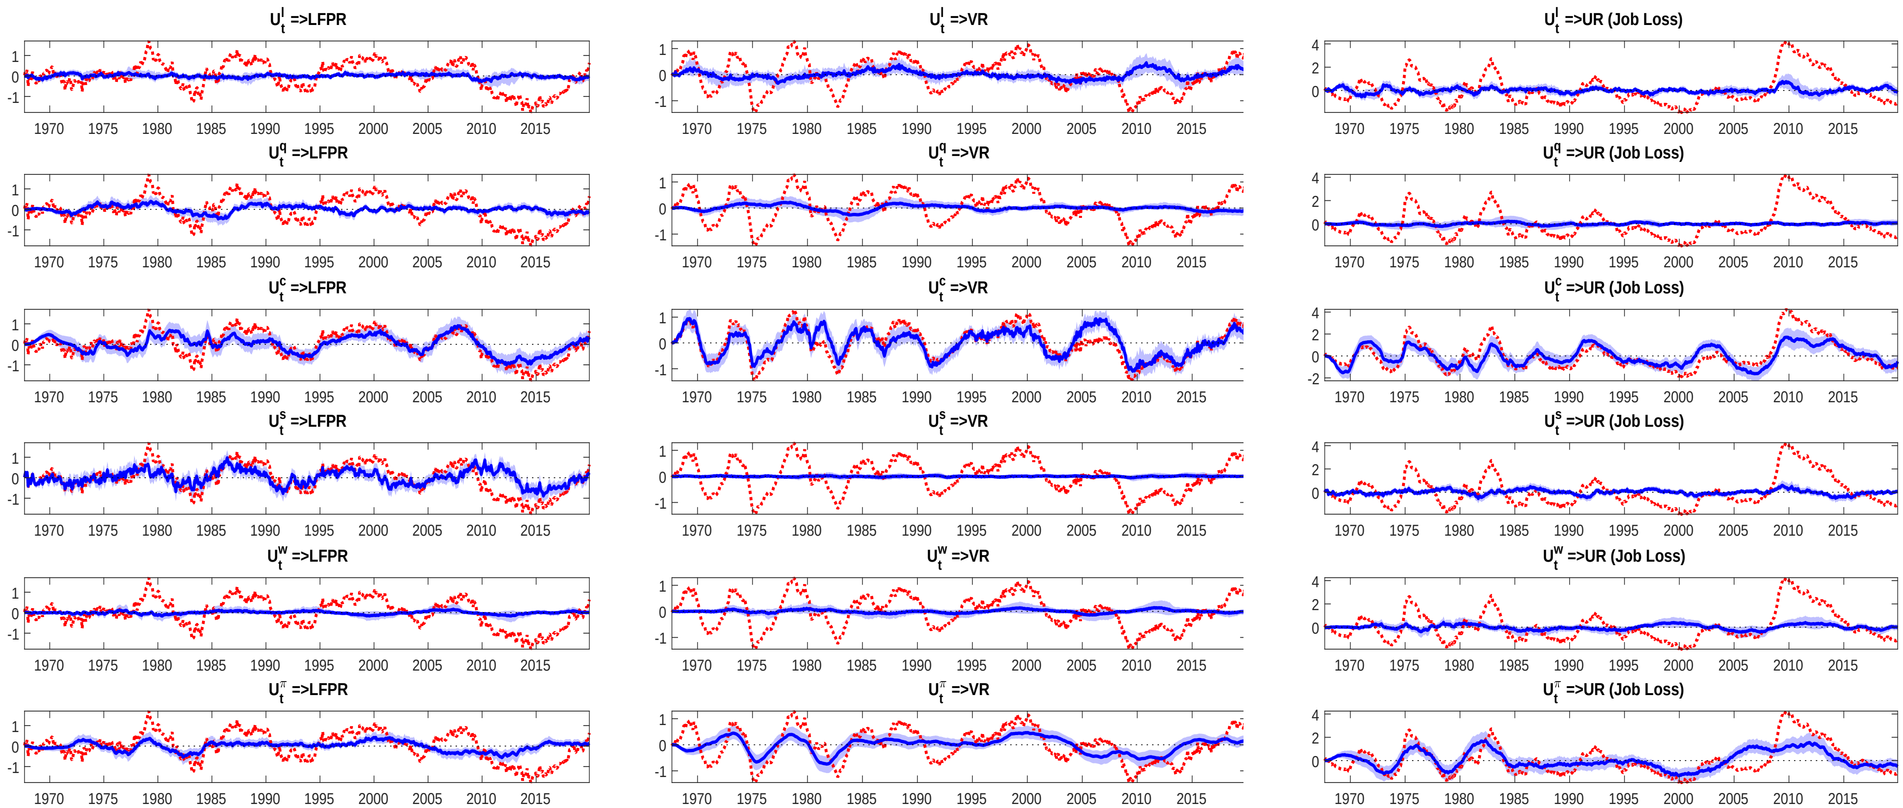

Notes: This figure displays the estimated historical decompositions for each variable (X-axis: calendar time; Y-axis: the magnitude of variations in the endogenous variable). The posterior median is shown as a solid line; the shaded area gives the 68 percent credible set. *LFPR* : the labor force participation rate; *VR* : the vacancy rate; *UR (Job Loss)* : the unemployment rate of job losers and entrants to the labor force; *UR (Quits)* : the unemployment rate of job quitters; *AHE* : Average hourly earnings; *Mkt-Core(adj)* : Adjusted market core inflation. U_t^l: job-loss reallocation shock; U_t^q: job-quits reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock; U_t^w: wage-specific shock; U_t^π: inflation-specific shock.

Source: Authors' calculations.

Figure 4: HISTORICAL DECOMPOSITION (2)

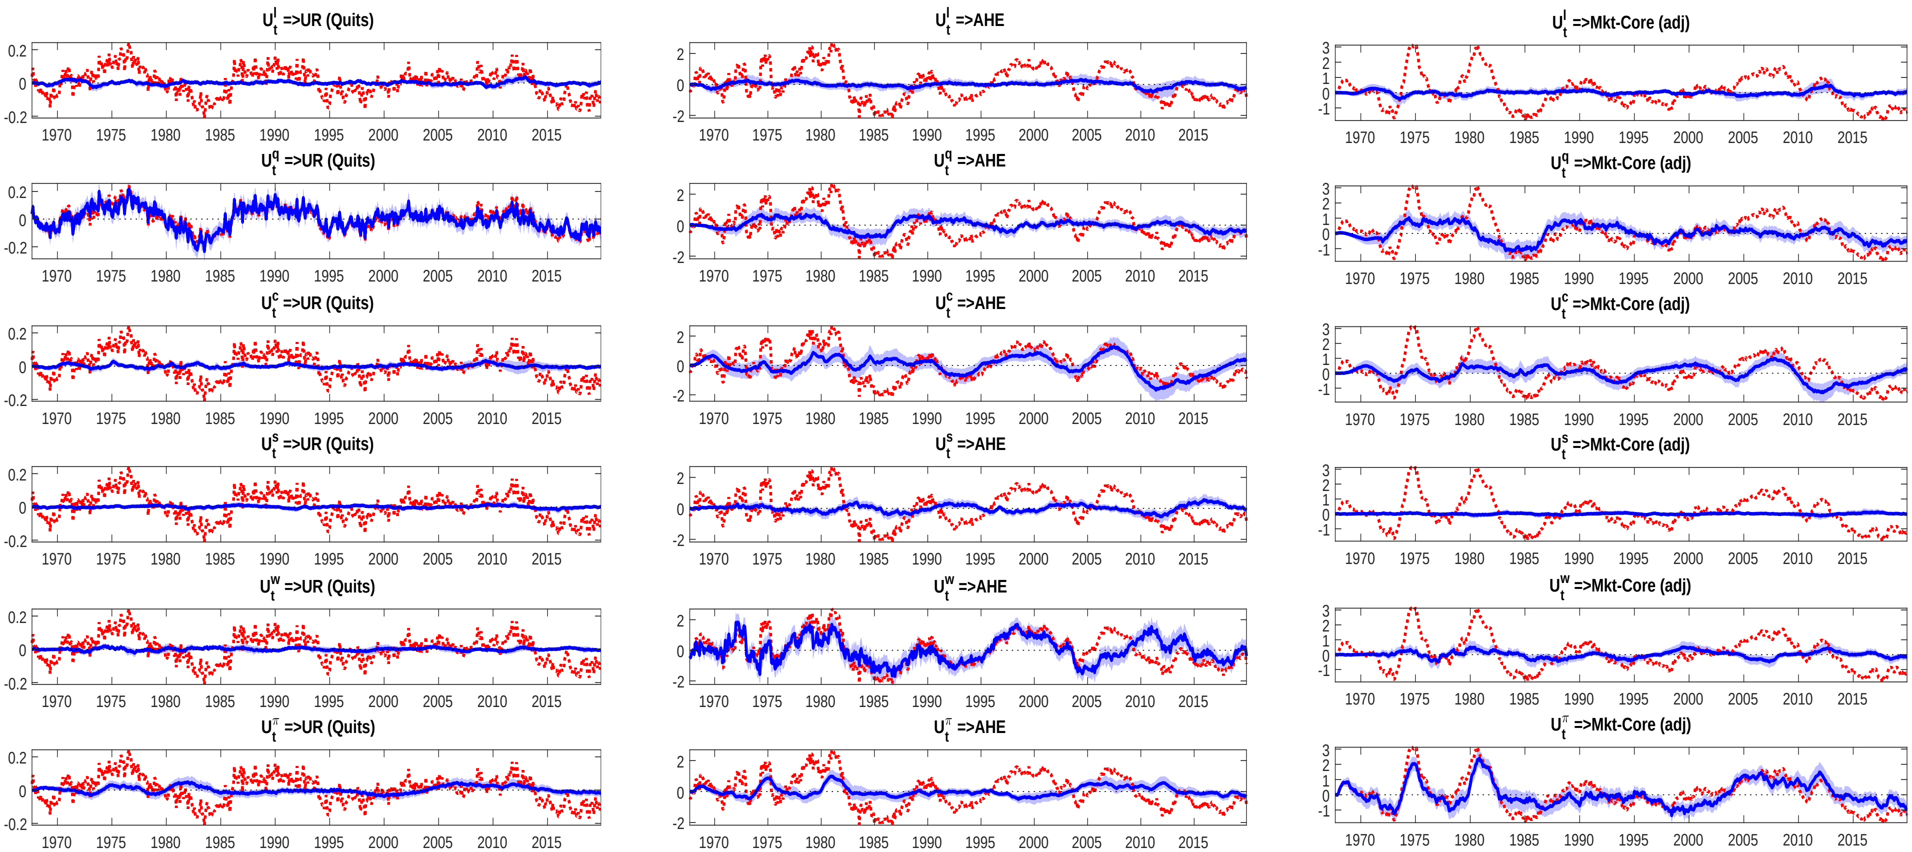

Notes: This figure displays the estimated historical decompositions for each variable (X-axis: calendar time; Y-axis: the magnitude of variations in the endogenous variable). The posterior median is shown as a solid line; the shaded area gives the 68 percent credible set. *LFPR* : the labor force participation rate; *VR* : the vacancy rate; *UR (Job Loss)* : the unemployment rate of job losers and entrants to the labor force; *UR (Quits)* : the unemployment rate of job quitters; *AHE* : Average hourly earnings; *Mkt-Core(adj)* : Adjusted market core inflation. U_t^l: job-loss reallocation shock; U_t^q: job-quits reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock; U_t^w: wage-specific shock; U_t^π: inflation-specific shock.

Source: Authors' calculation.

5 The labor market during the Covid-19 pandemic

This section focuses on the period that extends from the start of the Covid-19 pandemic up to the end of 2023. Section 5.1 discusses how we handle this period econometrically, and Section 5.2 describes the historical decompositions that we obtain over this period.

5.1 How we handle the pandemic period

The extremely large swings in labor-market and inflation indicators that were induced by the shutdown and restart of the economy represent influential outliers that will tend to dominate and distort any model estimates that include these observations. So, rather than using observations from the pandemic period to fit our model, we instead interpret the pandemic-induced swings as resulting from large transitory (structural) shocks that had no material effect on the underlying dynamics of the economy. We therefore use the model parameters estimated from the pre-pandemic period to compute historical decompositions over the pandemic period; in other words, the historical decompositions are extended through the pandemic period by using the pre-pandemic parameter values to back out the implied structural innovations and the dynamics through which they propagate.[45] This approach is in line with several recent studies that consider how to deal with observations from the pandemic period when fitting time-series models (e.g., Lenza and Primiceri, 2020; Ng, 2021; Cascaldi-Garcia, 2022).[46]

5.2 Historical decompositions for the pandemic period and after

Figure 5 reports the historical decompositions from our model over the period 2020–2023. In contrast to earlier periods, the reallocation and labor supply shocks make important contributions to the dramatic swings in labor market indicators that were observed during the pandemic period. Figure 5 shows the portions of the two unemployment rates, the vacancy rate, the labor force participation rate, wage growth, and price inflation that are explained by each structural shock, with the December 2019 values normalized to be zero.

[45] One of several potential issues with this approach is that it could cause the credible sets for the historical decompositions to be too wide during the pandemic period. In principle, this problem could be alleviated if we allowed for time variation in the SVAR's covariance matrix; in practice, we have too few observations (during and after the pandemic) to produce reliable covariance estimates.

[46] Specifically, Ng (2021) excludes the unusual observations owing to the Covid-19 shock, while Cascaldi-Garcia (2022) uses dummy variables as a way to minimize effects of pandemic outliers on the parameter estimates and on the resulting impulse responses.

Figure 5: HISTORICAL DECOMPOSITION : COVID-19 PANDEMIC

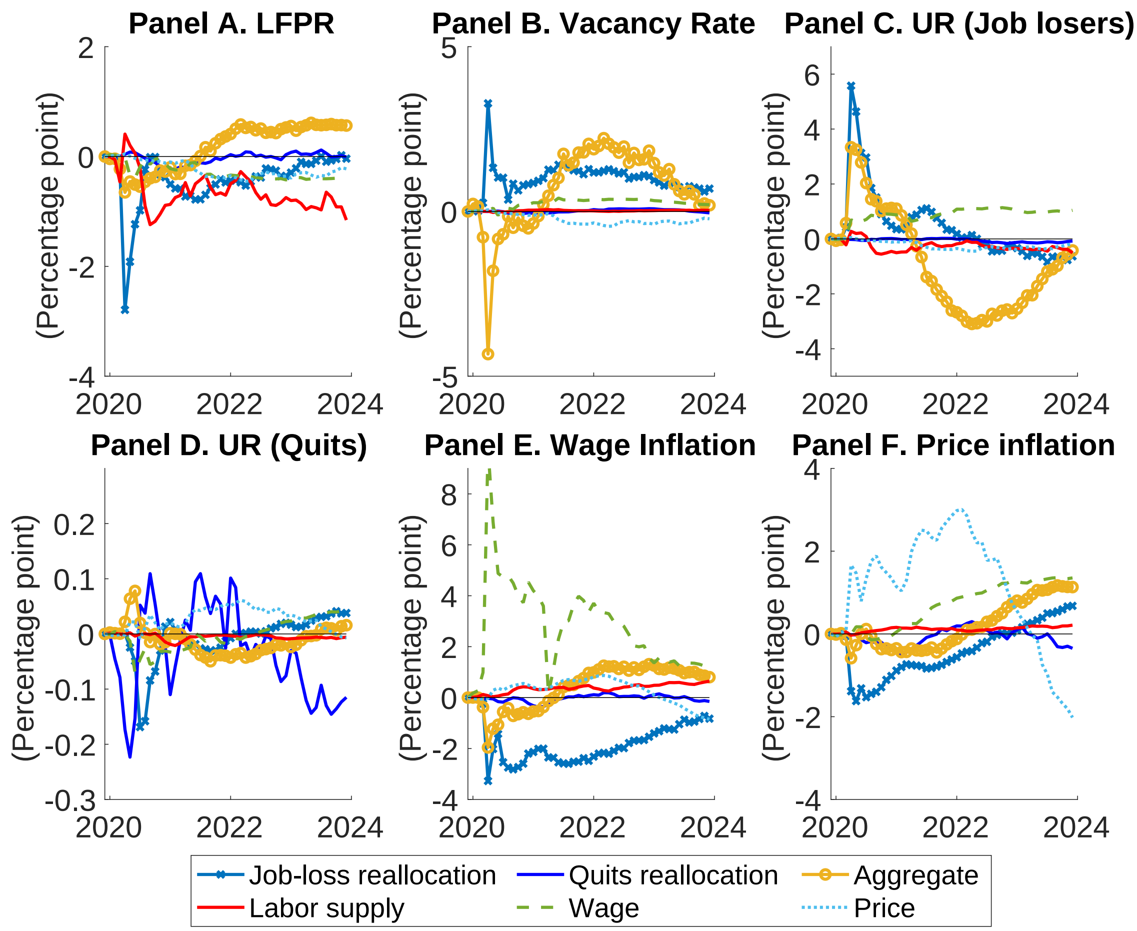

Notes : This figure displays the posterior medians of the historical decompositions of the variables in our VAR system during the Covid-19 pandemic (X-axis: calendar month; Y-axis: the magnitude of variations in the endogenous variable). *LFPR*: the labor force participation rate; *UR (Job losers)*: the unemployment rate of job losers and entrants to the labor force; *UR (Quits)*: the unemployment rate of job quitters; *Wage inflation*: Changes in average hourly earnings; *Price inflation*: Changes in adjusted market-based core PCE prices.
Source: Authors' calculations.

A. Labor force participation rate: Job-loss reallocation shocks are responsible for most of the 2020 swing in the participation rate (blue line with markers), including the sharp drop in the second quarter and subsequent rebound. From the second half of 2020 through the end of 2023, job-loss reallocation shocks continue to put downward pressure on the participation rate. Our interpretation of these movements in the participation rate is that they reflect job losers who, because of a deterioration in worker–job matching efficiency, experienced or anticipated a prolonged spell of joblessness that led them to exit the labor force. Additionally, negative labor supply shocks—which likely reflect the effects of school closures, increased care-giving responsibilities, fears of getting sick, and early retirements—depress the participation rate over much of this period (red line); this reduction in the labor force participation rate might have been facilitated by the unusually large stocks of savings that households held during this period.[47]

The recovery of the participation rate is attributable to a string of positive shocks to aggregate activity (yellow line). Hence, two offsetting factors govern the path of the participation rate during the pandemic: Participation is weighed down by both negative labor supply shocks and reallocative disturbances, while strong labor demand acts to push up the participation rate. Notably, the portion of the participation rate that is driven by aggregate activity shocks appears to have recovered to its pre-pandemic level in 2022.[48]

B. Vacancy rate: After dropping sharply in the pandemic's early phase, the vacancy rate moves higher over the remainder of the sample period. In the first half of 2020, job-loss reallocation shocks raise vacancies (blue line), but are more than offset by the effect of negative shocks to aggregate activity (yellow line). This positive effect of the reallocation shocks on job postings is consistent with findings by Barrero et al. (2020), as well as with Haltiwanger (2021a), who shows that new business formation spiked in the early phase of the pandemic. Hence, the job-loss reallocation shock might be raising the number of unemployed job losers in some sectors or firms, while at the same time creating new job postings elsewhere. From the second half of 2020, both the job-loss reallocation and aggregate activity shocks contribute to a persistent increase in the vacancy rate. The duration of the effect of the job-loss shocks is again consistent with the conclusions of Barrero et al. (2021) that Covid-19 acted like a *persistent* reallocation shock.

C and D. Unemployment rates: As shown in Panel C, job-loss reallocation shocks (blue line) cause the unemployment rate of job losers to rise sharply in April 2020 and to then decline rapidly over the second half of that year. From 2021-on, the aggregate activity shock lowers the unemployment rate of job losers, while the job-loss reallocation shock and the wage-specific shock (dashed green line) act to raise it. The upward pressure from the wage-specific shock might reflect the large increase in reservation wages that appears to have occurred during this period; this increase might have itself stemmed from households' large holdings of liquid savings (which would permit an extended period of search), elevated

[47] The initial (and short-lived) positive spike in the contribution of labor supply shocks to the participation rate probably doesn't reflect a structural shock to labor supply, but may instead capture the sorts of unusual measurement issues (such as misclassification during this period) that tend to specifically affect the participation rate (Ahn and Hamilton, 2022).

[48] This result aligns with findings by Hornstein, Kudlyak, Meisenbacher, and Ramachandran (2023). According to these authors, the labor force participation rate exceeded their estimate of its trend in 2022 mainly as a result of cyclical tightness in the labor market.

wage expectations, or a rise in their perceived disutility of work.[49] In this context, the portion of the unemployment rate of job losers that is attributable to job-loss reallocation shocks and wage-specific shocks likely implies changes in the non-cyclical rate of unemployment (NCRU). This boost to the NCRU gradually dissipates from mid-2021 through the end of 2023, primarily driven by the unwinding of the job-loss reallocation shocks. Meanwhile, the portion of the job losers' unemployment rate that is driven by aggregate demand shocks is about two percentage points lower than its pre-pandemic level at the end of 2022, suggesting that labor market conditions in that year were tighter than those that prevailed just prior to the pandemic. By the end of 2023, the portion of the unemployment rate that reflects aggregate activity shocks had largely returned to its pre-pandemic level; in addition, the effects of the job-loss reallocation shocks largely unwound as well, bringing the portion of unemployment rate driven by job-loss reallocation a little below its pre-pandemic level.[50]

Panel D shows the shocks' contributions to the unemployment rate of job quitters. The quits reallocation shocks raise the unemployment rate of job quitters in 2021 (solid blue line). This estimation result is consistent with the anecdotes of a "Great Resignation" that occurred in the recovery phase that followed the pandemic recession.[51] However, the effect on the total unemployment rate is small.

The negative labor supply shock lowers the unemployment rate of both groups by a small amount (red line). One interpretation is that the pandemic or policy responses to it (e.g., large transfer payments) might have reduced the intensity of job search by unemployed workers and motivated unemployed job seekers to leave the labor force (both would imply additional downward pressure on the NCRU). Once again, though, the magnitudes involved are not large.

E. AHE wage inflation: The evolution of wage inflation over the pandemic period is dominated by large swings in the contribution of own-shocks (the dashed green line in panel E), which in turn reflects a rapid change in the composition of employment: In the early phase of the pandemic, job losses were disproportionately concentrated among low-wage workers, which caused measured average wages to rise quickly. The sharp downward movement in wage growth in 2021 is the result of base

[49] This interpretation is consistent with the contribution of the wage-specific shock to the participation rate. The wage-specific shock lowered the participation rate by a small amount from 2021. The New York Fed's Survey of Consumer Expectations reports that households' self-reported reservation wages rose during the pandemic and stayed elevated through 2023.

[50] Consistent with our estimate, Crump et al. (2024) also find that their estimate of the unemployment rate gap is below its pre-pandemic level in 2022. However, in 2023 our estimate suggests less tightness in the labor market than what these authors find (their estimate for the gap is a few percentage points below its pre-pandemic level).

[51] See, for example, `https://www.bls.gov/opub/mlr/2022/article/empirical-evidence-for-the-great-resignation.htm`.

effects (recall that these are 12-month changes), as the earlier upward spike in wage growth drops out of the 12-month moving window. These composition effects have largely played out by the end of 2022; thereafter, the contribution of these shocks declines steadily. Elsewhere, the aggregate activity shock lowers wage inflation in 2020 and 2021, but pushes it up from 2022-on (yellow line). Meanwhile, the job-loss reallocation shock puts persistent downward pressure on wage growth (blue line with markers).

F. Market-based Core Inflation: Panel F indicates that own-shocks to price inflation (dotted light blue line) are the dominant force driving price inflation over this period. As is the case for wage inflation, the importance of the own-shocks reflects the fact that typically only a modest fraction of the variability in price and wage inflation can be explained by changes in resource utilization. In this episode, the own-shocks to price inflation likely reflect the effects of influences such as import price growth and the presence of supply–demand imbalances in product markets, neither of which is explicitly captured by the model's inflation equation. Own-shocks cause a reduction in inflation from 2022-on, which likely captures the unwinding of the supply bottlenecks that emerged when the world economy was reopened. Even so, the influence of at least some fundamentals is clearly apparent: Aggregate activity (yellow line) puts downward pressure on inflation until 2021; in addition, own-shocks to wage inflation (green dashed line) contribute about 1.5 percentage points to core inflation by 2023, implying a non-negligible amount of pass-through from wages to prices.[52] The contribution of job-reallocation shocks to price inflation is mostly negative until 2022 but turns positive in 2023, perhaps reflecting the re-employment of job losers after prolonged job searches.

6 Discussion

We conclude with a discussion of the implications of our results for the measurement of labor market slack and what they suggest about the nature of the inflation process, and suggest some directions for future research.

[52]The contribution of aggregate activity shocks to price inflation is somewhat smaller than their effect on wage inflation (note the difference in the y-axis scales), which is once again in line with studies that find wage Phillips curves to be steeper than price Phillips curves.

6.1 The anatomy of the Beveridge curve

An outward shift in the Beveridge curve has often been interpreted as indicating an increase in the noncyclical rate of unemployment. However, cyclical factors such as inflows to unemployment can also shift the Beveridge curve, which makes such an interpretation problematic (see Barnichon and Figura, 2010; Ahn and Crane, 2020; and Christiano, Eichenbaum, and Trabandt, 2015). In this context, our empirical decomposition of Beveridge curve dynamics into the contributions of different structural shocks can be used to isolate the noncyclical movements in the curve.

Figure 6 displays estimated Beveridge relations conditional on each structural shock; the vacancy rate (V) is plotted on the y-axis and the unemployment rate (U) is plotted on the x-axis, and the data shown span the pre-pandemic period. The aggregate activity shock yields a downward-sloping V-U relation whose slope is about -0.56 (Panel A), while job-loss reallocation shocks yield an *upward*-sloping V-U relationship with a slope of about $+0.21$ (Panel B). By contrast, the quits-reallocation shock plays a limited role in observed V-U movements because job quitters represent a small share of the unemployment pool (Panel C). Finally, labor supply shocks move the Beveridge curve sideways (Panel D).

While shocks to wage inflation do not create a meaningful trade-off between the unemployment and vacancy rates (panel E), shocks to price inflation produce a downward-sloping V-U relation with a slope that is slightly flatter (about -0.37) than what the aggregate activity shock generates (Panel F). The impulse responses suggest that these inflation shocks behave like negative disturbances to the labor market, with effects on unemployment and job vacancies that are qualitatively similar to—but nevertheless distinct from—what we observe following an aggregate activity shock. Such a finding is consistent with studies that examine the consequences for labor market outcomes of a shock to the price markup (as compared with a productivity-driven change in aggregate activity)—see Gertler and Trigari (2009) and Christiano, Eichenbaum, and Trabandt (2016).[53]

All told, the structural labor market shocks generate distinct sets of V-U dynamics, suggesting that we should consider the contribution of each shock when using the Beveridge curve to evaluate the cyclical state of the labor market.

[53]As before, our results also seem consistent with an interpretation in which nominal shocks induce a monetary policy response that takes time to have its full effect on real activity. In response to the increase in inflation, the monetary authority likely implements a contractionary policy that results in an increase in the unemployment rate and a reduction in the vacancy rate. (We would emphasize that these own shocks are not themselves monetary policy shocks.)

6.2 The Beveridge curve during the Covid-19 pandemic

Figure 7 includes the results for the pandemic period. The solid dots denote the pandemic observations, while the red open circles reproduce the pre-pandemic observations from Figure 6. The effects of the realized shocks in 2020 are captured with the dark blue dots, the subsequent realizations with lighter blue dots, and the most-recent realizations with yellow dots. The dark-red dashed line traces out the trajectory induced by each shock. Finally, the black lines show the pre-pandemic (linear) V-U relationships obtained from Figure 6.

Our estimates demonstrate just how unprecedented the pandemic-induced shocks to the labor market actually were. Not only did this period see an extraordinarily large change in aggregate activity (Panel A), it also saw a rise in job-loss reallocation that induced an outward shift in the Beveridge curve that was larger and more persistent than anything seen before (Panel B). By contrast, the effects of the labor-supply shocks (Panel D) were more in line with previous episodes, as were the effects of the shocks to price inflation (Panel F).[54]

Our anatomy of the movements in the Beveridge curve during and after the pandemic directly bears on the recent debate over the likelihood of a so-called soft landing in the labor market. Specifically, Blanchard et al. (2022) and Figura and Waller (2022) have expressed diametrically opposed views about the likely evolution of the labor market as the recovery matures. The former argue that the observed movements in vacancies and unemployment during the pandemic indicate that the underlying Beveridge curve is flat, which implies that a reduction in labor demand will lower the vacancy rate but induce a very large increase in the unemployment rate, making the likelihood of a soft landing during this business cycle essentially nil. Meanwhile, Figura and Waller argue that the observed changes in vacancies and unemployment reflected a high number of job separations (unemployment inflows) in the early phase of the pandemic, but that the underlying Beveridge curve is highly convex; as a result, a reduction in job vacancies (labor demand) can occur without a large increase in the unemployment rate.

We can evaluate this debate through the lens of our formal statistical framework. Essentially, what Blanchard et al. (2022) and Figura and Waller (2022) disagree about is the slope of the structural Beveridge curve—the locus of vacancy and unemployment rates that results from shocks to labor demand or the

[54] The effects of the wage-growth shocks are interesting inasmuch as they deviate noticeably from past experience; that said, the importance of composition effects make this period's wage behavior difficult to interpret.

aggregate activity shocks.[55] However, isolating this slope requires us to control for the effects of other structural shocks. Our decomposition does exactly that, and so provides a useful way to judge the merits of the two views.

Looking at the historical decomposition in Figure 5, both the job-loss reallocation and aggregate activity shocks exert comparable amounts of upward pressure on the vacancy rate in the second half of 2021. It is important to note that both the unwinding of job-loss reallocation shocks and cyclical slackening reduced the vacancy rate while pushing the unemployment rate in opposite directions. The job-loss reallocation shocks steadily lowered the unemployment rate from the second half of 2021 through the end of 2023, while the aggregate activity shock raised the unemployment rate over the same period of time. Owing to these offsetting factors, the decline in the vacancy rate in 2022 and 2023 was not accompanied by a large increase in the unemployment rate, even though a reduction in aggregate activity (labor demand) made an important contribution to the joint dynamics of the unemployment rate and vacancy rate over this period.

Taken as a whole, our analysis provides a description of labor-market dynamics during this period that is more nuanced than the analysis provided by either Blanchard, et al. or Figura and Waller. Rather than hinging on a particular shape of the Beveridge curve (or, equivalently, a particular parameterization of an hypothesized aggregate matching function), the prospects for a soft landing in the labor market are instead seen to depend on the size and persistence of a set of underlying structural shocks.[56] It is also important to keep in mind that different realizations of the structural shocks will induce different comovements of the unemployment rate and the vacancy rate; in particular, even if the underlying structural Beveridge curve is linear, the particular structural shocks that occur in a given period can give rise to a nonlinear reduced-form relationship. Hence, our findings provide an explanation for any observed convexity in the empirical or reduced-form Beveridge curve relations; by extension, our results underscore why assessing the likely future evolution of the labor market (including the probable effects of additional fiscal or monetary policy actions on unemployment) requires us to do more than just look for

[55]We use the terminology 'the structural Beveridge curve' to distinguish it from the empirical Beveridge curve—the observed locus of vacancy rate and unemployment rate. Although a negative aggregate activity shock can also raise inflows to unemployment, there won't be an outward shift of the Beveridge curve if the resulting reduction in labor demand reduces vacancy postings as well (it will just represent a movement along the curve). Only a burst of job losers that sufficiently outpaces the drop in the vacancy rate will be captured as a reallocation flow.

[56]See Rudd (2024) for a discussion of an alternative theoretical characterization of the Beveridge curve that does not involve a matching function.

suggestive patterns in a scatterplot.[57]

6.3 The anatomy of the reduced-form Phillips curve

Changes in aggregate demand should induce movements along the Phillips curve rather than shifts in the curve. We can therefore get an idea of the slope of the Phillips curve by quantifying the movements in the unemployment rate and inflation that obtain following an aggregate activity shock.[58] As our results demonstrate, one reason to take this approach is that other structural shocks can influence the size and direction of co-movements in the unemployment rate and price inflation that can in turn influence the reduced-form correlation between these variables.

To illustrate this point, we compute the correlation of the unemployment rate and market-based core inflation conditional on each structural shock. Using the impulse responses from the baseline model (Figure 2), we compute the ratio of the cumulative response of price inflation and the cumulative response of the unemployment rate (we combine both types of unemployment) following each structural shock. We compute the cumulative responses over h periods, so these ratios (which are the same thing as an integral multiplier) give the average change in the unemployment rate that is associated with a given change in price inflation through horizon h following a structural shock. For the aggregate activity shock, this ratio can be thought of as a measure of the inflation–unemployment tradeoff—loosely, the Phillips curve slope as it is typically defined. We also report analogous estimates for wage growth.[59]

Table 4 reports estimates of these "Phillips ratios" for the four real-side labor market shocks at various horizons.[60] Unsurprisingly, the aggregate activity shock induces a large and statistically significant tradeoff between inflation and the unemployment rate (Panel C), with a steeper tradeoff for wages than for prices. This is, of course, the same correlation that we recover in standard Phillips curve specifications over periods where changes in labor- and product-market slack largely reflect cyclical factors. Similarly, the larger wage response again echoes studies that find wage Phillips curves to be steeper than price Phillips curves.[61] We find a bit larger Phillips ratios for job reallocation shocks (Panel A); one interpretation is

[57] Our main focus is not the convexity of Beveridge curve. However, our empirical findings suggest that observed trajectories of vacancies and unemployment can spuriously suggest a convex relation.

[58] Barnichon and Mesters (2020) recover the slope of a "structural" Phillips curve based on a similar approach.

[59] One pitfall of such a measure is that if the unemployment response is close to zero following a shock, the integral multiplier becomes very large and will be imprecisely measured. For the structural shocks we consider, this is not much of an issue.

[60] The full set of ratios is plotted in Figure A.10.

[61] Using a time-varying parameter VAR (TVP–VAR) similar to the one used by Peneva and Rudd (2017), we find integral multipliers for the effect of unemployment on core market-based inflation on the order of −0.15 for the post-2000 sample, which is similar in size to what we find for an aggregate activity shock. This result obtains because the aggregate activity shock is the

Figure 6: The Anatomy of the Beveridge Curve (Pre-pandemic)

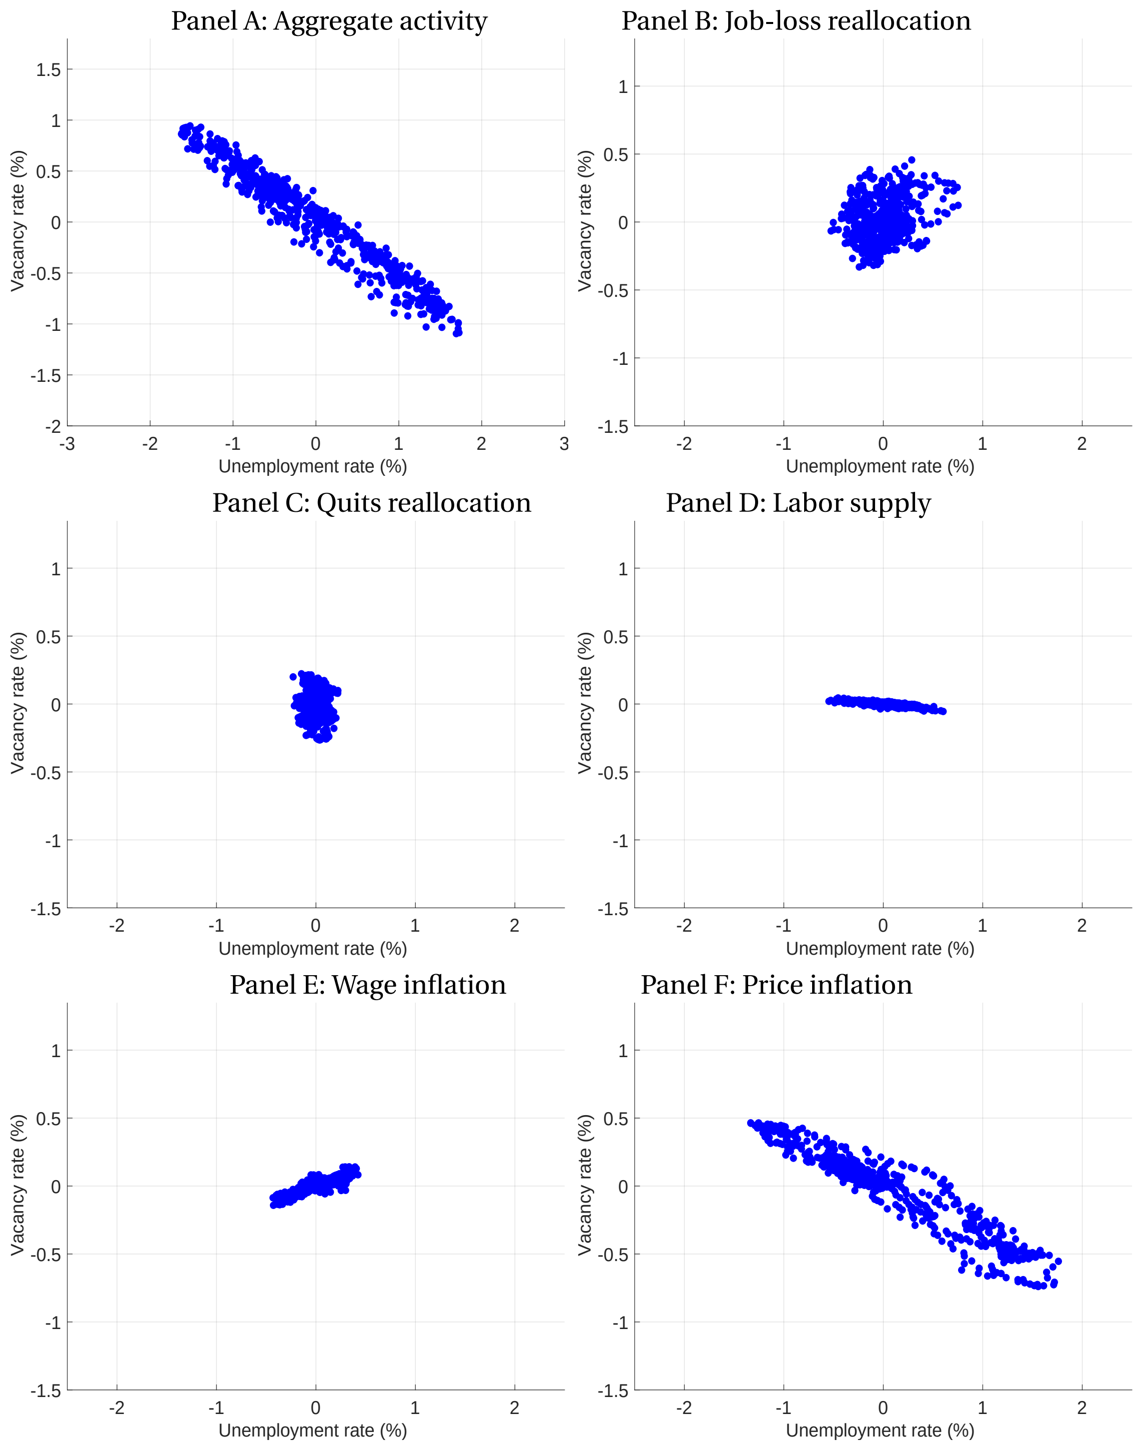

Notes: This figure displays the coordinates of vacancy and unemployment rates attributable to the six structural shocks (blue dots) before 2020. The red lines in Panels A, B, and F capture the linear relationship between unemployment and vacancy rates driven by aggregate activity, job-loss reallocation, and price-inflation specific shocks, respectively. The X-axis is the portion of unemployment rate driven by each structural shock (%). The Y-axis is the portion of vacancy rate driven by each structural shock (%).

Source: Authors' calculations.

Figure 7: THE ANATOMY OF THE BEVERIDGE CURVE DURING THE PANDEMIC

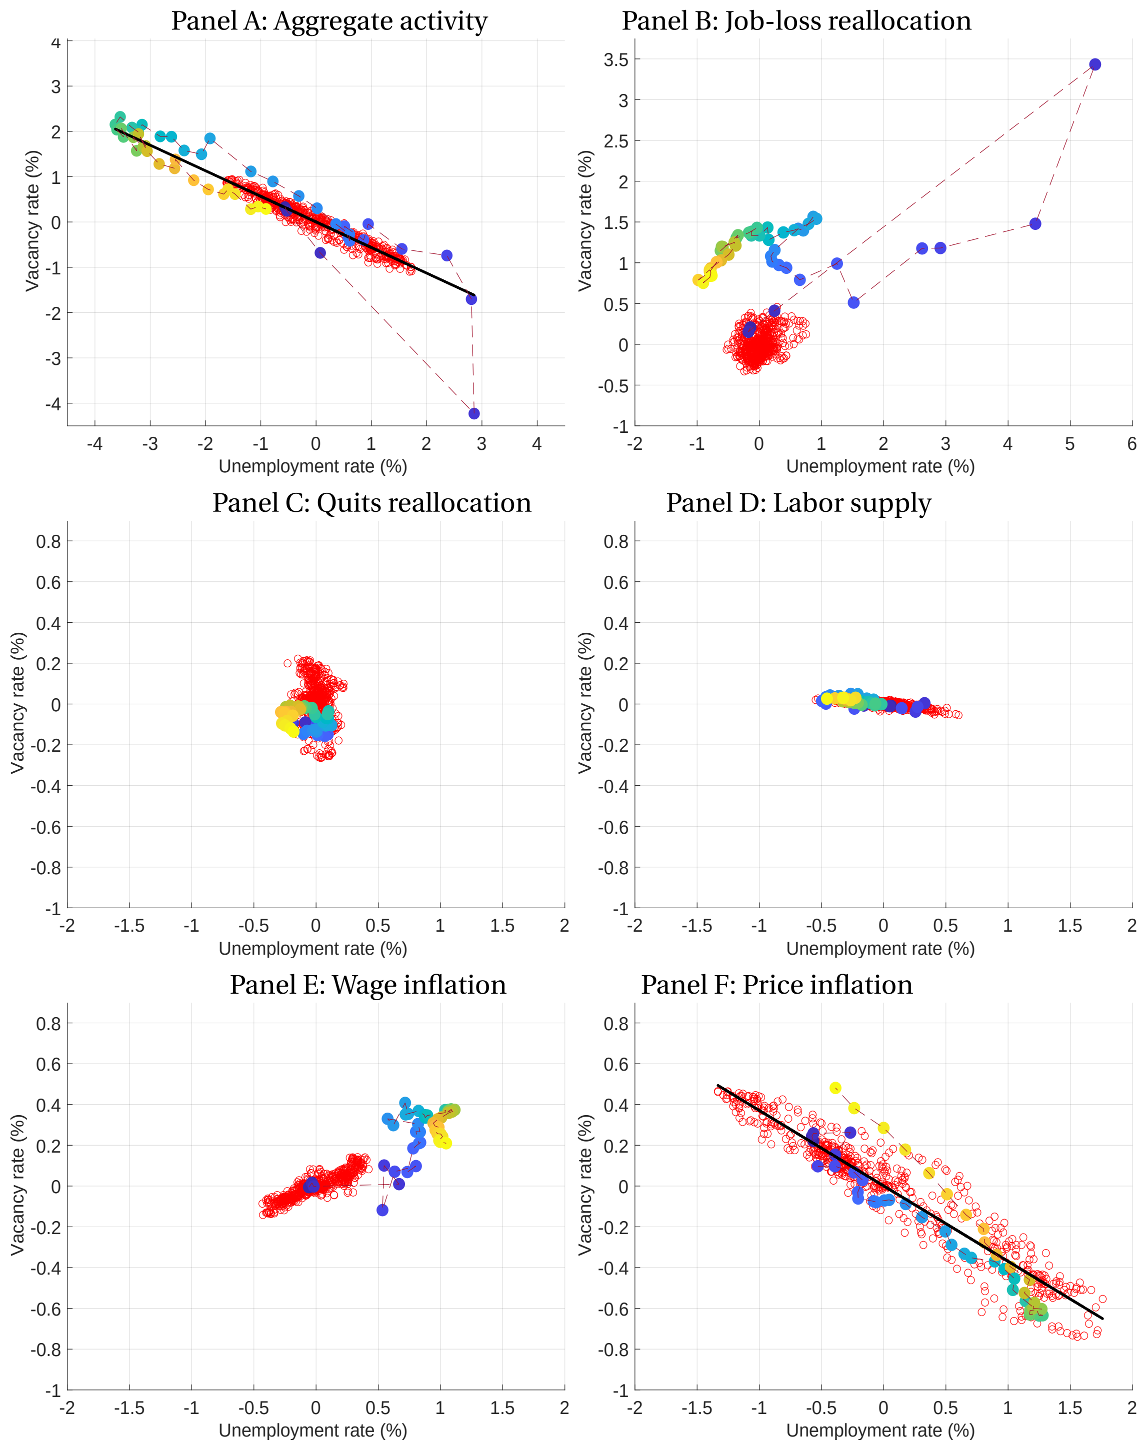

Notes: This figure displays the coordinates of vacancy and unemployment rates attributable to the six structural shocks. The red circles denote pre-pandemic observations; pandemic-period observations start as dark-blue dots and end as yellow dots. (Note the changes in scale between the first two panels and the remaining charts.) The black lines in Panels A, B, and F capture the linear relationship between unemployment and vacancy rates driven by aggregate activity, job-loss reallocation, and price-inflation specific shocks, respectively. The X-axis is the portion of unemployment rate driven by each structural shock (%). The Y-axis is the portion of vacancy rate driven by each structural shock (%).

Source: Authors' calculations.

that the increased job loss that results from this type of shock lowers the bargaining power of workers, thereby lowering wage growth and putting downward pressure on price inflation. Of note, the price Phillips ratio for the job-loss reallocation shock is short-lived relative to that of the aggregate activity shock, likely reflecting the acute and short-lived nature of job-loss shocks. The story is different in an interesting way for quits reallocation shocks (Panel B) in that a rise in unemployment associated with increased quits is actually *inflationary* (so the resulting inflation–unemployment tradeoff is *positively* sloped). Intuitively, we would expect increased quits to yield higher wage inflation as employers try to hold on to their remaining workers.[62] Increases in labor supply (panel D) have the expected (negative) effect on wage growth; their effect on price inflation is also negative and statistically significant.

The qualitative and quantitative patterns of the inflation–unemployment responses that arise from these shocks carry several implications for how we can assess the state of the labor market or measure the non-cyclical rate of unemployment (NCRU). First, because the effects of activity, labor supply, and job-loss reallocation shocks on unemployment and price inflation are very similar, it will not be possible to infer the cyclical position of the labor market simply by considering the behavior of inflation and unemployment in isolation (that is, by using a standard price Phillips curve); by extension, estimates of the NCRU using these sorts of Phillips curves are apt to be misleading. This problem is also present if we instead rely on a wage Phillips curve for these purposes: Even though the wage Phillips curve is "steeper" than the price Phillips curve, which would ordinarily permit tighter inference about labor-market slack and the NCRU, the observed correlation between wage growth and unemployment will also reflect a mix of true aggregate demand shocks and (observationally equivalent) shocks to job-loss reallocation. It is also worth pointing out that only quits-related reallocation shocks induce the pattern usually associated with an increase in the NCRU; namely, a rise in wage and price inflation accompanied by an increase in the unemployment rate.

This analysis also implies that we need to be cautious as to how we interpret observed changes in the slope of the Phillips curve. The reduced-form Phillips curve over a particular span of time will depend on the average mix or composition of the structural shocks that are realized over that period. Depending on that set of realizations, therefore, the slope of the reduced-form Phillips curve can change over time in ways that are unconnected with the actual tradeoff between real activity and inflation. As an illustration,

main source of variations in the unemployment rate, as is the own-shock to the unemployment rate in the TVP–VAR model.

[62] The higher price inflation that we see might reflect the pass-through of this faster wage growth, though the overall price response seems too large to be explained through this channel alone.

Table 4: CUMULATIVE WAGE AND PRICE PHILLIPS CORRELATION BY STRUCTURAL SHOCK

	Wage Phillips correlation		**Price Phillips correlation**	
	Median	68% posterior interval	Median	68% posterior interval
Panel A: Job-loss reallocation				
$h = 0$	-0.57	(-1.00, -0.35)	-0.27	(-0.40, -0.14)
$h = 6$	-0.79	(-1.34, -0.45)	-0.48	(-1.03, -0.23)
$h = 12$	-1.08	(-2.16, -0.57)	-0.55	(-1.44, -0.17)
$h = 18$	-1.22	(-2.78, -0.29)	-0.54	(-1.51, 0.20)
Panel B: Quits reallocation				
$h = 0$	0.10	(0.03, 0.25)	0.09	(0.03, 0.19)
$h = 6$	0.57	(0.04, 1.36)	0.72	(0.33, 1.34)
$h = 12$	0.82	(-0.39, 2.94)	1.43	(0.33, 3.82)
$h = 18$	0.65	(-3.30, 3.74)	1.52	(-4.13, 5.44)
Panel C: Aggregate activity				
$h = 0$	-0.56	(-0.83, -0.36)	-0.17	(-0.28, -0.08)
$h = 6$	-0.45	(-0.56, -0.34)	-0.10	(-0.19, -0.01)
$h = 12$	-0.48	(-0.59, -0.37)	-0.15	(-0.25, -0.06)
$h = 18$	-0.53	(-0.65, -0.41)	-0.23	(-0.34, -0.12)
Panel D: Labor supply				
$h = 0$	-0.21	(-0.38, -0.08)	-0.16	(-0.24, -0.08)
$h = 6$	-0.49	(-0.76, -0.27)	-0.23	(-0.44, -0.05)
$h = 12$	-0.61	(-0.95, -0.35)	-0.29	(-0.59, -0.04)
$h = 18$	-0.74	(-1.16, -0.44)	-0.32	(-0.71, -0.04)

Notes: This table reports the dynamic wage and price Phillips correlations conditional on the structural shocks. Panels A, B, C, and D report the wage and price Phillips correlations driven by job-loss reallocation, quits reallocation, aggregate activity and labor supply shocks. Rows in each panel (h) capture the magnitude of correlation h-period after the shock's impact.

Source: Authors' calculations.

after the pandemic recession the observed rise in price and wage inflation given conventional estimates of product- and labor-market slack led some researchers to conclude that the inflation–activity tradeoff had worsened.[63] Our analysis provides an alternative interpretation: The apparent steepening or non-linearity in the reduced-form Phillips curve in recent years might actually reflect the greater size and prevalence of job-loss reallocation and labor supply shocks that accompanied the unprecedentedly large aggregate activity and price-specific shocks of this period.[64]

[63] See, for example, Higgins (2021), Hall (2023), and Smith et al. (2023).

[64] The results are robust to alternative model specifications—see Appendix A.5.

References

Abraham, K. and L. F. Katz (1986). Cyclical unemployment: Sectoral shifts or aggregate disturbances? *Journal of Political Economy 94*, 507–522.

Ahn, H. J. (2023). The role of observed and unobserved heterogeneity in the duration of unemployment. *Journal of Applied Econometrics 38*(1), 3–23.

Ahn, H. J. and L. D. Crane (2020, March). Dynamic Beveridge Curve Accounting. Finance and Economics Discussion Series 2020-027, Board of Governors of the Federal Reserve System (U.S.).

Ahn, H. J. and J. D. Hamilton (2022). Measuring labor-force participation and the incidence and duration of unemployment. *Review of Economic Dynamics 44*, 1–32.

Ahn, H. J., B. Hobijn, and A. Şahin (2023, May). The Dual U.S. Labor Market Uncovered. Finance and Economics Discussion Series 2023-031, Board of Governors of the Federal Reserve System (U.S.).

Ahn, H. J. and S. Smith (2024). Inflation dynamics during the covid era: A high-frequency approach.

Atkeson, A. and L. E. Ohanian (2001). Are Phillips curves useful for forecasting inflation? *Quarterly Review 25*(Win), 2–11.

Bagga, S., L. Mann, A. Şahin, and G. Violante (2023, July). Job amenity shocks and labor reallocation. Technical report.

Barlevy, G., R. J. Faberman, B. Hobijn, and A. Şahin (2023, October). The shifting reasons for beveridge-curve shifts. Working Paper 31783, National Bureau of Economic Research.

Barnichon, R. (2010, December). Building a composite Help-Wanted Index. *Economics Letters 109*(3), 175–178.

Barnichon, R. and A. Figura (2010). What drives movements in the unemployment rate? a decomposition of the Beveridge curve. Finance and Economics Discussion Series 2010-48, Board of Governors of the Federal Reserve System (U.S.).

Barnichon, R. and A. Figura (2015). Labor market heterogeneity and the aggregate matching function. *American Economic Journal: Macroeconomics 7*(4), 222–249.

Barnichon, R. and G. Mesters (2020, 06). Identifying Modern Macro Equations with Old Shocks*. *The Quarterly Journal of Economics 135*(4), 2255–2298.

Barrero, J. M., N. Bloom, and S. J. Davis (2020, May). Covid-19 is also a reallocation shock. Working Paper 27137, National Bureau of Economic Research.

Barrero, J. M., N. Bloom, S. J. Davis, and B. H. Meyer (2021, May). Covid-19 is a persistent reallocation shock. *AEA Papers and Proceedings 111*, 287–91.

Baumeister, C. and J. Hamilton (2015, September). Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information. *Econometrica 83*(5), 1963–1999.

Baumeister, C. and J. Hamilton (2018). Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations. *Journal of Monetary Economics 100*(C), 48–65.

Baumeister, C. and J. Hamilton (2019, May). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. *American Economic Review 109*(5), 1873–1910.

Baumeister, C. and J. Hamilton (2021). Structural vector autoregressions with imperfect identifying information. *Working paper*.

Baumeister, C. and G. Peersman (2013, October). Time-varying effects of oil supply shocks on the us economy. *American Economic Journal: Macroeconomics 5*(4), 1–28.

Blanchard, O. J. and P. Diamond (1989). The Beveridge Curve. *Brookings Papers on Economic Activity 20*(1), 1–76.

Blanchard, O. J., A. Domash, and L. H. Summers (2022, July). Bad news for the Fed from the Beveridge space. Policy Briefs PB22-7, Peterson Institute for International Economics.

Blinder, A. S. and J. B. Rudd (2013, June). The supply-shock explanation of the Great Stagflation revisited. In *The Great Inflation: The Rebirth of Modern Central Banking*, pp. 119–175. University of Chicago Press.

Bloom, N., P. Bunn, P. Mizen, P. Smietanka, and G. Thwaites (2020, December). The Impact of Covid-19 on Productivity. NBER Working Papers 28233, National Bureau of Economic Research, Inc.

Cascaldi-Garcia, D. (2022, August). Pandemic Priors. International Finance Discussion Papers 1352, Board of Governors of the Federal Reserve System (U.S.).

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2015, January). Understanding the great recession. *American Economic Journal: Macroeconomics 7*(1), 110–67.

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2016). Unemployment and business cycles. *Econometrica 84*(4), 1523–1569.

Cochrane, J. H. (2005, January). Time series for macroeconomics and finance. Lecture note, University of Chicago.

Coibion, O. and Y. Gorodnichenko (2015, January). Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics 7*(1), 197–232.

Constâncio, V. (2015). Understanding inflation dynamics and monetary policy in a low inflation environment. Speech at the at the ECB Conference on “Challenges for Macroeconomic Policy in a Low Inflation Environment”, Frankfurt, November 5, 2015.

Crump, R. K., S. Eusepi, M. Giannoni, and A. Şahin (2024, March). The unemployment-inflation trade-off revisited: The phillips curve in covid times. (1086).

Davis, S. J. (1987). Fluctuations in the pace of labor reallocation. *Carnegie-Rochester Conference Series on Public Policy 27*, 335–402.

Davis, S. J. and J. Haltiwanger (1999a, December). On the driving forces behind cyclical movements in employment and job reallocation. *American Economic Review 89*(5), 1234–1258.

Davis, S. J. and J. Haltiwanger (1999b, December). On the driving forces behind cyclical movements in employment and job reallocation. *American Economic Review 89*(5), 1234–1258.

Doan, T., R. Litterman, and C. Sims (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric Reviews 3*(1), 1–100.

Elsby, M. W., B. Hobijn, and A. Şahin" (2015). On the importance of the participation margin for labor market fluctuations. *Journal of Monetary Economics 72*, 64–82.

Elsby, M. W. L., R. Michaels, and D. Ratner (2015, September). The beveridge curve: A survey. *Journal of Economic Literature 53*(3), 571–630.

Faberman, R. J. and A. Justiniano (2015). Job Switching and Wage Growth. *Chicago Fed Letter*.

Faccini, R. and L. Melosi (2021, February). Job amenity shocks and labor reallocation. Technical Report 2020-09.

Faccini, R. and L. Melosi (2023, January). Job-to-Job Mobility and Inflation. Working Paper Series WP 2023-03, Federal Reserve Bank of Chicago.

Ferrante, F., S. Graves, and M. Iacoviello (2023). The inflationary effects of sectoral reallocation. *Journal of Monetary Economics*.

Figura, A. and C. J. Waller (2022, July). What does the Beveridge curve tell us about the likelihood of a soft landing? FEDS Notes 2022-07-29, Board of Governors of the Federal Reserve System (U.S.).

Fitzgerald, T. J., C. Jones, M. Kulish, and J. P. Nicolini (2020, October). Is There a Stable Relationship between Unemployment and Future Inflation? Staff Report 614, Federal Reserve Bank of Minneapolis.

Friedman, M. (1968). The role of monetary policy. *American Economic Review 58*(1), 1–17.

Gallant, J., K. Kroft, F. Lange, and M. J. Notowidigdo (2020, October). Temporary unemployment and labor market dynamics during the covid-19 recession. Working Paper 27924, National Bureau of Economic Research.

Garcia-Cabo, J., A. Lipinska, and G. Navarro (2023). Sectoral shocks, reallocation, and labor market policies. *European Economic Review 156*, 104494.

Gertler, M. and A. Trigari (2009). Unemployment fluctuations with staggered nash wage bargaining. *Journal of Political Economy 117*(1), 38–86.

Hall, R. E. (2023, March). A major shock makes prices more flexible and may result in a burst of inflation or deflation. Working Paper 31025, National Bureau of Economic Research.

Haltiwanger, J. C. (2021a, June). Entrepreneurship during the covid-19 pandemic: Evidence from the business formation statistics. Working Paper 28912, National Bureau of Economic Research.

Haltiwanger, J. C. (2021b, June). Entrepreneurship During the COVID-19 Pandemic: Evidence from the Business Formation Statistics. NBER Working Papers 28912, National Bureau of Economic Research, Inc.

Hamilton, J. D. (1988, June). A Neoclassical Model of Unemployment and the Business Cycle. *Journal of Political Economy 96*(3), 593–617.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2020, October). The Slope of the Phillips Curve: Evidence from U.S. States. NBER Working Papers 28005, National Bureau of Economic Research, Inc.

Higgins, P. C. (2021, September). The Phillips Curve during the Pandemic: Bringing Regional Data to Bear. Policy Hub 94156, Federal Reserve Bank of Atlanta.

Hobijn, B., R. Miles, J. Royal, and J. Zhang (2023, January). The Recent Steepening of Phillips Curves. *Chicago Fed Letter 0*.

Hornstein, A., M. Kudlyak, B. C. Meisenbacher, and D. Ramachandran (2023, August). How Far Is Labor Force Participation from Its Trend? *FRBSF Economic Letter 2023*(20), 1–5.

Lenza, M. and G. E. Primiceri (2020, September). How to estimate a var after march 2020. Working Paper 27771, National Bureau of Economic Research.

Lilien, D. M. (1982, August). Sectoral Shifts and Cyclical Unemployment. *Journal of Political Economy 90*(4), 777–793.

Michaillat, P. and E. Saez (2022, July). u* = uv. NBER Working Papers 30211, National Bureau of Economic Research, Inc.

Montag, H. and D. Villar (2022). Price-Setting During the Covid Era. Technical report.

Moscarini, G. and F. Postel-Vinay (2017, May). The relative power of employment-to-employment reallocation and unemployment exits in predicting wage growth. *American Economic Review 107*(5), 364–68.

Moscarini, G. and F. Postel-Vinay (2023, July). The job ladder: Inflation vs. reallocation. Working Paper 31466, National Bureau of Economic Research.

Murphy, K. M. and R. H. Topel (1987, February). The Evolution of Unemployment in the United States: 1968–1985. In *NBER Macroeconomics Annual 1987, Volume 2*, NBER Chapters, pp. 11–68. National Bureau of Economic Research, Inc.

Ng, S. (2021, July). Modeling macroeconomic variations after covid-19. Working Paper 29060, National Bureau of Economic Research.

Nguyen, L. (2022). Bayesian inference in structural vector autoregression with sign restrictions and external instruments. Technical report.

Nguyen, L. (2024). Regime-switching structural vector autoregression identified by sign restrictions: Asymmetric effects of monetary policy revisited. Technical report.

Peneva, E. V. and J. B. Rudd (2017). The passthrough of labor costs to price inflation. *Journal of Money, Credit and Banking 49*(8), 1777–1802.

Ratner, D. and J. W. Sim (2022, May). Who Killed the Phillips Curve? A Murder Mystery. Finance and Economics Discussion Series 2022-028, Board of Governors of the Federal Reserve System (U.S.).

Rudd, J. B. (2024). *A Practical Guide to Macroeconomics.* Cambridge University Press.

Smith, S., A. Timmermann, and J. H. Wright (2023, April). Breaks in the phillips curve: Evidence from panel data. Working Paper 31153, National Bureau of Economic Research.

A Appendix

Section A.1 provides a more-detailed review of the relevant literature. Section A.2 details the modified estimation algorithm. Section A.3 uses our approach to estimate the original Blanchard–Diamond specification, and Section A.4 provides a number of robustness checks in the form of estimation results from various alternative models and alternative specifications.

A.1 Literature review

Our paper contributes to the literature on labor-market reallocation as well as the Phillips curve literature. We discuss our contributions and the implications of our results for the measurement of labor-market slack.

1. Labor-market reallocation

A large body of literature has treated aggregate activity shocks and reallocation shocks as providing competing explanations for the cyclical dynamics of labor markets (e.g., Lilien, 1982, Abraham and Katz, 1986). Interest in labor-market reallocation reemerged during the Covid-19 pandemic, as the pandemic hit some sectors and workers harder than others and resulted in significant disruptions to the labor market (see Bloom, Bunn, Mizen, Smietanka, and Thwaites, 2020, Barrero et al., 2020, and Barrero et al., 2021 for three examples). However, the measurement of labor-market reallocation is challenging, and much of the literature has focused on how best to do so. For instance, studies like Lilien (1982), Abraham and Katz (1986), and Murphy and Topel (1987) infer the intensity of reallocation activities from the dispersion of sectoral employment growth; Blanchard and Diamond (1989) estimate the effects of reallocation shocks from the positive comovement of unemployment and vacancy rates; and Davis and Haltiwanger (1999b), Barrero et al. (2020), and Barrero et al. (2021) measure reallocation intensity from the positive comovement of job creation and destruction flows. In contrast to these previous studies, Bagga, Mann, Şahin, and Violante (2023) claim that during the pandemic a labor reallocation shock occurred in the form of a disturbance that increased the value of non-pecuniary job amenities, and that this shock was an important source of Beveridge curve dynamics during the pandemic.[65]

[65] We discuss the literature on the Beveridge curve in the introduction; to avoid redundancy, we do not review it again here.

2. The Phillips curve and labor-market reallocation during the pandemic

Our paper also speaks to a literature that examines the effects of labor-market reallocation on inflation and on the estimated parameters of the Phillips curve.[66] Studies such as Faberman and Justiniano (2015) and Moscarini and Postel-Vinay (2017) argue that job quits or job switching can raise wage and price inflation. Moscarini and Postel-Vinay (2023) claim that job-to-job transitions are inflationary as they result in workers' climbing up the "job ladder," and Faccini and Melosi (2021) argue that on-the-job search can raise wages through increased competition for workers among firms. Relatedly, Faccini and Melosi (2023) develop a theory-based indicator of inter-firm wage competition and argue that inter-firm wage competition raised wage inflation by about a percentage point during the so-called Great Resignation. Quite differently, Ratner and Sim (2022) argue that the decreased bargaining power of trade unions weakened the Phillips correlation over a broader span of time (they propose an alternative model that they label a "Kaleckian" Phillips curve).[67]

Accounting for the role of reallocation shocks has direct implications for what we might infer about the NCRU or labor-market slack. Crump et al. (2024) estimate the natural rate of unemployment (*u-star*, henceforth) and the trend unemployment rate using a structural DSGE model with a new-Keynesian Phillips curve and a statistical model for the trend-cycle decomposition. These authors find their estimate of *u-star* increased during the pandemic and stayed elevated relative to its pre-pandemic level through the end of 2023. It is notable that in our model, which takes into account both wage and price inflation, changes in unemployment rate driven by job-loss reallocation are broadly in line with the *u-star* estimates from Crump et al. (2024). One important quantitative difference, though, is that our estimate has largely returned to its pre-pandemic level by the end of the sample, while the Crump et al. (2024) estimate has not. Hence, our estimates point to a smaller degree of labor market tightness around that time.

[66] Numerous studies argue that the Phillips correlation—the empirical association between inflation and real activity—weakened substantially in recent decades and almost disappeared (e.g., Atkeson and Ohanian, 2001; Coibion and Gorodnichenko, 2015; Constâncio, 2015; Haltiwanger, 2021b). Studies also disagree about the empirical magnitude of the Phillips correlation: Fitzgerald, Jones, Kulish, and Nicolini (2020) argue that the correlation is large and statistically significant once the correct specification is considered, while Hazell, Herreno, Nakamura, and Steinsson (2020) use results from state-level data to conclude that the Phillips correlation is weak. In a very different approach to the problem, Smith et al. (2023) use a panel-break framework to present evidence that the slope of the Phillips curve has in fact changed over time. Our own contribution is to show that a clear Phillips correlation emerges once one controls for the various underlying structural shocks that drive labor-market dynamics.

[67] Other related studies include Higgins (2021) (which is based on regional data) and Ahn and Smith (2024) (which uses high-frequency data on real sales and prices); both find empirical evidence for a steepening of the Phillips curve during the Covid-19 pandemic. In addition, Hall (2023) argues that inflation can change rapidly when two determinants of prices—input costs and productivity—become highly volatile, resulting in a steeper Phillips curve.

A.2 Modified estimation algorithm

As in Baumeister and Hamilton (2015), the algorithm used to estimate the modified characterization of the posterior distribution involves: [1] setting parameters for the prior distribution; [2] computing the target function; and [3] computing the posterior distribution.[68]

Step 1. Setting parameters for the prior distribution

In this step, we specify prior beliefs for $\mathbf{H}$, $\mathbf{D}$, and $\mathbf{B}$.

A. Specify $p(\mathbf{H})$

Specify the prior distribution $p(\mathbf{H})$ as a product of truncated Student t densities. For concreteness, suppose that there are two structural parameters in $\mathbf{H}$, h_α and h_β. Let $\mathbf{h}$ be defined as $\mathbf{h} = (h_\alpha, h_\beta)'$. Then $p(\mathbf{h})$ is given as

$$p(\mathbf{h}) = \begin{cases} \dfrac{f(h_\alpha; c_\alpha, \sigma, \nu)}{1 - F(0; c_\alpha, \sigma, \nu)} \dfrac{f(h_\beta; c_\beta, \sigma, \nu)}{F(h_\beta; c_\beta, \sigma, \nu)} & \text{if} \quad h_\alpha \geq 0 \quad \text{and} \quad h_\beta \leq 0 \\ 0 \qquad \text{otherwise,} & \end{cases} \tag{A.1}$$

where $f(x; c_\alpha, \sigma, \nu)$ is the probability density function of a Student t distribution with location c_α, scale σ, and degrees of freedom ν evaluated at x, and $F(\cdot)$ is the cumulative distribution function $F(x; c_\alpha, \sigma, \nu) = \int_{-\infty}^{x} f(z; c_\alpha, \sigma, \nu) dz$.

B. Specify $p(d_{ii}|\mathbf{H})$

Prior beliefs on structural variances are set so as to partly reflect the scale of the underlying data. Let $\hat{e}_{it}$ be the residuals from an eighth-order univariate autoregression fit to series i, and let $\hat{\mathbf{S}}$ be the sample variance of these residuals. The ith diagonal element of $(\mathbf{H^{-1}})\hat{\mathbf{S}}(\mathbf{H^{-1}})'$ is chosen for the prior mean of d_{ii}^{-1}. Because the prior mean of d_{ii}^{-1} is κ_i/τ_i, τ_i is a function of $\mathbf{h}$. If $\mathbf{a}_i$ denotes the ith row of $\mathbf{H}^{-1}$, the parameter τ_i can be expressed as $\tau_i(\mathbf{h}) = \kappa_i \mathbf{a}_i' \hat{\mathbf{S}} \mathbf{a}_i$.

C. Specify $p(b_i|\mathbf{H}, \mathbf{D})$

Doan et al. (1984) argue that an analyst should be more confident that the coefficients on higher lags are zero. In line with this recommendation, the prior mean of the lagged structural coefficients is set to be $\mathbf{m}_i(\mathbf{h}) = \boldsymbol{\eta}' \mathbf{a}_i$, where

$$\underbrace{\boldsymbol{\eta}}_{(n\times k)} = [\ \underbrace{\mathbf{I}_n}_{(n\times n)} \quad \underbrace{\mathbf{0}}_{[n\times(k-n)]}\ ],$$

[68]The following discussion follows Baumeister and Hamilton (2015) extremely closely, except for the alterations needed to describe our modification to their method.

for $k = mn + 1$.

Let s_{ii} denote the estimated variance of the residuals from a univariate eighth-order autoregression fit to variable i. Define $\mathbf{v}'_{1(1\times m)}$ and $\mathbf{v}'_{2(1\times n)}$ as follows:

$$\begin{aligned}
\mathbf{v}_{1(1\times m)} &= (1/(1^{2\lambda_1}), 1/(2^{2\lambda_1}), \cdots, 1/(m^{2\lambda_1}))' \\
\mathbf{v}_{2(1\times n)} &= (s_{11}^{-1}, s_{22}^{-1}, \cdots, s_{nn}^{-1})'.
\end{aligned}$$

The matrix $\mathbf{v}_3$ is defined to be

$$\mathbf{v}_3 = \lambda_0^2 \begin{bmatrix} \mathbf{v}_1 \otimes \mathbf{v}_2 \\ \lambda_3^2 \end{bmatrix}. \tag{A.2}$$

The matrix $\mathbf{M}_i$ is a diagonal matrix whose rth diagonal element is the rth element of $\mathbf{v}_3$. As in Baumeister and Hamilton (2015), we set $\lambda_0 = 0.2$, $\lambda_1 = 1$, and $\lambda_3 = 100$.[69]

Step 2. Computing the target function and setting initial values

For any $\mathbf{h}$, the target function $q(\mathbf{h})$ is calculated as:

$$\begin{aligned}
q(\mathbf{h}) \quad = \quad & \log p(\mathbf{h}) + (T/2)\log\{\det[(\mathbf{H}(\mathbf{h})^{-1})\hat{\Omega}_T(\mathbf{H}(\mathbf{h})^{-1})']\} \\
& - \sum_{i=1}^{2} (\kappa_i + T/2)\log\{[2\tau_i(\mathbf{h})/T] + [\zeta^*(\mathbf{h})/T]\} + \sum_{i=1}^{2} \kappa_i \log\tau_i(\mathbf{h}).
\end{aligned} \tag{A.3}$$

By numerically maximizing this target function, an initial guess of $\mathbf{h}$ and the associated Cholesky factor that provides the scale of the posterior distribution are obtained from $\hat{\mathbf{P}}_\Lambda \hat{\mathbf{P}}'_\Lambda = \hat{\Lambda} = -\frac{d^2 q(\mathbf{h})}{d\mathbf{h}d\mathbf{h}'}\Big|_{\mathbf{h}=\hat{\mathbf{h}}}$.

Step 3. Generate draws from the posterior distribution $p(\mathbf{H}, \mathbf{D}, \mathbf{B}|\mathbf{Y}_t)$

We take the following three steps to generate draws from the posterior distribution.

A. Random-walk Metropolis-Hastings algorithm for $p(\mathbf{H}|\mathbf{Y}_T)$

First set $\mathbf{h}^{(1)} = \hat{\mathbf{h}}$, and generate a candidate draw from the following:

$$\tilde{\mathbf{h}}^{(l+1)} = \tilde{\mathbf{h}}^{(l)} + \xi(\mathbf{P}_\Lambda^{-1})'\mathbf{v}_{l+1},$$

where $\mathbf{v}_{l+1}$ is a (2×1) vector of independent standard Student t variables with 2 degrees of freedom and ξ is

[69]The parameter λ_0 captures our overall confidence in the prior; λ_1 determines how quickly the prior for the lagged coefficients shrinks to zero as the lag length m increases; and the value of λ_3 that we choose ensures that the prior for the constant term will be essentially irrelevant.

a scalar tuning parameter that ensures an acceptance rate of about 30 percent. Next, if $q(\tilde{\mathbf{h}}^{(l+1)}) < q(\mathbf{h}^{(l)})$, set $\mathbf{h}^{(l+1)} = \mathbf{h}^{(l)}$ with probability $1 - \exp[q(\tilde{\mathbf{h}}^{(l+1)}) - q(\mathbf{h}^{(l)})]$; otherwise, set $\mathbf{h}^{(l+1)} = \tilde{\mathbf{h}}^{(l+1)}$. The first D draws are burn-in draws that are discarded.[70]

B. Generate draws from $p(\mathbf{D}|\mathbf{H},\mathbf{Y}_T)$

For each $\mathbf{h}^{(l)}$ with l from $D+1$, generate a draw of $\delta_{ii}^{(l)}$ as follows:

$$\delta_{ii}^{(l)} \sim \Gamma(\kappa_i + T/2, \tau_i(\mathbf{h}^{(l)}) + \zeta_i^*(\mathbf{h}^{(l)})/2)$$

and set the variance of the structural shocks, $d_{ii}^{(l)}$, equal to $1/\delta_{ii}^{(l)}$ for $i = 1, 2, \cdots, n$. These are the (i,i) elements of the diagonal matrix $\mathbf{D}^{(l)}$.

C. Generate draws from $p(\mathbf{B}|\mathbf{H},\mathbf{D},\mathbf{Y}_T)$

Take draws for the elements in $\mathbf{B}$, the matrix whose ith row is given by $\mathbf{b}_i'$, from

$$\begin{aligned}
\mathbf{b}_i^{(l)} &\sim N(\mathbf{m}_i^*(\mathbf{h}^{(l)}), d_{ii}^{(l)}\mathbf{M}_i^*) \\
\mathbf{m}_i^*(\mathbf{h}^{(l)}) &= (\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i')^{-1}(\tilde{\mathbf{X}}_i'\tilde{\mathbf{Y}}_i(\mathbf{h}^{(l)})) \\
\mathbf{M}_i^* &= (\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i')^{-1}.
\end{aligned}$$

[70]In our paper, we set $D = 1$ million.

A.3 Three-variable model

If we normalize $\gamma = 1$, the $\mathbf{H}$ matrix of Equation (2.1) becomes:

$$\underbrace{\begin{bmatrix} -\theta & \delta & 1 \\ \beta & 1 & 0 \\ 1 & -\epsilon & \omega \end{bmatrix}}_{\mathbf{H}}. \tag{A.4}$$

Figures A.1, A.2, and A.3 report this model's prior and posterior distributions for the structural parameters, the estimated impulse response functions, and the resulting historical decompositions, respectively.

Note that Blanchard and Diamond essentially calibrated their model with a set of pre-determined parameters and then estimated the full dynamic model using the method of moments. Despite this methodological difference, the estimated impulse responses that we obtain for the effects of each shock on each variable are qualitatively similar.[71] (The only differences worth noting are that the reallocation shock in our model has a more persistent positive effect on the unemployment and vacancy rates, and that the labor supply shock has a more persistent positive effect on the unemployment rate than in their model.)

[71] Note that Figure 9 in their paper and our Figure A.2 are comparable. The first column of Figure 9 in BD corresponds to the third column of our Figure A.2, and the third column of Figure 9 in BD corresponds to the first column of Figure A.2.

Figure A.1: PRIOR AND POSTERIOR DISTRIBUTIONS FOR STRUCTURAL PARAMETERS: THREE-VARIABLE MODEL

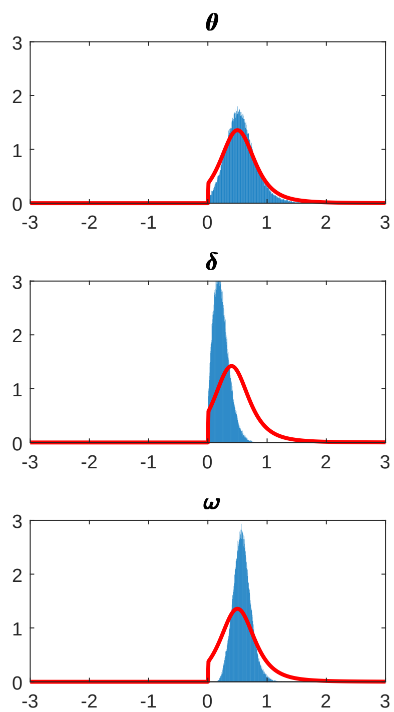

Note: This figure shows the prior and posterior distributions of structural parameters in **H**. Red lines are prior distributions; blue regions are posterior distributions.
Source: authors' calculation.

Figure A.2: Impulse responses (three-variable VAR)

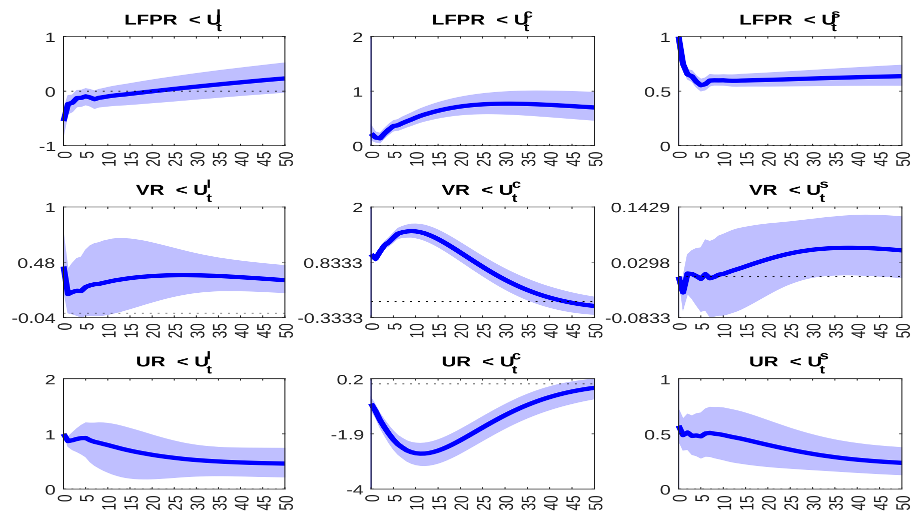

Notes: Posterior median (solid curve) and 68 percent credible sets (shaded regions). The X-axis denotes the number of months since the shock's impact. We use the following acronyms: *LFPR*: labor force participation rate; *VR*: vacancy rate; *UR*: unemployment rate. U_t^l: reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock. In the panel titles, $x < y$ denotes the effect of shock y on variable x.

Source: Authors' calculations.

Figure A.3: HISTORICAL DECOMPOSITION: THREE-VARIABLE VAR

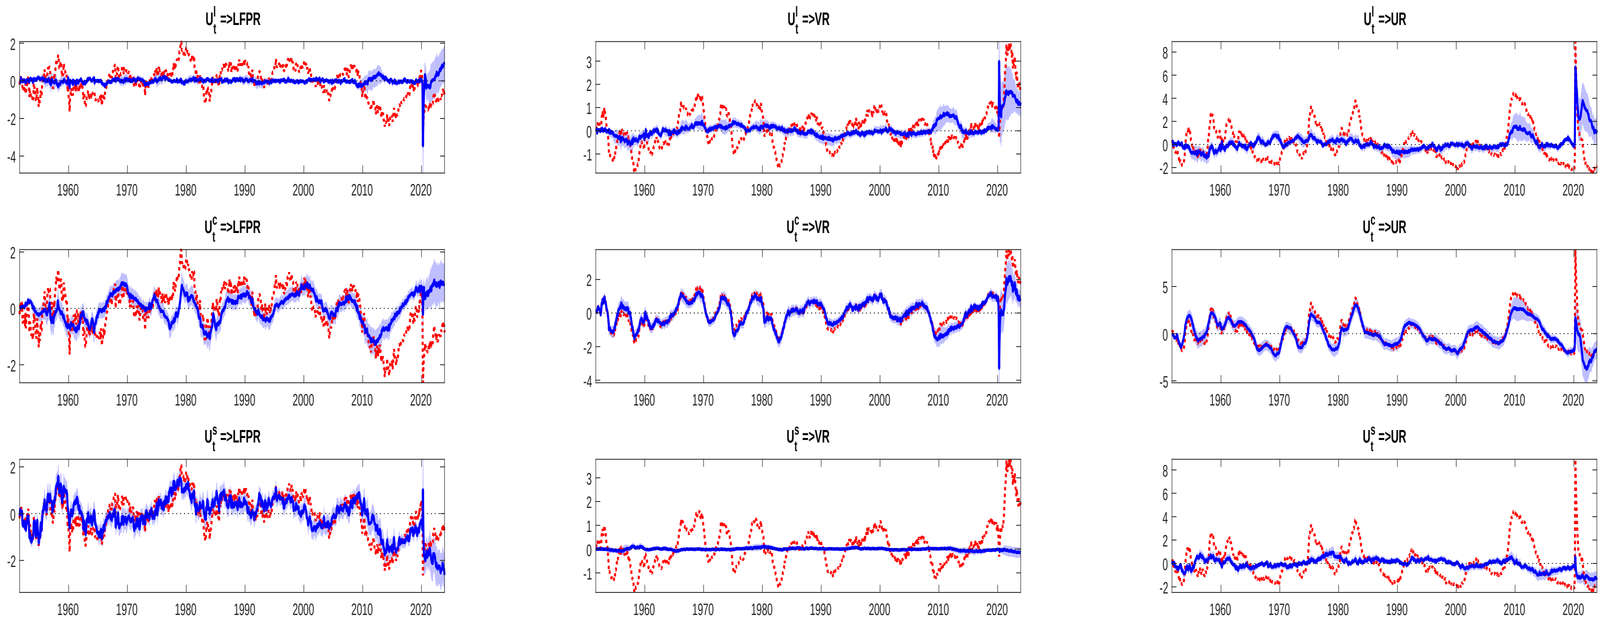

Notes: The posterior median is shown as a solid line; the shaded area gives the 68 percent credible set. We use the following acronyms: *LFPR* : the labor force participation rate; *VR* : the vacancy rate; *UR* : the unemployment rate. U_t^l: reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock.
Source: Authors' calculations.

A.4 Robustness checks

We examine the robustness of our estimates along the following three dimensions. First, we use market-based core inflation without an adjustment for the effects of the 1970s price controls in order to see whether this adjustment influences our results. Second, we relax the sign restrictions imposed on the impact effects of quits reallocation shocks on wage and price inflation. This is because some evidence suggests that workers who are prone to involuntary separations or who have unstable jobs are likely to quit their jobs (Ahn et al., 2023); these workers may not necessarily experience wage increases after quitting their jobs. Finally, we consider the case where there are additional zero restrictions on the impact effects that own shocks to wages and prices have on the unemployment rate of job losers and the impact effects that labor supply shocks have on wage and price inflation. The results from these exercises are broadly similar to what we obtain from our baseline specification.

A.4.1 Estimation results with unadjusted market-based core inflation

Figure A.4 reports the impulse responses and Figures A.5, and A.6 report the historical decompositions.

A.4.2 Estimation results under alternative restrictions

We first consider the case where the impact effects of the quits-reallocation shock on wage and price inflation are not required to be positive. Figures A.7 reports the impulse responses, and Figures A.8, and A.9 display the historical decompositions.

Figure A.4: Impulse responses with AHE and unadjusted market-based core inflation

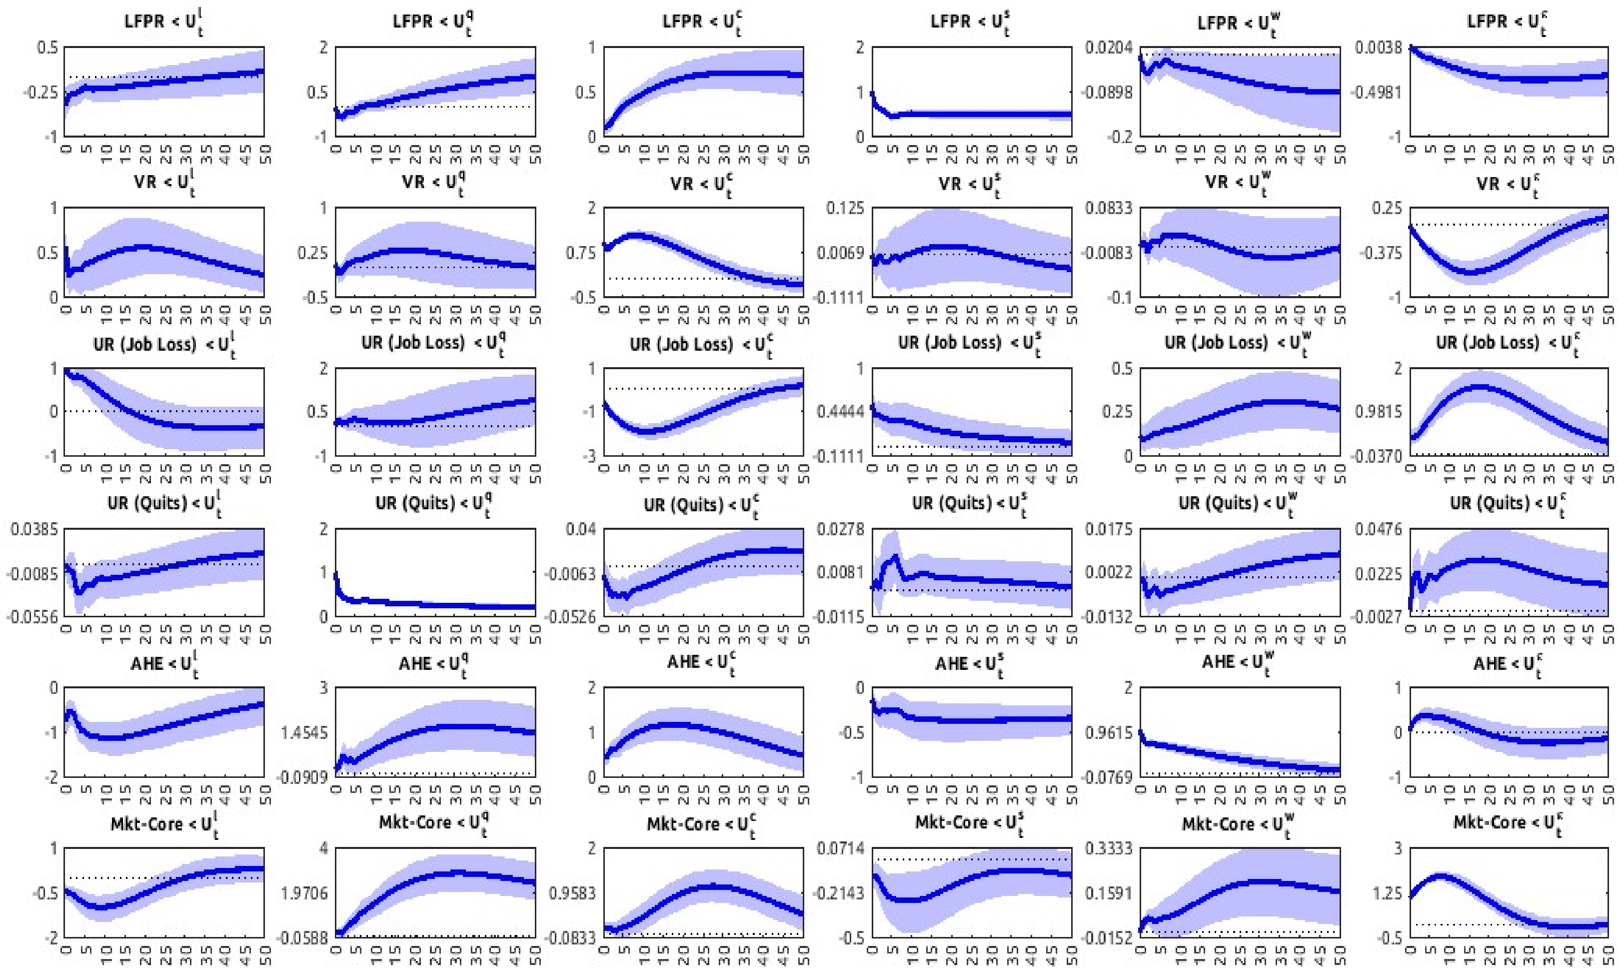

Notes: Posterior median (solid curve) and 68 percent credible set (shaded regions). The X-axis denotes months after the shock's impact. We use the following acronyms: *LFPR*: Labor force participation rate; *VR*: Vacancy rate; *UR (Job Loss)*: Unemployment rate of job losers and entrants to the labor force; *UR (Quits)*: Unemployment rate of job quitters; *AHE*: Average hourly earnings; *Mkt core*: Market-based core inflation; U_t^l: Job-loss reallocation shock; U_t^q: Job-quits reallocation shock; U_t^c: Aggregate activity shock; U_t^s: Labor supply shock; U_t^w: Wage-specific shock; U_t^{π}: Price-specific shock. In the panel titles, $x < y$ denotes the effect of shock y on variable x.
Source: Authors' calculations.

Figure A.5: Historical decomposition: Model with unadjusted market-based core inflation (1)

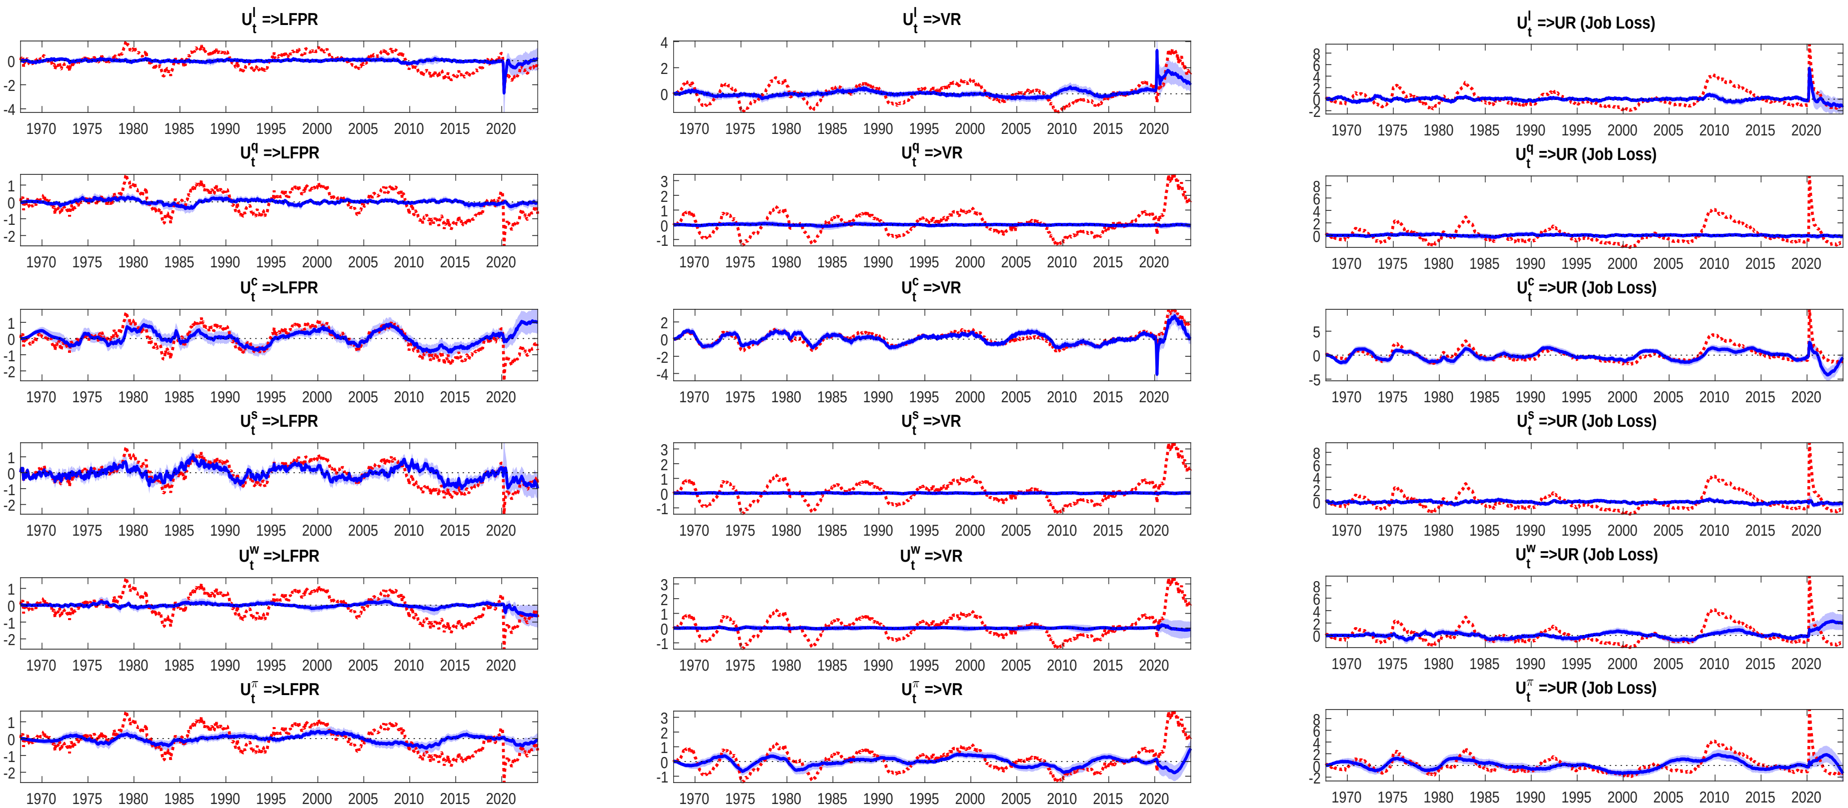

Notes: The posterior median is shown as a solid line; the shaded area gives the 68 percent credible set. *LFPR*: the labor force participation rate; *VR*: the vacancy rate; *UR (Job loss)*: the unemployment rate of job losers and entrants to the labor force; *UR (Quits)*: the unemployment rate of job quitters; *AHE*: Average hourly earnings; *Mkt*: Market-based core inflation. U_t^l: job-loss reallocation shock; U_t^q: job-quits reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock; U_t^w: wage-specific shock; U_t^π: price-specific shock.

Source: Authors' calculations.

Figure A.6: Historical decomposition: Model with unadjusted market-based core inflation (2)

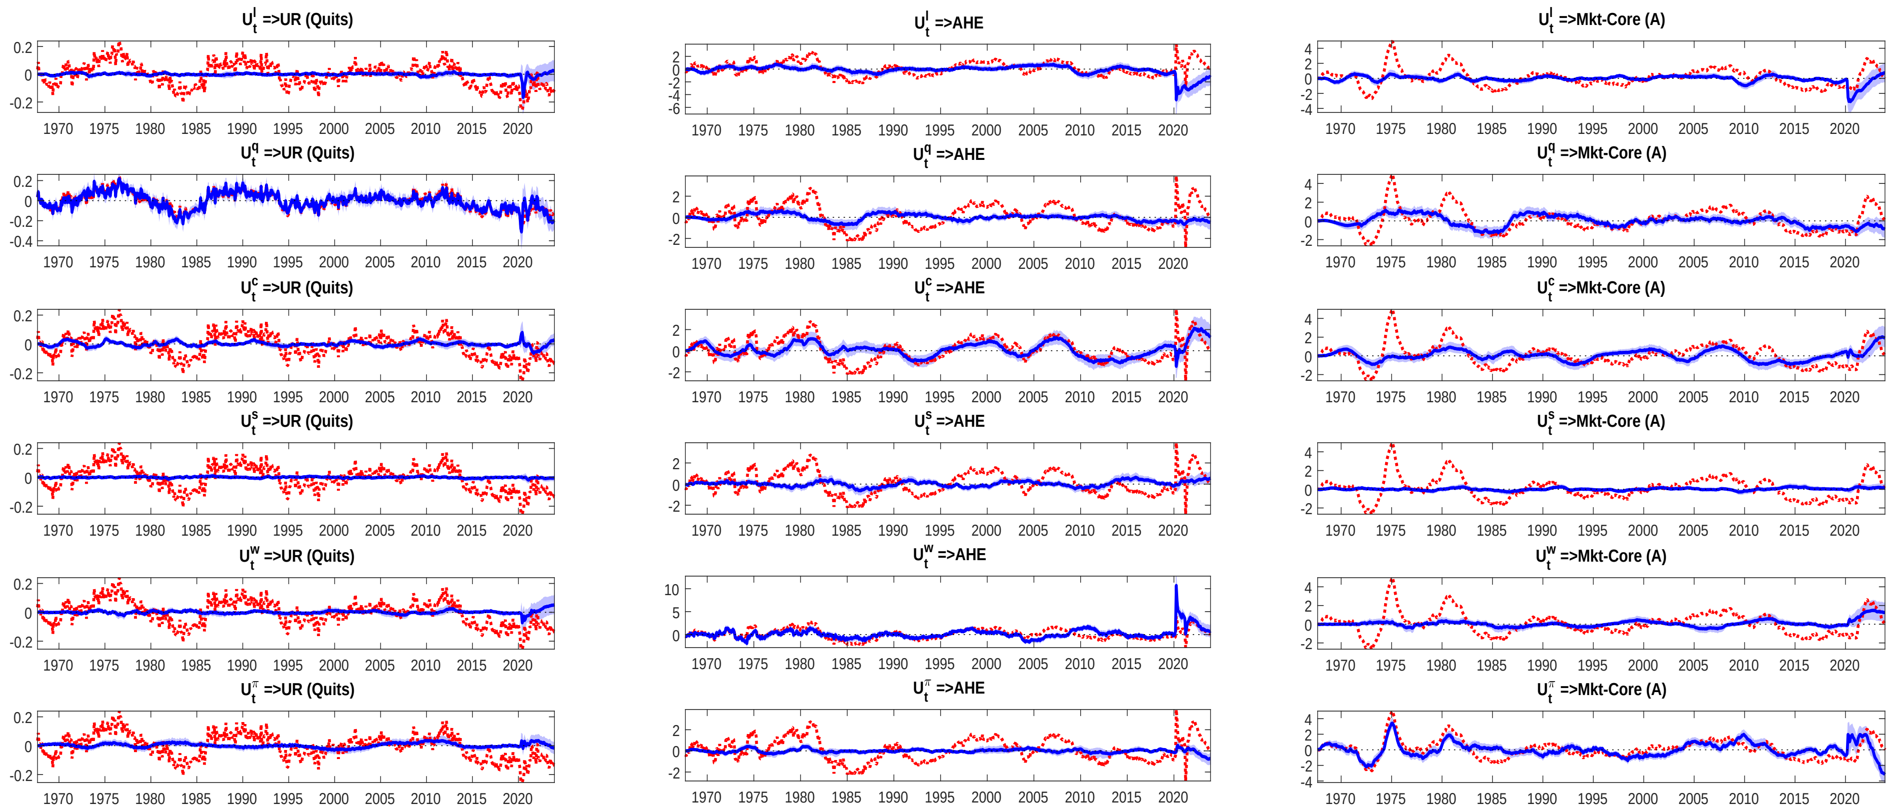

Notes: The posterior median is shown as a solid line; the shaded area gives the 68 percent credible set. *LFPR*: the labor force participation rate; *VR*: the vacancy rate; *UR (Job loss)*: the unemployment rate of job losers and entrants to the labor force; *UR (Quits)*: the unemployment rate of job quitters; *AHE*: Average hourly earnings; *Mkt-core*: Market-based core inflation. U_t^l: job-loss reallocation shock; U_t^q: job-quits reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock; U_t^w: wage-specific shock; U_t^π: price-specific shock.

Source: Authors' calculations.

Figure A.7: IMPULSE RESPONSES: NO SIGN RESTRICTIONS ON THE IMPACT EFFECTS OF QUITS REALLOCATION SHOCK ON WAGE AND PRICE INFLATION

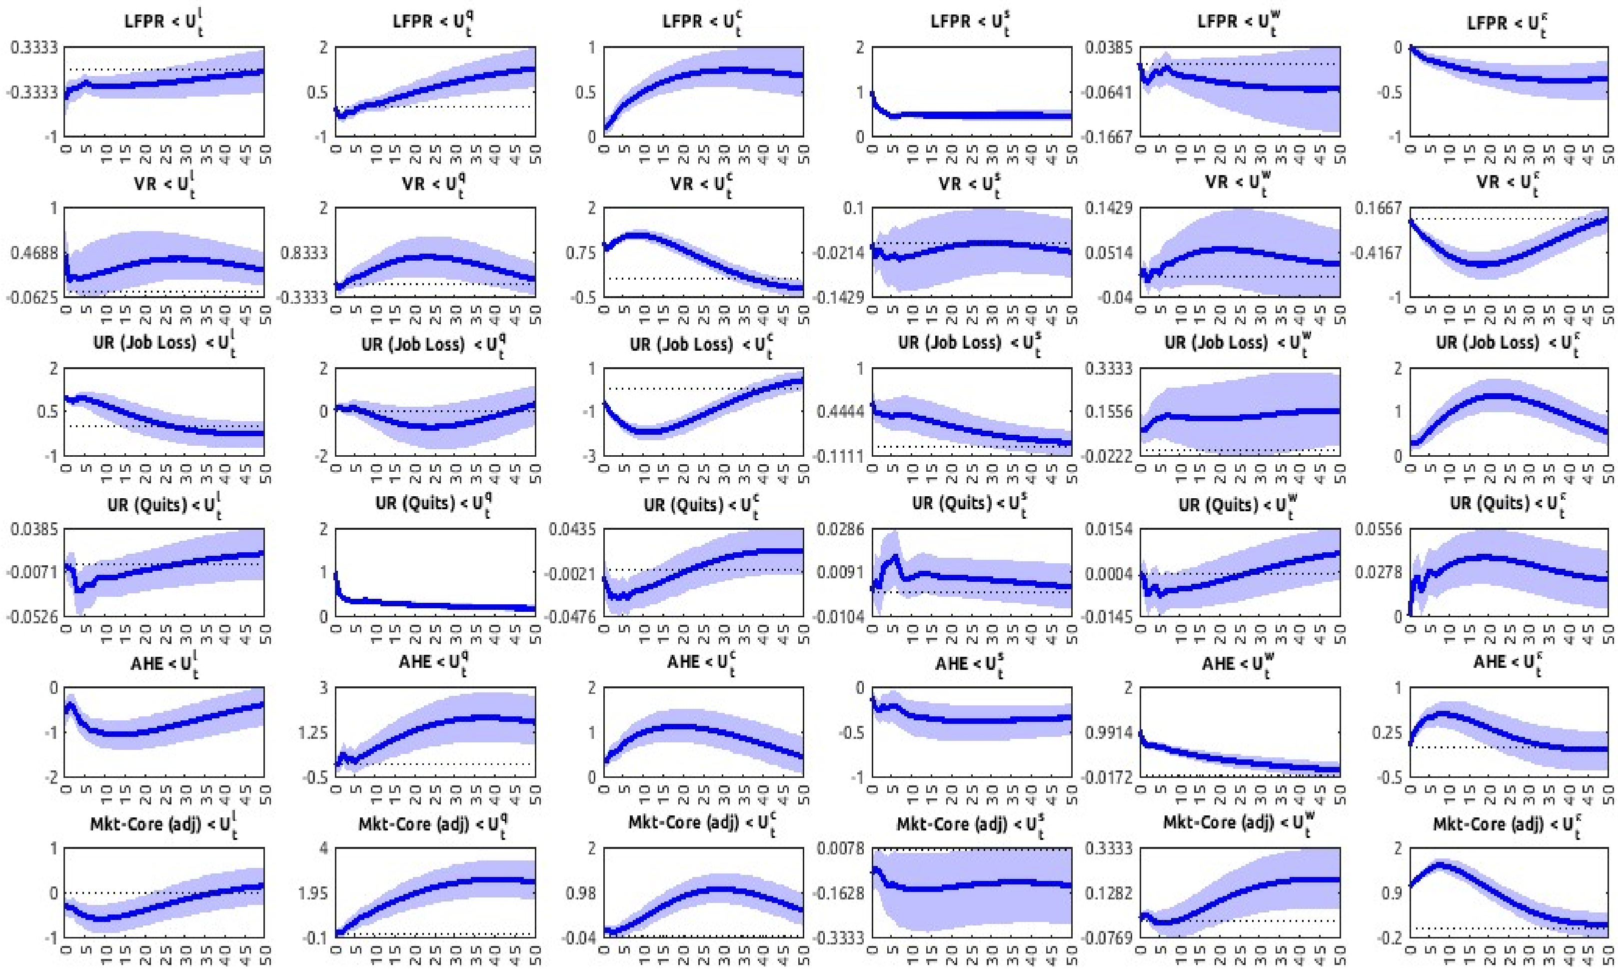

Notes: Posterior median (solid curve) and 68 percent credible set (shaded regions). The X-axis denotes months after the shock's impact. We use the following acronyms: *LFPR*: Labor force participation rate; *VR*: Vacancy rate; *UR (Job Loss)*: Unemployment rate of job losers and entrants to the labor force; *UR (Quits)*: Unemployment rate of job quitters; *AHE*: Average hourly earnings; *Mkt-Core (adj)*: Adjusted market-based core inflation; U^l_t: Job-loss reallocation shock; U^q_t: Job-quits reallocation shock; U^c_t: Aggregate activity shock; U^s_t: Labor supply shock; U^w_t: Wage-specific shock; U^π_t: Price-specific shock. In the panel titles, $x < y$ denotes the effect of shock y on variable x.
Source: Authors' calculations.

Figure A.8: HISTORICAL DECOMPOSITION (1): NO SIGN RESTRICTIONS ON THE IMPACT EFFECTS OF QUITS REALLOCATION SHOCK ON WAGE AND PRICE INFLATION

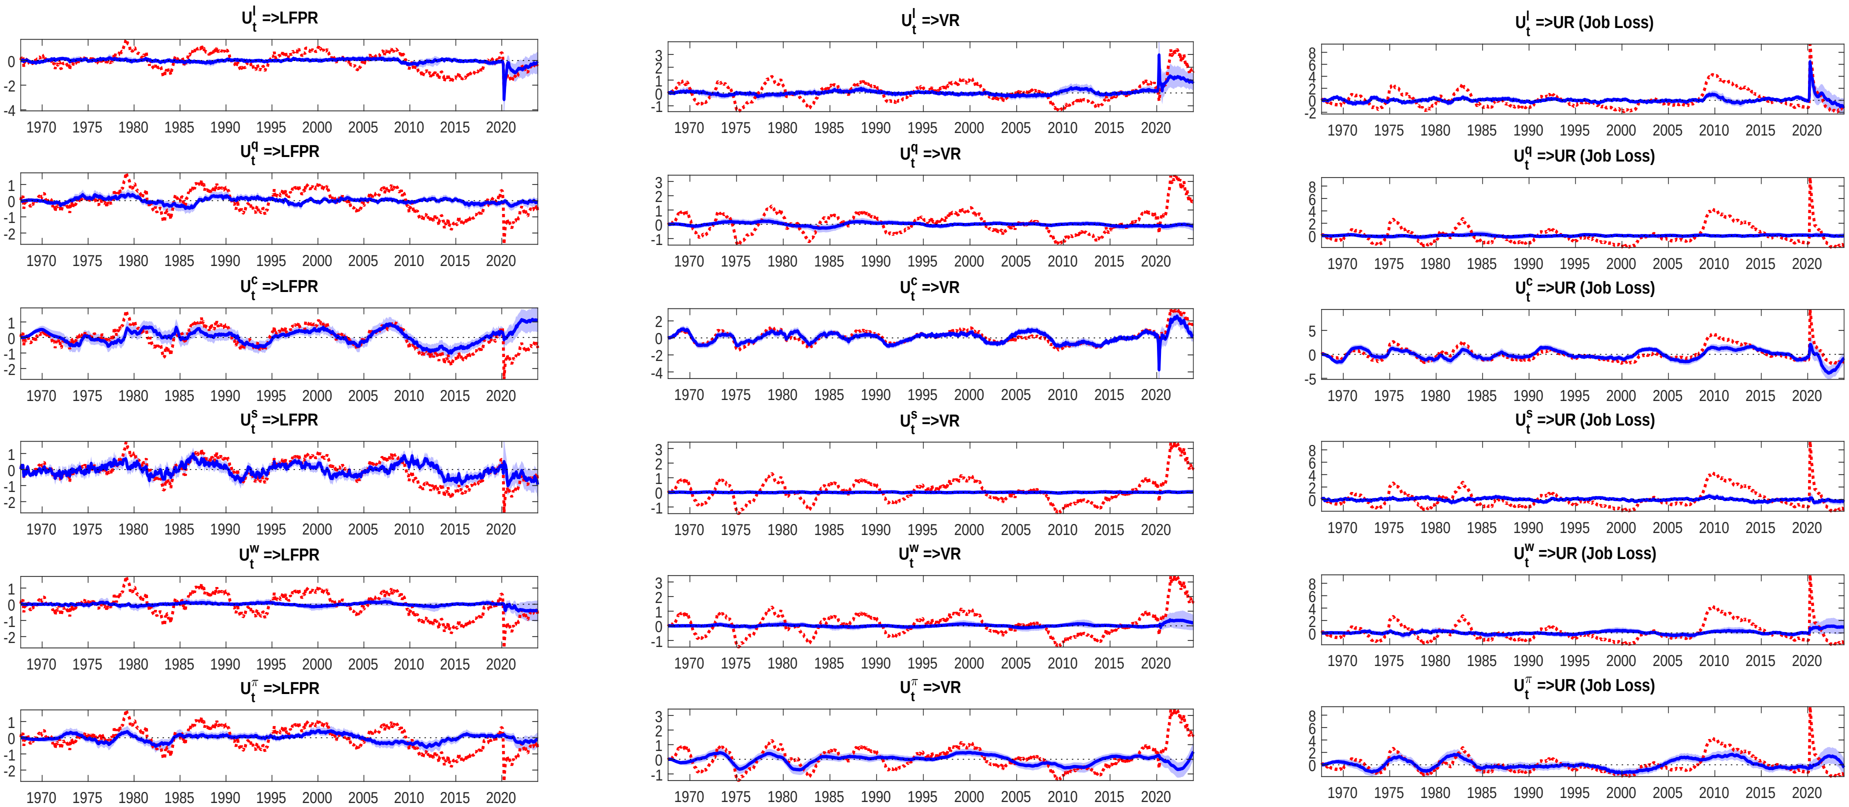

Notes: The posterior median is shown as a solid line; the shaded area gives the 68 percent credible set. *LFPR*: the labor force participation rate; *VR*: the vacancy rate; *UR (Job loss)*: the unemployment rate of job losers and entrants to the labor force; *UR (Quits)*: the unemployment rate of job quitters; *AHE*: Average hourly earnings; *Mkt-Core (A)*: Adjusted market-based core inflation. U_t^l: job-loss reallocation shock; U_t^q: job-quits reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock; U_t^w: wage-specific shock; U_t^π: price-specific shock.

Source: Authors' calculations.

Figure A.9: Historical decomposition (2): No sign restrictions on the impact effects of quits reallocation shock on wage and price inflation

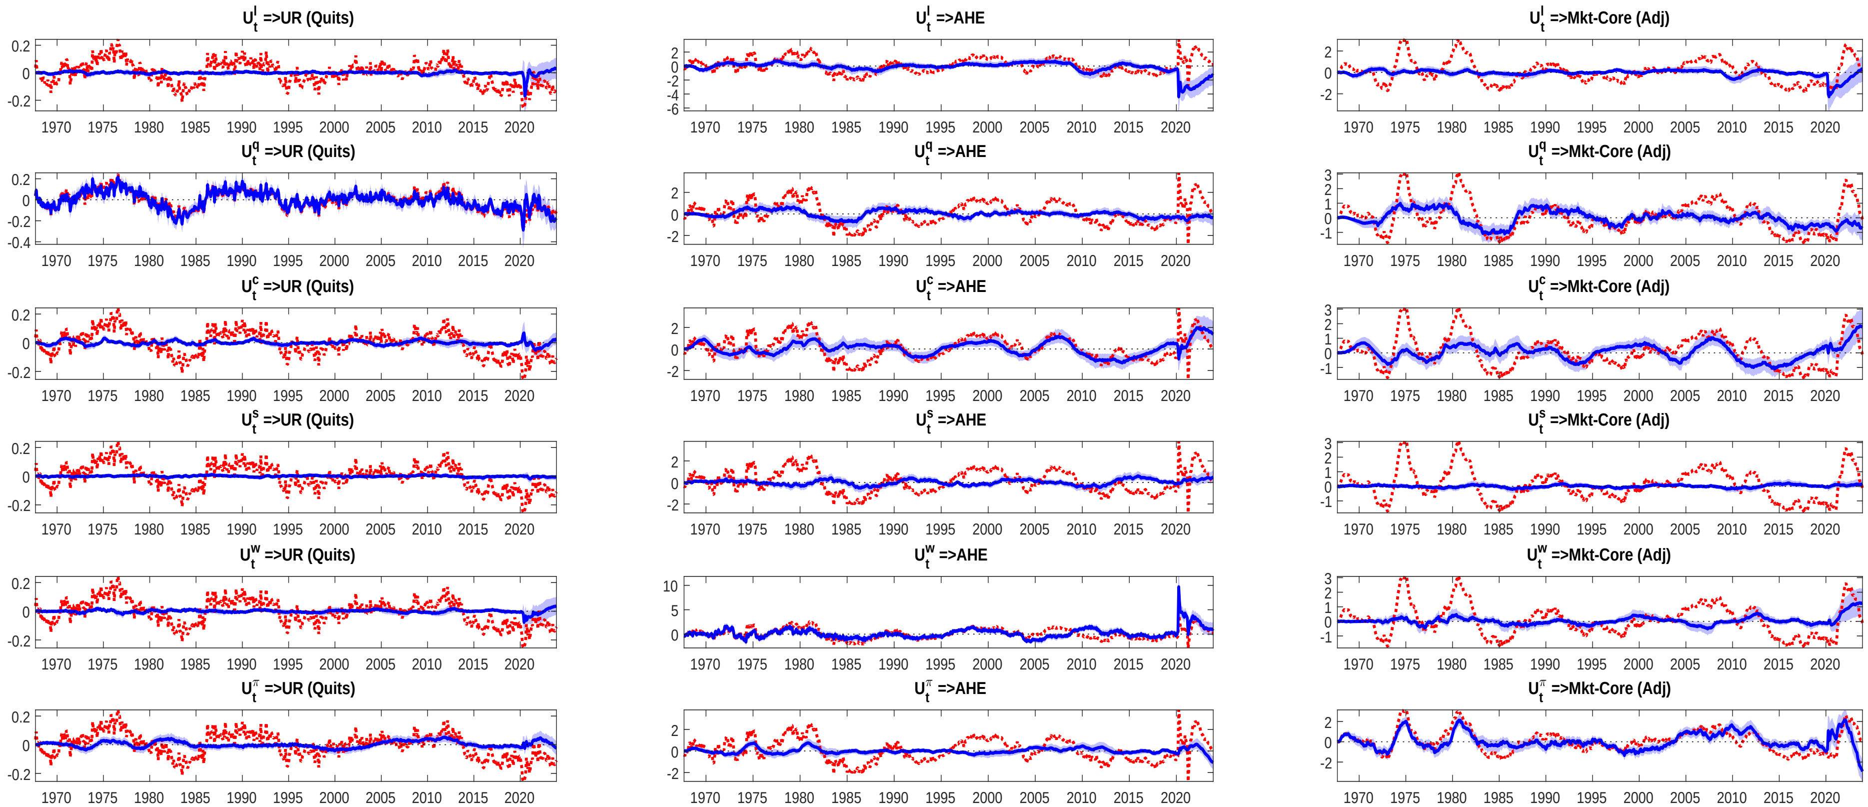

Notes: The posterior median is shown as a solid line; the shaded area gives the 68 percent credible set. *LFPR*: the labor force participation rate; *VR*: the vacancy rate; *UR (Job loss)*: the unemployment rate of job losers and entrants to the labor force; *UR (Quits)*: the unemployment rate of job quitters; *AHE*: Average hourly earnings; *Mkt-core (A)*: Adjusted market-based core inflation. U_t^l: job-loss reallocation shock; U_t^q: job-quits reallocation shock; U_t^c: aggregate activity shock; U_t^s: labor supply shock; U_t^w: wage-specific shock; U_t^π: price-specific shock.
Source: Authors' calculations.

A.5 Phillips correlations from the models with alternative restrictions on the impact matrix

This section reports additional estimates of Phillips correlations. First, Figure A.10 plots the full set of baseline dynamic Phillips correlations that are reported in Table 4 for specific horizons. Second, we report implied wage and price Phillips correlations for two alternative models: (1) the model with unadjusted market-based core inflation (Figure A.11); and (2) the model where the impact effects of quits reallocation shock on wage and price inflation do not have explicit sign restrictions (Figure A.12). The results are similar to those obtained from our baseline specification.

Figure A.10: CUMULATIVE WAGE AND PRICE PHILLIPS CORRELATION BY STRUCTURAL SHOCK

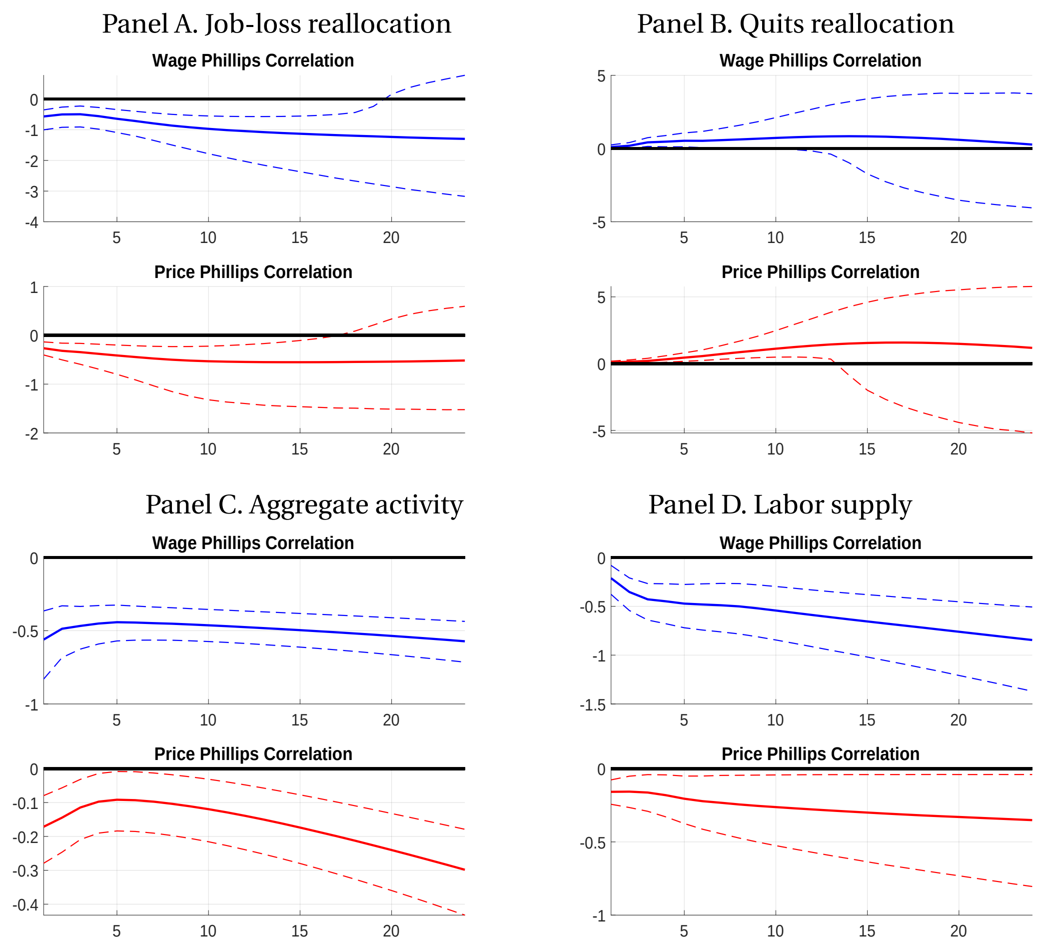

Notes: This figure displays the dynamic wage and price Phillips correlations conditional on the structural shocks. The dashed lines denote 68 percent posterior intervals. The X-axis is the horizon of response after the impact of shock. The Y-axis is the ratio of cumulative response of wage or price inflation to that of unemployment rate by each structural shock. Rows in each panel (h) capture the magnitude of correlation h-period after the shock's impact.
Source: Authors' calculations.

Figure A.11: WAGE AND PRICE PHILLIPS CORRELATIONS FOR EACH STRUCTURAL SHOCK USING UNADJUSTED MARKET-BASED CORE INFLATION

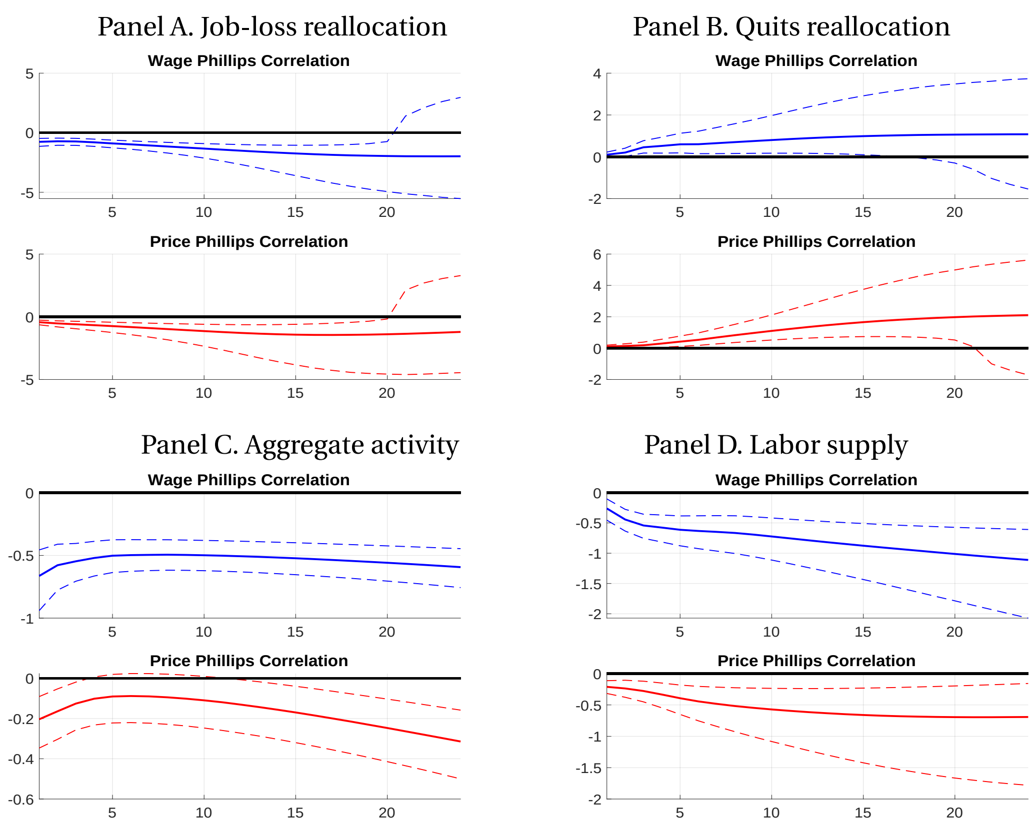

Notes: This figure displays the dynamic wage and price Phillips correlations conditional on the structural shocks. The dashed lines denote 68 percent posterior intervals. The X-axis is the horizon of response after the impact of shock. The Y-axis is the ratio of cumulative response of wage or price inflation to that of unemployment rate by each structural shock. Rows in each panel (h) capture the magnitude of correlation h-period after the shock's impact.
Source: Authors' calculations.

Figure A.12: WAGE AND PRICE PHILLIPS CORRELATIONS FOR EACH STRUCTURAL SHOCK (NO SIGN RESTRICTIONS ON IMPACT EFFECTS OF QUITS REALLOCATION SHOCK ON WAGE AND PRICE INFLATION)

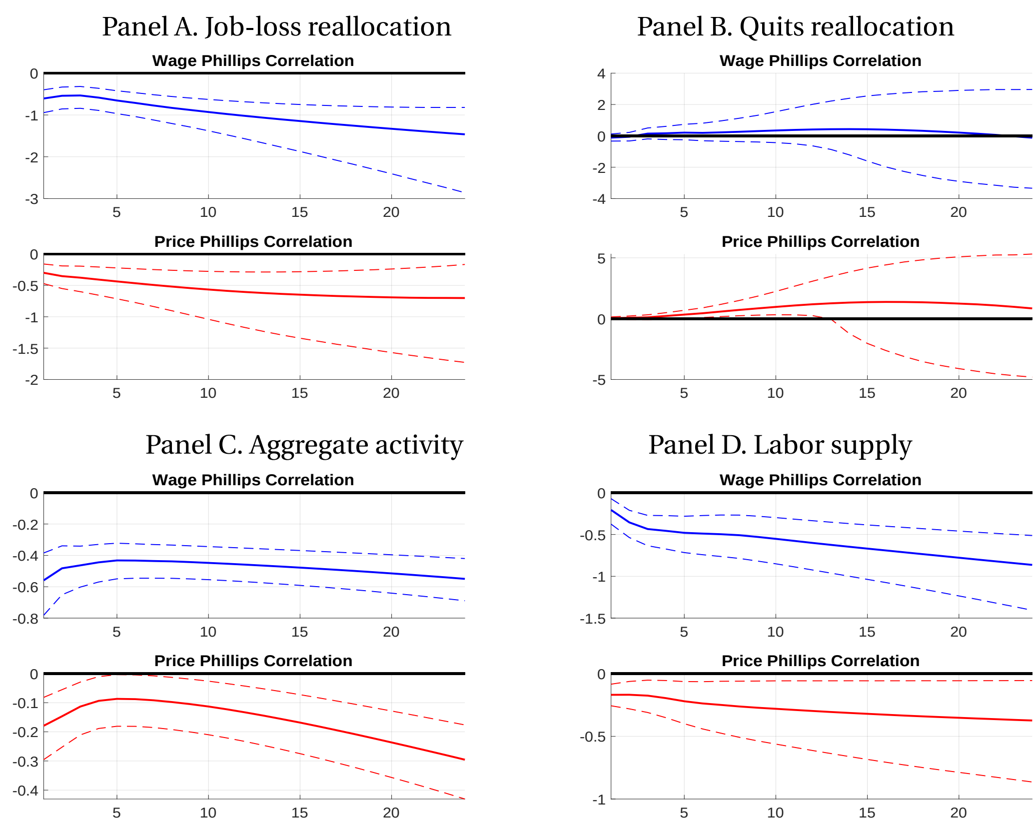

Notes: This figure displays the dynamic wage and price Phillips correlations conditional on the structural shocks. The dashed lines denote 68 percent posterior intervals. The X-axis is the horizon of response after the impact of shock. The Y-axis is the ratio of cumulative response of wage or price inflation to that of unemployment rate by each structural shock. Rows in each panel (h) capture the magnitude of correlation h-period after the shock's impact.
Source: Authors' calculations.