WARWICK

ECONOMICS

Search Frictions and Evolving Labour Market Dynamics

Michael Ellington, Chris Martin & Bingsong Wang

May 2019 No: 1195

Warwick Economics Research Papers

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

Search Frictions and Evolving Labour Market Dynamics

Michael Ellington[*]
University of Liverpool

Chris Martin[†]
University of Bath

Bingsong Wang[‡]
University of Warwick

May 20, 2019

Abstract

This paper puts search frictions models under novel empirical scrutiny and tests their ability to match empirical observations. To capture changing dynamics we fit an extended Bayesian time-varying parameter VAR to US labour market data from 1962–2016. Our results indicate that these models are unable to match the responses of key variables to identified structural shocks. We document substantial changes in the transmission, and economic importance, of shocks throughout time that this framework is unable to explain. In particular, although search frictions models use productivity and job separations shocks to explain the labour market, these shocks only explain at best 50% of fluctuations in key labour market variables.

Keywords: time-varying parameter model, real wages, search frictions,

JEL Classification: *E23, E32, J23, J30, J64*

Acknowledgements

We would like to thank Gary Koop and Marco Fogoni for their insightful comments. We are also grateful to seminar participants at: The University of Bath; Brunel University; The University of Strathclyde; and the 12$^{\text{th}}$ International Conference on Computational and Financial Econometrics held at the University of Pisa in December 2018. Declarations of interest: none.

1 Introduction

Macroeconomic models of the labour market look to explain cyclical, long-run, and secular relationships among key variables; namely unemployment, job vacancies and wages. Search and matching models, developed by Diamond (1982), Mortensen and Pissarides (1994) and Pissarides (2000), are the workhorse of modern labour economics. This framework examines the incentives of firms to post vacancies, how unemployed workers find a job match, and the resulting wage of a successful job match. They also provide an explanation for the underlying structural dynamics of the labour market, and historically, are successful in assessing the welfare implications of labour market policies. Their success stems from their ability to match key empirical regularities in the data, such as the negative link between unemployment and vacancies.

[*]m.ellington@;liverpool.ac.uk; Management School, University of Liverpool, Liverpool L69 7ZH UK.
[†]c.i.martin@bath.ac.uk; Department of Economics, University of Bath, Bath BA2 7AY UK.
[‡]b.wang.9@warwick.ac.uk; Department of Economics, University of Warwick, Coventry CV4 7AL, UK

In this paper we subject the search frictions model to novel empirical scrutiny. We find that the model has three key weaknesses. First, the model is not able to match the responses of key labour market variables to identified structural shocks. This suggests that the search frictions approach does not capture dynamics adequately. In particular, the vacancy surge required by search frictions models in order to match the response of unemployment to a productivity shock is not observed in the data. Second, the most prominent search frictions models in the literature require a negative relationship between the volatilities of wages and unemployment in order to generate a large volatility of unemployment. This negative relationship is not present in the data. Third, search frictions models use productivity and job destruction shocks to explain movements in labour market variables; but these shocks explain at best 50% of total variation.

One of the main drawbacks of the search and matching approach is that structural relationships are assumed to be constant throughout time. A growing empirical literature using models accounting for parameter and volatility variation within labour markets casts doubt on this (see e.g. Benati and Lubik (2014); Mumtaz and Zanetti (2015); Guglielminetti and Pouraghdam (2017))[1]. These empirical models do not include the real wage in their specifications. We build on this literature in two ways. First, we fit a time-varying parameter VAR model (TVP VAR) to US labour market data from 1952–2016 that includes the real wage. We present tests that establish that support our use of a TVP VAR model and reject simpler empirical models. Using the model, we find evidence of statistically significant and economically meaningful parameter and volatility change. We then further develop the literature by using our estimated TVP VAR model to evaluate the transmission mechanism central to the search frictions model. In particular, our model allows us to investigate changes in the correlation between the volatilities of wages and unemployment and changes in estimated impulse response functions. Simpler empirical models would not enable us to do this.

Our work is pertinent to two main areas of literature; the first being the extensive debate on unemployment volatility. Shimer (2005) finds that a calibrated search frictions model, based on Nash wage bargaining, is unable to replicate the large volatility of unemployment in the data. His conjecture is that this is due to the response of wages to shocks dampening vacancy creation. In response, a large literature (eg, Hall (2005), Hagedorn and Manovskii (2008) and Hall and Milgrom (2008)) has developed alternative models that reduce the impact of labour market conditions on the wage and allows the impact of the shock to largely fall on vacancy creation. This enables the model to match the volatility of unemployment in the data.

These models have two testable features. First, a productivity shock leads to a large surge in vacancy creation and hence a reduction in unemployment. Second, since a stronger response of unemployment requires a more stable wage rate, a negative correlation between unemployment volatility and wage volatility arises. Our results suggest that neither of the above receives empirical support. More specifically, there is a strong response of vacancies to productivity shocks during the first half of our sample; but this is due to the large volatility of these shocks. The response is smaller in the second half of our sample, due to a lower shock volatility. But at no point do we observe the large surge of vacancy creation predicted by the Shimer hypothesis. Furthermore, our empirical estimates permit a direct test on the correlation between unemployment and wage variability using time-varying volatilities. These indicate that the negative link between unemployment and wage volatility is present only during the 1970s and 1980s; with strong positive correlations in the other decades.

Our work also relates to the literature that seeks to match key features of the data with calibrated simulations of the labour market. This literature seeks calibrations of a search frictions model that replicate

[1] Benati and Lubik (2014) find variation in the position and slope of the Beveridge Curve over time. Mumtaz and Zanetti (2015) find substantial time-variation in the response of key variables to study the response of key labour market variables to technology shocks.Guglielminetti and Pouraghdam (2017) find marked differences in the job creation process over time.

the volatilities and correlations in key labour market variables (see e.g. Yashiv (2006), Hall (2005) and Hagedorn and Manovskii (2008))[2]. Although this literature assesses whether search frictions models can match summary statistics of the data, it does not address the question of whether the key transmission mechanisms of the search frictions model are present in the data. In order to address this, we use a novel approach. Transmission mechanisms directly relate to impulse response functions rather than volatilities and correlations. We can therefore test the transmission mechanisms of the search frictions model by examining the match between estimated and simulated impulse response functions. A close match provides evidence that the transmission mechanisms are indeed present in the data. To do this, we calibrate a theoretical search frictions model to match one of the estimated impulse responses for each shock at a given date, and investigate how closely the other simulated responses match the estimated impulse responses.

Our results cast doubt on the search frictions framework for a number of reasons. First, a search frictions model calibrated to match the empirical response of unemployment to productivity shocks is unable to match the responses of vacancies and wages to these shocks. Second, simulated responses of all variables to job separations shocks are unable to match what we observe in the data. Third, our estimates provide substantial evidence against the assumption that, by construction, fluctuations in labour market variables can be entirely explained by productivity and job separations shocks. We find that these identified shocks explain at best half of the variance in labour market variables throughout our sample. The implication here is that there are other important shocks present in the data that the search and matching approach does not account for.

The immediate policy implications of this paper are twofold: first, policies should evolve with labour market dynamics. For instance, the large volatility of unemployment we document in the 1970s implies that the adverse effects of business cycles fell mainly on unemployed workers. Thereby suggesting a policy focus on protecting employment. However, in recent decades unemployment has become less volatile, so that business cycles are increasingly affecting employed workers, suggesting a shift in policy focus toward stabilising their incomes. Second, and more widely, policy should not rely upon a single theoretical framework. Our critical scrutiny of search frictions models suggests it may be a potentially misleading framework for policy evaluation.

The structure of the remainder of this paper is as follows. Section 2) describes data and outlines the econometric model; Section 3) presents reduced form results. In Section 4), we explain how we identify structural shocks and present our structural estimates. Section 5) outlines our theoretical search frictions model, present our results on the fit between estimated and simulated impulse responses and draws conclusions from these. Section 6) concludes and outlines areas for future research.

2 Data Description and Econometric Model

We use quarterly US data from 1952 to 2016, a period that contains 10 NBER recessions, ensuring that we are able to detect business cycle effects. Our choice of sample is entirely driven by available data, since the

[2]There are two main alternative methodologies in the literature. The first is to examine whether simulations of a calibrated search frictions model are consistent with key features of the data. Yashiv (2006) uses indirect inference (Gourieroux et al. (1993)) to test whether the autocorrelations and volatilities of unemployment, the hiring rate and the labour share as well as the correlations between these variables derived from a simulated search frictions model match the same statistics derived from an estimated reduced form VAR model. Less formally, a large stream of work (see e.g. Hall (2005), and Hagedorn and Manovskii (2008)) investigates whether results from simulations of a search frictions model match selected features of the labour market, most prominently the volatility of unemployment, but also the correlations between unemployment, vacancies and labour market tightness. The second is to estimate a structural search frictions model and compare the estimated structural parameters with widely-used values in the literature(see e.g. Lubik (2009); Faccini et al. (2013)).

vacancy data in Barnichon (2010*a*) ends in December 2016. Our measure of US productivity is constructed by dividing data on GDP, drawn from the Federal Reserve Bank of St. Louis database, by a measure of total hours worked collected by the Bureau of Labor Statistics (BLS); see Hall (2007) for a discussion of why output per worker hour is the appropriate measure of productivity in this context. We use data on unemployment, again collected by the BLS. For vacancies, we use the composite Help Wanted index proposed by Barnichon (2010*a*). For real wages, we take the logarithmic difference between the "Compensation Per Hour of the Nonfarm Business Sector"taken from the FRED Economic Database, and the US Consumer Price Index [3]. Labour productivity and real wages are converted into annual growth rates as 100 multiplied by the logarithmic difference of each respective series. We plot our US labour market data in Figure 1.

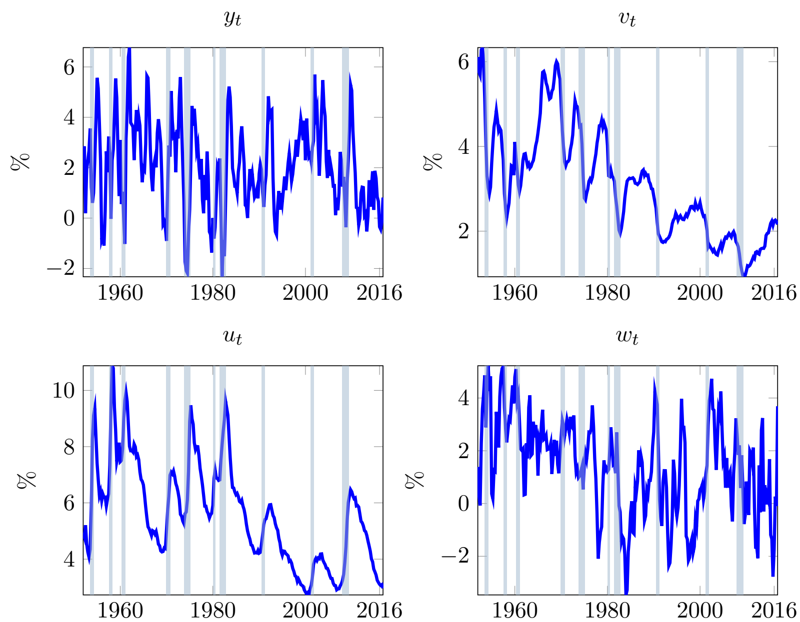

Figure 1: **US Macroeconomic data from 1952 to 2016**
Notes: This figure plots quarterly growth rates of US macroeconomic data from 1952Q1–2016Q4. The top left panel plots the annual growth rate of labour productivity, y_t; the top right panel plots the vacancy rate, v_t; the bottom left panel plots the unemployment rate, u_t; and the bottom right panel plots the annual growth rate of real wages, w_t. Grey bars indicate NBER recession dates.

This study works with the following TVP VAR model with $p = 2$ lags and $N = 4$ variables:

$$Y_t = \beta_{0,t} + \beta_{1,t}Y_{t-1} + \cdots + \beta_{p,t}Y_{t-2} + \epsilon_t \equiv X_t^{'}\theta_t + \epsilon_t \tag{1}$$

where $Y_t \equiv [y_t,\, v_t,\, u_t,\, w_t]^{'}$ is a vector of endogenous variables. Here y_t is labour productivity, v_t is the vacancy rate, u_t is the unemployment rate, and w_t is real wages. $X_t^{'}$ contains lagged values of Y_t and a constant. The VAR's time-varying parameters are collected in θ_t and evolve as

$$p(\theta_t|\theta_{t-1}, Q) = I(\theta_t)f(\theta_t|\theta_{t-1}, Q_t) \tag{2}$$

[3] We have also computed an alternative wage series in by using "Compensation of Employees, Received: Wages and Salary Disbursements", scaled by labour force participation. Specifically, we use series A576RC1 ("Compensation of Employees, Received: Wages and Salary Disbursements"), taken from the FRED Economic Database and scaled by the labour force participation, series CLF16OV. Real wages are then constructed as the logarithmic difference between this and the Consumer Price Index.

where $I(\theta_t)$ is an indicator function that rejects unstable draws, thereby imposing a stability constraint on the VAR where, conditional on the roots of the VAR polynomial lying outside the unit circle, $f(\theta_t|\theta_{t-1}, Q_t)$ follows a random walk. Adding an indicator function that rejects draws for the coefficient matrices in every t truncates and renormalises the prior. This stability constraint imposes a belief, a priori, that explosive representations of the model are implausible.

$$\theta_t = \theta_{t-1} + \gamma_t \tag{3}$$

with $\gamma_t \equiv [\gamma_{1,t}, : \gamma_{2,t}, ..., : \gamma_{N\cdot(Np+1),t}]'$, where $\gamma_t \backsim N(0, Q_t)$. Q_t is diagonal, and collecting these elements in the vector $q_t \equiv [q_{1,t}, : q_{2,t}, ..., : q_{N\cdot(Np+1),t}]'$, they evolve as geometric random walks

$$\ln q_{i,t} = \ln q_{i,t-1} + \kappa_t \tag{4}$$

with $\kappa_t \backsim N(0, Z_q)$. The innovations in (1) follow $\epsilon_t \backsim N(0, \Omega_t)$. Ω_t is the time–varying covariance matrix which is factored as

$$\Omega_t = A_t^{-1} H_t (A_t^{-1})' \tag{5}$$

with A_t being a lower triangular matrix with ones along the main diagonal, and the elements below the diagonal contain the contemporaneous relations. H_t is a diagonal matrix containing the stochastic volatility innovations. Collecting the diagonal elements of H_t and the non-unit non-zero elements of A_t in the vectors $h_t \equiv [h_{1,t}, : h_{2,t}, ..., h_{N,t}]'$, $\alpha_t \equiv [\alpha_{21,t}, : \alpha_{31,t}, \ldots, \alpha_{NN-1,t}]'$ respectively, they evolve as

$$\begin{aligned} \ln h_{i,t} &= \ln h_{i,t-1} + \eta_t && (6) \\ \alpha_t &= \alpha_{t-1} + \zeta_t && (7) \end{aligned}$$

where $\eta_t \backsim N(0, Z_h)$, and $\zeta_t \backsim N(0, S)$. The innovations in the model are jointly Normal, and the structural shocks, ψ_t are such that $\epsilon_t \equiv A_t^{-1} H_t^{\frac{1}{2}} \psi_t$. Similar to Primiceri (2005), S is a block diagonal matrix that implies the non-zero and non-unit elements of A_t evolve independently. The prior specification of our model are similar to Baumeister and Benati (2013). To calibrate the initial conditions of the model, we use the point estimates of a time-invariant VAR model estimated using the first 10 years of data. We estimate the model using Bayesian methods allowing for 20,000 runs of the Gibbs sampler. Upon discarding the initial 10,000 iterations as burn-in, we sample every 10^{th} draw to reduce autocorrelation. Details of our prior specification, and an outline of the posterior simulation algorithm is provided in the Online Appendix.

3 Reduced Form Estimates

Before presenting empirical results, it is necessary to evaluate the fit of our baseline model. We use the Bayesian deviance information criterion (DIC) proposed in Spiegelhalter et al. (2002). The DIC consists of two terms, one evaluating the fit of the model, and the other a penalty term for model complexity. Specifically, the DIC is given by

$$\text{DIC} = \bar{\text{D}} + \text{pD} \tag{8}$$

where $\bar{\mathrm{D}} = -2\mathbb{E}(\ln \mathrm{L}(\mathbf{\Lambda}_i))$, the measure of fit, is equal to minus two multiplied by the expected value of the log likelihood evaluated over the draws of the MCMC, and $\mathrm{pD} = \bar{\mathrm{D}} + 2\ln \mathrm{L}(\mathbb{E}(\mathbf{\Lambda}_i))$, is the measure of model complexity; with $\ln \mathrm{L}(\mathbb{E}(\mathbf{\Lambda}_i))$ being the log likelihood evaluated at the posterior mean of parameter draws. The lower the DIC, the better the model fit. For time-varying coefficient VARs with stochastic volatility, the DIC is estimated using a particle filter that evaluates the likelihood function to deal with the non-linear interaction of the stochastic volatilities (Mumtaz and Sunder-Plassmann, 2013). Restricted variants of the time-varying coefficient models include: a conventional Bayesian VAR; a time-invariant coefficient VAR with stochastic volatility; and a time-varying coefficient VAR with constant covariance matrix[45].

Table 1 reports the estimated DIC statistics, for competing models. It is clear that our time-varying coefficient VAR model with stochastic volatility fit the data best; relative to restricted variants. Based on these results, we proceed by reporting results from our TVP VAR model[6].

Table 1: **Bayesian DIC Statistics for Competing VAR Models**
Notes: This table reports the DIC statistics from a battery of competing Bayesian VAR models. The row highlighted in bold font indicates the model with the lowest DIC, and therefore the model that best fits the data.

	DIC
TVP VAR time-varying covariance matrix	**67.66**
TVP VAR constant covariance matrix	807.42
Bayesian VAR stochastic volatility	411.5
Linear Bayesian VAR	865.53

The upper panel of Figure 2 plots the posterior median and 80% highest posterior density intervals for the logarithmic determinant of the time-varying covariance matrices. As in Guglielminetti and Pouraghdam (2017), this proxies total prediction variation in the model, and is characterised as the amount of 'noise hitting the system'. This increases from the mid-1960s to a peak in the early 1980s. It then falls sharply before rising again in recent years. This pattern is reflected in the stochastic volatilities of US labour market variables, presented in the lower panel of Figure 2. The volatilities of productivity growth and vacancies fall in the early 1980s and account for the fall in the overall volatility in our model in that period. By contrast, the volatility of wages is gradually increasing throughout our sample, especially in the post-2008 period; this accounts for the increase in the estimated overall volatility in recent years[7].

Adding to this, Table 2 reports contemporaneous correlations between unemployment and wage volatilities implied by our TVP VAR. More specifically, we report correlations using posterior median estimates of stochastic volatilities: i) over decades; ii) by splitting the sample in half; and iii) using the full sample of 1962Q1–2016Q4. As we can see the only decades in which the correlation between unemployment and wage volatility is negative are the 1970s and 1980s; with correlations of -0.89 and -0.91 respectively. However, when looking at later decades, there is a stark change in the correlation between unemployment and wage

[4] These models were all estimated with standard priors within the literature. In particular, BVARs were estimated with a Minnesota prior on the coefficients, models with constant covariance matrices were assumed to have inverse-Wishart priors (see e.g. Koop and Korobilis (2010)), and those with time-varying parameters or stochastic volatility were estimated using analogous priors to the time-varying coefficient VAR models as outlined in the Appendix.

[5] We choose restricted variants of our extended TVP VAR as we do not wish to presume that periods of economic boom and recession can be represented by just two (or possibly three) sets of parameters; like regime-switching models impose.

[6] Available on request are results from estimated rolling VAR models. The dynamics of reduced form results from these simple rolling VARs are consistent with those provided in the main text.

[7] Reduced form volatilities have been investigated in the previous literature, although with different specifications of the VAR. Our estimate of overall model volatility and the volatility of vacancies is similar to Guglielminetti and Pouraghdam (2017). Our estimated volatility of unemployment is more stable than Mumtaz and Zanetti (2015). The previous literature has not modelled wages or productivity growth.

volatility with this being estimated as 0.97 from 2001Q1–2009Q4. Overall, our stochastic volatility estimates suggest that the argument of a negative link between unemployment and wage volatility is weak at best.

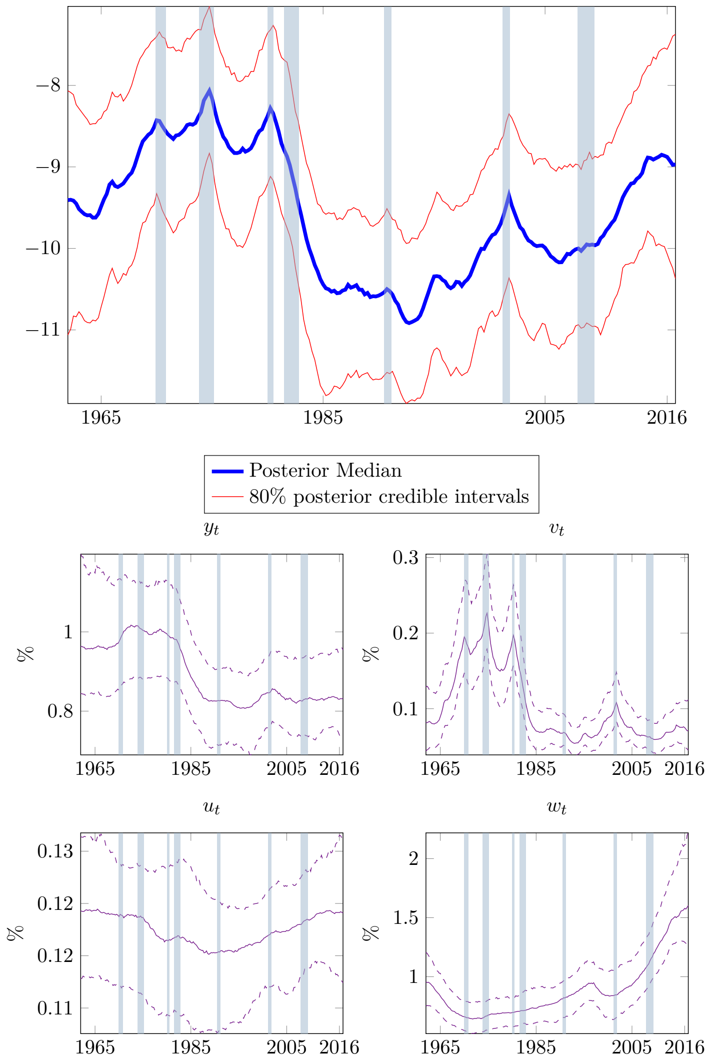

Figure 2: **Total Prediction Variation,** $\ln|\Omega_{t|T}|$**, from 1962 to 2016**
Notes: The upper panel plots the posterior median, and 80% posterior credible intervals of logarithmic determinant of the time-varying reduced-form covariance matrices, $\ln|\Omega_{t|T}|$, from 1962Q1–2016Q4. The lower panel plots the posterior median, and 80% posterior credible intervals of the reduced-from stochastic volatility innovations of labour productivity growth, y_t (top left panel); the vacancy rate, v_t (top right panel); the unemployment rate, u_t (bottom left panel); and real wage growth, w_t (bottom right panel) from 1962Q1–2016Q4. Grey bars indicate NBER recession dates.

In Figure 3, we report the estimated time-varying pairwise correlations between our variables. Note that our model is able to reproduce the switch in the correlation between productivity growth and unemployment growth from negative to positive in the early-1980s that was first shown by Barnichon (2010*b*). By contrast, there is no switch in the correlation between productivity growth and vacancies growth. The Beveridge Curve correlation between vacancies and unemployment is negative throughout our sample; yet varies markedly. In particular the Beveridge Curve correlation is highly negative in late 1960s to the early 1980s, but substantially muted thereafter. This switch in correlation may indicate that productivity shocks are not a major driving force of labour market dynamics, since these shocks imply a negative relationship between productivity and unemployment. The correlations between wages and the other variables in the model are small and never significant. The lack of correlation between the real wage and the other variables suggests that labour market conditions may not have had a strong impact on the real wage. It also suggests that the increase in wage volatility since 1980s has been independent of the other variables and so the search frictions framework may be unable to explain this increased volatility.

Table 2: **Model-Implied Correlations between Unemployment and Wage Volatility**
Notes: This table reports contemporaneous correlations between the posterior median stochastic volatility estimates of unemployment and wage volatility, $\hat{\rho}(\sigma_t^u, \sigma_t^w)$. The table reports correlations by decade, a sample split, and using the full sample of 1962Q1–2016Q4

Sample	$\hat{\rho}(\sigma_t^u, \sigma_t^w)$	Sample	$\hat{\rho}(\sigma_t^u, \sigma_t^w)$
1962Q1-1969Q4	0.84	1990Q1-1999Q4	0.33
1970Q1-1979Q4	-0.89	2000Q1-2009Q4	0.97
1980Q1-1989Q4	-0.91	2010Q1-2016Q4	0.76
1962Q1-1989Q2	0.11	1989Q3-2016Q4	0.87
1962Q1-2016Q4	0.34		

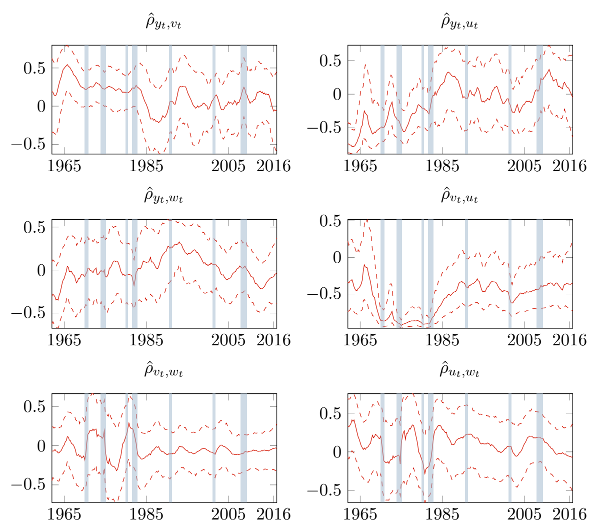

Figure 3: **Reduced-form correlations from 1962 to 2016**
Notes: This figure plots the posterior median, and 80% posterior credible intervals of the reduced-from model implied correlations of variables within the TVP VAR model from 1962Q1–2016Q4. $\hat{\rho}_{i_t,j_t}$ denotes the model implied correlation of variable i and j at time t respectively. y_t, v_t, u_t, w_t denote labour productivity growth, the vacancy rate, the unemployment rate, and real wage growth, respectively. Grey bars indicate NBER recession dates.

4 Structural Estimates

We partially identify the structural model using contemporaneous sign restrictions following a variant of Algorithm 1 in Arias et al. (2018) that is combined with the method proposed in Rubio-Ramirez et al. (2010); further details are provided in the Online Appendix as well the algorithm we use to compute impulse response functions (Koop et al., 1996). We identify a productivity and job separation shock using the sign restrictions outlined in Table 3.

Table 3: **Contemporaneous Impact of Identified Shocks on Labour Market Variables**
Notes: This table shows the contemporaneous sign restrictions imposed on variable $x = \{y_t,\ v_t,\ u_t,\ w_t\}$ to a productivity shock, ψ_t^{Prod}; and a job separation shock, ψ_t^{JS} respectively. y_t is the annual growth rate of labour productivity; v_t is the vacancy rate; u_t is the unemployment rate; and w_t is the annual growth in real wages.

	y_t	v_t	u_t	w_t
ψ_t^{Prod}	$\geq$	$\geq$	$\leq$	$\geq$
ψ_t^{JS}	$\leq$	$\geq$	$\geq$	$\leq$

These restrictions are consistent with a simple search frictions model. A positive productivity shock increases the value of a productive job match; this leads firms to post increased vacancies. This results in more job matches and so unemployment falls. The impact on wages depends on how wages are determined. The most widely-used forms of wage determination in the literature are either worker-firm Nash bargaining or the alternating offer bargaining protocol of Hall and Milgrom (2008). In either case, a positive productivity shock will increase the wage. A shock that increases the job separation rate will increase unemployment, reduce employment thereby reducing output and productivity. The reduction in employment leads firms to post more vacancies, and the real wage falls due to a reduction in labour market tightness.

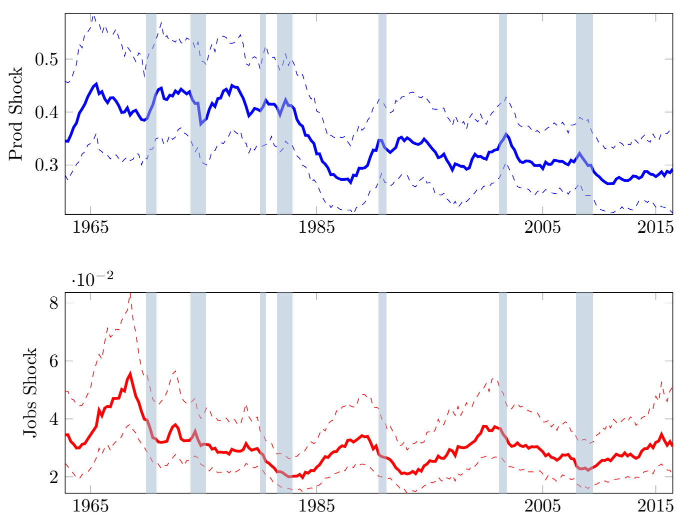

Figure 4: **Volatility of Structural Productivity and Job Separation Shocks from 1962 to 2016**
Notes: This figure plots the posterior median and 80% equal-tailed point-wise posterior probability bands for the volatility of identified productivity and job separations shocks from 1962Q1–2016Q4. Grey bars indicate NBER recession dates.

Figure 4 presents the estimated volatilities of the identified structural productivity and job separation shocks. There is a reduction in the volatility of productivity shocks from the early 1980s, which is consistent with the reduction in the volatilities of unemployment and vacancies in the reduced form estimates. There is no increase in the volatility of either structural shock over the past two decades. This suggests that neither structural shock can account for the rise in the volatility of wages in the reduced form model.

Figure 5 reports the estimated impulse response functions of US labour market data with respect to a one standard deviation productivity shock. Impulse response functions are normalised, such that the productivity shock causes labour productivity growth to increase by 1%. In the four leftmost quadrants, we report the posterior median impulse response functions of our variables throughout time, over a 20 quarter horizon. In the rightmost four quadrants, we examine the statistical credibility of productivity shocks by showing the evolution of the posterior median and 80% equal-tailed point-wise probability bands of impulse response functions four quarters after the shock. In the lower panel we plot the impulse response functions for 1974Q2, 1980Q2, 1982Q1 and 2001Q4, all relative to 2008Q4. These dates correspond to the mid-point of recessions as defined by the NBER, so these impulse responses correspond to similar points in the business cycle.

Our estimated impulse responses have similar features across our sample. Vacancies and wages increase in response to a positive productivity shock, while unemployment falls. Unemployment and vacancies respond gradually; vacancies reach their peak response within 1-2 quarters while the unemployment response reaches its peak within 2-3 quarters. After reaching the peak, the response of both variables gradually dies away. By contrast, the response of wages immediately jumps to its peak response and then dies away rapidly.

The estimated impulse responses also show evidence of substantial change. Impulse responses for unemployment and vacancies were large and countercyclical in the 1970s and early 1980s. From around 1985 onwards, the responses are relatively more subdued, and less cyclical. There are clear differences between these impulse responses during recessions in the 1970s and the recession of 2007-8. We also find that the response of wages to productivity shocks becomes stronger over time. The lower panel confirms this, showing that the response in the recent recession was much larger than in previous recessions. The Online Appendix provides statistical support for these conclusions, showing a clear and usually statistically significant difference between the impulse responses for unemployment and vacancies in 1974Q2 and 1980Q2 and the responses in 2008Q4. The gradual strengthening in the response of wages to a productivity shock results in a statistically significant difference between the responses in the 1974Q2 and 1980Q2 and 2008Q4. The upper panel of Figure 5) shows that after 1985, the response of wages to a productivity shock becomes stronger but there is no significant change in the response of unemployment to a productivity shock. This is consistent with the reduced form results above and provides further evidence for our second finding, that increased unemployment volatility does not require reduced wage volatility. In summary, the role of labour productivity in driving labour market dynamics has been diminishing since the late 1980's, both because productivity shocks have become less volatile (as shown in Figure 4) and because the response of unemployment and vacancies to those shocks has weakened (as shown in Figure 5).

Figure 6) reports the posterior median impulse response functions of US labour market data with respect to a job separation shock. The distribution of these impulses have been normalised such that the shock causes the unemployment rate to increase by 1% on impact. The four leftmost quadrants show the posterior median impulse response functions of our variables throughout time, over a 20 quarter horizon, and the rightmost four quadrants examine the statistical credibility of these impulse responses by reporting the posterior median and 80% equal-tailed point-wise probability bands of impulse response functions four quarters after the shock. In the lower panel, we plot the impulse response functions for 1974Q2, 1980Q2, 1982Q1 and 2001Q4, all relative to 2008Q4. The impulse responses for unemployment and vacancies are generally counter-cyclical in the 1970s and 1980s but display no apparent trend. By contrast, the impulse response for wages increases over time. We show in the Online Appendix that there are statistically significant changes in the responses of wages in 2008Q4 relative to the first three recessions in our sample; yet no statistical change in the response of unemployment and vacancies. Similar to Figure 5), this figure the increased response of wages to a separations shock is not associated with a reduced response of unemployment; our finding that increased

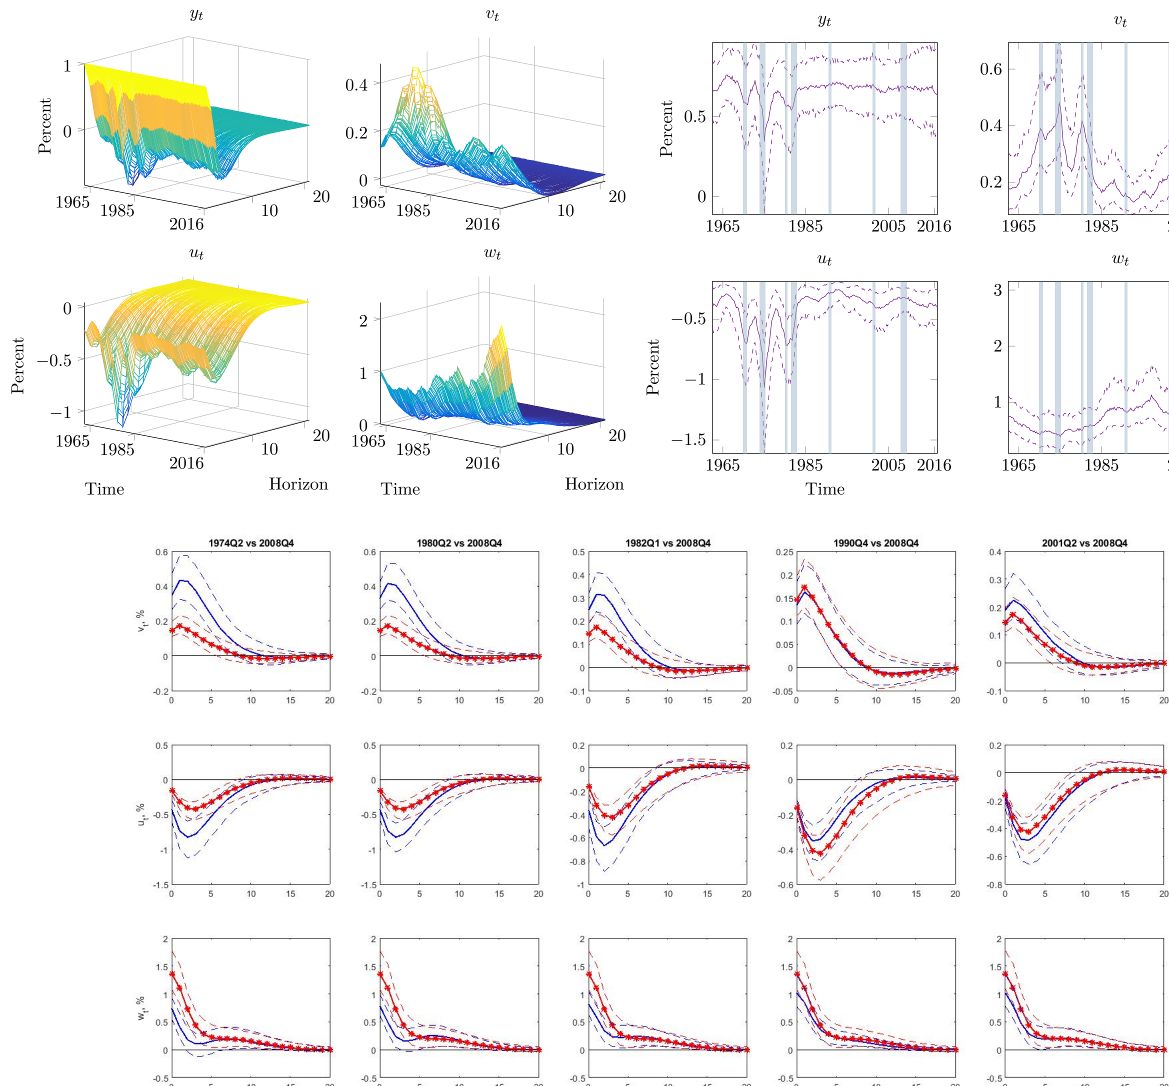

Figure 5: **Impulse Response Functions with Respect to a Productivity Shock from 1962 to 2016**
Notes: The top left hand side of this figure plots the posterior median generalised impulse response functions of US labour market data with respect to a one standard deviation productivity shock from 1962Q1 to 2016Q4. y_t, v_t, u_t, w_t denote the response of labour productivity growth, the vacancy rate, the unemployment rate, and real wage growth respectively. Impulse responses are computed for a 20 quarter horizon and normalised such that the shock causes labour productivity to increase by 1%. The top right hand side of this figure reports the posterior median and 80% equal-tailed point-wise posterior probability bands for the responses, at a 4 quarter horizon, of US labour market data with respect to a one standard deviation productivity shock from 1962Q1 to 2016Q4. y_t, v_t, u_t, w_t denote the response of labour productivity growth, the vacancy rate, the unemployment rate, and real wage growth respectively. Grey bars indicate NBER recession dates. The lower panel shows the posterior median generalised impulse response functions of US labour market data with respect to a one standard deviation productivity shock at selcted dates with 80% equal-tailed point-wise posterior probability bands

unemployment volatility does not require reduced wage volatility is not restricted to productivity shocks [8].

Figure 7) plots the posterior median and 80% equal-tailed point-wise probability bands of the percent share of forecast error variance attributable to productivity and job separation shocks at a 20 quarter horizon. We note that our structural shocks are only able to explain around half of the variance in our variables across our sample. This suggests that the data reflect the impact of structural shocks that are not considered by the search frictions model. The upper panel shows that forecast error variances for vacancies and unemployment with respect to productivity shocks are larger and more volatile in the 1970s and 1980s, similar to the impulse responses reported in Figure 5), but are generally procyclical rather than counter cyclical. The reduction in the contribution of these shocks after the early 1980s reflects the reduced volatility of this shock in this period shown in Figure 4). There is no clear trend in the forecast error variance for wages, in contrast to the estimated impulse responses[9]. The lower panel shows that job separation shocks had a much smaller impact on the volatility of unemployment and vacancies in the 1970s and 1980s. The Online Appendix shows that these differences are statistically significant relative to recessions occurring in the latter half of our sample. The forecast error variance for wages remains relatively constant, with no significant differences over time. The increasing strength of the impulse response of wages following a job separations shock is not reflected in an increased forecast error variance for this shock.

[8] Our estimates show a negative response of productivity in response to separations shocks. In the theoretical model outlined below, separations shocks do not affect productivity. Imposing this as a restriction on our estimates does not affect our main results.

[9] The Online Appendix shows no significant differences in forecast error variances across the different recessions in our sample.

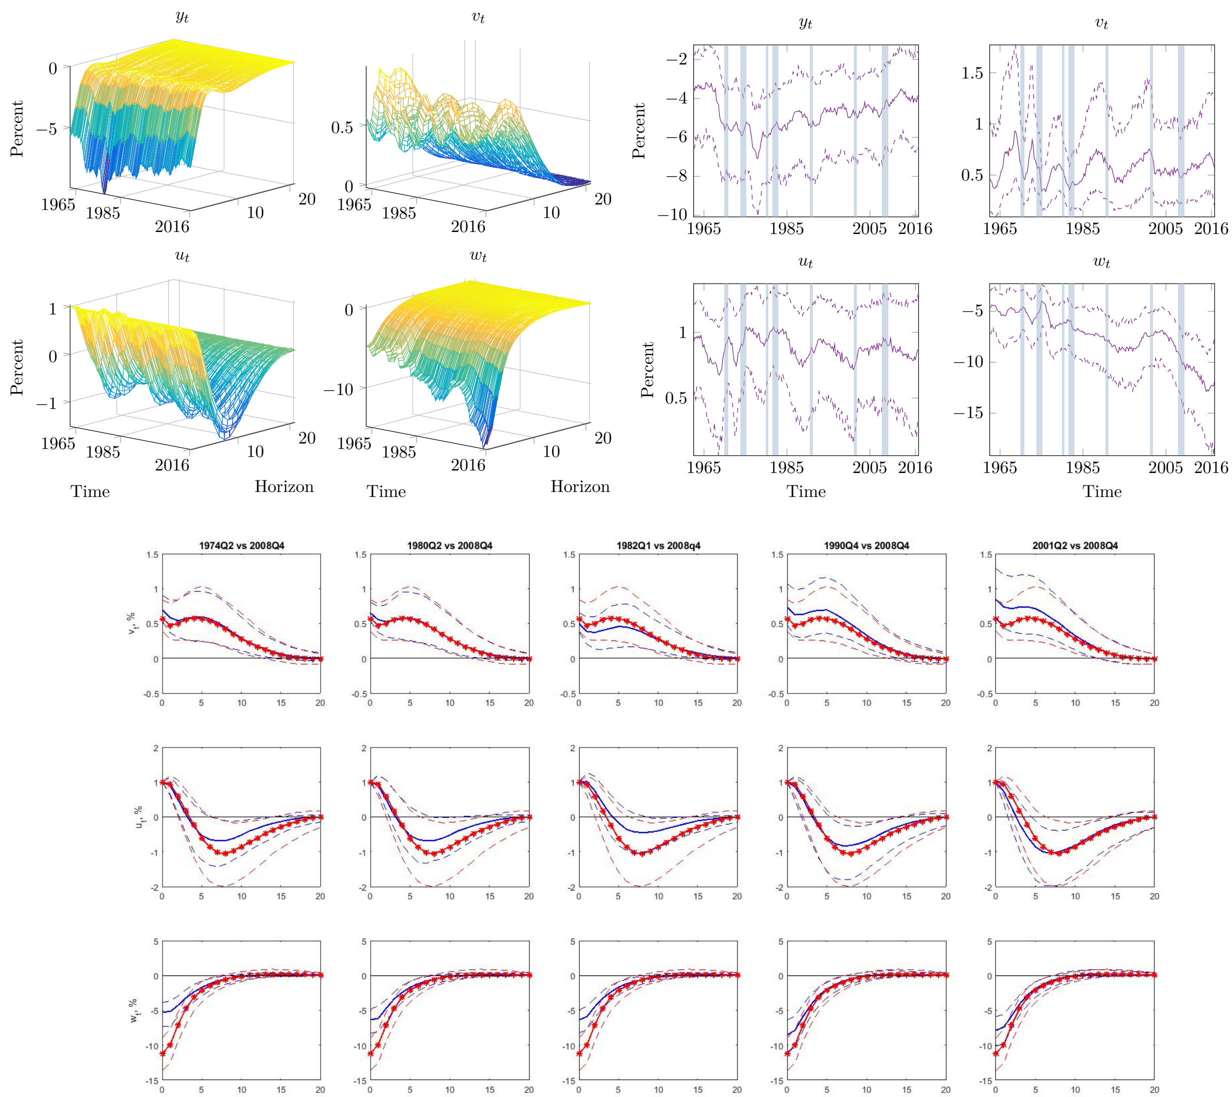

Figure 6: **Impulse Response Functions with Respect to a Job Separation Shock from 1962 to 2016**

Notes: The top left hand side of this figure plots the posterior median generalised impulse response functions of US labour market data with respect to a one standard deviation job separation shock from 1962Q1 to 2016Q4. y_t, v_t, u_t, w_t denote the response of labour productivity growth, the vacancy rate, the unemployment rate, and real wage growth respectively. Impulse responses are computed for a 20 quarter horizon and normalised such that the shock causes unemployment to increase by 1%. The top right hand side of this figure reports the posterior median and 80% equal-tailed point-wise posterior probability bands for the responses, at a 4 quarter horizon, of US labour market data with respect to a one standard deviation job separation shock from 1962Q1 to 2016Q4. y_t, v_t, u_t, w_t denote the response of labour productivity growth, the vacancy rate, the unemployment rate, and real wage growth respectively. Grey bars indicate NBER recession dates. The lower panel shows the posterior median generalised impulse response functions of US labour market data with respect to a one standard deviation job separation shock at selcted dates with 80% equal-tailed point-wise posterior probability bands

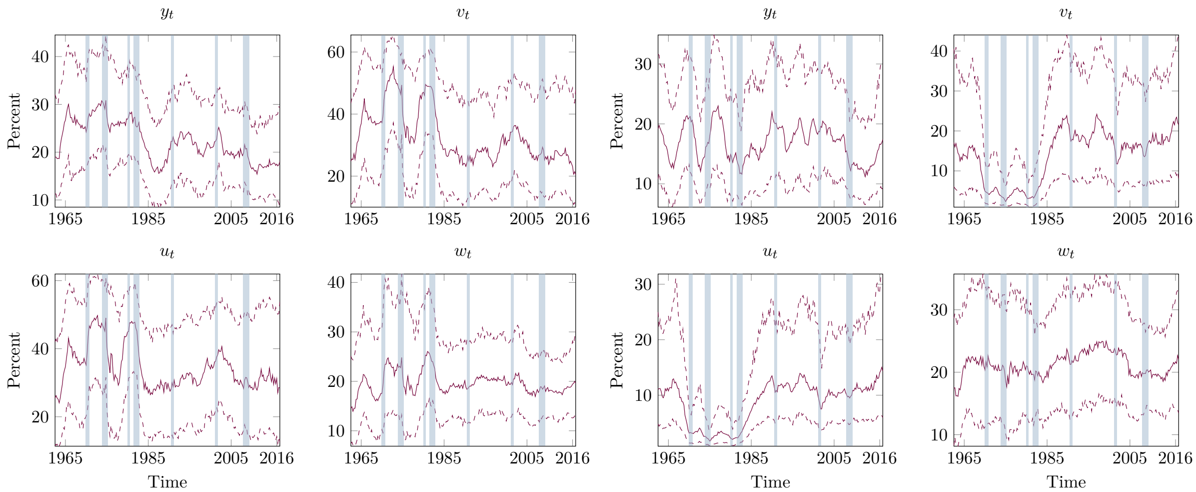

Figure 7: **Forecast Error Variance Decomposition of Productivity and Job Separation Shocks from 1962 to 2016**
Notes: This figure plots the posterior median, and 80% posterior credible intervals of the percent share of variance attributable, at a 20 quarter horizon, to productivity (four leftmost quadrants) and job separation shocks (four rightmost quadrants) for labour productivity growth, y_t; the vacancy rate, v_t; the unemployment rate, u_t; and real wage growth, w_t from 1962Q1–2016Q4. Grey bars indicate NBER recession dates.

5 Matching Theory to the Evidence

In this section, we calibrate and simulate a simple model with search frictions to investigate if it can match the evidence of changes to impulse responses documented above. The labour market is characterised by search frictions. Aggregate hiring is determined by the matching function

$$h_t = m u_t^{\alpha} v_t^{1-\alpha} \tag{9}$$

where h_t is the number of workers hired, u_t is unemployment and v_t are vacancies. m and α are parameters characterising the matching function. Defining labour market tightness as $\theta_t = \frac{v_t}{u_t}$ and the vacancy filling rate, the probability of a firm filling a vacancy, as $q_t = \frac{h_t}{v_t}$ and the job finding rate of unemployed workers as $f_t = \frac{h_t}{u_t}$, we can write $q(\theta_t) = m\theta_t^{-\alpha}$ and $f(\theta_t) = m\theta_t^{1-\alpha}$. Existing job matches dissolve at the end of the period with exogenous but time-varying probability τ_t; we assume that $\tau_t = \tau e^{\varepsilon_t^{\tau}}$ where $\varepsilon_t^{\tau} = \rho^{\tau}\varepsilon_{t-1}^{\tau} + \eta_t^{\tau}$ where η_t^{τ} is distributed as $N(0, \sigma_{\tau}^2)$. The evolution of unemployment is given by

$$1 - u_t = (1 - \tau_t)(1 - u_{t-1}) + h_t \tag{10}$$

There is a continuum of identical workers on the unit interval. The value function for an employed worker is

$$L_t = w_t + \frac{1}{1+r} E_t[(1-\tau_t)L_{t+1} + \tau_t U_{t+1}] \tag{11}$$

The worker earns (and consumes) real wage w_t. The value function for an unemployed worker is

$$U_t = z + \frac{1}{1+r} E_t[f_{t+1}L_{t+1} + (1-f_{t+1})U_{t+1}] \tag{12}$$

where z denotes the opportunity cost of employment. If unemployed, an individual finds a job and is employed in the next period with endogenous probability f. There is a continuum of identical firms on the unit interval. Each firm can hire up to one worker and a firm with an employed and non-shirking worker produces

$$y_t = s_t \tag{13}$$

where $s_t = e^{\varepsilon_t^s}$; we assume $\varepsilon_t^s = \rho^s \varepsilon_{t-1}^s + \eta_t^s$ where η_t^s is distributed as $N(0, \sigma_s^2)$. The value function for a productive job match is

$$J_t = y_t - w_t + \frac{1}{1+r} E_t[\tau_t V_{t+1} + (1 - \tau_t) J_{t+1}] \tag{14}$$

Firms must pay a real cost of γ to post a vacancy. Vacancies are then filled at the start of the next period with endogenous probability q. The value function for a vacant job is therefore

$$V_t = -\gamma + \frac{1}{1+r} E_t[q_{t+1} J_{t+1} + (1 - q_{t+1}) V_{t+1}] \tag{15}$$

We follow the timing convention of Gertler and Trigari (2009) and assume that new job matches become productive immediately.

The simple search frictions model assumes free entry of firms, so $V_t = 0$. This implies that the value function for vacancies simplifies to

$$J_t = (1+r)\frac{\gamma}{q_t} \tag{16}$$

and so the value function for a filled job becomes

$$(1+r)\frac{\gamma}{q_t} = y_t - w_t + (1 - \tau_t)\frac{\gamma}{E_t q_{t+1}} \tag{17}$$

From this, we obtain the job creation condition

$$y_t = w_t + \lambda_t \tag{18}$$

where $\lambda_t = \gamma \left(\frac{(1+r)}{q_t} - \frac{(1-\tau_t)}{E_t q_{t+1}} \right)$ is the real cost of hiring a worker.

We assume that wages are set through worker-firm unilateral Nash bargaining. Defining the surplus from a productive job match as $S = J + (L - U)$, the wage bargain divides the surplus as $J = (1 - \phi)S$ and $L - U = \phi)S$, where ϕ is the worker's relative bargaining power. The resultant wage is

$$w_t = (1 - \phi)z + \phi(y_t + \gamma\theta_t) \tag{19}$$

Most calibrations of search frictions models are designed to assess whether the model fits the data; the focus is on whether simulations of the calibrated model can generate volatilities of unemployment and vacancies as well as the correlations between them, that match values of these statistics in the data. There is little controversy around the values of observable structural parameters, such as the average job separations rate (τ) and the discount rate (r). Precise values of less observable parameters, such as the opportunity cost of employment (z), the elasticity of matching with respect to unemployment (α) and the cost of posting vacancies γ, are less certain and investigators usually select values within a range that is seen as plausible. These parameters are chosen to match the average values of unemployment and other variables, as well as, in some cases, the volatility of unemployment. The parameters of the processes generating productivity and job separation shocks are calibrated to match the time series properties of productivity and job separations.

In this paper, we use a modification of this approach to reflect the different research question addressed in our paper. Unlike most of the literature, we do not ask whether the search frictions model can generate volatilities and correlations that match statistics in the data; we accept the finding in the literature (eg Hall (2005) and Hagedorn and Manovskii (2008)) that it can. Instead, we ask whether the transmission mechanism embedded in the model is consistent with the data. To do this, we calibrate our model so that simulations of the calibrated model match the estimated impulse response function of unemployment to the productivity shock at a chosen date as closely as possible. Given these values, we then calibrate the volatility of the separations shock to match the estimated impulse response function of unemployment to the separations shock at the same date as closely as possible. We then assess whether the calibrated model can also match the other four estimated impulse response functions at that date. Evidence that they can provides evidence in favour of the transmission mechanism. We then re-calibrate the model for another date and repeat the process.

Our calibrated parameter values are summarised in Table 4. We normalize a time period to be one month. We first calibrate the model for 1974Q2. We set $r = 0.004$, equivalent to an annual discount rate of 5%. We also calibrate the average monthly job separation rate as $\tau = 0.033$. For the matching function, we set $\alpha = 0.5$; this is consistent with the range of estimates obtained by Petrongolo and Pissarides (2001). The opportunity cost of employment is $z = 0.66$; this is slightly below the value used by Hall and Milgrom (2008) ($z = 0.71$) but is close to the mid-point of the range of alternative estimates based on alternative specifications of the flow value of non-work reported by Chodorow-Reich and Karabarbounis (2016), $0.47 - 0.96$. Worker bargaining power is set as $\phi = 0.6$, a slightly higher value than the more usual assumption of $\phi = 0.5$. We have two free parameters, m, matching efficiency, and γ, the cost of posting a vacancy. We set the value of these two parameters to match the unemployment rate in 1974Q2 and the response of unemployment to a productivity shock. Doing so gives m=0.8 and γ=0.42. We follow the standard calibration starting from Shimer (2005) to calibrate the parameters governing the productivity shock, setting $\sigma_s = 0.01$ and $\rho_s = 0.878$. Given these parameters, we then calibrate the separation shock to match the response of unemployment to a separations shock as closely as possible. To do this we use $\rho_\tau = 0.733$ from Shimer (2005) but calibrate a smaller volatility, $\sigma_\tau = 0.01$.

We then re-calibrate the model for 2008Q4. Our three calibration targets, the unemployment rate and impulse responses of unemployment, have changed. To respond to this, we reduce the volatilities of the shocks to $\sigma_s = 0.005$, and $\sigma_\tau = 0.005$, reduce the efficiency of matching to $m = 0.64$ and retain the other parameter values. There are alternative strategies: we might also match the changing targets by increasing worker bargaining power or by increasing the opportunity cost of employment. These alternative calibrations do not improve the fit between estimated and simulated impulse responses in 2008Q4.

Figure 10) presents simulated and estimated impulse responses for 1974Q2 and 2008Q4. Although the estimated and simulated responses of unemployment to a productivity shock are similar by design, the other simulated impulse responses are markedly different from their corresponding estimated impulse responses. We particularly note four findings. First, the simulated impulse responses for vacancies in response to a productivity shock are much larger than the estimated impulse responses, for both 1974Q2 and 2008Q4. Second, the simulated impulse response functions for vacancies following a productivity shock do not match the behaviour of the estimated impulse responses. The estimated impulse responses increase for 3-6 months before declining, whereas the simulated impulse responses immediately jump to their maximum values and then decline. Third, although the second calibration enables the simulated model to match the weaker

Table 4: **Calibrated Parameter Values: The Search Frictions Model**
This table reports the values of the calibrations for 1974Q2 and 2008Q4.

Parameter	Interpretation	1974Q2	2008Q4
τ	Ave separation rate	0.033	0.033
r	Discount rate	0.004	0.004
α	Elasticity of matching function	0.50	0.50
b	Opportunity cost of unemployment	0.65	0.65
γ	Vacancy posting cost	0.42	0.42
m	Matching coefficient	0.80	0.64
ϕ	Bargaining power	0.60	0.60
ρ^s	Persistence of productivity shock	0.878	0.878
σ^s	Volatility of productivity shock	0.01	0.005
ρ^τ	Persistence of job separation shock	0.733	0.733
σ^τ	Volatility of job separation shock	0.01	0.005

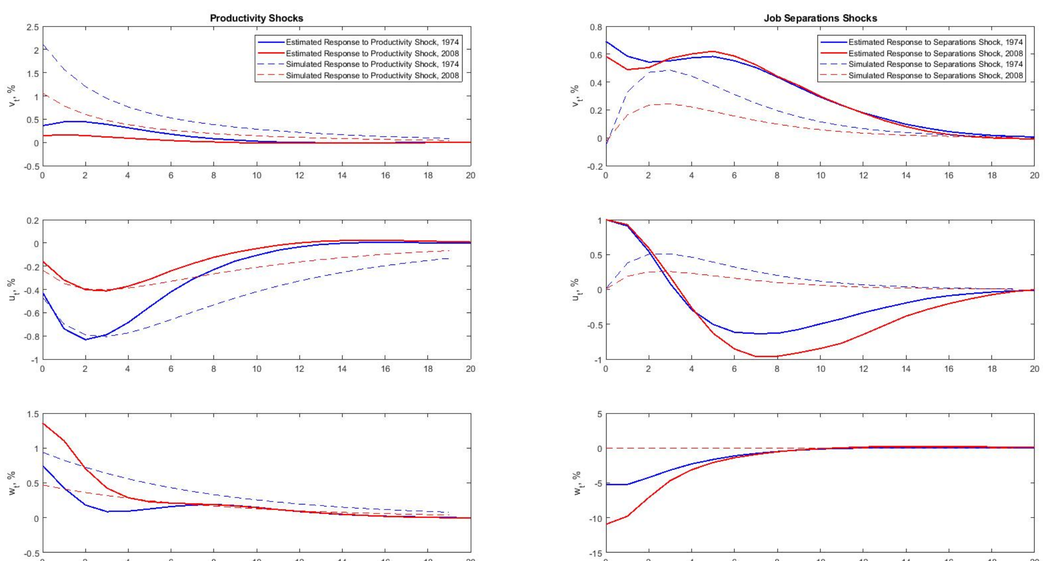

Figure 8: **Estimated and Simulated Impulse Response Functions for the Search Frictions Model**
Notes: The dashed lines on this figure plots the impulse response functions of vacancies (top row), the unemployment rate (middle row), and wages (bottom row) with respect to a productivity shock (LHS column) and a job separations shock (RHS column) from two calibrated search frictions models. The first uses data from 1974Q2 (blue dashed lines), and the second uses data from 2008Q4 (red dahsed lines). Along with these impulse response functions, we provide analogous posterior median estimates from our TVP VAR model which are represented by solid lines.

impulse response of unemployment to a productivity shock in 2008Q4 compared to 1974Q2, it also implies a weaker impulse response of wages to these shocks, contrary to the empirical evidence[10]. Fourth, the simulated impulse response functions for a job separation shock are markedly different from the estimated impulse responses. The estimated impulse responses for wages, unemployment and vacancies are much larger than the simulated responses.

This evidence highlights several potential weaknesses of the simple search frictions model. The central mechanism that enables the search frictions model to match the large empirical volatilities of unemployment and vacancies relies on a surge of vacancy creation by firms following a positive productivity shock. The first finding suggests that this mechanism is empirically rather weak [11][12]. Although the estimates are consistent with the prediction from search frictions model that a reduction in the response of the wage to a productivity shock leads to an increase in the response of unemployment and vacancies, the estimated responses of unemployment and vacancies to a productivity shock are less sensitive to the wage response than in the simulated model. This casts doubt on research that seeks to generate a large volatility of unemployment by adjusting the search frictions model to make wages less sensitive to productivity shocks (eg Hall (2005) and Hall and Milgrom (2008)). It may be that the response to other, non-identified, shocks may provide an alternative explanation of movements in unemployment across the business cycle[13].

Our second finding indicates the presence of dynamic adjustment of the stock of vacancies, similar to that for unemployment. This raises issues with the way that vacancies are modeled in the simple search frictions, as it suggests that not all unfilled vacancies are immediately destroyed. Our final finding suggests that the search frictions model does not capture the channels through which job separations shocks affect the economy. In the model, separations shocks affect the wage through their impact on labour market tightness; this effect is small as separations shocks increase unemployment and vacancies by similar amounts. Our estimates suggest that wages are actually highly sensitive to separations shocks; the effect is strong enough to lead to a reduction in unemployment as part of the adjustment process.

[10]This finding also holds if we increase worker bargaining power rather than reducing matching efficiency in our calibration for 2008Q4. As (19) shows, bargaining power affects the response of the wage to productivity shocks in two ways; there is a direct effect and an indirect effect operating through labour market tightness. By increasing bargaining power, our calibration for 2008Q4 makes the wage more sensitive to productivity in our model through the direct effect; but it also makes the wage less responsive through the indirect effect as it reduces the sensitivity of labour market tightness to productivity. The indirect effect dominates and so the wage is less responsive to a productivity shock.

[11]It is possible that some vacancies are created and filled within a quarter but are not detected due to the granularity of our data. If this is the case, the impulse response for vacancies in the first quarter following the shock would understate the true response. However, in order for the unobserved true response of vacancies to match the simulated response, the number of unobserved vacancies would have to exceed the measured level by a factor of 4-5. Such a large value does not seem plausible.

[12]We note that a one percentage point increase in vacancies leads to more than one percentage point decrease in unemployment. This may suggest that some vacancies are opened and then filled but are not measured. That is plausible since vacancies are hard to measure and we know that some workers get jobs through informal channels. If this is the case, a given level of measured vacancies leads to a higher number of job matches and so a higher number of workers leaving unemployment than would be the case if all vacancies were observed. That would be consistent with evidence that the job filling rate exceeds 1, so that the number of job matches exceeds the number of vacancies (eg Hall and Schulhofer-Wohl (2018)). However, this does not imply that there is a large enough vacancy surge in the data to match our simulated surge. If there are unmeasured vacancies, the impulse response of total vacancies will be larger than the impulse response of measured vacancies. To get an idea of how much larger this response is, we note that the response of unemployment to a productivity shock in Figure 8) is up to 50% larger than response of vacancies. Using the matching function, this suggests that the impulse response of total vacancies is at most twice as large as the impulse response of measured vacancies. As a result, the impulse response of measured vacancies would therefore still be much smaller than the simulated impulse response in Figure 8).

[13]On this note, Buchheim et al. (2019) offer an explanation that investment promotes higher wage growth in tight rather than slack labour markets. Incorporating investment, and indeed labour market conditions, within the search and matching framework is an interesting avenue for future research.

5.1 Robustness and Alternative Approaches

We assess the robustness of our findings by considering the effect of alternative treatments of our data. The search frictions literature often derives stylised facts from key statistics of data that have been de-trended using the Hodrick-Prescott (HP) filter. We do not report this, as recent evidence suggests that the HP filter can distort the data, leading to unreliable results. Hamilton (2018) shows how using the HP filter to de-trend data can induce spurious dynamics that are inconsistent with the data generating process. In the Online Appendix, we show our main results are preserved if we use the a one-sided Kalman filter[14].

We also adopt an alternative identification strategy where the impact response of productivity with respect to a job separation shock is zero. Then, following Baumeister and Benati (2013), we draw a sequence of productivity shocks that neutralises the response of productivity throughout the impulse horizon. This technique is free from the Lucas critique since we are manipulating only structural shocks; not the structural parameters themselves. Fro ma theoretical perspective, this approach is arguably more consistent with the theoretical calibrations we present in Section 5.1. Results from our alternative identification strategy, which are available on request, are consistent with those we present in the main text.

The evidence above suggests that the simple search frictions model does not provide an adequate description of the dynamic response of vacancies to shocks. Fujita and Ramey (2007), Leduc and Liu (2016) and Coles and Moghaddasi-Kelishomi (2018) develop an alternative model in which some matching opportunities persistent beyond the end of each period, so not all unfilled vacancies are immediately destroyed. This introduces dynamics for vacancies similar to that for unemployment. We calibrate these alternative search frictions models to examine whether they provide a better match to our estimated impulse responses; results are available upon request. Our results from alternative calibrations are qualitatively similar to those we present in the main text. Although this approach is better able to capture the initial increase in vacancies following a productivity shock, the match between estimated and simulated responses to a productivity shock is generally worse, while the responses to a job separations shock are no closer.

6 Concluding Remarks

This paper puts search and matching models, the workhorse of modern labour market marcoeconomics, under novel empirical scrutiny. Using state-of-the-art Bayesian estimation techniques, we fit an extended TVP VAR to US labour market data from 1962–2016. We depart from existing literature (see e.g. Yashiv (2006); Faccini et al. (2013); Hall (2005); Hagedorn and Manovskii (2008); Lubik (2009)) by proposing a simple and intuitive method to test calibrated model impulse response functions against empirical estimates. Our results provide three main messages. First, search frictions models are unable to match the responses of key labour market variables to structural shocks. Second, we find evidence against the hypothesis outlined in Shimer (2005) that unemployment and wage volatility possess a negative relationship. Third, empirical estimates show that the key shocks underpinning search and matching models explain, at best, 50% of total variation.

Our results and conclusions hold for a variety of alternative specifications. More specifically, our results are robust to both alternative transformations of data, and extended search and matching models. Therefore for economists, our work calls for a refinement of search and matching models in order to match what we observe in the data. For policymakers, we note that policy should evolve with labour market dynamics. For

[14] Available on request are results using a standard HP filter. It is important to note that our main conclusions are unchanged using this type of transformation.

example, our analysis suggests that the burden of business cycle troughs lies on unemployed workers during the 1970s and 1980s. However, due to relatively tranquil unemployment volatility and surging wage volatility, more recent business cycle troughs lie on the employed. Therefore policy focus should predominantly lie on income stabilisation.

References

Arias, J., Rubio-Ramirez, J. and Waggoner, D. (2018), 'Inference Based on SVARS Identified with Sign and Zero Restrictions: Theory and Applications', *Econometrica* **86**(2), 685–720.

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

Barnichon, R. (2010*b*), 'Productivity and Unemployment over the Business Cycle', *Journal of Monetary Economics* **57**(8), 1013–1025.

Baumeister, C. and Benati, L. (2013), 'Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound', *International Journal of Central Banking* **9**(2), 165–212.

Benati, L. and Lubik, T. A. (2014), The Time-varying Beveridge Curve, *in* 'Advances in Non-linear Economic Modeling', Springer, pp. 167–204.

Buchheim, L., Watzinger, M. and Wilhelm, M. (2019), 'Job Creation in Tight and Slack Labor Markets', *Journal of Monetary Economics (Forthcoming)* .

Chodorow-Reich, G. and Karabarbounis, L. (2016), 'The Cyclicality of the Opportunity Cost of Employment', *Journal of Political Economy* **124**(6), 1563–1618.

Coles, M. and Moghaddasi-Kelishomi, A. (2018), 'Do Job Destruction Shocks Matter in the Theory of Unemployment', *American Economic Journal: Macroeconomics.* **10**(3), 118–136.

Diamond, P. A. (1982), 'Aggregate Demand Management in Search Equilibrium', *Journal of Political Economy* **90**(5), 881–894.

Faccini, R., Millard, S. and Zanetti, F. (2013), 'Wage Rigidities in an Estimated Dynamic Stochastic General Equilibrium Model of the UK Labour Market', *The Manchester School* .

Fujita, S. and Ramey, G. (2007), 'Job Matching and Propagation', *Journal of Economic Dynamics and Control* **31**(11), 3671–3698.

Gertler, M. and Trigari, A. (2009), 'Unemployment Fluctuations with Staggered Nash Wage Bargaining', *Journal of Political Economy* **117**(1), 38–86.

Gourieroux, C., Monfort, A. and Renault, E. (1993), 'Indirect Inference using Simulation Techniques', *Journal of Applied Econometrics* **8**.

Guglielminetti, E. and Pouraghdam, M. (2017), 'Time-varying Job Creation and Macroeconomic Shocks', *Labour Economics* **50**, 156–179.

Hagedorn and Manovskii (2008), 'The Cyclical Behavior of Unemployment and Vacancies Revisited', *American Economic Review* .

Hall, R. (2005), 'Employment Fluctuations with Equilibrium Wage Stickiness', *American Economic Review* **98**(4).

Hall, R. E. (2007), 'Sources and Mechanisms of Cyclical Fluctuations in the Labor Market', *Manuscript, Stanford University, www.Standford.edu/-rehall* .

Hall, R. E. and Milgrom, P. R. (2008), 'The Limited Influence of Unemployment on the Wage Bargain', *American Economic Review* **98**(4), 1653–74.

Hall, R. and Schulhofer-Wohl, S. (2018), 'Measuring job-finding rates and matching efficiency with heterogeneous job-seekers', *American Economic Journal: Macroeconomics* .

Hamilton, J. D. (2018), 'Why You Should Never use the Hodrick-Prescott Filter', *Review of Economics and Statistics* **100**(5), 831–843.

Koop, G. and Korobilis, D. (2010), *Bayesian Multivariate Time Series Methods for Empirical Macroeconomics*, Now Publishers Inc.

Koop, G., Pesaran, M. H. and Potter, S. M. (1996), 'Impulse Response Analysis in Nonlinear Multivariate Models', *Journal of Econometrics* **74**(1), 119–147.

Leduc, S. and Liu, Z. (2016), 'The Slow Job Recovery in a Macro Model of Search and Recruiting Intensity', *Federal Reserve Bank of San Francisco* .

Lubik, T. A. (2009), 'Estimating a Search and Matching Model of the Aggregate Labor Market', *Economic Quarterly* **95**(2), 101–120.

Mortensen, D. T. and Pissarides, C. A. (1994), 'Job Creation and Job Destruction in the Theory of Unemployment', *Review of Economic Studies* **61**(3), 397–415.

Mumtaz, H. and Sunder-Plassmann, L. (2013), 'Time-varying Dynamics of the Real Exchange Rate: An Empirical Analysis', *Journal of Applied Econometrics* **28**(3), 498–525.

Mumtaz, H. and Zanetti, F. (2015), 'Labor Market Dynamics: A Time-varying Analysis', *Oxford Bulletin of Economics and Statistics* **77**(3), 319–338.

Petrongolo, B. and Pissarides, C. A. (2001), 'Looking into the Black Box: A Survey of the Matching Function', *Journal of Economic Literature* **39**(2), 390–431.

Pissarides, C. A. (2000), *Equilibrium Unemployment Theory*, MIT press.

Primiceri, G. E. (2005), 'Time-varying Structural Vector Autoregressions and Monetary Policy', *Review of Economic Studies* **72**(3), 821–852.

Rubio-Ramirez, J. F., Waggoner, D. F. and Zha, T. (2010), 'Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference', *Review of Economic Studies* **77**(2), 665–696.

Shimer, R. (2005), 'The Cyclical Behavior of Equilibrium Unemployment and Vacancies', *American Economic Review* **95**(1), 25–49.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P. and Van Der Linde, A. (2002), 'Bayesian Measures of Model Complexity and Fit', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64**(4), 583–639.

Yashiv, E. (2006), 'Evaluating the Performance of the Search and Matching Model', *European Economic Review* .