Dynamic Beveridge Curve Accounting

Hie Joo Ahn and Leland D. Crane

Dynamic Beveridge Curve Accounting[*]

Hie Joo Ahn   Leland D. Crane[†]

March 25, 2020

Abstract

We develop a dynamic decomposition of the empirical Beveridge curve, i.e., the level of
vacancies conditional on unemployment. Using a standard model, we show that three
factors can shift the Beveridge curve: reduced-form matching efficiency, changes in the
job separation rate, and out-of-steady-state dynamics. We find that the shift in the Beveridge curve during and after the Great Recession was due to all three factors, and each
factor taken separately had a large effect. Comparing the pre-2010 period to the post-
2010 period, a fall in matching efficiency and out-of-steady-state dynamics both pushed
the curve upward, while the changes in the separation rate pushed the curve downward.
The net effect was the observed upward shift in vacancies given unemployment. In previous recessions changes in matching efficiency were relatively unimportant, while dynamics and the separation rate had more impact. Thus, the unusual feature of the Great
Recession was the deterioration in matching efficiency, while separations and dynamics
have played significant, partially offsetting roles in most downturns. The importance
of these latter two margins contrasts with much of the literature, which abstracts from
one or both of them. We show that these factors affect the slope of the empirical Beveridge curve, an important feature in recent welfare analyses estimating the natural rate
of unemployment.

1 Introduction

The empirical Beveridge curve—the level of vacancies conditional on unemployment—has long been of interest to economists and policy makers. Interest intensified in the wake of the Great Recession, as the curve appeared to shift upwards (see Figure 1), fueling concerns about the functioning of the labor market. There is not currently consensus on the cause of this shift (or historical Beveridge curve shifts). Many papers have attributed the shift to falling matching efficiency (whether due to mismatch, duration dependence, recruiting intensity, heterogeneity, or other causes.) Other researchers have argued that mechanical out-of-steady state dynamics can account for the apparent shift. Finally, it has also been noted that variation in the employment separation rate can also produce shifts in the Beveridge curve. Each of these threads of the literature has taken a slightly different modelling approach. Some authors use steady-state approximations, while others assume a constant job separation rate.

In this paper we provide a new, unified accounting model for the Beveridge curve and a related decomposition method. In our baseline model, where the labor-force status is either employed or unemployed, there are three main factors that matter for the position of the Beveridge curve: (1) matching efficiency, (2) the job-separation probability, and (3) out-of-steady-state dynamics. We analyze how much each of these factors shifted the Beveridge curve. The model allows us to estimate how the contribution of each factor changed in different recessionary and recovery episodes. We extend our model to include the labor-force participation margin, to see how important labor-supply factors are in the dynamics of Beveridge curve.

We find that matching efficiency, job separations and out-of-steady-state dynamics are *all* important in understanding the shifts of the Beveridge curve over business cycles, particularly in the Great Recession. Out-of-steady state dynamics (defined below) produced a net upward shift in the Beveridge curve during and after the Great Recession, as suggested
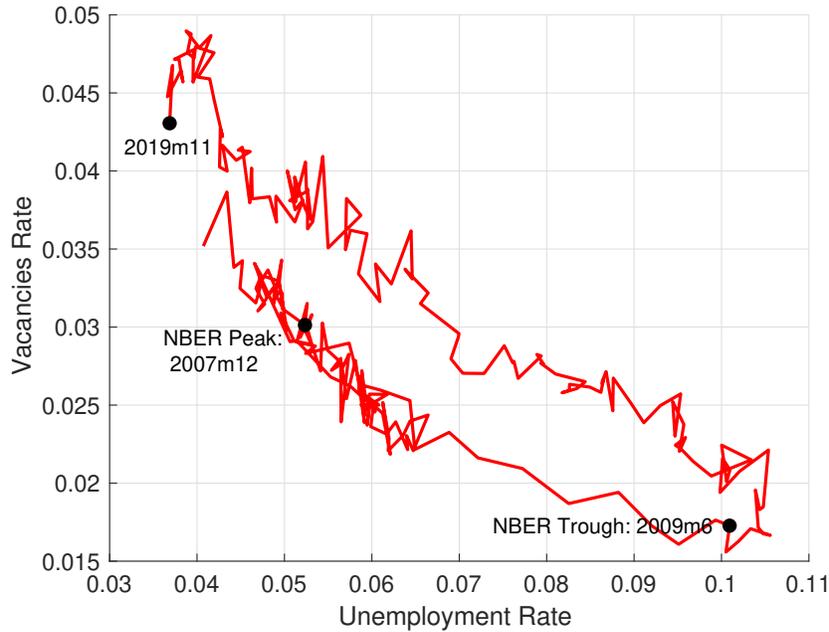
1

Figure 1: The Beveridge curve

by Christiano et al. (2015) and Furlanetto and Groshenny (2016).[1] Those papers assume a constant job separation rate (i.e., the rate at which workers lose their jobs and enter unemployment is constant over the business cycle), but we find that changes in the job separation rate shifted the Beveridge curve sharply *down* on net in the later part of the Great Recession. This downward shift of the Beveridge curve partially offset the combined upward shifts from out-of-steady-state dynamics and matching efficiency. In fact, changes in the separation rate were the largest single factor moving the Beveridge curve.

Our accounting exercise is conditional on the observed path of unemployment. When considering, say, a counterfactual path for matching efficiency, we calculate the level of vacancies in each period that is consistent with the *actual* path of unemployment unemployment under counterfactual matching efficiency. Intuitively, conditional on a path for unemployment, lower matching efficiency require more vacancies to offset lower hiring rates

---

[1]See also Eichenbaum (2015) for related discussion.

conditional on tightness. Thus, lower matching efficiency shifts the level of vacancies (the Beveridge curve) up.

Perhaps less intuitively, a higher job separation rate will shift the Beveridge curve up. Higher separations implies more inflows to unemployment. To keep unemployment at the observed levels, vacancies must be higher, to absorb the extra workers. The job separation probability was high in the downswing of the Great Recession, and it later fell back to more normal levels in the recovery. This had the effect of shifting the Beveridge curve up in the downswing and down in upswing. Elsby et al. (2015) documented a similar point, though they did not quantify the extent of the shift or compare it to the other shifters.[2] We also find that matching efficiency fell significantly during and after the Great Recession, which pushed the Beveridge curve up. This result is consistent with, e.g., Barnichon and Figura (2015).[3]

Analyses which ignore one or more of these shifters will either fail to match the data or will risk making mistaken inferences. This leads to several concrete conclusions and recommendations: First, the importance of out-of-steady-state dynamics implies that the usual flow steady-state approximations are *not* appropriate for studying the Great Recession, or similar periods of rapid change in the unemployment rate. Flow steady-state approximations have become a fundamental tool for simplifying and understanding the labor market (see, for example, Fujita and Ramey (2009), Elsby et al. (2009), Shimer (2012), Barnichon et al. (2012), Elsby et al. (2015).) Unfortunately, in the Great Recession unemployment was consistently far from the steady-state value implied by inflows and outflows, thus the approximation is poor during this period. Similarly, We find a large role for out-of-steady-state dynamics in some previous recessions.

Second, time-variation in the job separation probability is critical for understanding the Beveridge curve, and indeed was the single largest shifter of the Beveridge curve in the

---

[2]Hall and Schulhofer-Wohl (2018) also noted the unemployment inflow rate complicates the behavior of the Beveridge curve.

[3]See also Barnichon and Figura (2010) and Barnichon et al. (2012) for more on matching efficiency.

Great Recession. Thus, the common simplifying assumption of a constant separation rate (made in, e.g., Christiano et al. (2015)) is not appropriate when trying to model the Beveridge curve. In fact, we find that variation in the separation rate was an important shifter of the Beveridge curve in many previous recessions as well, and this variation also affects the slope of the empirical curve. Our analysis *does not* speak directly to the debate over the relative importance of the separations versus the job findings for the evolution of unemployment (see, e.g., Fujita and Ramey (2009), Elsby et al. (2009), Shimer (2012), Ahn and Hamilton (2019)). Rather, we simply point out that the Beveridge curve cannot be properly understood without this ingredient.

Third, we confirm that there was a clear fall in reduced-form matching efficiency in the Great Recession, as has been documented in several other papers (see Elsby et al. (2010), Barnichon and Figura (2015)). We show that this drop in matching efficiency shifted the Beveridge curve substantially and persistently upward in the Great Recession (though the other shifters partially obscure this effect.) In this paper we do not attempt to explain *why* matching efficiency fell, instead we seek to quantify the effects on the Beveridge curve and the interactions with other factors.[4]

Though all three of these factors are crucial in understanding the Beveridge curve, we also find that the relative importance of each factor differed across recessionary episodes.[5] We find that the 1990's recession was similar to the Great Recession in that matching efficiency was the key factor to the persistent outward shift of Beveridge curve. However, in the other recessions in the 1970's, 1980's and 2001, the job separation probability and out-of-steady-state dynamics played more important roles than matching efficiency.

In addition to clarifying the source of loops in the Beveridge curve, we show that these shifters affect the *slope* of the empirical Beveridge curve. This occurs because the curve is being shifted while labor market upswings and downswings progress, not just at peaks and

---

[4]Many papers have offered explanations for the fall in reduced-form matching efficiency among them Davis et al. (2013), Sahin et al. (2014), Elsby et al. (2015), Barnichon and Figura (2015), Kroft et al. (2016), Ahn and Hamilton (2019), and Hall and Schulhofer-Wohl (2018).

[5]Daly et al. (2011) and Diamond and Sahin (2015) document historical Beveridge curve shifts.

troughs. Thus the slope of the steady-state Beveridge curve under constant separations and constant matching efficiency is very different from the empirical slope. This has direct implications for the work of Michaillat and Saez (2019), who exploit the slope of the Beveridge curve to estimate the efficient level of unemployment and the unemployment gap. A back of the envelope exercise shows that using an arguably more appropriate slope cuts the estimated unemployment gap in half, relative to Michaillat and Saez (2019). We view this as evidence that more work is needed to understand how time-varying factors affect the slope of the empirical Beveridge curve.

For our baseline results, we work with a log-linearized Beveridge curve, which expresses the vacancy rate a linear function of various factors. This first-order approximation matches the observed Beveridge curve quite well, and the factors and their associated coefficients are easily interpretable. This analytical tool makes it easy to trace out the contributions of factors to the shifts in the Beveridge curve, and trace out counterfactual curves that hold various factors constant.

One possible concern is that results based on a Taylor series approximation can be inaccurate. In addition, under an approximate Beveridge curve the implied paths of vacancies will not be exactly consistent with the matching function and the law of motion for unemployment. To address this concern we perform similar decompositions, holding various factors constant, using the actual, non-linear Beveridge curve relation, and show that the results are nearly unchanged. Of course, when using the non-linear version the exact contributions of each margin depend on the ordering of the variables in the decomposition. But the results are qualitatively consistent across all orderings.

The next section introduces the basic model. Section 3 discusses the data. Section 4 linearizes the model and presents the results for the Great Recession. Previous recessions are covered in Section 5, and the results of a three-state model are discussed in Section 6. Section 7 concludes. Appendix A addresses the robustness of the linearized results by calculating exact non-linear decompositions.

2 Model

This section derives a version of the simple Beveridge curve framework used in Christiano et al. (2015) (hereafter CET) and Eichenbaum (2015), which is nearly identical to that of Elsby et al. (2015). We do not close the model by making assumptions about the job creation process, wage determination, or other fundamentals. Instead we focus on deriving conclusions that must hold for any general equilibrium model whose labor market is described by (1) the standard law of motion for unemployment and (2) the usual matching function relationship.

Let U_t be the unemployment rate in month t. Similarly, let V_t and H_t denote the vacancy and hires rates both of which are normalized by the labor force. There is no on-the-job search, no participation margin, and the size of the labor force is constant and normalized to unity. At the beginning of each period the unemployed search for jobs, and those that find matches are hired. At the end of the period a fraction of continuing (and not new) matches are destroyed. The flow of new hires in month t, H_t, is given by the standard Cobb-Douglas matching function:

$$H_t = \sigma_t U_t^{1-\alpha} V_t^{\alpha} \tag{1}$$

where α is the elasticity of the matching function and σ_t is matching efficiency, which can vary over time. Then the job-finding probability is given by

$$f_t = \sigma_t (V_t / U_t)^{\alpha}. \tag{2}$$

The law of motion for unemployment is

$$U_{t+1} = s_t (1 - U_t) - f_t U_t + U_t \tag{3}$$

where s_t is the job separation probability. Substituting equation (2) into (3) and rearranging

we arrive at

$$V_t = \left[\frac{s_t(1 - U_t) - \Delta U_{t+1}}{\sigma_t U_t^{1-\alpha}} \right]^{1/\alpha} \tag{4}$$

where $\Delta U_{t+1} = U_{t+1} - U_t$. This is a slight generalization of CET equation 5.2. Whereas CET assume that s_t and σ_t are constants, we permit time-variation in these parameters. Note that if s_t is set to its observed values and σ_t is chosen to verify equation (1), then equation (4) is an identity.

Equation (4) is at the core of our analysis. To understand it better, consider the case where s_t, σ_t and U_t are constants:

$$V = \left[\frac{s(1 - U)}{\sigma U^{1-\alpha}} \right]^{1/\alpha}. \tag{5}$$

This is the steady state Beveridge curve relationship at the core of textbook search models (see Pissarides (2000)): a steady-state with low U must have high V, and vice-versa. Taking equation (5) as the reference point, variation in s_t, σ_t and ΔU_{t+1} changes the level of V_t given U_t. Thus, with a slight abuse of terminology, we will refer to these factors as *shifters*.[6]

Given a path for unemployment and hypothesized, possibly counterfactual, values of the parameters (α, s_t, σ_t), one can calculate the implied path of vacancies from equation (4) and compare it to the true path of vacancies. This is the essence of our exercises in Section 4.

3   Data

We require data on all the variables and parameters in equations (3) and (4). We use the standard approaches, based mostly on Shimer (2012) and Barnichon and Figura (2015). We set U_t as the number of unemployed divided by the labor force, as measured in the Current Population Survey (CPS). We set V_t equal to the count of vacancies from Job Openings and

---

[6]We use shifters to mean factors that change V_t given U_t. Note s and σ also shift the steady-state Beveridge curve (5), while ΔU_{t+1} does not. The dynamics captured by ΔU_{t+1} produce loops *around* the steady-state Beveridge curve, but do not change that model-based relationship.

Labor Turnover Survey (JOLTS), divided by the size of the labor force. Figure 2 plots the two series.
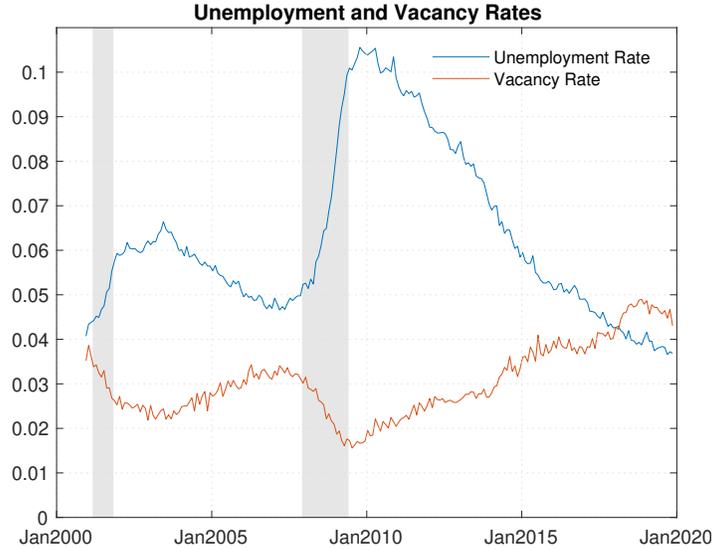
Unemployment and Vacancy Rates

Note: Monthly data, 2000-2019. NBER recessions shaded in gray.
Source: Current Population Survey (U.S. Census Bureau), Job Openings and Labor Turnover Survey (U.S. Bureau of Labor Statistics).

Figure 2: Unemployment and Vacancy Rates

We set the monthly job-finding probability, f_t as in Shimer (2012), using data on the number of short-term unemployed each month.[7] We then choose s_t to satisfy the law of motion (3) exactly.[8]

Figure 3 shows the job finding and separation probabilities. It is notable that the job finding probability fell by about 50 percent in the Great Recession and the separation probability increased by about 50 percent.[9] This suggests that both margins may have played a

---

[7]That is, we set $f_t = 1 - \frac{U_{t+1} - U_{t+1}^s}{U_t}$, where U_{t+1}^s is the number of workers unemployed for less than five weeks in month $t+1$. Thus f_t is the probability that a worker unemployed in month t finds a job by $t+1$. In the data it is possible for such a worker to both find and lose a job (or multiple jobs) before $t+1$, but the discrete-time model we use rules out this possibility.

[8]In both our setup and the continuous time formulation of Shimer (2012), job separation flows are set so as to make the observed sequence of stocks consistent with the flows. In the three-state model of Section 6 the transition rates are taken directly from the data and adjusted to be consistent with the stocks.

[9]Christiano et al. (2015) note that the job separation rate, as measured by JOLTS, *fell* in the Great Recession. The JOLTS separation rate includes job-to-job flows, which are known to be highly procyclical, as well as flows to nonparticipation. Their model, like ours, does not allow for job-to-job flows. The JOLTS separation rate is likely

significant role in the evolution of unemployment. We will confirm this impression in what follows.
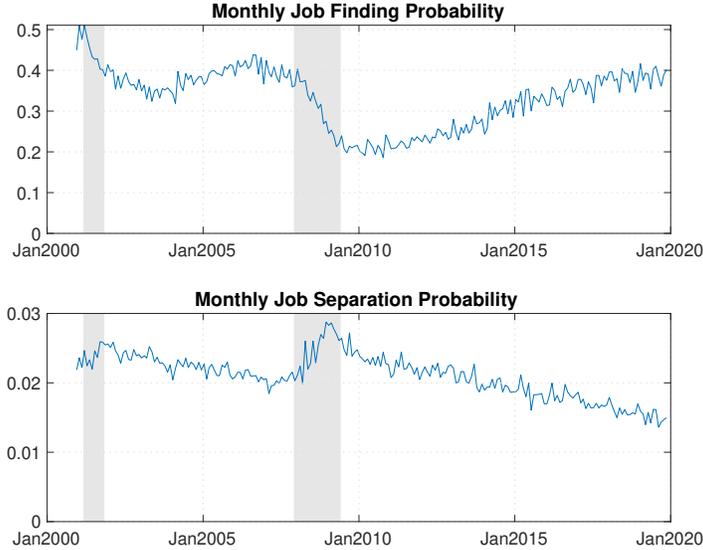
Note: Monthly data, 2000-2019. NBER recessions shaded in gray.
Source: Current Population Survey (U.S. Census Bureau), Job Openings and Labor Turnover Survey (U.S. Bureau of Labor Statistics).

Figure 3: Observed Transition Probabilities

Measurement of α and σ_t require estimation of the matching function. We run the usual regression

$$\ln f_t = \ln \overline{\sigma} + \alpha \ln \left(\frac{V_t}{U_t} \right) + \varepsilon_t \tag{6}$$

where ε_t is the mean-zero error term, $\sigma_t = \overline{\sigma} \exp(\varepsilon_t)$ is time-varying matching efficiency, and $\overline{\sigma}$ is interpreted as average matching efficiency.

Figure 4 plots the log job finding probability against the log V-U ratio. The data for different periods are plotted in different colors. It is evident that matching efficiency deteriorated significantly post-2008. Any change in the matching elasticity α was minor by comparison,

---

the correct measure when considering the firm's problem, since it gives the expected duration of the match. But when considering the evolution of unemployment it is better to use the inflow to unemployment, rather than including job-to-job flows.

so we will continue assuming that α is a constant throughout the paper (as is standard in the literature).
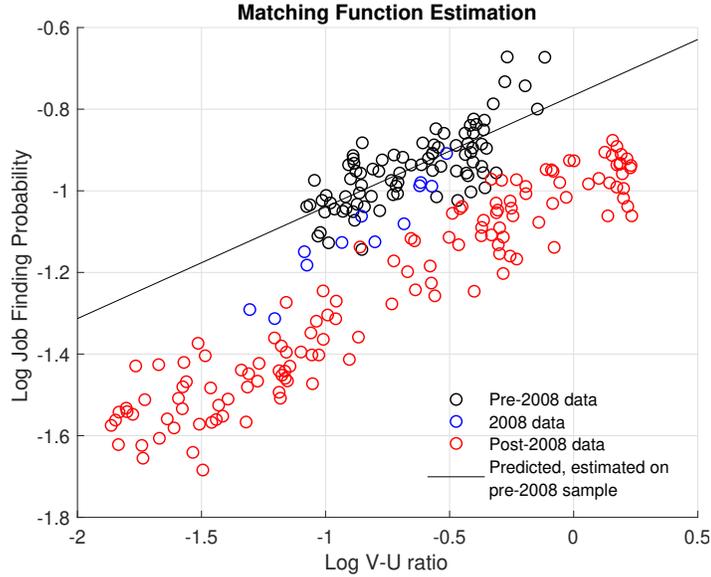
Figure 4: Matching Function Estimation

We run equation (6) on a sample starting in 2000 (when the JOLTS series begins) and ending in 2007, a period where it is plausible that σ was indeed constant. We also run the regression on a post-2008 sample. Table 1 presents the results. The point estimates put α near 0.3, very similar to the estimates of Shimer (2005) and Barnichon and Figura (2015), who use longer time series. It is evident that average matching efficiency fell about 25% between the two samples.

4 Linearization and Results

In order to simplify the discussion, we log-linearize equation (4). In particular, we take the first order Taylor approximation around a point $(U_t, s_t, \sigma_t, \Delta U_{t+1}) = (\overline{U}, \overline{s}, \overline{\sigma}, 0)$. The result is the following expression

	(1)	(2)
	Pre-2008 Sample	Post-2008 Sample
$\ln\overline{\sigma}$	-0.77^{***}	-1.00^{***}
	(0.02)	(0.01)
α	0.27^{***}	0.34^{***}
	(0.03)	(0.01)

Notes: OLS estimates of average matching efficiency ($\ln\overline{\sigma}$) and the matching function elasticity (α). *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively. Standard errors are in parentheses.

Table 1: Matching Function Estimates

$$\ln V_t \approx \ln\overline{V} - \left(\frac{\overline{U}}{\alpha\left(1-\overline{U}\right)} + \frac{1-\alpha}{\alpha}\right)\left(\ln U_t - \ln\overline{U}\right)$$

$$-\underbrace{\frac{\overline{U}}{\alpha\overline{s}(1-\overline{U})}\Delta\ln U_{t+1}}_{\substack{\text{Shift due to}\\\text{Dynamics}}} + \underbrace{\frac{1}{\alpha(1-\overline{U})}\left(\ln s_t - \ln\overline{s}\right)}_{\substack{\text{Shift due to}\\\text{Separations}}} - \underbrace{\frac{1}{\alpha}\left(\ln\sigma_t - \ln\overline{\sigma}\right)}_{\substack{\text{Shift due to}\\\text{Matching Efficiency}}} \qquad (7)$$

where $\overline{V}$ is equation (4) evaluated at $\left(\overline{U}, \overline{s}, \overline{\sigma}, 0\right)$.

The first line of equation (7) is the (approximate) steady-state Beveridge curve. The second line contains the "shifters". Treating $\ln V_t$ as a linear function of $\ln U_t$, these shifters move the y-intercept of the steady-state curve up and down. For example, we can see that when unemployment is rising ($\Delta\ln U_{t+1} > 0$) then $\ln V_t$ will be lower than the steady state curve. This is because, all else equal, rising unemployment implies low finding and thus low $\ln V_t$, which is the out-of-steady-state dynamics mechanism outlined in Pissarides (2000).

While increasing in $\Delta\ln U_{t+1}$ shifts $\ln V_t$ down, increases in the job separation probability s_t shift the curve up. The intuition is that a higher job-separation probability, conditional on a fixed value of $\Delta\ln U_{t+1}$, requires more equilibrium vacancies to absorb the unemployment inflows. Increases in matching efficiency σ_t obviously shift the curve down, as fewer vacancies are needed to rationalize the observed value of $\Delta\ln U_{t+1}$.
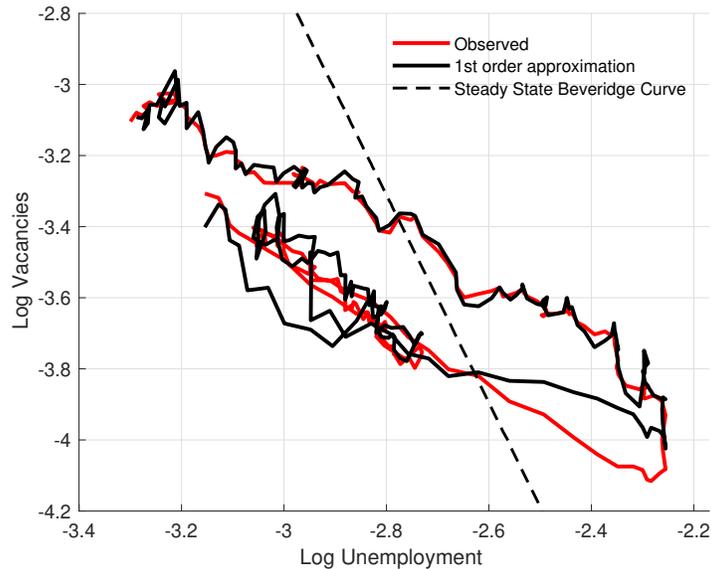
11

We are interested in approximating the Beveridge curve around the Great Recession. To that end, we center the Taylor approximation around post-2007 averages. This yields $\overline{U} = 0.068$, $\bar{s} = 0.020$, and $\bar{\sigma} = 0.359$. We set $\Delta \ln U_{t+1} = 0$ at the approximation point, which is close to its post-2007 average anyway.

4.1 Results

Figure 5 plots the (log) observed Beveridge curve, the first order approximation, and the steady-state Beveridge curve. The approximate Beveridge curve, which includes all the (first order) effects of the shifters, follows the actual curve closely, aside from a brief period near the trough of the Great Recession. Most importantly, the approximate curve shows nearly the same shift (between recession downswing and recovery) as the observed curve. The good fit of the linearized curve gives us confidence that our decomposition of the linearized curve will also be accurate for the exact curve. Appendix A addresses any lingering concerns about the accuracy of the linearized results by calculating a series of nonlinear decompositions on the exact Beveridge curve.

Both the actual Beveridge curve and the approximate curve are significantly flatter than the steady state curve. In log space, the slope of the steady state curve is roughly $-\frac{1-\alpha}{\alpha} = -2.66$, while the slope of the empirical curve is near unity. The difference in slopes is due to slow variation in the shifters, which pushed vacancies up as the Great Recession took hold, and then pushed vacancies down in the recovery. Figure 6 plots the time paths of the three shifter terms in equation (7), along with the net shift (the black line), all normalized to be zero in April 2007. The blue line shows the shift in the Beveridge curve attributable to out-of-steady-state dynamics (that is, $-\frac{\overline{U}}{\alpha \bar{s}(1-\overline{U})} \Delta \ln U_{t+1}$.) The red and yellow lines similarly show the shifts due to separations and matching efficiency.

Relative to the pre-Great Recession period (say, 2007), the net effect of the shifters was to move vacancies sharply upward during the recession. This effect then dissipated very slowly, with the shifters returning to their pre-recession net value only in 2017. This com-

Figure 5: Beveridge Curves

bined effect explains why the slope of the empirical Beveridge curve is so much flatter than the steady state curve. We return to this point in Section 4.2.

Turning to each shifter separately, contribution of each factor is complicated and time-varying. Out-of-steady-state dynamics pushed the Beveridge curve intercept sharply down in the recession, and modestly up in the recovery, more or less the way Pissarides (2000) describes. The contribution of separations is roughly the opposite, raising the intercept sharply, especially late in the recession, and then eventually pushing the intercept down. Finally, the deterioration in matching efficiency raised the intercept during and after the recession.

Figure 6 cannot clearly tell us which factors are responsible for the shift in the *empirical* Beveridge curve between the downswing and the upswing of the Great Recession. To understand that, we need to condition on a level of unemployment and examine the vertical shift evident in Figure 5.

Note: 3 month moving averages of monthly data, 2000-2019. NBER recessions shaded in gray. Shifters are relative to April 2007 values.
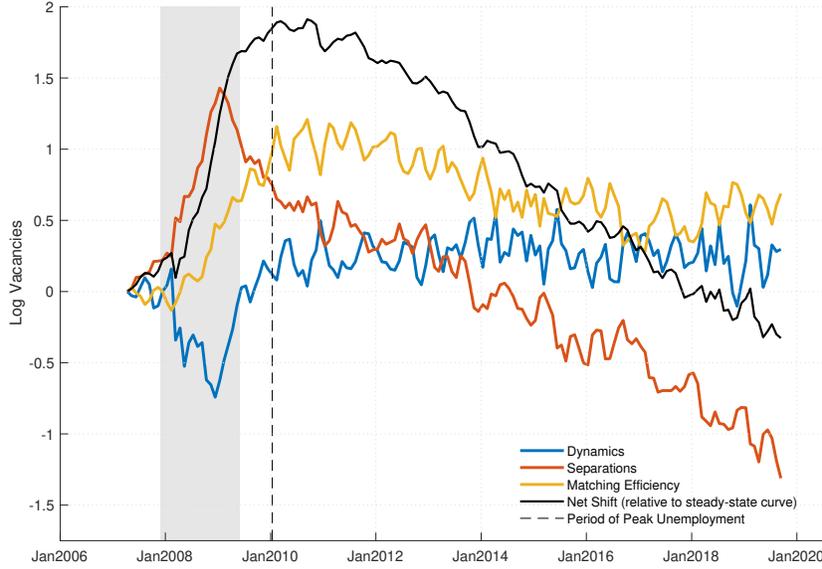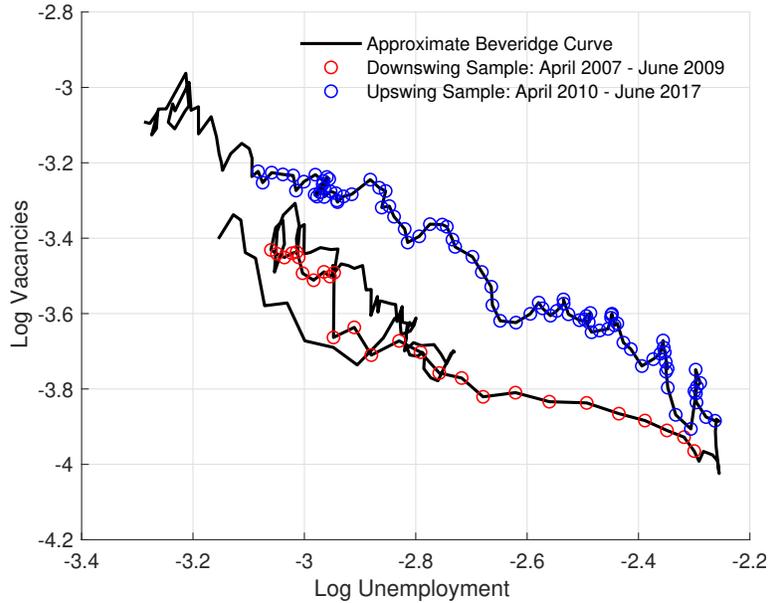*Source:* Current Population Survey (U.S. Census Bureau), Job Openings and Labor Turnover Survey (U.S. Bureau of Labor Statistics).

Figure 6: Shifters of the Approximate Beveridge Curve

Say that there were two months, t and t', where observed unemployment rates were exactly equal, $U_t = U_{t'}$. Then using equation (7) we could decompose the (approximate) difference in vacancies, $\ln V_{t'} - \ln V_t$, as follows:

$$\ln V_{t'} - \ln V_t \approx$$

$$\underbrace{-\frac{\overline{U}}{\alpha \overline{s}(1-\overline{U})} \left(\Delta \ln U_{t'+1} - \Delta \ln U_{t+1} \right)}_{\substack{\text{Shift due to} \\ \text{Dynamics}}} + \underbrace{\frac{1}{\alpha(1-\overline{U})} \left(\ln s_{t'} - \ln s_t \right)}_{\substack{\text{Shift due to} \\ \text{Separations}}} - \underbrace{\frac{1}{\alpha} \left(\ln \sigma_{t'} - \ln \sigma_t \right)}_{\substack{\text{Shift due to} \\ \text{Matching Efficiency}}} \quad (8)$$

Equation (8) provides an additive decomposition of the vertical shift in the Beveridge curve. The portion of $\ln V_{t'} - \ln V_t$ due to, say, differences in matching efficiency between t and t' is just the log difference in matching efficiency, $\ln \sigma_{t'} - \ln \sigma_t$, multiplied by $1/\alpha$. The shifts due to dynamics and separations are similar. The only wrinkle in implementing equation (8) is that we never observe two months with exactly the same unemployment rate, so

14

Figure 7: Downswing and Upswing Samples

we linearly interpolate all relevant series.

As the reference points, we select the unemployment rates observed between April 2007 and June 2009. These are highlighted in red in Figure 7 (the "downswing sample"). We compare the downswing sample to the upswing sample, which begins in April 2010 (highlighted in blue). For each of the downswing points, we calculate the vertical distance between observed vacancies and the (linearly interpolated) upswing vacancy levels. We also calculate each of the terms in equation (8).

The result is Figure 8. The x-axis is the unemployment rate. For each unemployment rate, the black line shows the vertical distance between the upswing and downswing samples, as measured in log vacancies. This is the shift in the Beveridge curve we are trying to explain. The black line is the sum of the other three lines, which are the contributions in equation (8). There are several striking results. First, the job-separation probability is responsible for a large shift *down* in the Beveridge curve. This is because separations rose early
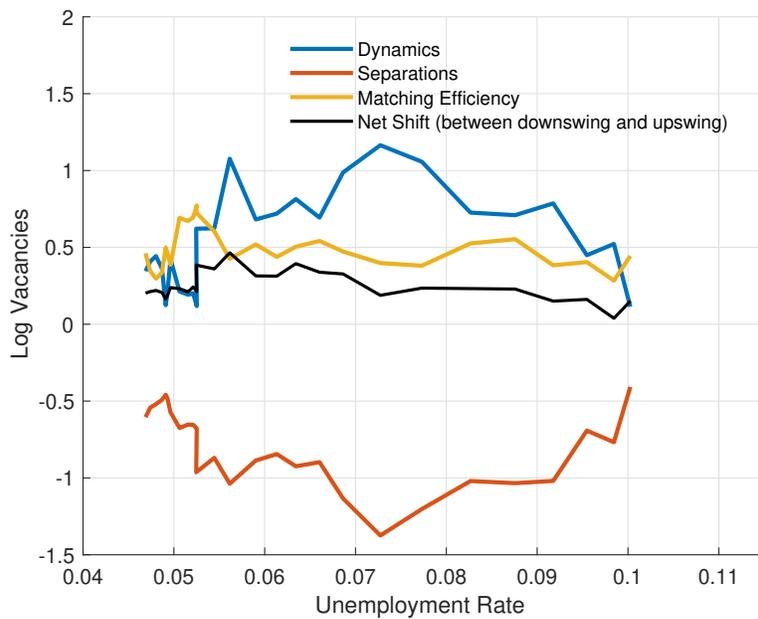
15

Figure 8: Accounting for the Vertical Shift

in the recession. Our accounting exercise is conditional on the path of unemployment, so rising separations implies a higher level of vacancies to keep unemployment at its observed values. Separations later fell, resulting in a lower path of vacancies during the recovery. This net shift is offset by the combined effects of dynamics and matching efficiency, which both pushed the curve up (on net.)

Interestingly, out-of-steady-state dynamics played a prominent role, with a contribution larger than that of matching efficiency over much of the range. This specific result is consistent with Christiano et al. (2015)'s argument that, because the Great Recession was so large and so sudden, dynamics can produce a realistic loop in the Beveridge curve. However, their analysis ignores the separation probability and matching efficiency, which are at least as important for understanding what happened. In particular, matching efficiency more than accounts for the net shift across most of the range, so without a change in matching efficiency the Beveridge curve would have shifted *down*, not up.

To summarize, all three of the factors we consider shifted the Beveridge curve in nontrivial ways. The vertical shift in the empirical Beveridge curve is the net result of out-of-steady state dynamics and matching efficiency both shifting the curve up, an effect which is partially offset by a large negative contribution from the separation probability. The time paths of these shifters are complicated and non-monotonic, leading the slope of the empirical Beveridge curve to differ from the model-implied steady-state curve. We now turn to this result in more detail.

4.2 The Slope of the Beveridge Curve

Recent innovative work by Michaillat and Saez (2019) (MS) has emphasized the importance of the Beveridge curve slope for welfare and the natural rate of unemployment. In this section we show how our measurement methods relate to their results.

In many models with a matching function (e.g., Shimer (2005)), the Beveridge curve describes the possible steady-state values of vacancies and unemployment. In short, an

17

economy that sustains a lower level of unemployment *must* have more vacancies in equilibrium, and vice versa. MS point out that this relationship can be used to estimate the welfare-maximizing level of unemployment in a particularly simple and general way. They note that a social planner will seek to equalize the costs of additional vacancies to the costs of additional unemployment. In other words, the social planner will seek the location on the Beveridge curve where the marginal cost of additional unemployment equals the social value of the resulting reduction in vacancies. This point then defines the natural rate of unemployment, and the difference between observed unemployment and natural rate is the unemployment gap. MS use estimates of the costs of vacancies, the costs of unemployment, and the slope of the Beveridge curve to make their calculations.

MS measure the slope of the Beveridge curve by estimating regressions of V_t on U_t in periods where the Beveridge curve appeared stable (dropping the troughs of recessions, for example.) As we show above, these observed slopes reflect both (1) movements along a stable Beveridge curve (changes in U_t for fixed separations, matching efficiency and dynamics) and (2) time variation in the shifters. This second factor can distort the empirical Beveridge curve relative to the planner-relevant, steady-state curve. For example, consider a bare bones model where the separation probability and matching efficiency are exogenous processes, possibly correlated with the aggregate productivity shock. Such a model fits in our framework (and that of MS), and could produce the observed data, including the empirical Beveridge curve and the paths of the shifts. However, a planner, facing such an economy, would not look to the empirical Beveridge curve to estimate the unemployment-vacancy tradeoff. The reason is that the empirical curve include the effects of the (purely cyclical) shifters, while the planner is interested in long-run, steady state relationships. The correct slope for the planner comes from the linearized curve (7), which treats the shifters as fixed:

$$-\left(\frac{\overline{U}}{\alpha\left(1-\overline{U}\right)}+\frac{1-\alpha}{\alpha}\right) \approx -\frac{1-\alpha}{\alpha} \tag{9}$$

and is determined by the shape of the matching function. The planner would make decisions based on the steady-state curve in Figure 5, not the empirical curve. Thus, in this toy example the empirical Beveridge curve does not directly give us the planner-relevant, long-run relationship we seek.

The key question is whether the planner should incorporate the effects of the shifters when making a choice about the long-run level of unemployment. Clearly, out-of-steady-state dynamics are fundamentally transitory, so the planner should always purge the Beveridge curve of their effect. However, it is possible that the separation probability and matching efficiency are, to some extent, functions of the long-run level of unemployment (unlike in the toy example above). In this case the planner should not remove (all of) their influence when calculating the vacancy-unemployment tradeoff.

Determining the exact nature of the variation in separations and matching efficiency is well beyond the scope of this paper. Instead, we provide an example to demonstrate that these issues can have an economically meaningful impact on welfare calculations. From equation (9), the slope of the steady-state curve (treating the shifters as fixed) is very close to $-\frac{1-\alpha}{\alpha}$. Averaging together the two estimates of α in Table 1, we set set $\alpha = 0.3$, implying a Beveridge curve slope of -2.33. This is far steeper than the estimates of MS, which are around -0.9 for the same period.

We can calculate the efficient levels of unemployment using equation (5) from MS, based on our two estimates of the Beveridge slope (-2.33 and -0.9). In both cases we use MS's preferred values for the costs of vacancies and unemployment. Figure 9 shows the results (this figure is comparable to Figure 3 Panel D in Michaillat and Saez (2019).) The blue line is the actual unemployment rate. The red line shows the efficient level of unemployment according to MS's calibration, with a Beveridge slope of -0.9. The black line shows the
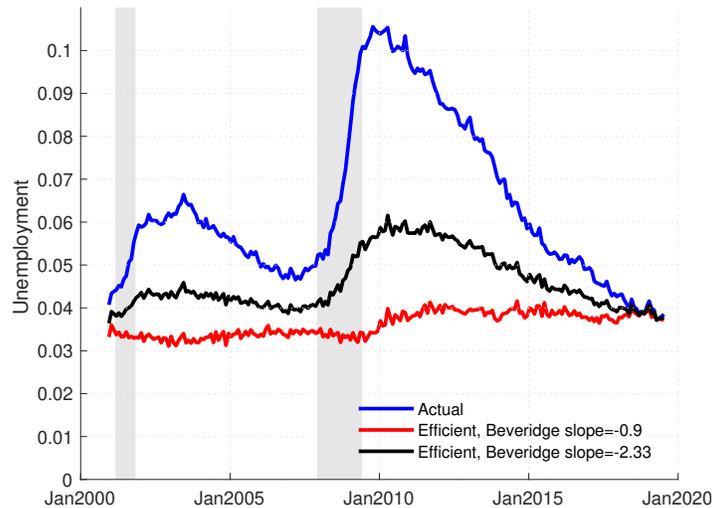
Figure 9: Efficient Unemployment Based on the Beveridge Tradeoff

efficient level of unemployment using our preferred Beveridge slope of -2.33. It is evident that the steeper Beveridge curve significantly raises the efficient level of unemployment, as reducing unemployment with a steep Beveridge curve is more costly in terms of vacancies. In our calibration the natural rate of unemployment fluctuates between about 4 percent and 6 percent, near the range of other estimates including the Congressional Budget Office (CBO)'s short-term natural rate of unemployment. Notably, our calibrated estimate moves very similarly to the CBO's estimate during the post Great-recession period.

Our results suggest that careful work is needed to disentangle which features of the Beveridge curve the planner should care about. These choices have real consequences for the measurement of efficiency, as Figure 9 shows. One approach is to specify a more complete model, which explicitly links separations and matching efficiency to the rest of the economy. With such a model in hand, one could determine the planner-relevant Beveridge curve slope.

5 Previous Recessions

We can also use our framework to analyze recessions prior to the Great Recession. In terms of data, the only change is that up through 2016 we use the composite vacancy series from Barnichon (2010) instead of JOLTS. After 2016 we continue the series by splicing on the JOLTS series. For four historical labor market downturns, we calculate the log-linearized Beveridge curve, as in Section 4. For each episode the curve is linearized around the local mean, to ensure a good fit. Figure 10 compared the observed and linearized Beveridge curves. The fit is generally good, although some of the linearized Beveridge curves show less of a shift, or counter-clockwise loop, than their observed counterparts. We view this as a topic for further investigation

With the linearized Beveridge curves in hand, we can read off the implied contribution of each factor to the shift in the curve at every point in time. Figure 11 is analogous to Figure 6 for each recession: the net shift in the Beveridge curve, and the contributions, as functions of time. It is apparent that in each recession the Beveridge curve intercept began shifting up at the onset of the recession, and slowly drifted down once unemployment began falling. Rising separations usually drove this upward shift, partially offset by out-of-steady-state dynamics.

It can be seen that in all recessions, out-of-steady-state dynamics shifted the Beveridge curve significantly down in the initial stages and generally up in the recovery. Interestingly, this shift is partially offset by the contribution of separations, which (as in the Great Recession) tend to push Beveridge curve sharply upward in the initial stages of a recession and more moderately upward afterward. Thus the changes in the job-separation probability tend to flatten the observed Beveridge curve, and cancel out some of the counter-clockwise loop that out-of-steady-state dynamics induce.

In most previous recessions, changes in matching efficiency had little impact, and were swamped by changes in the other factors. The 1990 recession appears to be an exception here. During the 1990 recession and the recovery period, the deterioration in matching effi-

1973 Recession

1981 Recession

(a)

(b)

1990 Recession

2001 Recession

(c)

(d)
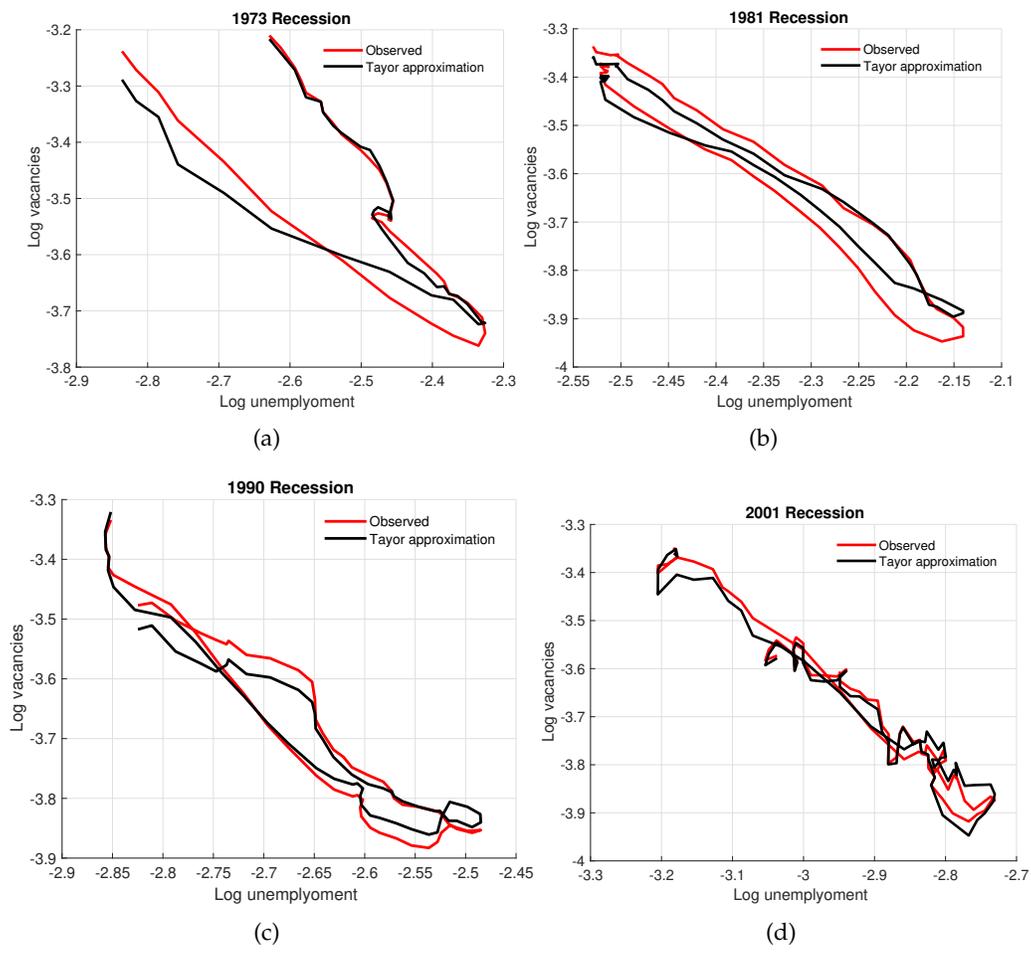
Note: Observed and approximated Beveridge curves for historical downturns. Log-linear Taylor approximations are taken about averages for the periods plotted.
Source: Current Population Survey (U.S. Census Bureau), Barnichon (2010), and authors' calculations.

Figure 10: Observed and Approximate Beveridge Curves

Figure 11: Decompositions of Approximate Beveridge Curves
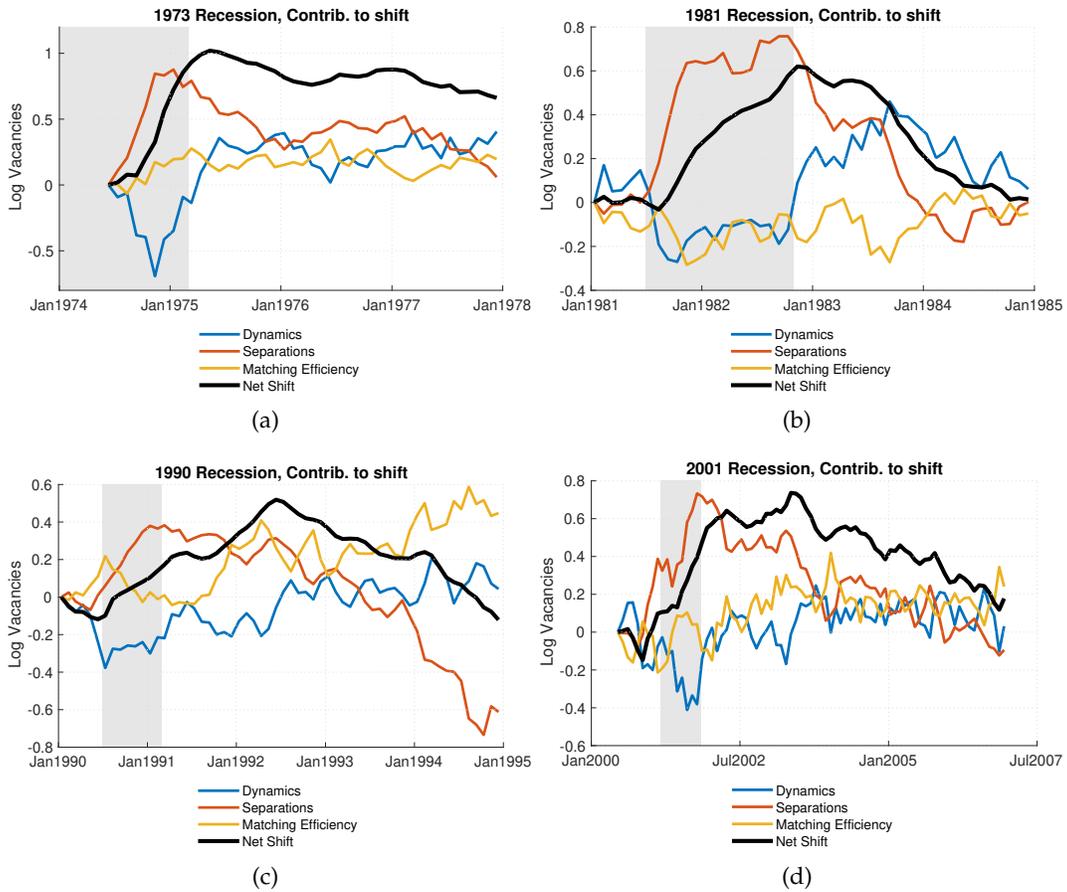
Note: Solid black line shows the time path of the net shift of the (log-linearized) Beveridge curve intercept. NBER recessions shaded in gray.

Source: Current Population Survey (U.S. Census Bureau), Barnichon (2010), and authors' calculations.

ciency continued to push the Beveridge curve up, which is quite similar to what happened in the Great Recession. In fact, the two recessions are similar to each other in a sense that long-term unemployment continued to increase substantially after the recession was over. This suggests that mismatch or related factors might have been an important driver in the rise of long-term unemployment in the two recession episodes. We view this line of reasoning as a topic for future research.

The tentative conclusion is that the Great Recession was exceptional, insofar as the drop in matching efficiency had first-order effects on the Beveridge curve (with the possible exception of the early 1990s recession). In previous recessions matching efficiency usually played little role. However, the modest counter-clockwise loops in previous recession were not simply the product of modest out-of-steady-state dynamics, but were the net result dramatic dynamics being offset by large contributions from the separations margin. Out-of-steady-state dynamics and the separations margin played critical roles in all the recessions examined here.

6   Three State Model

The results so far have assumed that all workers are either employed or unemployed. This is a simplification, since empirically flows into and out of the labor force are important for understanding total hires and evolution of unemployment. In this section we add the participation margin and discuss the robustness of our results in the expanded model.

6.1   Model

The population is still normalized to unity, but we add a nonparticipation state. Let N_t be the stock of nonparticipants, so that $E_t + U_t + N_t = 1$. Consider the law of motion for

unemployment when workers can move into and out of the labor force:

$$\Delta U_{t+1} = E_t + N_t n u_t - U_t u n_t - U_t u e_t \tag{10}$$

The transition rate from nonparticipation to unemployment in month t is nu_t. The terms un_t and ue_t are similarly defined, with eu_t replacing s_t for consistency. The law of motion for nonparticipation is symmetric:

$$\Delta N_{t+1} = E_t e n_t + U_t u n_t - N_t n e_t - N_t n u_t \tag{11}$$

Summing equations (10) and (11) yields an expression involving total hires ($H_t = N_t n e_t + U_t u e_t$)

$$\Delta U_{t+1} + \Delta N_{t+1} = E_t e u_t + E_t e n_t - H_t \tag{12}$$

where the flows between unemployment and nonparticipation have canceled.

We can write the matching function as

$$H_t = \sigma_t (U_t + \xi_t^N N_t)^{1-\alpha} V_t^\alpha \tag{13}$$

where ξ_t^N is the search effort of the nonparticipants relative to the unemployed. Thus the effective mass of searchers is $U_t + \xi_t^n N_t$ and σ_t continues to represent reduced-form matching efficiency.

Combining equations (12) and (13), and assuming balanced matching (that is, hires from unemployment are a share $\frac{U_t}{U_t + \xi_t^N N_t}$ of total hires), we have the following expression for vacancies:

$$V_t = \left[\frac{(1 - U_t - N_t)(eu_t + en_t) - \Delta U_{t+1} - \Delta N_{t+1}}{\sigma_t (U_t + \xi_t^N N_t)^{1-\alpha}} \right]^{1/\alpha} \tag{14}$$

When the non-employed can participate in job search, it is more sensible to think of a Bev-

25

eridge curve which relates vacancies to *searchers* (both unemployed and nonparticipants) instead of unemployment. To this end, we define two new groups. First, we define the pool of searchers S_t as

$$S_t = U_t + \xi_t^N N_t. \tag{15}$$

Second, we define the pool of "true nonparticipants" as

$$\tilde{N}_t = \left(1 - \xi_t^N\right) N_t. \tag{16}$$

While we take no stand on whether ξ_t^N is the fraction of nonparticipants who search or the search effort of each nonparticipant relative to the unemployed, the former interpretation is convenient here. Note that if $\xi_t^N = 1$ all the nonparticipants search and $\tilde{N}_t = 0$. Using S_t and $\tilde{N}_t$, we can write (14) as

$$V_t = \left[\frac{\left(1 - S_t - \tilde{N}_t\right) x_t - \Delta S_{t+1} - \Delta \tilde{N}_{t+1}}{\sigma_t S_t^{1-\alpha}} \right]^{\frac{1}{\alpha}} \tag{17}$$

where $x_t = eu_t + en_t$ is the total job-separation probability. Log-linearizing yields

$$
\begin{aligned}
\ln V_t = &-\frac{1}{\alpha} \left[\ln \sigma_t - \ln \sigma_0\right] \\
&- \left\{ \frac{(1-\alpha)}{\alpha} + \frac{1}{\alpha} \frac{S_0}{(1 - S_0 - \tilde{N}_0)} \right\} \left[\ln S_t - \ln S_0\right] \\
&- \left\{ \frac{1}{\alpha} \frac{S_0}{(1 - S_0 - \tilde{N}_0) x_0} \right\} \left[\Delta \ln S_{t+1}\right] \\
&- \left\{ \frac{1}{\alpha} \frac{\tilde{N}_0}{(1 - S_0 - \tilde{N}_0)} \right\} \left[\ln \tilde{N}_t - \ln \tilde{N}_0\right] \\
&- \left\{ \frac{1}{\alpha} \frac{\tilde{N}_0}{(1 - S_0 - \tilde{N}_0) x_0} \right\} \left[\Delta \ln \tilde{N}_{t+1}\right] \\
&+ \frac{1}{\alpha} \left[\ln x_t - \ln x_0\right]
\end{aligned}
\tag{18}
$$

26

Like equation (4), equation (17) can be used to analyze the Beveridge curve. This decomposition, naturally, has more shifters than the two-state model. In this model movements along the Beveridge curve are captured by the $\ln S_t - \ln S_0$ term, since the curve is defined in terms of searchers, not merely the unemployed. The effects of matching efficiency and separations still appear, on the first and last lines of equation (17) respectively. Finally, there are now two out-of-steady state terms, $\Delta \ln S_{t+1}$ and $\Delta \ln \tilde{N}_{t+1}$, as well as a term capturing the level of non-searchers, $\ln \tilde{N}_t - \ln \tilde{N}_0$. Not all of these terms have a transparent interpretation, but as we shall see below, many of them are not quantitatively important either.
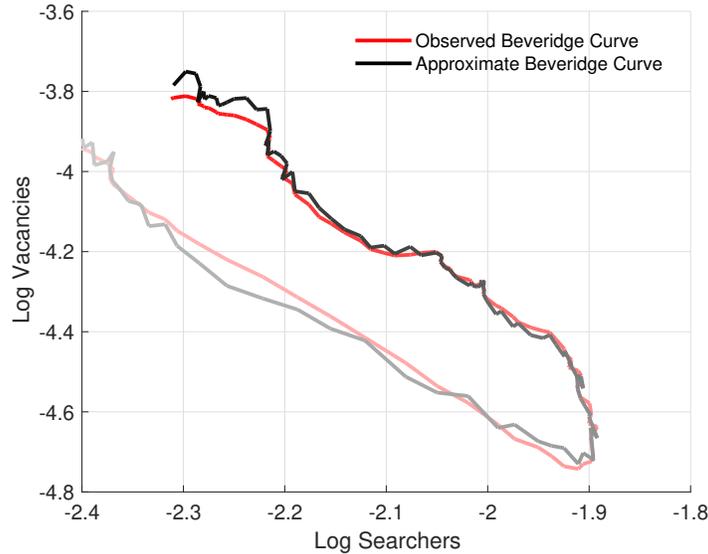
6.2 Data

To implement the three state model, we need data on the terms appearing in equation (14). We obtain the stocks of employed, unemployed, and nonparticipants from the CPS labor force status flows.[10] We normalize these stocks to satisfy $E_t + U_t + N_t = 1$ in all periods. The transition rates eu_t, nu_t, un_t, ue_t are also taken from the labor force status flows. These transition rates are not exactly consistent with the stocks, due to missing month-to-month linkages and sample rotation. We adjust for the inconsistency by raking the rates until they are consistent with the stocks.[11] This results in very small adjustments to the transition rates.

Under the assumption of balanced matching, ξ_t^N can be identified by the ratio of transition rates to employment:

$$\xi_t^N = \frac{ne_t}{ue_t}$$

Finally, α and σ_t can be identified by the matching function regression, using $U_t + \xi_t^N N_t$ as the population of effective searchers.

Figure 12: Three State Approximate Beveridge Curve

6.3   Results

Figure 12 shows that, as with the two state model, the three state approximate Beveridge curve is a good approximation of the observed curve. Here "searchers" are the pool of actively searching workers, $U_t + \xi_t^N N_t$. To show the direction of time, more recent periods are shaded darker.

Figure 13 shows the shifters as a function of time, similar to Figure 6 the story is similar to the two-state model. Matching efficiency slowly and steadily pushed the Beveridge curve upwards during and after the Great Recession. The separation probability, x_t, pushed the Beveridge curve up during the recession, but this was short-lived. The out-of-steady-state dynamics terms, on net, pushed the curve down, though interestingly the $\Delta \tilde{N}_{t+1}$ term partially offsets the $\Delta \tilde{S}_{t+1}$ term. Strikingly, there is no shift in the Beveridge curve under

---

[10]Accessible at `https://www.bls.gov/webapps/legacy/cpsflowstab.htm`.

[11]This is also called iterative proportional fitting: alternately scaling each row and column of the transition matrix to match the stocks until they converge.
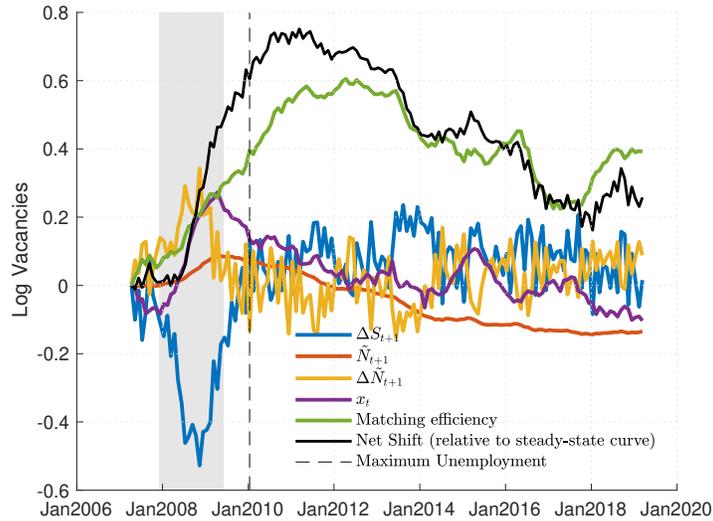
Figure 13: Three State Model - Shifters of the Approximate Beveridge Curve

constant matching efficiency. This confirms the results from the two state model (and much of the literature) that the decline in matching efficiency was an important contributor to the loop in the Beveridge curve.

7 Conclusion

The empirical Beveridge curve is easy to calculate, as it only requires data on the stocks of unemployed workers and job openings. This ease of measurement may help explain the attention it has received. Unfortunately, the Beveridge curve is (even in a simple model) the product of multiple factors, and can be difficult to interpret. Our hope is that our results help clarify the behavior of the Beveridge curve and reconcile some conflicting ideas in the literature.

We have shown that reduced-form matching efficiency, changes in the separation probability, and out-of-steady-state dynamics all played important roles in the recent shift of the

Beveridge curve. Comparing the pre-2010 period to the post-2010 period, out-of-steady-state dynamics and a fall in matching efficiency both pushed the curve upward, while the changes in the separation probability pushed the curve downward. The net effect was the observed upward shift in the empirical Beveridge curve. Our results are largely unchanged when we include a nonparticipation margin. One area for more research is the effect of on-the-job search, which would affect the measurement of matching efficiency.

A realistic model of the Great Recession therefore needs, (1) a mechanism for reduced-form matching efficiency to fall during and after the recession, (2) a non-constant separation probability, which can generate an increase in job losses towards the end of the recession. Furthermore, models should not be evaluated using steady-state approximations, since the rapid changes in the labor market around the Great Recession made out-of-steady-state dynamics a first-order issue.

We reach similar conclusions regarding earlier recessions, though the role of matching efficiency is generally smaller. Importantly, the relatively small Beveridge curve loops in earlier recessions were the product of changes in the separation probability nearly offsetting out-of-steady-state dynamics. We find that these shifters move the intercept of the Beveridge curve continuously, not just at business cycle peaks and troughs. As a result, the slope of the empirical Beveridge curve is distinct from the slope of the implied (constant separation probability, constant matching efficiency) steady-state curve.

References

Ahn, Hie Joo and James D. Hamilton, "Heterogeneity and Unemployment Dynamics," *Journal of Business and Economic Statistics*, 2019.

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

__ **and Andrew Figura**, "What drives movements in the unemployment rate? a decomposition of the Beveridge curve," Finance and Economics Discussion Series 2010-48, Board of Governors of the Federal Reserve System (U.S.) 9 2010.

__ **and** __ , "Labor Market Heterogeneity and the Aggregate Matching Function," *American Economic Journal: Macroeconomics*, October 2015, *7* (4), 222–49.

__ **, Michael Elsby, Bart Hobijn, and Aysegul Sahin**, "Which industries are shifting the Beveridge curve?," *Monthly Labor Review*, 6 2012, *135* (6), 25–37.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt, "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 110–67.

Daly, Mary C., Bart Hobijn, and Robert G. Valletta, "The Recent Evolution of the Natural Rate of Unemployment," IZA Discussion Papers 5832, Institute for the Study of Labor (IZA) July 2011.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger, "The Establishment-Level Behavior of Vacancies and Hiring," *The Quarterly Journal of Economics*, 2013, *128* (2), 581–622.

Diamond, Peter A. and Aysegul Sahin, "Shifts in the Beveridge curve," *Research in Economics*, 2015, *69* (1), 18–25.

Eichenbaum, Martin S., "Comment," *NBER Macroeconomics Annual*, 2015, *29* (1), 129–145.

Elsby, Michael W. L., Bart Hobijn, and Aysegul Sahin, "The Labor Market in the Great Recession," *Brookings Papers on Economic Activity*, 2010, *41* (1 (Spring), 1–69.

_ , Ryan Michaels, and David Ratner, "The Beveridge Curve: A Survey," *Journal of Economic Literature*, September 2015, *53* (3), 571–630.

_ , _ , and Gary Solon, "The Ins and Outs of Cyclical Unemployment," *American Economic Journal: Macroeconomics*, January 2009, *1* (1), 84–110.

Fujita, Shigeru and Garey Ramey, "The Cyclicality Of Separation And Job Finding Rates," *International Economic Review*, May 2009, *50* (2), 415–430.

Furlanetto, Francesco and Nicolas Groshenny, "Mismatch Shocks and Unemployment During the Great Recession," *Journal of Applied Econometrics*, November 2016, *31* (7), 1197–1214.

Hall, Robert E. and Sam Schulhofer-Wohl, "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers," *American Economic Journal: Macroeconomics*, January 2018, *10* (1), 1–32.

Kroft, Kory, Fabian Lange, Matthew J. Notowidigdo, and Lawrence F. Katz, "Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Nonparticipation," *Journal of Labor Economics*, 2016, *34* (S1), S7–S54.

Michaillat, Pascal and Emmanuel Saez, "Beveridgean Unemployment Gap," Working Paper 26474, National Bureau of Economic Research November 2019.

Pissarides, Christopher A., *Equilibrium Unemployment Theory, 2nd Edition*, Vol. 1 of *MIT Press Books*, The MIT Press, January 2000.

Sahin, Aysegul, Joseph Song, Giorgio Topa, and Giovanni L. Violante, "Mismatch Unemployment," *American Economic Review*, November 2014, *104* (11), 3529–3564.

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

— , "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, April 2012, *15* (2), 127–148.

A Full Decompositions

One may be concerned that results based on the Taylor approximation are not robust. While the fit of the approximate Beveridge curve is strikingly good, it is not perfect. Therefore, there is some room for non-linearities to affect the results. A related issues is that the log-linearized Beveridge curve is not dynamically consistent: If we plug implied vacancies into the matching function and the unemployment law of motion, we generally won't get the observed U_{t+1} back.

In this section we decompose the shift in the empirical Beveridge curve using the exact vacancy equation rather than the log-linearized version. Again, the goal is to measure the contributions to the shift due due to out-of-steady-state dynamics, changes in the separation probability, and changes in matching efficiency.

The starting point of our decomposition is the standard, steady-state Beveridge curve, with constant matching efficiency and separations:

$$V_t^{\bar{s},\bar{\sigma},\overline{\Delta U}} = \left[\frac{\bar{s}(1 - U_t)}{\bar{\sigma}U_t^{1-\alpha}} \right]^{1/\alpha} \tag{19}$$

The steady state Beveridge curve sets $\Delta U_{t+1} = 0$. It therefore the level of vacancies that would prevail after many months of constant $\bar{s}$ and $\bar{\sigma}$.

Let t^{down} be a month from the downswing sample, and let t^{up} be the corresponding (interpolated) period from the upswing with the same level of unemployment. Then the observed vertical shift in the Beveridge curve is $V_{up} - V_{down}$. The steady-state Beveridge curve (19) obviously entails no shift, so $V_{up}^{\bar{s},\bar{\sigma},\overline{\Delta U}} - V_{down}^{\bar{s},\bar{\sigma},\overline{\Delta U}} = 0$.

We can define other counterfactual vacancy series. We use superscripts with bars to denote that the margin is being held constant. Thus, for example,

$$V_t^{\overline{\sigma}} = \left[\frac{s_t(1 - U_t) - \Delta U_{t+1}}{\overline{\sigma} U_t^{1-\alpha}} \right]^{1/\alpha} \tag{20}$$

$$V_t^{\overline{\sigma}, \overline{\Delta U}} = \left[\frac{s_t(1 - U_t)}{\overline{\sigma} U_t^{1-\alpha}} \right]^{1/\alpha} \tag{21}$$

with $V_t^{\overline{s}}$, $V_t^{\overline{s},\overline{\sigma}}$, $V_t^{\overline{s},\overline{\Delta U}}$, and $V_t^{\overline{\Delta U}}$ defined similarly.

Next, consider the accounting identity

$$
\begin{aligned}
V_{up} - V_{down} = {} & \left(V_{up} - V_{down} \right) - \left(V_{up}^{\overline{\sigma}} - V_{down}^{\overline{\sigma}} \right) \\
& + \left(V_{up}^{\overline{\sigma}} - V_{down}^{\overline{\sigma}} \right) - \left(V_{up}^{\overline{s},\overline{\sigma}} - V_{down}^{\overline{s},\overline{\sigma}} \right) \\
& + \left(V_{up}^{\overline{s},\overline{\sigma}} - V_{down}^{\overline{s},\overline{\sigma}} \right) - \left(V_{up}^{\overline{s},\overline{\sigma},\overline{\Delta U}} - V_{down}^{\overline{s},\overline{\sigma},\overline{\Delta U}} \right).
\end{aligned} \tag{22}
$$

This writes $V_{up} - V_{down}$ as three double differences. The terms on the right hand side have the following interpretation:

- $\left(V_{up} - V_{down} \right) - \left(V_{up}^{\overline{\sigma}} - V_{down}^{\overline{\sigma}} \right)$: The shift in the Beveridge curve accounted for by the time-variation in matching efficiency, conditional on having s_t and ΔU_t at their observed values.

- $\left(V_{up}^{\overline{\sigma}} - V_{down}^{\overline{\sigma}} \right) - \left(V_{up}^{\overline{s},\overline{\sigma}} - V_{down}^{\overline{s},\overline{\sigma}} \right)$: The shift accounted for by time-variation in the separation probability, conditional on having ΔU_t at its observed values and σ held constant.

- $\left(V_{up}^{\overline{s},\overline{\sigma}} - V_{down}^{\overline{s},\overline{\sigma}} \right) - \left(V_{up}^{\overline{s},\overline{\sigma},\overline{\Delta U}} - V_{down}^{\overline{s},\overline{\sigma},\overline{\Delta U}} \right)$: The shift accounted for by time-variation in ΔU_{t+1}, conditional on having matching efficiency and the separation probability held constant. Note that $V_{up}^{\overline{s},\overline{\sigma},\overline{\Delta U}} - V_{down}^{\overline{s},\overline{\sigma},\overline{\Delta U}} = 0$ by construction.

Thus, we can interpret equation (22) as moving us from the steady-state Beveridge curve

Ordering	Dynamics	Separations	Matching
$\Delta U_{t+1}, s, \sigma$	115.26	-177.18	161.92
$\Delta U_{t+1}, \sigma, s$	115.26	-444.82	429.55
$s, \Delta U_{t+1}, \sigma$	121.46	-183.38	161.92
$s, \sigma, \Delta U_{t+1}$	213.19	-183.38	70.19
$\sigma, \Delta U_{t+1}, s$	205.80	-444.82	339.01
$\sigma, s, \Delta U_{t+1}$	213.19	-452.21	339.01

Notes: Percentage point contributions to the vertical shift in the Beveridge curve, averaged over the "downswing" sample points discussed earlier. "Ordering" column shows the order in which margins are set to their observed values. For example, the row $\Delta U_{t+1}, s, \sigma$ starts with the steady-state curve, then adds observed ΔU_{t+1}, then adds observed s_t, and finally adds the observed σ_t.

Table 2: Contributions to the Shift in the Beveridge Curve

(which cannot shift by construction) to the observed shift, by successively adding the observed time-variation in margins. Equation (22) first adds observed dynamics, then adds observed the separation probability, then adds observed matching efficiency. With three margins there are six possible orderings, and the results will, in general, depend on the ordering.

Table 2 shows the results of all six orderings. The results are remarkably consistent. In all versions, separations push the Beveridge curve down during the upswing period, relative to the downswing period. Both dynamics and matching efficiency have the opposite effect, contributing to the counter-clockwise loop in the observed Beveridge curve. Generally, the contribution of matching efficiency is larger than that of dynamics, sometimes dramatically so. The only outlier is the fourth row. However, we believe that the first two rows are the most important, because they put ΔU_{t+1} first in the ordering, which ensures dynamic consistency.

Nearly all of the contributions in Table 2 are well above 100 percent. This shows just how important all three margins are in understanding the shift of the Beveridge curve. The shift we observe empirically is relatively small, when compared to the effects of the shifters taken separately.