

Beveridge Curve Shifts and Time-Varying Parameter VARs

Thomas A. Lubik, Christian Matthes, and Andrew Owens

At first glance, many macroeconomic time series exhibit some form of nonlinearity. For instance, output growth and inflation show less volatility in the 1980s and 1990s than during the Great Inflation period of the 1970s, an observation that has been labeled the Great Moderation. Over the business cycle, the unemployment rate exhibits an asymmetric sawtooth pattern whereby it rises rapidly during downturns and declines only gradually during a recovery. Many price variables, such as exchange rates or commodity prices, appear stable for a long period followed by sudden level shifts. The literature has studied various specific forms of nonlinearity—such as structural breaks, time-varying volatility, or business cycle asymmetries—using sophisticated time-series methods ranging from threshold and Markov switching to vector-autoregressions (VARs) with time-varying parameters and stochastic volatility. The result as to whether there is nonlinearity in the data has been mixed.¹ A key issue in this

■ John Bailey Jones, Daniel Ober-Reynolds, Zhu Wang, and John Weinberg provided helpful comments that improved the quality and exposition of the article. We also wish to thank conference participants at the 2015 Meeting of the Society for Computational Economics in Taipei and the 2015 Symposium of the Society for Nonlinear Dynamics and Econometrics in Oslo for useful feedback and suggestions. The views expressed in this paper are those of the authors and should not necessarily be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System. Lubik: Research Department, Federal Reserve Bank of Richmond. P.O. Box 27622, Richmond, VA 23261. Email: thomas.lubik@rich.frb.org. Matthes: Research Department, Federal Reserve Bank of Richmond. P.O. Box 27622, Richmond, VA 23261. Email: christian.matthes@rich.frb.org. Owens: Department of Economics, University of Rochester. 280 Hutchison Road, Box 270156, Rochester, NY 14627. Email: owensaa@gmail.com.

¹ See, for instance, Hamilton (1989), Primiceri (2005), Sims and Zha (2006), or Amir-Ahmadi, Matthes, and Wang (2016), who use a wide range of empirical methodologies, data sets, and sample periods.

literature is that tests for nonlinearity tend to have low power against linear alternatives.


Against this background, time-varying parameter vector-autoregressions (TVP-VARs) with stochastic volatility have emerged as a promising framework to analyze a wide range of underlying nonlinearities.² In this class of models, the coefficients of the time-series representation for economic data are allowed to vary over time. The idea is that this feature approximates the underlying nonlinearity in the data-generating process to a satisfactory degree and in a parsimonious manner. For instance, a structural break in a deep parameter, or a switch in regimes, could be captured by a shock to the innovation in a random-walk VAR coefficient. Since TVP-VARs offer this flexibility, that is, since they can be understood as approximations to a wide range of underlying nonlinear behavior, they have become increasingly popular in recent years as empirical modeling devices.³ TVP-VARs are estimated almost exclusively using Bayesian methods. This is necessitated by the fact that, as with any model that features many parameters, the use of prior information is crucial to deliver sensible estimates. In TVP-VARs the choice of priors is of special importance because, with standard sample sizes, they have a substantial impact on how much of the variation in observables is attributed to stochastic volatility versus time variation in other coefficients.⁴ At the same time, there is a growing sense, e.g., Lubik and Matthes (2015), that the conclusions drawn from the TVP-VAR literature warrant skepticism. More specifically, TVP-VARs often find not much time variation in the lag coefficients. Instead, they attribute the variation seen in the data to movements in volatilities as the right incidence of shocks can in principle capture a range of time-series patterns.

The purpose of this article is to investigate the extent to which an inherently nonlinear TVP-VAR with stochastic volatility does, in fact, pick up nonlinear features in the underlying data. We do so by applying the TVP-VAR methodology to data generated from a simple (but nonlinear) search and matching model that is designed to generate endogenous shifts in parameters. We thus ask whether a TVP-VAR is capable of detecting the resulting nonlinearity in Beveridge curve dynamics. We follow standard procedure and prescriptions in the literature to specify the TVP-VAR and to choose the prior. The results from these benchmark exercises show that the concerns about proper

² See Lubik and Matthes (2015) for an introduction and survey of TVP-VARs.

³ See Canova, Ferroni, and Matthes (2015) for a discussion of these issues.

⁴ This point is demonstrated by means of a simple example in Lubik and Matthes (2015).

Figure 1 The Beveridge Curve over the Great Recession

attributions of the sources of nonlinearity are warranted. We attempt to resolve some of these concerns by means of an alternative strategy in choosing priors with only partial success. While these findings are largely negative and are also highly conditional on the chosen theoretical model environment, we argue that they serve as a cautionary tale when conducting and interpreting TVP-VAR studies.

Our chosen framework to analyze these issues is the labor market regularity captured by the so-called Beveridge curve. It describes the joint behavior of unemployment and vacancies over the business cycle and is often seen as indicative of the state of the labor market. The Beveridge curve depicts a negative relationship between these two variables, whereby movements along this curve reflect expansions and recessions. The behavior of the curve over the course of the Great Recession and its aftermath has attracted much interest in the literature (e.g., Barlevy 2011; Lubik 2013; or Şahin et al. 2014). Figure 1 shows the Beveridge curve for data over the Great Recession period. The unemployment-vacancies relationship is often represented by a scatter plot of the two series against each other, resulting in a downward-sloping curve. For purposes of illustration, in Figure 1 we fitted a regression line to data from 2001 up to September 2008, which

cluster tightly around the Beveridge curve. At the onset of the Great Recession, the unemployment rate rises rapidly and vacancies fall. In the graph, the data points start moving off the normal curve and appear to settle at a location above their normal, or expected, level. In other words, during the Great Recession, the Beveridge curve appears to have shifted outward in a discrete manner, which could be indicative of a structural break in a labor market parameter.

More generally, the Beveridge curve over the last sixty years reveals a substantial degree of nonlinearity (see Benati and Lubik 2014). There are discernible inward and outward shifts, tilts, and even the occasional slope reversal over short periods. This is not necessarily *prima facie* evidence of the presence of nonlinearities since these patterns can be rationalized through the right incidence of various shocks (e.g., Blanchard and Diamond 1989; Barlevy 2011; or Lubik 2013). It is nevertheless suggestive of underlying structural changes in the labor market. We take this observation as a starting point for our investigation into the practice of TVP-VAR estimation.

We develop a simple search and matching model of the labor market, where we allow for endogenous threshold switching in a key parameter, namely in the efficiency of a match between employer and job seeker. This match efficiency is captured by a level parameter in the matching function and summarizes the efficacy of the labor market. We assume that it can take on two values, which indicate different but parallel locations of the Beveridge curve. A high level of match efficiency translates into a location of the Beveridge curve closer to the origin, whereby a lower level shifts it outward. Under high efficiency, employers need to open fewer vacancies for the same number of job seekers to fill a desired number of positions. The economy switches between the two efficiency parameters when a threshold embedded in the model is crossed endogenously. We assume that this threshold is given by a low level of output that we associate with a weak labor market performance. This threshold can be reached with a sequence of bad and persistent productivity draws; that is, in this case, the labor market exhibits damage to the extent that the Beveridge curve shifts only when a recession is deep and drawn out. In terms of the behavior of the model, this threshold switch implies nonlinearity in the dynamics of the economic variables.

In order to study the implications of this specific form of nonlinearity for empirical modeling, we solve the full nonlinear model and simulate data on unemployment and vacancies. We then estimate a Bayesian TVP-VAR with stochastic volatility on these data and assess how well the nonlinear *atheoretical* time-series model captures the underlying nonlinearity in the model. Given a standard initialization and

choice of priors, the evidence suggests that the TVP-VAR attributes almost all of the changes in the simulated data to changes in the reduced-form innovation variances. We argue that this raises doubts as to the validity of TVP-VAR models with standard priors in detecting shifts. In order to address this shortcoming, we suggest an approach that tries to elicit priors for the TVP-VAR, but it is only moderately successful. In order to better capture the time variation in parameters, researchers will need to adapt the priors to the question at hand in more sophisticated ways. One possibility that delivers better performance is to estimate the hyperparameters associated with the parameters governing the amount of time variation in the model (Amir-Ahmadi, Matthes, and Wang 2017).

The contribution of this article is twofold. First, using simulated data, we study to what extent a generic TVP-VAR with stochastic volatility deals with a specific form of nonlinearity in these underlying data. Our results suggest that some findings of this literature should be regarded with skepticism since they attribute too much of this nonlinearity to time variation in the shocks rather than to structural breaks in the underlying model parameters. Second, and of independent interest, we demonstrate how Beveridge curve shifts can be explained conceptually via an endogenous mechanism that moves the economy between a high-performing and a low-performing labor market. This mechanism can thus be used to address issues like hysteresis, where temporary shocks, such as business cycle shocks, can have permanent effects.

The article is structured as follows. In the next section, we lay out our simple modeling framework of the standard search and matching model and describe how we introduce the threshold-switching mechanism that leads to nonlinearity in the model. The second section describes how we calibrate the model. In this section, we also describe simulation results from the model and discuss the TVP-VAR that we use to estimate the simulated data. Section 3 presents the estimation results and details the shortcomings of the TVP-VAR approach in this environment, while Section 4 introduces an alternative method to elicit priors for the TVP-VAR. The final section concludes.

1. A STRUCTURAL MODEL OF BEVERIDGE CURVE SHIFTS

We now describe the simple structural labor market framework that we use to model the Beveridge curve. We hereby draw heavily from the existing literature, most prominently Shimer (2005). The specification of the model follows Lubik (2013). Our working assumption is that

the Beveridge curve has experienced a structural shift, as seen by the evolution of unemployment and vacancies in Figure 1. We model the structural break in terms of a threshold-switching process: when a target variable, aggregate output in our case, hits a threshold in terms of deviations from its long-run level, it triggers a shift in a structural labor market parameter. The idea is to capture the observation that Beveridge curve shifts appear to occur during strong and deep recessions and expansions (see Benati and Lubik 2014).

We assume that in our model economy time is discrete and the time period is a quarter. The labor market in this economy is characterized by search and matching frictions, which help rationalize the existence of equilibrium unemployment. Specifically, a job, that is, a relationship between a worker and a firm for the purpose of engaging in production, is the outcome of a matching process. New jobs M are generated by combining unemployed job seekers U with job openings (vacancies) V . This process can be represented by a constant returns matching function, $M_t = m(s_t)U_t^\xi V_t^{1-\xi}$, where $0 < \xi < 1$ is the match elasticity. $m(s_t) > 0$ is the match efficiency that captures the ease with which the unemployed are transformed into workers.

We assume that match efficiency is subject to structural shifts. Specifically, the level parameter in the matching function $m(s_t)$ can switch between two values, $s_t \in \{s_H, s_L\}$, with $m(s_L) < m(s_H)$. In our framework, the switch is generated endogenously by a trigger mechanism, in contrast to the exogenous regime changes in Markov-switching models. We implement this trigger by tying it to the severity of the business cycles. Whenever GDP deviates too much from its current target level, the labor market experiences a structural shift in terms of a change in the matching efficiency. As Lubik (2013) argues, Beveridge curve shifts are most parsimoniously and plausibly modeled by a change in this one parameter. More specifically, one can show that declines in match efficiency are associated with outward shifts of the curve.

For the purposes of capturing Beveridge curve dynamics, we assume that the threshold mechanism is attached to aggregate output. More specifically, we assume that match efficiency $m_t = m(s_t)$ follows a threshold process:

$$m_t = \begin{cases} m(s_H) & \text{if } Y_t \geq \underline{Y} \\ m(s_L) & \text{if } Y_t < \underline{Y} \end{cases}, \text{ where } m(s_H) > m(s_L). \quad (1)$$

Y_t is aggregate output and \underline{Y} is the threshold at which the labor market experiences a structural shift. In the simple search and matching framework, we assume linear production so that Y_t is given by:

$$Y_t = A_t N_t, \quad (2)$$

where N_t is the stock of employed workers, and A_t is an aggregate productivity process that obeys the law of motion:

$$\log A_t = (1 - \rho_A) \log \bar{A} + \rho_A \log A_{t-1} + \varepsilon_{A,t}, \quad (3)$$

where $0 \leq \rho_A < 1$ and $\varepsilon_{A,t} \sim \mathcal{N}(0, \sigma_A^2)$. We normalize the mean of the process \bar{A} to a value of unity without loss of generality.

The dynamics of the model are such that sequences of low and persistent productivity draws—in other words, a recession—will occasionally move aggregate output below the threshold \underline{Y} . This damages the labor market in the sense that match efficiency declines and the Beveridge curve shifts outward. This shift is persistent because of the persistence in the productivity process and the inherent persistence of employment in the search and matching framework. Once the recession abates, the labor market recovers in terms of a switch back to a “normal” level of match efficiency. In that sense, our framework shares similarities with the “plucking” model of recessions, where the economy is plucked away occasionally from its normal evolution due to a deep recession but then transitions back over time.

The dynamics of employment are governed by the following relationship:

$$N_t = (1 - \chi_t) \left[N_{t-1} + m(s_{t-1}) U_{t-1}^\xi V_{t-1}^{1-\xi} \right]. \quad (4)$$

This is a stock-flow identity that relates the stock of employed workers N to the flow of new hires, $M = mU^\xi V^{1-\xi}$, into employment. The timing assumption is such that variations in match efficiency do not affect employment contemporaneously. Unemployment is defined as:

$$U_t = 1 - N_t, \quad (5)$$

where the labor force is normalized to 1. Inflows to unemployment arise from exogenous job destruction at rate $0 < \chi < 1$. We assume that the separation rate χ follows the process:

$$\log \chi_t = (1 - \rho_\chi) \log \bar{\chi} + \rho_\chi \log \chi_{t-1} + \varepsilon_{\chi,t}, \quad (6)$$

where $0 < \rho_\chi < 1$ and $\varepsilon_{\chi,t} \sim \mathcal{N}(0, \sigma_\chi^2)$.

The matching function can be used to define the job-matching rate, i.e., the probability that a firm is matched with a worker:

$$q(\theta_t) = \frac{M_t}{V_t} = m_t \theta_t^{-\xi}, \quad (7)$$

where $\theta_t = V_t/U_t$ is labor market tightness. From the perspective of an individual firm, the aggregate match probability $q(\theta_t)$ is exogenous, and hence new hires are linear in number of vacancies posted for individual firms: $M_{it} = q(\theta_t) V_{it}$.

A firm chooses the optimal number of vacancies V_t to be posted and its employment level N_t by maximizing the intertemporal profit function:⁵

$$E_0 \sum_{t=0}^{\infty} \beta^t [A_t N_t - w_t N_t - \kappa V_t], \quad (8)$$

subject to the employment accumulation equation (4). Profits are discounted at rate $0 < \beta < 1$. Wages paid to the workers are w , while $\kappa > 0$ is a firm's fixed cost of opening a vacancy. The first order conditions are:

$$N_t : \quad \mu_t = A_t - w_t + \beta E_t(1 - \chi_{t+1})\mu_{t+1}, \quad (9)$$

$$V_t : \quad \kappa = q(\theta_t)\beta E_t(1 - \chi_{t+1})\mu_{t+1}, \quad (10)$$

where μ_t is the multiplier on the employment equation. Combining these two first-order conditions results in the *job-creation* condition:

$$\frac{\kappa}{q(\theta_t)} = \beta E_t \left[(1 - \chi_{t+1}) \left(A_{t+1} - w_{t+1} + \frac{\kappa}{q(\theta_{t+1})} \right) \right], \quad (11)$$

which captures the trade-off faced by the firm. The marginal, effective cost of posting a vacancy, $\frac{\kappa}{q(\theta_t)}$, that is, the per-vacancy cost κ adjusted for the probability that the position is filled, is weighed against the discounted benefit from the match. The latter consists of the surplus generated by the production process net of wage payments to the workers plus the benefit of not having to post a vacancy again in the next period.

Wages are determined based on the Nash bargaining solution: surpluses accruing to the matched parties are split according to a rule that maximizes the weighted average of the respective surpluses. We relegate the full discussion of the derivation to the Appendix (see also, Lubik 2013). The resulting wage equation is:

$$w_t = \eta (A_t + \kappa\theta_t) + (1 - \eta)b. \quad (12)$$

Wage payments are a weighted average of the worker's marginal product A_t , of which the worker can appropriate a fraction η , and the outside option b , of which the firm obtains the portion $(1 - \eta)$. Moreover, the presence of fixed vacancy posting costs leads to a hold-up problem where the worker extracts an additional $\eta\kappa\theta_t$ from the firm.

⁵ For ease of exposition and notation, we will drop the firm-specific subscripts and discuss the problem of a representative optimizing firm with the understanding that firms are ex-ante heterogeneous in this framework.

We can substitute the wage equation and the job-matching rate into the job-creation condition to obtain:

$$\frac{\kappa}{m_t} \theta_t^\xi = \beta E_t \left[(1 - \chi_{t+1}) \left\{ (1 - \eta) (A_{t+1} - b) - \eta \kappa \theta_{t+1} + \frac{\kappa}{m_{t+1}} \theta_{t+1}^\xi \right\} \right]. \quad (13)$$

Firms are more willing to post vacancies if productivity shocks increase the wedge to the outside option of the worker; they are less willing if there are expected separations as this will reduce the present value of a hired worker.

In our simulation and empirical analysis, we make use of the simple structure of the model. The dynamics can be fully described by two equations, the employment accumulation equation (4) and the job-creation condition (13), after convenient substitutions. Intuition for why an outward shift of the Beveridge curve is generated by a fall in match efficiency can be gleaned from equation (4) and the logic of the matching function. At any given unemployment rate, firms would need to post more vacancies to achieve a target hiring quota since the matching process is now less efficient.⁶ However, there is also a countervailing effect, namely through the influence of match efficiency on firms' vacancy posting decisions. A fall in m raises effective vacancy posting costs as captured by the left-hand side of the job-creation condition (13). This implies that vacancies are increasing in match efficiency. The overall effect of a change in m therefore depends on the interaction of these two margins. As Lubik (2013) shows, the stock-flow identity of the law of motion has to hold in equilibrium, so that the first effect via the matching function dominates and shifts the Beveridge curve outward for a smaller m , whereas the effect via the job-creation margin generates movements along this equilibrium relationship. We now turn to a discussion of our solution and simulation approach.

2. SIMULATION AND ESTIMATION

Calibration

We calibrate our model to representative parameter values in the literature. Our benchmark calibration rests on the parameter estimates

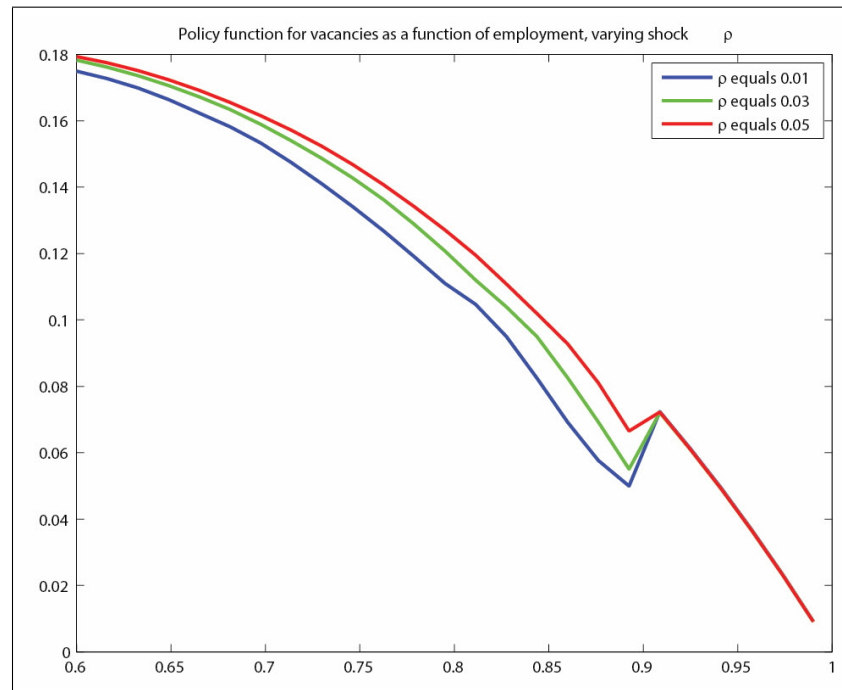
⁶ Formally, this can also be seen from the steady-state representation of the employment equation (4), which describes an equilibrium locus of combinations of U and V such that inflows and outflows to (un)employment are balanced:

$$V = \left(\frac{1 - \chi}{m} \frac{\chi}{1 - \chi} \right)^{\frac{1}{1-\xi}} \left(\frac{1 - U}{U} \right)^{\frac{1}{1-\xi}} U.$$

Table 1 Calibration

| Parameter | Value | Source |
|---------------------------------|-------|---|
| Separation Rate $\bar{\chi}$ | 0.036 | Shimer (2005); Monthly JOLTS Data |
| Match Elasticity ξ | 0.49 | Beveridge Curve Estimation: Lubik (2013) |
| Match Efficiency m_H | 0.90 | Beveridge Curve Estimation: Lubik (2013) |
| Match Efficiency m_L | 0.70 | Beveridge Curve Estimation: Lubik (2013) |
| Benefit b | 0.90 | Hagedorn and Manovskii (2008) |
| Bargaining η | 0.49 | Hosios-Condition: $\eta = \xi$ |
| Job Creation Cost κ | 0.18 | Imputed from Steady-State Sample Means $\bar{V} = 2.6\%$ and $\bar{U} = 5.2\%$. |
| Discount Factor β | 0.99 | Annual Real Interest Rate |
| Productivity \bar{A} | 1.00 | Normalized |
| Threshold Value \underline{Y} | 0.91 | Cumulative Decline in U.S. GDP 2008–10 |
| AR(1) Coefficient ρ_A | 0.95 | Standard Value |
| AR(1) Coefficient ρ_s | 0.95 | Standard Value |
| StD Productivity σ_A | 0.01 | Standard Value |
| StD Separation Rate σ_s | 0.01 | Standard Value |

in Lubik (2013) for the period 2000–08, after which a potential shift in the Beveridge curve appears evident from the data (see Figure 1). The calibrated parameters are reported in Table 1. We set the mean of the separation rate to a value of 0.036. This follows the value reported in Shimer (2005) for monthly data. We choose the match efficiency in the high state $m_H = 0.90$ and in the low state $m_L = 0.70$ based on the estimate in Lubik (2013). The match elasticity is set to $\xi = 0.49$. These values broadly determine the slope and the location of the Beveridge curve in a scatter plot of vacancies and unemployment. We set the discount factor $\beta = 0.99$ and choose the bargaining parameter by imposing the Hosios-condition for social efficiency, $\eta = \xi = 0.49$. As mentioned before, we normalize the mean of the level of productivity to $\bar{A} = 1$. Next, we assume that the outside option of the worker makes up 90 percent of the productivity level, $b/\bar{A} = 0.9$. The calibration is therefore close to that of Hagedorn and Manovskii (2008), who argue that a high outside option for the worker is needed to match the cyclical properties of the data. The job-creation condition can then be used to back out the cost parameter κ for a given level of unemployment and vacancies. We compute these from the sample averages for the period 2000–08, $\bar{V} = 2.6$ percent and $\bar{U} = 5.2$ percent. This implies $\kappa = 0.18$. Finally, we set the threshold value for $\underline{Y} = 0.91$ to approximate the cumulative decline in U.S. GDP over the course of the Great

Figure 2 Policy Functions

Recession of 2007–09. We set the persistence parameter of the technology process and the separation rate to 0.95 and the standard deviations of the respective innovations to 0.01.

Model Simulation and Discussion

We solve the search and matching model with threshold switching in match efficiency fully nonlinearly by means of the monotone mapping algorithm. The algorithm computes an approximation of the firm's decision rule, which determines the number of vacancy postings given the economy's state variables: employment N_t , the exogenous productivity shock A_t , and the separation rate process χ_t . The algorithm is detailed in the Appendix.

In order to understand the underlying dynamics of the model before we turn to the estimation exercise, we compute the policy functions under the baseline calibration for given realizations of the shocks. The key driving force behind the shifts in the Beveridge curve are movements

in productivity. An adverse enough realization of productivity A_t can drive output below the threshold value, which then generates a switch to a lower match efficiency. However, equilibrium outcomes across the threshold and within the distinct regions depend on the subtle interplay between the state variables. To give a sense of the nonlinearities present in our model, we plot the policy function for vacancies in Figure 2. For the purposes of this exercise, we hold the productivity shock fixed at its unconditional mean $\bar{A} = 1$. Vacancies are graphed against the model's sole endogenous state variable, namely the level of employment. We plot this relationship for different realizations of the separation rate $\bar{\chi}$.

Figure 2 shows the key aspect of the model. The policy function has two distinct regions that coincide with the two distinct states of the labor market. For given productivity, the policy function is discontinuous at the implied threshold level, $\bar{N} = \bar{Y} = 0.91$. To the right of the threshold, the labor market is in its normal state with match efficiency $m(s_H)$, and to the left, it has suffered from a deterioration of the latter. We also note that vacancies are decreasing in employment. When employment is high (and unemployment low), few vacancies are being posted because the vacancy-unemployment ratio θ is high and the labor market is tight. That is, the firm's probability of finding a worker is low relative to the costs of hiring him. When employment is low, the labor market is awash with job seekers, so firms can more easily recoup the implicit hiring costs. We also note that the vacancy policy function is increasing in the productivity shock.

The policy function for the high match efficiency case tends to lie to the right and above the respective function in the low efficiency regime. Other things being equal, a lower match efficiency reduces the firm's hiring probability and thereby the incentive to post vacancies relative to the high efficiency scenario. An interesting pattern emerges when we additionally vary the policy function across separation rates. We find that the higher the separation rate, the higher the vacancy postings for given productivity and employment. More separations mean higher churn, so for given employment, more vacancies need to be posted. The differences between the policy functions, however, are quantitatively small for the low efficiency case and almost nonexistent under high efficiency.⁷ What is interesting is that the relationship between the separation rate and levels of match efficiency appears nonlinear in its effect on vacancy postings.

⁷ This is consistent with the empirical finding in Lubik (2009) and the assumption and interpretation in Shimer (2005) that movements in the separation rate are not key drivers of labor market fluctuations.

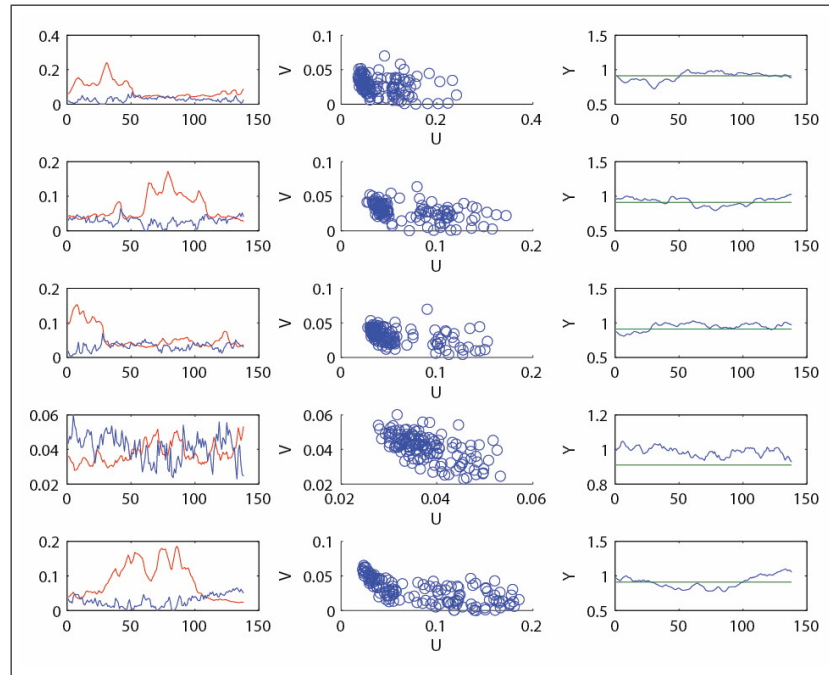
Table 2 Selected Moments

| | $\sigma(V)$ | $\sigma(U)$ | $\sigma(V/U)$ | $\rho(V, U)$ |
|----------|-------------|-------------|---------------|--------------|
| Sample 1 | 0.53 | 1.88 | 0.32 | -0.35 |
| Sample 2 | 0.49 | 2.10 | 0.39 | -0.51 |
| Sample 3 | 0.43 | 1.22 | 0.41 | -0.47 |
| Sample 4 | 0.39 | 0.57 | 0.41 | -0.43 |
| Sample 5 | 0.35 | 1.09 | 0.66 | -0.69 |

We now simulate the model for 590 periods under the benchmark calibration. We discard the first 450 periods as burn-in. We are thus left with a sample of size 140, of which we will use the first forty observations as a training sample in the estimation of the VAR. This leaves us with 100 periods, or twenty-five years, of data for the actual estimation. Table 2 reports moments for five representative samples. We present this as a first pass for whether our regime-switching framework can potentially capture salient labor market facts. The last column shows the correlation between unemployment and vacancies, which is considerably negative and ranges from -0.35 to -0.69 but is below the correlation found in U.S. data. Nevertheless, the model can replicate to some extent the strongly negative comovement between these two labor market variables.

The model is less successful in terms of volatilities. The first two columns of Table 2 report the standard deviations of vacancies and unemployment relative to the standard deviation of (labor) productivity $A_t = Y_t/N_t$ as in Shimer (2005). Vacancies are roughly half as volatile as productivity, while unemployment is twice as volatile for samples 1 and 2. The standard deviation drops considerably in sample 4, while samples 3 and 5 show the volatilities of the driving process and the endogenous variables as roughly equal. The low volatility of vacancies is also reflected in that of labor market tightness V/U . Our framework thus falls prey to the critique espoused in Shimer (2005), namely that the basic search and matching model has difficulty replicating the observed high volatility of unemployment and vacancies. As Lubik (2009) shows, this can be remedied by additional shocks to the model such as the exogenous variations in the separation rate, but this comes at the price of reducing the correlation between U and V since a shock to separations moves unemployment and vacancies in the same direction.

Figure 3 shows data plots of the five simulation samples, including the training sample, in the same order as presented in Table 2. Each row in the graph represents one simulation. The panels on the left show time series plots of unemployment (in red) and vacancies (in blue). The middle column shows the same data as a scatter plot in order to

Figure 3 Simulated Data

highlight shifts in the Beveridge curve that are potentially induced by the mechanism in our framework. The last column shows aggregate output Y_t for each simulation along with its threshold value for the regime switch. The graphs confirm that the simulated model reproduces the negative correlation between unemployment and vacancies; that is, the model generates a Beveridge curve. What is notable visually from the middle column of Figure 3 is that there are generally two separate clusters of data (with the exception of the sample in the fourth row). On the face of it, this lends support to the mechanism in our framework as it can replicate the shift patterns seen in actual data. This outcome is not preordained, however, as is evident from sample 4. In this simulation, the threshold is never reached despite values of output persistently below its mean for extended periods. The model economy suffers from a recession, but not one that is deep enough to do damage to the labor market.⁸ Consequently, a stable Beveridge curve

⁸ This is consistent with the interpretation of Benati and Lubik (2014) that most shifts of the Beveridge curve during recessions are too small to be plausibly and statis-

pattern arises over the full simulation period. We also note that this sample stands out in Table 2 because of the low volatilities of unemployment and vacancies.⁹

In the other sample paths, output falls below the threshold for lengthy periods. For instance, in the first row, the initial productivity draw pushes output below the threshold and keeps it there for fifty periods. During this period, there are two opposing forces at play. First, the productivity process is mean-reverting; that is, eventually there will be enough positive innovations to push productivity above its mean and thereby drag output back above the threshold.¹⁰ The strength of this effect depends on the degree of persistence in productivity. If it is high enough, very large negative draws can have staying power to keep the economy below the threshold. Low persistence, on the other hand, leads to faster mean reversion. The second, endogenous force works against this pattern. If the economy is below the threshold, vacancy postings are lower than they otherwise would be (see Figure 2 and the discussion above). Consequently, matching with lower match efficiency reinforces the threshold switch. The observed Beveridge curve shift would thus be consistent with the hypothesis that prolonged periods of high unemployment are generated by mismatch in the labor market (see Şahin et al. 2014).

To summarize, we show that the simple model with an endogenous threshold switch can qualitatively and, with some qualifications, quantitatively replicate the business cycle patterns of key labor market variables. More importantly for our purposes, we demonstrate that our model can generate structural shifts in the Beveridge curve. This raises two questions. First, are these shifts large enough to be statistically different from a standard adjustment pattern, such as a counterclockwise loop as discussed in Blanchard and Diamond (1989)? Using different methodologies and sample periods, Lubik (2013) and Benati and Lubik (2014) answered this question in the negative. In this paper, we ask a second question, namely whether the shifts are even detectable as such in a flexible time-series framework.

tically judged as structural. They are thus more consistent with the counterclockwise loop identified by Blanchard and Diamond (1989). Yet, a few recessions, notably the most recent one, fall outside this pattern.

⁹ What drives this pattern, that is, the lack of a Beveridge curve shift, is the combination of not-large-enough random shocks in the simulation and a lack of adjustment dynamics to the new (conditional) steady state associated with lower match efficiency.

¹⁰ An alternative specification would have productivity also obey a threshold switch, so that the effect of very bad recessions would be much more protracted. This would render the model closer to the implications of a Markov-switching model such as Hamilton (1989).

A TVP-VAR for the Simulated Data

Given the simulated data, we now turn to assessing whether statistical approaches can uncover the underlying shifts in the Beveridge curve. For this purpose, we rely on a TVP-VAR with stochastic volatility, which has proved to be a flexible and useful tool to study nonlinear behavior in aggregate time series. It has recently been applied to the question of Beveridge curve shifts by Benati and Lubik (2014). Our specific time-series model builds on Cogley and Sargent (2005) and Primiceri (2005). The exposition below follows Lubik and Matthes (2015), who provide further details on the implementation.

We stack the unemployment rate U_t and the vacancy rate V_t in a column vector y_t , which we assume is determined by the following law of motion:

$$y_t = \mu_t + \sum_{j=1}^L A_{j,t} y_{t-j} + e_t. \quad (14)$$

μ_t is a drift term that can contain deterministic and stochastic components. The $A_{j,t}$ are conformable coefficient matrices that contain time-varying parameters. e_t is a vector of residuals. Most of the literature on TVP-VARs that use quarterly data pick the lag length in the reduced-form specification as $L = 2$. We follow this convention since we use a quarterly calibration for our matching model. We define $X_t' \equiv I \otimes (1, y_{t-1}', \dots, y_{t-L}')$ to provide a concise representation of the dynamics of y_t . We thus rewrite equation (14) as:

$$y_t = X_t' \theta_t + e_t. \quad (15)$$

We assume that the law of motion for the time-varying parameters in the coefficient matrices $A_{j,t}$ is given by:

$$\theta_t = \theta_{t-1} + u_t, \quad (16)$$

where u_t is a zero mean i.i.d. Gaussian process. To characterize stochastic volatility, we assume that the covariance matrix of the one-step-ahead forecast error e_t can be decomposed using two matrices such that:

$$e_t = \Lambda_t^{-1} \Sigma_t \varepsilon_t, \quad (17)$$

where the standardized residuals are distributed as $\varepsilon_t \sim N(0, I)$. Λ_t is a lower triangular matrix with ones on the main diagonal and representative nonfixed element λ_t^i . Σ_t is a diagonal matrix with representative nonfixed element σ_t^j . The dynamics of the nonfixed elements of Λ_t and Σ_t are given by:

$$\lambda_t^i = \lambda_{t-1}^i + \zeta_t^i. \quad (18)$$

$$\log \sigma_t^j = \log \sigma_{t-1}^j + \eta_t^j. \quad (19)$$

We assume that all these innovations are normally distributed with covariance matrix V . In order to provide some structure for the estimation, we restrict the joint behavior of the innovations as follows (see Primiceri 2005):

$$V = \text{Var} \left[\begin{pmatrix} \varepsilon_t \\ u_t \\ \zeta_t \\ \eta_t \end{pmatrix} \right] = \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix}. \quad (20)$$

S is further restricted to be block diagonal, which simplifies inference. We use a Gibbs-sampling algorithm to generate draws from the posterior. The implementation of the Gibbs-sampling approach used for Bayesian inference follows Del Negro and Primiceri (2013).

A key choice for TVP-VAR modeling is how to set the prior. In order to achieve sharp inference, given the multiple sources of variation in TVP-VAR models, a researcher needs to impose restrictions on the relationship between the covariance matrices of the parameters. The trade-off, however, is that a too restrictive prior may not leave room for the time variation to appear. In our benchmark, we impose a typical choice of prior as recommended in, for instance, Primiceri (2005). Specifically, we assume the following:

$$Q \sim IW(\kappa_Q^2 * 40 * V(\theta_{OLS}), 40), \quad (21)$$

$$W \sim IW(\kappa_W^2 * 2 * I, 2), \quad (22)$$

$$S \sim IW(\kappa_S^2 * 2 * V(\Lambda_{OLS}), 2), \quad (23)$$

where IW denotes the Inverted Wishart distribution priors for all other parameters are the same as in Primiceri (2005). For the prior hyperparameters κ_Q , κ_W , and κ_S , we use the values $\kappa_Q = 0.01$, $\kappa_W = 0.01$, and $\kappa_S = 0.1$. We will discuss alternative prior choices below.

3. ESTIMATION RESULTS

We report estimation results for our benchmark TVP-VAR on simulated unemployment and vacancies data in Figures 4 and 5. In each figure, we report posterior mean estimates from the five representative data samples discussed in the previous section. Since we specify a two-variable VAR with two lags, we report eight series overall for the lag coefficients, two series for the variances, and one for the covariance.


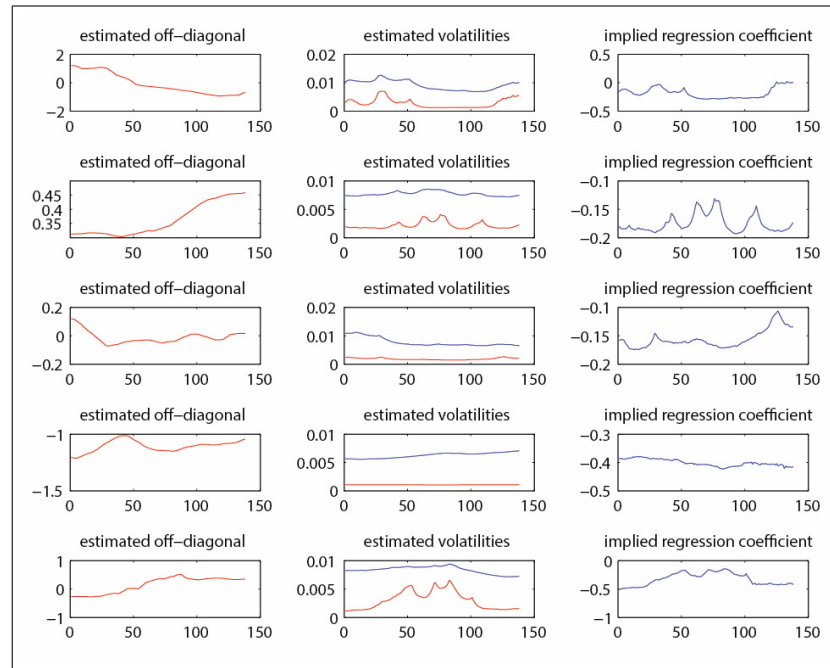
Figure 4 Posterior Means of VAR Coefficients

Figure 4 shows the median posterior estimates of the coefficients in the lagged matrices $A_{j,t}$ in (14) for each sample and over the entire sampling horizon. Figure 5 shows additional estimated statistics. The left column of Figure 5 reports the estimated off-diagonal elements of the covariance matrix of the one-step-ahead forecast errors, while the middle column depicts the posterior means of the diagonal elements, that is, the variances. We also report the implied regression coefficients of a period-by-period population regression of unemployment on vacancies for each sample.

Figure 5 Summary of Benchmark Results: Estimated Posterior Means



The results are almost unequivocal. Across all simulations, the TVP-VAR attributes the shifts in the simulated Beveridge curve to changes in the forecast error variance only. While both volatilities and contemporaneous correlations change with shifts in the underlying series, *all* lag coefficients are estimated to be unvarying and effectively constant (see Figure 4). The estimates for the individual samples show that when there appears to be a shift in the Beveridge curve it is associated with a gradual drift in the coefficients of the variance-covariance matrix. Consider as a baseline case the simulated sample in the fourth row of Figures 3 and 5. As discussed before, this sample path includes declines in output that never cross the threshold and therefore do not lead to Beveridge curve shifts. The TVP-VAR produces essentially constant variances of the shock innovations and an implied population regression coefficient (i.e., a Beveridge curve slope) that is fairly constant at -0.4. There is some variation in the covariance, which rises from -1.2 to -1.0 before retreating again. This seems commensurate with the increase in unemployment and the fall in vacancy postings as

the economy enters a downturn in the first half of the simulated sample. The resulting pattern is that of a movement *along* the Beveridge curve but not a shift.

These patterns are noticeably different when we consider sample paths that include movements of output across the threshold. First, the population regression coefficients exhibit more variation and are smaller (in absolute value) over the full sample period compared to those of a sample path that does include a switch. Along a given Beveridge curve, unemployment and vacancies move in opposite directions. But in the transition between the two Beveridge curves, unemployment and vacancies tend to move in the same direction as vacancy postings rise in order to counteract the lower match efficiency. Shifts in the Beveridge curve are associated with shifts in the elements of the covariance matrix. In particular, periods of high volatility and positive covariation are associated with unemployment-vacancy combinations arising from low match efficiency. As discussed above, because of the constant mean and the mean-reversion of the productivity process, large and persistent enough negative shocks are required to push output below the threshold. These shocks also induce high volatility in unemployment and vacancies. The TVP-VAR then attributes this increased volatility to time-variation in the innovation covariance matrix. The positive correlation in the innovations thus mirrors the lower implied regression coefficient.

To summarize our findings, we posit that an econometrician who attempts to discover shifts in the Beveridge curve using a standard TVP-VAR would come to an erroneous conclusion. What appears in the data as a parallel shift in the curve is interpreted by the TVP-VAR as the outcome of time-variation in the variance-covariance matrix of the shocks. Large shocks drive the labor market variables away from their present location. Given the inherent persistence in the search and matching model, this would then cluster temporally close data points in a pattern that indicated a shift.¹¹ In the logic of the search and matching model, this outcome would be consistent with a higher incidence and severity of shocks that primarily affect the matching process and transitional labor dynamics as captured in equation (4) (see Barlevy 2011; Lubik 2013).

¹¹ Incidentally, this reasoning is consistent with the argument in Lubik (2013) that the degree of estimation, parameter, and model uncertainty in the empirical model is large enough that it would be difficult to distinguish statistically between competing hypotheses, especially when compared to the relatively short time span of a Beveridge curve cycle in the data. On the other hand, Benati and Lubik (2014) impose further restrictions and utilize longer samples to show that a few Beveridge curve cycles do allow for sharper inference, including the Great Recession.

However, and to reiterate this point, the underlying data are generated from a model where the presence of a structural shift for lengthy periods of time is quite noticeable. The TVP-VAR thus attributes these shifts erroneously to changes in volatility. This observation is consistent with many studies using TVP-VARs that tend not to find substantial changes in the lag coefficient matrices, but rather apportion excess volatility and breaks in behavior to stochastic volatility. Our finding is also reminiscent of the critique by Benati and Surico (2009) of Sims and Zha's (2006) argument that the switch from the Great Inflation of the 1970s to the Great Moderation of the 1980s and beyond was not driven by a break in policy but by a decline in the volatility of the shocks. By means of a simulation study, Benati and Surico (2009) show that a regime-switching VAR cannot recover a break in policy coefficients in the underlying model. Instead, it erroneously attributes the change in reduced-form behavior to changes in the innovation variance, in a manner similar to our results.

This naturally leads to the deeper question of why the TVP-VAR we use is not capable of picking up these shifts seen in the theoretical model. TVP-VARs are a very flexible modeling framework that, in theory, can certainly capture substantial shifts in parameters. At the same time, they also possess many moving parts, and the contribution of each to the ultimate estimation result is not trivial to disentangle. One important aspect is certainly the length of the sample over which the model is estimated. It is well-known that inference under heteroskedasticity (or time variation in the innovation covariance matrix) is quite problematic in short sample (e.g., Toyoda 1974). For that reason, TVP-VARs generally perform better in longer samples, as in Amir-Ahmadi, Matthes, and Wang (2016). A second aspect is that in models with many parameters, the choice of priors can be very important. In particular, priors in TVP-VARs encode a particular view of how much of the variation in the data is due to changes in parameters, changes in volatilities, or additive shocks. The following section shows one alternative to the standard practice that could be used to elicit priors. With standard priors, we would need drastic and sudden shifts in the data to have the estimated coefficients move substantially. Our search and matching model can generate those shifts, but they would arguably not be regarded as realistic for many developed economies.

4. ELICITING PRIORS FOR A TVP-VAR


A key element of TVP-VAR modeling is the choice of the prior on the time-varying components. In our benchmark specification, we follow the generally accepted practice in the literature going back to Primiceri

(2005). However, for data with considerably different properties than those commonly used in the literature or to which TVP-VAR models have been applied, our chosen values might not capture a researcher's prior view on time variation in the data set at hand. More specifically, the prior on the lag coefficient matrices $A_{j,t}$ may be too tight in our framework, so the true underlying time variation in the reduced-form coefficients is instead forced into the covariance matrix. We therefore consider an alternative that is based on a prior predictive analysis.¹²

Our alternative approach proceeds as follows. We first estimate fixed-coefficient VARs on rolling samples of the same length as our training sample (forty periods) to get paths for the time-varying coefficients and volatilities. In a separate exercise, we then simulate paths for those coefficients based on the benchmark priors described above. The hyperparameters of the alternative prior are chosen to match a set of moments from the paths of the time-varying coefficients and volatilities obtained from the rolling window estimation. We choose the average volatilities of the three sets of time-varying coefficients and volatilities. For each set of κ values that govern the tightness of the prior distribution on the covariance matrix, we run twenty-five simulations to generate paths of the same length as the paths from our rolling window estimation and average over the moments obtained in those simulations. We then pick the vector of κ coefficients that minimizes the quadratic distance between the moments from the simulations and the rolling window estimation. The difference in the moments obtained by simulation and the rolling window estimation is rescaled by their value obtained in the rolling window estimation. This avoids one set of moments dominating our calculation since the coefficients have different scales. We use a grid of values for the κ parameters. As lower bounds for the grid, we impose the values used by Primiceri (2005) since we are worried about not capturing enough time variation. Upper bounds are roughly ten times the values chosen by Primiceri (2005).

Figure 6 shows the resulting values for the prior hyperparameters for our full set of twenty-five simulated samples. The horizontal green line shows our benchmark values. As it turns out, there are, in fact, substantial differences between the values chosen by Primiceri (2005) and the values implied by our approach. In the case of the innovation variance in the law of motion for the time-varying parameters in the coefficient matrices $A_{j,t}$, equation (16), there are only two samples for which our prior choice deviates from the one chosen by this approach;

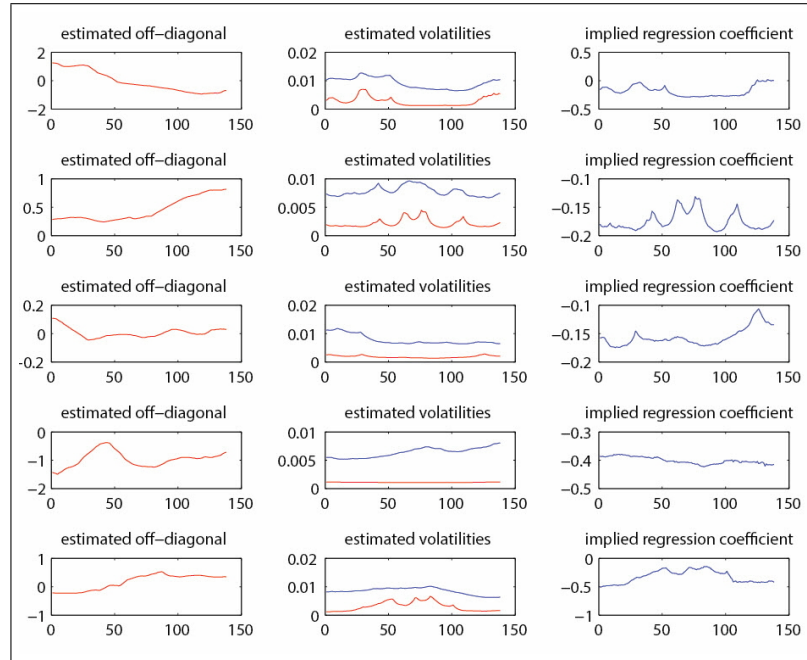
¹² An alternative would be to directly estimate the prior hyperparameters with the rest of the parameters of the model. A Gibbs sampler to do this is described in Amir-Ahmadi, Matthes, and Wang (2017).

Figure 6 Eliciting Priors: Values of Prior Hyperparameters

the deviation is only $\kappa_Q = 0.03$ when compared to our benchmark choice of $\kappa_Q = 0.01$. The hyperparameter scaling the variance of the innovation in the triangular decomposition matrix of the forecast-errors covariance matrix shows more deviations. They can be larger by a factor of up to five, but this is not consistent across each simulation. The largest difference to our benchmark choice can be found for the innovations on the process of the error variances. As can be seen from the bottom graph in Figure 6, the hyperparameter κ_W is larger by an order of magnitude.

This raises the question of whether this alternative prior choice has an effect on the implications derived from the TVP-VAR. We therefore reestimate our model with the much wider priors chosen by the procedure described above. The estimation results are reported in Figure 7, which can be compared directly with Figure 5. We find that the parameter estimates using the alternative prior are almost identical to those obtained using the benchmark specification. The estimated entries of the VAR companion form matrix (not reported) are also virtually identical to the benchmark case. Our conclusion that the Beveridge curve

Figure 7 Summary of Results: Estimated Posterior Means from Alternative Choice of Hyperparameters



shifts in the simulated data are erroneously attributed by the TVP-VAR to time variation in the covariance matrix of the one-step-ahead prediction errors therefore remains intact. In order to get substantial differences in estimated parameters, the prior hyperparameters need to be increased dramatically (e.g. $\kappa_Q = 1$). For our application, the benchmark values consistent with the existing literature therefore seem to be a good choice as far as a naive exercise—that is, without knowledge of the underlying dynamics—is concerned.

5. CONCLUSION

This article makes a simple point. TVP-VARs appear to be predisposed to capture time variation in the underlying data by means of changes in the innovation terms and not via movements in lag coefficients. We arrived at this conclusion by means of a simulation study where we generate a specific form of nonlinearity that would imply time variation in the data. This conclusion holds for a standard choice of priors as well

as an alternative set of priors that we obtain from a prior predictive analysis.

Naturally, the results derived in this article are model dependent and should therefore be taken with a grain of salt. As our model analysis shows, the degree of nonlinearity in the policy function or in the simulated data does not appear to be, heuristically speaking, large. It thus may very well be that the posterior sampler in the Bayesian estimation attributes this type of variation in the data to residual shocks, just as a fixed-coefficient VAR would. What supports this argument is that during times of economic upheaval, chiefly the Great Depression period, TVP-VARs do tend to exhibit considerable time variation in the lag coefficients (Benati and Lubik 2014; Amir-Ahmadi, Matthes, and Wang 2016). That said, we argue that the basic point still applies as to the interpretability of TVP-VAR results. At the very least, researchers should consider a more careful approach to prior selection.

In addition, and independently of the TVP-VAR angle, we propose in this article a modeling framework that conceptualizes structural changes in the labor market and links them to business cycle movements. The mechanism works via an endogenous regime shift in a key labor market parameter, whereby the shift is driven by the interaction of shocks and the intrinsic dynamics of the model. In the case of a simple labor market model, we show that a deep and long recession that originates in adverse productivity realizations can be prolonged by deterioration in the labor market matching process. This mechanism thus offers a convenient setup for studying the behavior of the labor market over the business cycle.

REFERENCES

- Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang. 2016. “Drifts and Volatilities under Measurement Error: Assessing Monetary Policy Shocks over the Last Century.” *Quantitative Economics* 7 (July): 591–611.
- Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang. 2017. “Estimating Prior Hyperparameters in VAR Models with Time-Varying Parameters and Stochastic Volatility.” Manuscript.
- Barlevy, Gadi. 2011. “Evaluating the Role of Labor Market Mismatch in Rising Unemployment.” Federal Reserve Bank of Chicago *Economic Perspectives* 35 (Third Quarter): 82–96.
- Benati, Luca, and Thomas A. Lubik. 2014. “The Time-Varying Beveridge Curve.” In *Advances in Non-linear Economic Modeling: Theory and Applications*, edited by Frauke Schleer-van Gellecom. Heidelberg: Springer Berlin Heidelberg, 167–204.
- Benati, Luca, and Paolo Surico. 2009. “VAR Analysis and the Great Moderation.” *American Economic Review* 99 (September): 1636–52.
- Blanchard, Olivier J., and Peter A. Diamond. 1989. “The Beveridge Curve.” *Brookings Papers on Economic Activity* (No. 1): 1–76.
- Canova, Fabio, Filippo Ferroni, and Christian Matthes. 2015. “Approximating Time Varying Structural Models With Time Invariant Structures.” Federal Reserve Bank of Richmond Working Paper 15-10 (September).
- Cogley, Timothy, and Thomas J. Sargent. 2005. “Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.” *Review of Economic Dynamics* 8 (April): 262–302.
- Del Negro, Marco, and Giorgio Primiceri. 2013. “Time-Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum.” Federal Reserve Bank of New York Staff Reports 619 (May).
- Hagedorn, Marcus, and Iourii Manovskii. 2008. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited.” *American Economic Review* 98 (September): 1692–706.
- Hamilton, James D. 1989. “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle.” *Econometrica* 57 (March): 357–84.

- Lubik, Thomas A. 2009. "Estimating a Search and Matching Model of the Aggregate Labor Market." Federal Reserve Bank of Richmond *Economic Quarterly* 95 (Spring): 101–20.
- Lubik, Thomas A. 2013. "The Shifting and Twisting Beveridge Curve: An Aggregate Perspective." Federal Reserve Bank of Richmond Working Paper 13-16 (October).
- Lubik, Thomas A., and Christian Matthes. 2015. "Time-Varying Parameter Vector Autoregressions: Specification, Estimation, and an Application." Federal Reserve Bank of Richmond *Economic Quarterly* 101 (Fourth Quarter): 323–52.
- Primiceri, Giorgio. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies* 72 (July): 821–52.
- Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *American Economic Review* 104 (November): 3529–64.
- Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (March): 25–49.
- Sims, Christopher A., and Tao Zha. 2006. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review* 96 (March): 54–81.
- Toyoda, Toshihisa. 1974. "Use of the Chow Test Under Heteroscedasticity." *Econometrica* 42 (May): 601–8.

APPENDIX: DERIVATION OF THE WAGE SCHEDULE

The wage that firms pay to workers is derived as the outcome of a Nash bargaining process. Denoting the workers' weight in the bargaining process as $\eta \in [0, 1]$, this implies the sharing rule:

$$\mathcal{W}_t - \mathcal{U}_t = \frac{\eta}{1 - \eta} \mathcal{J}_t, \quad (\text{A1})$$

where \mathcal{W}_t is the asset value of employment, \mathcal{U}_t is the value of being unemployed, and \mathcal{J}_t is, as before, the value of the marginal worker to the firm. In models with one-worker firms, the net surplus of a firm is given by $\mathcal{J}_t - \mathcal{V}_t$, with \mathcal{V}_t the value of a vacant job. By free entry, \mathcal{V}_t is then assumed to be driven to zero. The value of employment to a worker is described by the following Bellman equation:

$$\mathcal{W}_t = w_t + \beta E_t[(1 - \chi_{t+1})\mathcal{W}_{t+1} + \chi_{t+1}\mathcal{U}_{t+1}]. \quad (\text{A2})$$

Workers receive the wage w_t and transition into unemployment in the next period with probability s . The value of searching for a job, when currently unemployed, is:

$$\mathcal{U}_t = b + \beta E_t[\chi_t(1 - \chi_{t+1})\mathcal{W}_{t+1} + (1 - \chi_t(1 - \chi_{t+1}))\mathcal{U}_{t+1}]. \quad (\text{A3})$$

An unemployed searcher receives benefits b and transitions into employment with probability $\chi_t(1 - \chi_{t+1})$. It is adjusted for the probability that a completed match gets dissolved before production begins next period. Substituting the asset equations into the sharing rule (A1), results, after some algebra, in the wage equation found in the text:

$$W_t = \eta(A_t + \kappa\theta_t) + (1 - \eta)b. \quad (\text{A4})$$

APPENDIX: MODEL SOLUTION

We solve the simple search and matching model fully nonlinearly by means of the monotone mapping algorithm. The algorithm computes an approximation of the decision rule $\widehat{h}^V(N_t, A_t, \chi_t)$, which determines the number of vacancy postings given the state variables: employment N_t , the exogenous productivity shock A_t , and the separation rate process χ_t . The algorithm contains the following steps:

1. Specify a threshold switching value \underline{Y} and discretize the state space S . Formulate an initial guess for the decision rule: $\widehat{h}_0^V(N, A, \chi) \forall \{N, A, \chi\} \in S$.
2. Compute a residual function $\overline{R}(V_t; \{N_t, A_t, \chi_t\})$ based on the following:
 - (a) Y_t is calculated and m_t is given according to the threshold process (1).
 - (b) Calculate next period's employment from (4).
 - (c) Expected values of next-period values in the firm's first-order condition appear as:

$$E_t \left[\frac{(1 - \eta)(A_t^{\rho_A} e^{\varepsilon A_t} - b) - \eta \kappa \left(\frac{V_{t+1}}{1 - N_{t+1}} \right) + \left(\frac{\kappa}{m_{t+1}} \right) \left(\frac{V_{t+1}}{1 - N_{t+1}} \right) \xi}{Y_{t+1}} \right] = X_t,$$

which can be approximated with the truncated distribution:

$$\widehat{X}_t = \int_{\underline{\varepsilon}}^{\varepsilon^*} \phi(\varepsilon; \sigma_\varepsilon^2) \frac{\Phi(N_{t+1}, A_t)}{\Psi(N_{t+1}, A_t)} d\varepsilon + \int_{\varepsilon^*}^{\bar{\varepsilon}} \phi(\varepsilon; \sigma_\varepsilon^2) \frac{\Phi(N_{t+1}, A_t)}{\Psi(N_{t+1}, A_t)} d\varepsilon,$$

where:

$$\begin{aligned} \Phi(N_{t+1}, A_t) &= (1 - \eta)(A_t^{\rho_A} e^{\varepsilon A_t} - b) - \\ &\quad \eta \kappa \left(\frac{\widehat{h}_0^V(N_{t+1}, A_t^{\rho_A} e^{\varepsilon A_t})}{1 - N_{t+1}} \right) + \\ &\quad + \left(\frac{\kappa}{m_t} \right) \left(\frac{\widehat{h}_0^V(N_{t+1}, A_t^{\rho_A} e^{\varepsilon A_t})}{1 - N_{t+1}} \right) \xi, \end{aligned}$$

$$\Psi(N_{t+1}, A_t) = A_t^{\rho_A} e^{\varepsilon A_t} N_{t+1} - \kappa \widehat{h}_0^V(N_{t+1}, A_t^{\rho_A} e^{\varepsilon A_t}).$$

Estimate this expression with a trapezoid rule. Linear interpolation is used in the implementation of the decision rule.

(d) Given the expectations, the residual function is:

$$\bar{R}(V_t; \{N_t, A_t, \chi_t\}) = \left| (1 - \chi_t)\beta Y_t \hat{X} - \frac{\kappa}{m_t} \left(\frac{V_t}{1 - N_t} \right)^\xi \right|,$$

which can be interpreted as the absolute value of the difference between the right-hand side and left-hand side of the firm's first-order condition.

3. This residual function is minimized over V_t for every triple $\{N_i, A_j, \chi_k\}$ in S . The decision rule is then updated based on:

$$\hat{h}_2^V(N_i, A_j, \chi_k) = \arg \min \{ \bar{R}(V_t; \{N_i, A_j, \chi_k\}) \} \forall \{n_i, A_j, \chi_k\} \in S.$$

4. The algorithm is repeated until:

$$\max \left| \hat{h}_{k+1}^V(N_i, A_j, \chi_k) - \hat{h}_k^V(N_i, A_j, \chi_k) \right| < \varepsilon.$$