

# The Dual Beveridge Curve

Anton Cheremukhin\*

Paulina Restrepo-Echavarria†

February 14, 2024

## Abstract

The recent behavior of the Beveridge Curve significantly differs from past recessions and is hard to explain with traditional gradual changes in fundamentals. We propose a novel dual vacancy model where we acknowledge that not all vacancies are made equal—when firms post a vacancy they can hire from unemployment or they can poach a worker from another firm. Our dual vacancy model segments the labor market into separate search processes for unemployed and employed workers and provides a better fit to the data than traditional models assuming a homogeneous market. By analyzing labor market data from 2000 onwards, we estimate the proportions of the two types of vacancies and find a significant rise in poaching vacancies since the mid-2010s. The behavior of the share of poaching vacancies is explained by the residual hires to quits ratio and by an increasing trend in the profit-cost ratio of these positions. Once we adjust the Beveridge Curve to only include vacancies for the unemployed, the recent puzzling behavior disappears. These results imply that a slowdown in the demand for overall workers is likely to have a diminished effect on unemployment, affecting the implications of monetary policy for unemployment.

Keywords: Beveridge Curve, Vacancies, Unemployment.

JEL Codes: J23, J63, J64, E52

---

\*Federal Reserve Bank of Dallas, chertosha@gmail.com

†Federal Reserve Bank of St. Louis, paulinares@me.com

We would like to thank Serdar Birinci, Andreas Mueller, Gianluca Violante, and several participants of the seminar at the Dallas Fed and St. Louis Fed for their insightful comments. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Reserve Banks of Dallas and St. Louis or the Federal Reserve System.

# 1 Introduction

The negatively sloped relationship between the number of unemployed and the number of job openings over the business cycle in the U.S. is known as the Beveridge curve. Since its inception in the 1950s, this curve has been used by policymakers to assess the health of the labor market and measure the distance to full employment.

As shown in Figure 1, the recent behavior of the Beveridge curve is strikingly different from its behavior in previous recessionary episodes. Historically, we have observed movements along lines of similar slopes, with only gradual shifts of the intercept in between recessions. However, in the current episode, both the slope and the intercept seem to have shifted multiple times. This new behavior of the Beveridge curve is puzzling.

Economists have naturally turned to explanations that can account for the past gradual shifts of the curve, but don't seem to work for the current episode. In this paper, we propose a new and completely different explanation for the recent puzzling behavior of the Beveridge curve: a dual vacancy model.

We know that firms sometimes hire unemployed workers and sometimes poach workers from other firms. This choice depends on the various skill requirements for various types of jobs, including their location on the job ladder. Naturally, firms tailor their job postings as much as they can to attract the type of worker they are more interested in hiring; an unemployed worker or an already employed worker. Therefore, we think it is important to consider a model where there are two separate types of vacancies, those designed for unemployed workers and those designed to poach workers from existing positions at other firms, and two different matching functions.

Each type of vacancy has a different effect on the labor market. If a vacancy leads to a hire from the unemployment pool, it reduces the unemployment rate and increases the employment rate. In contrast, when a firm poaches an employee, a worker moves between two positions and potentially increases his or her wage in the process, but employment and unemployment are unaffected.

Our dual vacancy model takes the view that these two types of job postings operate in separate, segmented markets, breaking the overall search and matching process into two non-overlapping processes. In our model, unemployed workers search for and match only with vacancies intended for unemployed workers, while employed workers match only with the vacancies that are open for workers who are already employed. Since poaching vacancies (intended to be filled from the employed pool) do not affect employment and unemployment, in our model, the Beveridge curve relationship applies only to the first sub-market that matches unemployed workers with unemployment vacancies.

We use the dual vacancy model and available data on labor market stocks and flows to estimate the parameters as well as the numbers of both types of vacancies, for the U.S. economy, and its sub-sectors at a monthly frequency, starting from the year 2000. We find that there has been a large disproportionate increase in the number of poaching vacancies starting no later than mid-2010s,




Figure 1: Beveridge Curves over Business Cycles.

Source: BLS. Notes: Henderson moving averages of the unemployment and vacancy rates are shown.

and that their increase and cyclical conduct is largely explained by the behavior of residual hires (hires minus quits) and quits. Furthermore, if we adjust the Beveridge curve by considering only unemployment vacancies, the recent puzzling behavior of the Beveridge curve disappears. This result says that the rise in poaching vacancies since 2020 explains the unusual recent behavior of the Beveridge Curve.

Note that because of equal opportunity rights laws, firms can't explicitly state in a job posting whether they are intending to hire an unemployed worker or they rather hire an already employed worker. Hence, we can't provide direct evidence of the split of vacancies in the real world. However, we can evaluate the statistical significance of our model versus the standard random matching model that assumes a single homogeneous market with just one type of vacancy. When we run this statistical test, we find that the dual vacancy model matches the data substantially better than the standard random

matching model with a single homogeneous market. This result reflects the fact that the business cycle responsiveness of quits and residual hires to the vacancy rate differs substantially, making it hard to match both with a single matching function elasticity.

Furthermore, our assumption that markets are completely segmented is extreme, and we are aware of the fact that there is probably some mixing going on. Some unemployment vacancies can be filled by employed workers and some poaching vacancies can be filled by unemployed workers. In Section 6, we consider an extension of the model where we allow for matches across the two sub-markets. Our empirical estimates of the model suggest that this interaction is relatively small, and that our full segmentation assumption is not far from what is going on in reality.

But why have poaching vacancies increased so much, apart from what is explained by the behavior of residual hires to quits? To understand the possible underlying factors driving this result, we develop a simple model of the Dual Beveridge curve where vacancies for the unemployed are determined by the number of unemployed and their profit-cost ratio, and poaching vacancies are determined by the number of employed and their respective profit-cost ratio. We find that although there is a negative trend in unemployment, the profit-cost ratio of poaching vacancies has been trending upwards since 2010, explaining the expansion in the share of poaching vacancies. In Section 8 we discuss possible explanations for the increase in the profit-cost ratio of vacancies opened for employed workers.

Finally, it is important to mention that our results imply that because the share of poaching vacancies has been increasing, a slowdown in the demand for workers is likely to have a diminished effect on unemployment (as poaching vacancies will go down by more than unemployment vacancies). This has implications for monetary policy and its effects on unemployment. We discuss policy implications of our findings in more detail in Section 8.

We contribute to three strands of the literature. First, the extensive literature on the Beveridge curve, the business cycle relationship between the numbers of unemployed and vacancies was first noted by Beveridge (1944) and the Beveridge curve relationship first plotted by Dow and Dicks-Mireaux (1958). The interest in the curve has been summarized in surveys by Elsby, Michaels, and Ratner (2015) among others. The shifts in the Beveridge curve were analyzed both for the U.S. (see e.g. Ahn and Crane (2020), Diamond and Sahin (2014)) and for other developed countries (see Bonthuis et al. (2016) and Hobijn and Sahin (2012)).

The puzzling behavior that we have observed recently has led to both a lively discussion of its causes and a policy debate. Lubik (2021) has attributed the breakdown puzzle to a decline in matching efficiency due to sectoral shifts and a change in skill requirements. Rodgers and Kassens (2022) have attributed the flattening of the curve to changes in the cost of remaining unemployed and the larger-than-expected number of retirements. Another proposed explanation is that technological change has made it easier to search for a job but harder to convert a match into an offer. For the policy discussion, see Figura and Waller (2022a) and Blanchard et al. (2022).

Our paper contributes to this literature and discussion by enhancing our understanding of the medium-term behavior of the Beveridge curve, including the most recent episode. We propose a novel mechanism that drives the puzzling behavior and narrow down the list of possible fundamental explanations.

Second, we contribute to the understanding of the matching function. According to the survey by Petrongolo and Pissarides (2001), models used in the literature have traditionally incorporated job-to-job flows into the matching process by assuming a joint matching function that combines all workers searching for a job, both employed and unemployed, with the total number of vacancies. In this paper, we propose an alternative model with two separate processes for employed and unemployed workers and show that our approach fits the data much better than the traditional approach. We advance the measurement of the search effort of employed workers, which allows us to estimate the coefficients of both matching functions for the whole U.S. economy and its sub-sectors.

Third, we contribute to the emerging literature on segmented labor markets. Recent studies by Hall and Kudlyak (2020) and Ahn et al. (2022) have identified segments of the labor market that differ in behavior on the worker side. We analyze market segmentation on the firm side, proposing a split of job openings into those designed for different types of workers. An estimation of this split is new to the literature. Our findings also provide empirical evidence that can be used to design and better calibrate theoretical models where there are different types of vacancies used for unemployment-to-employment and job-to-job transitions. To our knowledge, the first (and only) study to introduce such a theoretical split of vacancies, in a calibrated directed search framework, was Menzio and Shi (2011).

## 2 A Simple Model

We propose a very simple model, where we assume that there are two separate matching functions in the labor market; one for unemployed workers ( $M_u$ ) and another for employed workers ( $M_e$ ). This means that there are two different types of vacancies: those that are open to hire from unemployment ( $V_u$ ), and those that are open to poach an already employed worker from another firm ( $V_e$ ).

More specifically, in the unemployed matching function, the unemployed ( $U$ ) search for unemployment vacancies ( $V_u$ ) and get hired according to a standard constant-returns-to-scale matching function:

$$M_u = B_u U^\alpha V_u^{1-\alpha}, \tag{1}$$

where  $M_u$  is the number of hires from the unemployment pool,  $\alpha \in [0, 1]$  is an elasticity, and  $B_u$  is a parameter characterizing the efficiency of the matching process.

On the other hand, we know that a subset of all employed workers ( $E_s$ ) engage in on-the-job search

and hence search for vacancies designed to poach them from their current positions ( $V_e$ ) and switch jobs according to a second matching function:

$$M_e = B_e E_s^\beta V_e^{1-\beta}, \quad (2)$$

where  $M_e$  is the number of workers who voluntarily quit their positions to join a new employer,  $\beta \in [0, 1]$  is an elasticity, and  $B_e$  is a parameter characterizing the efficiency of the matching process for already employed workers.

Finally, according to these assumptions, the two types of vacancies have to sum up to the total number of vacancies, such that,  $V_u + V_e = V$ .

The assumption of this simple model is that the two matching processes are completely separate, and unemployed workers never match with vacancies designed for the employed, and employed workers never get hired to positions designed for the unemployed. These admittedly extreme assumptions are very helpful for transparency in a simple model, and we shall fully relax them in Section 6.

### 3 Methodology

Our goal is to estimate the split of total vacancies ( $V$ ), into unemployment vacancies ( $V_u$ ) and poaching vacancies ( $V_e$ ), and estimate the matching efficiencies  $B_e$ ,  $B_u$ , plus the elasticities  $\alpha$ , and  $\beta$ . To do so, we use Equations (1) and (2), together with  $V_u + V_e = V$ , and observed data for  $M_e$ ,  $M_u$ ,  $U$ ,  $E_s$ , and  $V$ .

We approximate the number of hires from the employment pool ( $M_e$ ) by the number of quits in the Job Openings and Labor Turnover Survey (JOLTS) data, since the majority of voluntary separations are due to job switches. The number of hires from the unemployment pool ( $M_u$ ) is then equal to the difference between total hires and quits in the JOLTS data. Total vacancies ( $V$ ) are the total number of job openings from JOLTS.

The key question here is how to approximate the search effort of the unemployed and of the employed workers. Theoretically, a transition rate is calculated by dividing the total number of matches by the total number of searchers. However, this calculation may not be accurate if the total number of searchers is measured imprecisely or there are systematic factors that affect their search effort. In such cases, one can estimate this unobserved search effort by measuring the difference between the ratio of the number of matches to the number of searchers and the corresponding transition rate.

The search input of the unemployed,  $U$ , is traditionally approximated by the total number of unemployed persons as reported by the Bureau of Labor Statistics (BLS). This is consistent with both the BLS definition of an unemployed person as a person that actively searches for a job and survey

evidence that more than 99 percent of the unemployed spend time actively searching for a job.




Figure 2: Ratio of Hires to EE Transitions vs Adjusted Employment

One way to check this assumption is to use data on the unemployment-to-employment (UE) transition rate measured from the Current Population Survey (CPS). The ratio of the number of hires from the unemployment pool to the search input of the unemployed must equal the transition rate. This proportion is not true by construction, as the data on the number of hires, the number of unemployed and the transition rate are collected and computed separately. Therefore, the ratio of hires to the transition rate should give a measure of the search input by the unemployed. Indeed, this ratio matches the total number of unemployed very closely, which implies that the number of unemployed is a very good measure of their search effort.

We employ the same method to obtain a measure of search input of the employed. Fallick and Fleischman (2004) and then Moscarini and Postel-Vinay (2022) have used CPS data to measure an employment-to-employment (EE) transition rate. We obtain our measure of the search input of the employed by dividing the number of hires from the employed pool by the EE transition rate.

An alternative way to obtain this measure is to use observations from the Survey of Consumer Expectations (SCE). Using these data, Faberman et al. (2022) document that only a small fraction of the employed (22%) engage in active search, but those who do engage are much more efficient than the unemployed at finding new jobs. We subtract from the total number of employed workers a highly smoothed measure of trend employment scaled by a factor 0.78, representing the 78% of employed who do not engage in active search. With this method, we obtain a measure of search input of the

employed, which behaves very similarly to the ratio of hires to EE transitions, as shown in Figure 2. Although this method is somewhat less precise, we shall use this fact to study sectoral data for which EE transition rates are not available, as well as for modeling purposes.

We observe all of these data at a monthly frequency starting from December 2000. Thus, for a sector of the economy, or for the economy as a whole, we can measure the variables  $M_u, M_e, U, V, E_s$ . The remaining unknowns to be estimated are the split of vacancies into unemployment vacancies ( $V_u$ ) and poaching vacancies ( $V_e$ ), and the parameters  $B_u, B_e, \alpha$ , and  $\beta$ . For the estimation we assume random white-noise measurement errors on each of the matching functions (Equations (1) and (2)).

For the economy as a whole, and for its sectors, we estimate the parameters of the model jointly using Bayesian methods. We compute the likelihood of the data given the parameters and multiply it by a relatively uninformative prior for the parameters. We use Bayesian estimation because in some cases, some parameters may not be fully identified. This means that using pure likelihood maximization may result in multiple local maximums and relatively flat areas connecting them. As a result, finding a unique maximum can be challenging. By using a Bayesian approach we introduce additional curvature to the parameter space by multiplying the likelihood by a relatively flat prior. This approach helps explore the parameter space, improves convergence, and provides a diagnostic method to detect cases where the parameters are not well-identified. In our results, we can easily detect such cases by comparing the shapes of the prior and posterior distributions, which should be similar. We evaluate the posterior distribution using a Random Walk Metropolis (RWM) algorithm as described in An and Schorfheide (2007). We use multiple chains all starting from the posterior mode that amount to a total of 100,000 posterior draws and make sure that the acceptance rates remain in the range from 0.2 to 0.5.

We do not estimate the standard deviations of measurement errors together with other parameters. This is because doing so would make the likelihood function flat, regardless of the other parameter values. Instead, we compute the likelihood of the data conditional on the parameters and their posterior distributions while keeping the standard deviations of the shocks fixed at their sample means. The likelihood function weighs both equations' errors equally. Hence, it aims to make the standard errors of the two equations the same and, with  $T+4$  degrees of freedom for  $2 \cdot T$  observations, is able to achieve that goal. Therefore, through maximizing the likelihood, we arrive at an estimate of the vacancy split that implies equal standard deviations of errors in the two equations, as can be seen in Table 3.

## 4 Results

In this section we report the parameter values that we recover and the estimated split of total vacancies into unemployment vacancies and poaching vacancies. Our estimated parameters for the economy as






Figure 3: Total Vacancies: Poaching vs Unemployed

a whole and for its sectors are shown in Table 1. For the whole economy, we estimate  $\alpha = 0.2$  and  $\beta = 0.9$ , and the level shifters  $B_u$  and  $B_e$  simply reflect proper scale. We also estimate these parameters for sectors that combine 1) manufacturing and construction, 2) business services and retail trade and 3) education, health and leisure services. All the parameter estimates are reported in Table 1. Using the estimated parameters we are also able to split job openings for the U.S. economy, and for each subsector, into those designed for the unemployed and those designed for poaching. This breakdown, calculated for the period 2001 to 2022, at a monthly frequency, is shown in Figure 3.

There are two important observations one can make from Figure 3. First, the fraction of poaching vacancies has increased significantly since approximately 2015, compared with the preceding period. This suggests that the reason the Beveridge curve has shifted upward is due to the dramatic increase in poaching vacancies (because these don't contribute to reducing unemployment). Second, while the business cycle behavior of the two types of vacancies was similar in the period prior to 2015, both dropped during recessions and recovered during booms, it was dramatically different in the most recent recession episode. Although poaching vacancies fell in 2020, but quickly recovered soon after, the vacancies designed for the unemployed increased in the recession period.

We find that the increase in poaching vacancies over time and their cyclical behavior is identified largely by the behavior of the ratio of residual hires (hires minus quits) to quits. This is not surprising given that conditional on vacancy filling rates being equal across sectors, the ratio of poaching vacancies

Table 1: Parameter estimates of the dual vacancy model

Parameter	Prior			Posterior		
	mean	st.dev.	mode	mean	st. dev.	conf. int. [5-95]
Total private industries						
$\alpha$	0.5	0.2	0.22	0.20	0.04	[0.14, 0.28]
$\beta$	0.5	0.2	0.90	0.89	0.004	[0.89, 0.90]
$B_u$	0.2	0.1	0.84	0.75	0.16	[0.49, 1.01]
$B_e$	0.2	0.1	0.11	0.15	0.11	[0.03, 0.37]
Construction and Manufacturing						
$\alpha$	0.5	0.2	0.65	0.66	0.08	[0.48, 0.75]
$\beta$	0.5	0.2	0.74	0.71	0.09	[0.52, 0.80]
$B_u$	0.2	0.1	0.54	0.55	0.08	[0.39, 0.66]
$B_e$	0.2	0.1	0.14	0.15	0.05	[0.06, 0.20]
Business services and retail trade						
$\alpha$	0.5	0.2	0.29	0.30	0.07	[0.21, 0.43]
$\beta$	0.5	0.2	0.82	0.80	0.05	[0.72, 0.86]
$B_u$	0.2	0.1	0.80	0.77	0.15	[0.56, 1.05]
$B_e$	0.2	0.1	0.20	0.26	0.17	[0.10, 0.61]
Other services						
$\alpha$	0.5	0.2	0.60	0.55	0.10	[0.41, 0.71]
$\beta$	0.5	0.2	0.53	0.44	0.12	[0.26, 0.63]
$B_u$	0.2	0.1	0.46	0.41	0.12	[0.20, 0.61]
$B_e$	0.2	0.1	0.27	0.23	0.06	[0.13, 0.34]

*Notes:* The priors for  $\alpha$  and  $\beta$  were drawn from a beta distribution with support on the interval [0.1, 0.9], and priors for  $B_u$  and  $B_e$  were drawn from a gamma distribution with positive support.

to unemployment vacancies will equal the ratio of residual hires to quits. Furthermore, in Section 7 we show that the remaining increase in poaching vacancies that is not identified by the ratio of residual hires to quits, is explained by a trend in the ratio of vacancy filling rates.

To understand our results, we need to look at them through the lens of an adjusted Beveridge curve. Recall that only the vacancies designed for the unemployed match with unemployed workers and lead to increases in employment. Thus, the proper Beveridge curve relationship should only consider vacancies for the unemployed and disregard poaching vacancies. The adjusted Beveridge curves for the whole economy and for 3 broad sectors are shown in the bottom row of Figure 4 compared with the un-adjusted (or classical) Beveridge curves in the top row.

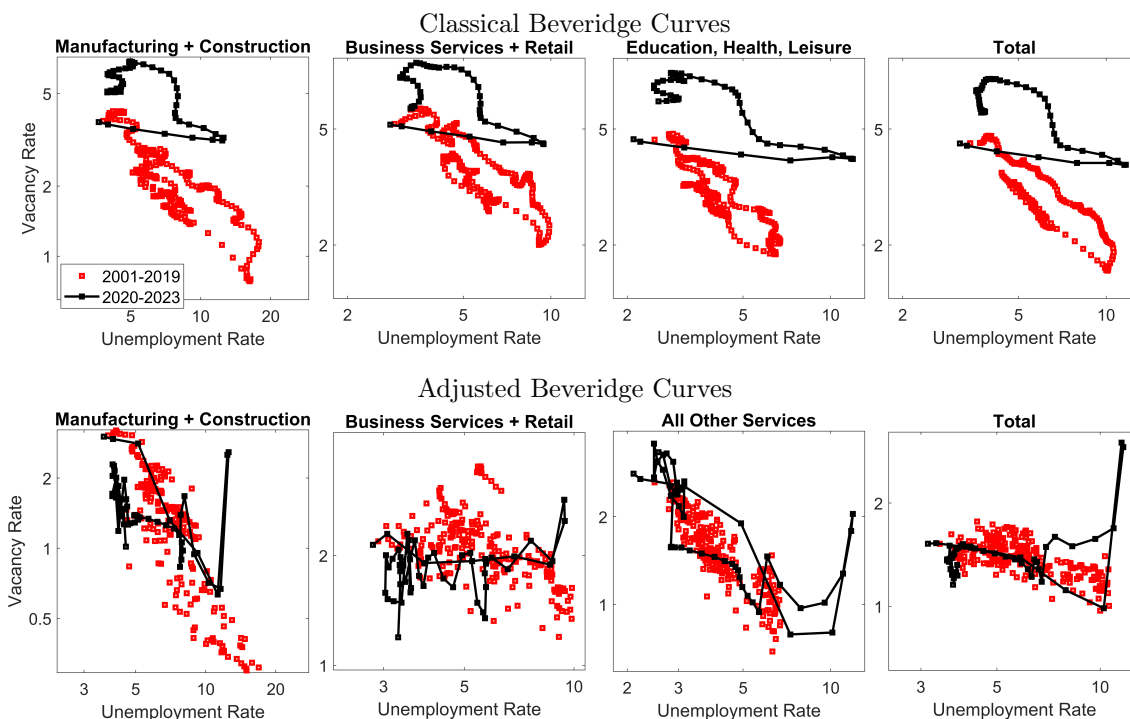


Figure 4: Classical and Adjusted Beveridge Curves

Source: BLS, authors' calculations. Notes: Henderson moving averages of the unemployment and vacancy rates are shown in the top four panels. Moving averages were not used in the estimation for the bottom four panels.

Figures (3) and (4) are illustrative of what happened in labor markets since the onset of the Covid pandemic. The first few months of the pandemic saw a decline in demand due to widespread social distancing, which increased unemployment and reduced poaching. In the next few months, mask and distancing mandates led to a separation shock where many more people were laid off than would be consistent with lower demand, so vacancies designed for the unemployed increased, and a lot of people were hired back from unemployment very quickly. After the spike in hires from unemployment ended, stimulative fiscal and monetary policy increased purchasing power and created strong excess demand for goods. The excess demand prompted firms to expand, but this excess demand for workers could not be met by hiring from the unemployment pool. Together with supply chain bottlenecks, the excess demand for goods led to a surge in inflation, and excess demand for workers led to an increase in poaching which then drove up nominal wage growth.

This interpretation provides us with two lessons. First, the (adjusted) Beveridge curve relationship between unemployed workers and unemployment vacancies has not changed, at either the aggregate or sectoral level. In other words, the current puzzling behavior of the Beveridge Curve disappears once we replace total vacancies with unemployment vacancies. Second, abnormalities in the classical Beveridge curve are due to a disproportional expansion of poaching vacancies after 2015. We explore

the underlying cause for this shift in Section 7.

## 5 Model Fit and Comparison to Standard Model

Since we do not have direct evidence on the split of vacancies into those designed for poaching and those designed for the unemployed, and it is not a given that such a clear split exists, it would be helpful to understand whether our new dual-vacancy model provides a better description of the data than existing models. In order to answer this question, we adopt the traditional model with a single matching function to fit our observables and estimate its parameters.

According to the standard model, a single constant-returns-to-scale matching function combines the total number of job seekers  $U + E_s$  with the total number of vacancies  $V$  to produce the total number of matches  $M_u + M_e$ . In order to give the model the chance of matching the data, we add extra flexibility to this overly restrictive model. We allow the proportion of total matches going to the unemployed to differ from their share of the search effort and estimate an additional parameter responsible for this split. Thus, our version of the traditional model consists of two equations:

$$M_u = B_u U \left( \frac{V}{U + E_s} \right)^{1-\theta}, \quad (3)$$

$$M_e = B_e E_s \left( \frac{V}{U + E_s} \right)^{1-\theta}. \quad (4)$$

We estimate this model using the same methods as the dual vacancy model. This allows us to compare fit because both models approximate the same set of data, with a different number of parameters. In particular, the traditional model has only one elasticity of the matching function,  $\theta$ , and combines vacancies into a single series, while the dual-vacancy model has two elasticities of the matching function,  $\alpha$  and  $\beta$ , and recovers a hidden variable, the split of the vacancies.

The parameter estimates for the traditional model are presented in Table 2. The estimates of the matching elasticity tend to hit the upper bound of the support range both for the whole economy and for its subsectors, while for the dual vacancy model it is common to have interior estimates of both elasticities (see Table 1).

In Table 3 we present measures of model fit. The dual-vacancy model fits the data uniformly better based on smaller estimated standard errors for each of the two equations of each model. This is because the business cycle responsiveness of quits and residual hires to the vacancy rate differs substantially, making it hard to match both with a single matching function elasticity. The dual vacancy model does a much better job at fitting both rates because it has two elasticity parameters rather than one, and also because it has the ability to split vacancies into two subsets - one for each matching rate. The superior fit of the dual vacancy model is also supported by the marginal data

Table 2: Parameter estimates of the model with a single matching function

Parameter	Prior			Posterior		
	mean	st.dev.	mode	mean	st. dev.	conf. int. [5-95]
Total private industries						
$\theta$	0.5	0.2	0.90	0.89	0.003	[0.89, 0.90]
$B_u$	0.2	0.1	0.22	0.22	0.08	[0.12, 0.37]
$B_e$	0.2	0.1	0.07	0.09	0.12	[0.003, 0.38]
Construction and Manufacturing						
$\theta$	0.5	0.2	0.90	0.89	0.004	[0.89, 0.90]
$B_u$	0.2	0.1	0.20	0.22	0.09	[0.08, 0.40]
$B_e$	0.2	0.1	0.06	0.07	0.11	[0.002, 0.35]
Business services and retail trade						
$\theta$	0.5	0.2	0.90	0.90	0.004	[0.89, 0.90]
$B_u$	0.2	0.1	0.33	0.31	0.09	[0.15, 0.46]
$B_e$	0.2	0.1	0.10	0.08	0.11	[0.003, 0.36]
Other services						
$\theta$	0.5	0.2	0.90	0.90	0.004	[0.89, 0.90]
$B_u$	0.2	0.1	0.25	0.21	0.11	[0.04, 0.41]
$B_e$	0.2	0.1	0.07	0.09	0.11	[0.003, 0.37]

*Notes:* The priors for  $\alpha$  were drawn from a beta distribution with support on the interval [0.1, 0.9], and priors for  $B_u$  and  $B_e$  were drawn from a gamma distribution with positive support.

densities, which we report further in Table 3. In all four cases, Bayes factors strongly favor the dual vacancy model.

## 6 Extensions and Robustness

In this section, we consider extensions that may affect the estimates of the matching functions and the vacancy split and substantially generalize our results. First, we use a simplified version of a targeted search model (see Cheremukhin, Restrepo-Echavarria, and Tutino (2020)) as an inspiration to generalize our matching function specifications to the case where both types of workers sometimes confuse the two vacancy types and therefore apply to the wrong type of vacancy, and also both types of jobs sometimes accept workers for which they were not originally designed. This confusion creates additional matches between the wrong types, and their numbers should have the following forms:  $M_{u+} = A_u U^\alpha V_e^{1-\alpha}$ , and  $M_{e+} = A_e E_s^\beta V_u^{1-\beta}$ . This would affect our measurement equations by adding

Table 3: Comparison of model fit

Sector	Standard Errors		Marginal Data Density		Bayes factor
	DV	SMF	DV	SMF	
Total private industries	0.06, 0.06	0.32, 0.07	-484.2	-545.2	exp(61.0)
Construction and manufacturing	0.09, 0.09	0.40, 0.18	-508.6	-563.8	exp(55.2)
Business services and retail trade	0.08, 0.08	0.32, 0.16	-500.3	-543.4	exp(43.1)
Other services	0.12, 0.12	0.33, 0.25	-517.9	-547.3	exp(29.4)

*Notes:* DVM stands for dual vacancy model, and SMF stands for single matching function model. Standard errors report two numbers - representing standard errors on each of the two equations of each model. The marginal data density was computed using the method of Chib and Jeliazkov (2001). Using Geweke's (1999) modified harmonic mean leads to similar results.

these matches to the observations. Relaxing some of the restrictions on parameters we postulate the following general functional forms:

$$M_u = B_u U^\alpha V_u^{1-\alpha} \left[ 1 + C_u \left( \frac{V_e}{V_u} \right)^\gamma \right], \quad (5)$$

$$M_e = B_e E_s^\beta V_e^{1-\beta} \left[ 1 + C_e \left( \frac{V_u}{V_e} \right)^\gamma \right], \quad (6)$$

where the mixing coefficients  $C_u$  and  $C_e$  are the amounts of unemployed workers that are able to get a job with a firm that intended to poach and of employed workers that take up jobs intended for the unemployed.

Second, we allow for loglinear time trends in unobservable matching efficiencies of the form  $B_u e^{\delta_u t}$  and  $B_e e^{\delta_e t}$ .<sup>1</sup> Using the same maximum likelihood maximization method as in the previous section, our estimates of parameters for total private industries are:  $\alpha = 0.2$ ,  $\beta = 0.704$ ,  $\gamma = 0.51$ ,  $C_u = 0.035$ ,  $C_e = 0.134$ ,  $B_u = 0.92$ ,  $B_e = 0.184$ ,  $\delta_u = -0.000504$ ,  $\delta_e = -0.001128$ . The extra parameters reduced the standard errors by a factor of 3 compared with Table 3. Another important difference is that  $\beta = 0.7$  is now not close to the upper bound of its range, which indicates that its parameter value is much better identified under this setup. Both of these changes are mainly due to allowing for time trends in matching efficiencies. Nevertheless, the vacancy split is nearly identical to that identified with the more restrictive model of Section 2.

The estimated number of "wrong" matches that occur, that is the number of matches between unemployed workers and poaching vacancies, and employed workers and unemployment vacancies is

<sup>1</sup>We added this assumption because the estimated measurement errors of the model in Section 2 are suggestive of a trend of this form.

relatively small, accounting on average for 4 percent and 12 percent of total matches respectively. The matching elasticity  $\gamma$  that characterizes the mixing up of workers and jobs is in between the estimated matching elasticities  $\alpha$  and  $\beta$  for the two main matching processes, as would be expected from theory. These results show that the estimates from the restricted model hold more generally, where we don't make the extreme assumption that the matching process for unemployed workers and the matching process for employed workers are fully independent.

As another robustness check, we evaluate the possibility that our estimates are affected by flows in and out of the labor force. The transition rates between employment (E), unemployment (U), and out of the labor force (N) are reported in the Current Population Survey at a monthly rate since 1990, covering our period of interest. The transition rates between unemployment and out of the labor force (UN and NU) did not respond at all to the pandemic, so they are unlikely to affect our estimates. The transition rates between employment and out of the labor force (NE and NE) did respond in a way similar to transition rates between employment and unemployment, explaining about one fifth of the decline in employment in 2020 and its subsequent recovery. This suggests that we should consider the possibility that workers out of the labor force are applying to (both types of) vacant positions, and filling some of those positions. Given that this is unlikely to affect quits directly, these matches will show up in residual hires. This means that we should add an extra term of the form  $M_{u++} = B_L N_s^\psi V^{1-\psi}$  to equation (5). As this new term is correlated with the two terms already present in the equation we need to calibrate  $\psi$ .

We calibrate it based on a regression of the log ratio of the number of NE transitions to the labor force, independently measured in the CPS, on the log ratio of total vacancies to the labor force. This regression has good fit and yields the value  $\psi = 0.75$ . We also assume, for lack of a better option, that the number of searchers not in the labor force,  $N_s$ , are a constant proportion of the labor force itself. Conditional on this calibration, we re-estimate the model and find that the contribution of NE transitions to equation (5) is quite significant, explaining more than one third of the flows. However, this re-estimation does not alter in a meaningful way any of the other estimated parameters, nor the estimated series for the vacancy split. Its effect is similar to a change in the intercept and a slight improvement in fit. Nevertheless, we think that these estimates of the properties of matching from out of the labor force could be of independent interest for calibrating theoretical models going forward.

## 7 Dual Beveridge Curve Model

As was shown in Section 4, there has been a disproportionate increase in the number of poaching vacancies at least since 2015. But what is the underlying cause of this increase? In this section, we write down a simple theoretical model of the dual Beveridge curve and use it to analyze the potential driving factors behind the expansion in poaching vacancies and the shift in the Beveridge curve.

We take the model described in Section 2 and add three equations. First, a stationary labor market implies that entry into and exit from the pool of unemployed must be balanced. More specifically, the number of workers entering unemployment must equal the number of new matches that the unemployed form. Second, there are two free entry conditions, one for each type of vacancy:

$$M_u = (LF - U) s, \quad (7)$$

$$1 = \frac{M_u}{V_u} y, \quad (8)$$

$$1 = \frac{M_e}{V_e} z, \quad (9)$$

where  $LF$  represents the labor force,  $s$  is the separation rate,  $y$  is the profit-cost ratio for vacancies designed for the unemployed, and  $z$  is the profit-cost ratio for vacancies designed for poaching.

Equations (1), (2), and (7)-(9), together with the fact that the search effort of employed workers can be expressed as a function of the number of unemployed,  $E_s = E - 0.78E^*$ , describe the whole system of equations, with  $\{M_u, M_e, V_u, V_e, U, \}$  being endogenous unknowns, and  $\{s, y, z\}$  being the exogenously given driving forces.<sup>2</sup>

We further simplify the model and notation by detrending by the labor force, for each variable  $X$  defining a lower case detrended analog  $x = X/LF$ , and by substituting the matching functions and the search effort of the employed to get:

$$(1 - u) s = B_u u \left( \frac{v_u}{u} \right)^{1-\alpha}, \quad (10)$$

$$\left( \frac{v_u}{u} \right)^\alpha = B_u y, \quad (11)$$

$$\left( \frac{v_e}{0.27 - u} \right)^\beta = B_e z. \quad (12)$$

where we omitted the mixing matching terms for simplicity. Although these three equations have three endogenous variables  $u, v_u, v_e$ , the first two equations could be solved separately with respect to  $u$  and  $v_u$  — who's relationship determines the adjusted Beveridge Curve. Poaching vacancies are then determined by the third equation, driven by fluctuations in their profitability  $z$  and the unemployment rate. The solution to the model then looks as follows:

$$u^* = \frac{1}{1 + \frac{B_u}{s} (B_u y)^{\frac{1}{\alpha} - 1}}, \quad (13)$$

$$v_u^* = u^* (B_u y)^{\frac{1}{\alpha}}, \quad (14)$$

$$v_e^* = (0.27 - u^*) (B_e z)^{\frac{1}{\beta}}. \quad (15)$$

---

<sup>2</sup>Note that based on our approximation in Section 3, search effort of the employed can be expressed as a function of the number of unemployed  $E_s = E - 0.78E^* = LF - U - 0.78(LF - U^*) = 0.22LF + 0.78U^* - U \approx 0.27LF - U$ .



Log-linearizing the model with respect to  $u^*$  and  $v_u^*$ , we can find that the slope of the adjusted Beveridge curve is  $-\frac{\alpha+u^*}{1-\alpha}$ . We further denote the "steady-state" share of poaching vacancies by  $x = \frac{v_e}{v_u+v_e}$ . To find the slope of the classical Beveridge curve, we need to understand the relationship between movements in  $y$  and  $z$  over the business cycle. In standard search models, movements in profitability of a match  $y$  reflect changes in productivity or demand driving the business cycle. It is natural to expect the profitability of poaching vacancies to be driven by similar factors. Therefore, we would expect  $y$  and  $z$  to have a common factor reflecting business-cycle fluctuations. We denote the elasticity of the co-movement between profitabilities by  $\phi$ , reflecting the ratio of their log standard deviations:  $\ln\left(\frac{z}{z^*}\right) \propto \phi \ln\left(\frac{y}{y^*}\right)$ . Then we can show that the slope of the classical Beveridge curve is  $-(1-x)\frac{\alpha+u^*}{1-\alpha} - x\left(\frac{u^*}{0.27-u^*} + \frac{\alpha}{1-\alpha}\frac{\phi}{\beta}\frac{1}{1-u^*}\right)$ . The first term reflects movement in vacancies designed for the unemployed in the adjusted Beveridge curve. The second term reflects the movements in search effort of the employed and the movements in the profitability of poaching vacancies over the business cycle.

To put some numbers to these slopes, we use the estimated parameters  $\alpha = 0.2$ ,  $\beta = 0.7$ , the estimated steady-state share of poaching vacancies  $x = 0.5$ ,<sup>3</sup> and the steady-state level of unemployment  $u^* = 0.06$ . For this calibration, the slope of the adjusted Beveridge curve is -0.33, consistent with Figure 4 for total private industries. To get the slope of the classical Beveridge curve to -1 to be consistent with Figure 4 for total private industries prior to 2020, we calibrate  $\phi = 3.5$ , thus assuming that the profitability of poaching vacancies is more sensitive to the business cycle than that of vacancies typically opened for the unemployed. If the steady-state level of poaching vacancies were to increase to 0.8, as we might have seen recently, then the Beveridge curve could have steepened to a slope of -1.4.

Now that we understand the slope of the Beveridge curve and how the model operates, we can analyze the behavior of the driving forces, if they have trends, if trends have changed over time and how the driving forces co-move at business cycle frequencies.

We use available data series to plot  $s = m_u/(1-u)$ ,  $y = (v_u/u)^\alpha$ , and  $z = (v_e/e_s)^\beta$  in the first three panels of Figure 5. The second panel shows that there is a downward trend in the separation rate,<sup>4</sup> consistent with the literature documenting a secular decline in labor market dynamism in the US, e.g. Molloy et al (2016). The downward trend in separations accounts for the decline in the trend of the unemployment rate over the past 25 years consistent with a downward trend in most existing measures of the natural rate of unemployment, see e.g. Crump et al (2019).

Substituting the separation rate into the model solution for the unemployment rate (13), while keeping other parameters constant, is enough to account for the secular downward movement in

<sup>3</sup>To take a conservative approach, we use the estimated vacancy split that we observe prior to 2010.

<sup>4</sup>It is important to highlight that the trend in our measured  $s$  is coming from  $m_u$  and not  $u$ . This finding can also be confirmed directly from empirical measures of separation rates.




Figure 5: Trends and Co-movement of Driving Forces

unemployment, and therefore, it is no surprise that we see no trend in the profit-cost ratio for vacancies designed for the unemployed, as shown in the first panel of Figure 5. Consequently, consistent with the model, vacancies designed for the unemployed have only a mild downward trend, similar to the unemployment rate, jointly implying a slow downward shift in the adjusted Beveridge curve over time.

In contrast, the profit-cost ratio for poaching vacancies demonstrates a clear break in trend. It is shown in the third panel of Figure 5, estimated under the assumption that the break occurred around 2009. The profit-cost ratio for poaching vacancies remained stable prior to 2009, but then started expanding on an upward trend. It is unclear if this expansion is over or will continue. Substituting this trend estimate into the equation for poaching vacancies, accounts for most of the expansion in poaching vacancies. Thus, the driving force for the recent anomalous behaviour of the Beveridge curve is the expansion of the profit-cost margin for poaching vacancies, while the adjusted Beveridge curve remained largely intact.

The fourth panel of Figure 5 shows the behavior of log deviations of  $y$  and  $z$  from their estimated trends (plotted in the first and third panel of Figure 5). As conjectured in our derivation of the slope of the Beveridge curve, these fluctuations of the driving forces are strongly correlated, jointly accounting for business cycle fluctuations in the model. We estimate the ratio of standard deviations of their first principal component to be 2.7, not far from our theoretical calibration of 3.5.

Figure 6 compares the joint behaviour of unemployment and vacancies with the predictions of our calibrated theoretical model for the dual Beveridge curve. Instead of parameter values for  $B_u$ ,  $B_e$ ,  $s$ ,  $z$ , and  $y$  we input their estimated linear trend values for 2007 and 2019 - two pre-recession

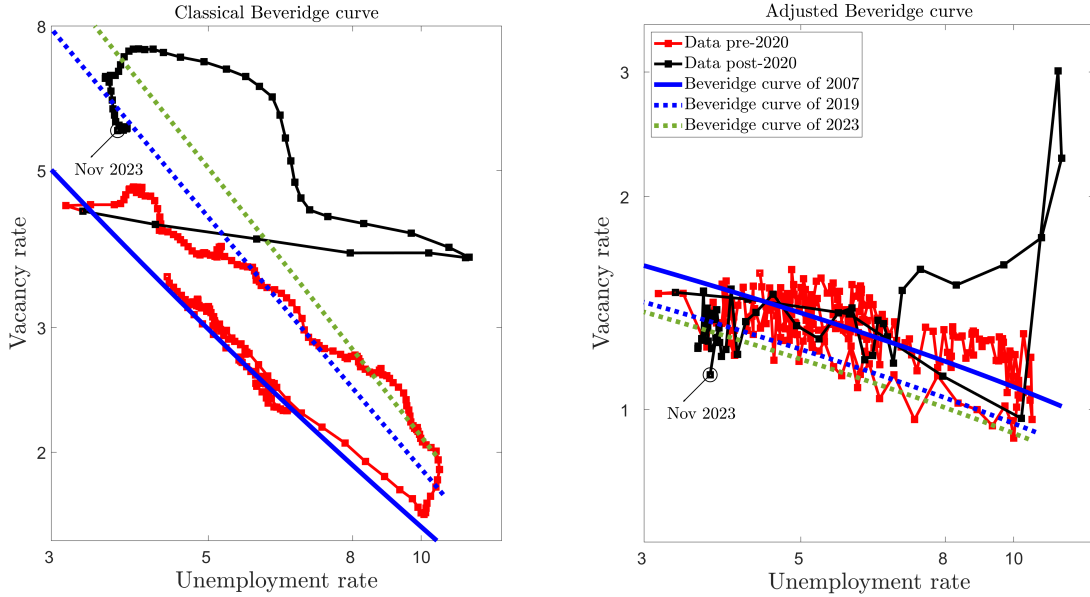


Figure 6: Theoretical Dual Beveridge Curves

peaks commonly used as reference points, and 2023 - the latest observation at hand. The adjusted Beveridge curve in the right panel shifted down only mildly due to the reduced labor market dynamism, as captured by the decline in the trend separation rate. The classical Beveridge curve in the left panel both shifted outwards and steepened its slope, due to the increase in steady-state profit-cost ratio  $z$  and the consequent expansion in the steady-state level of poaching vacancies. It expanded and steepened further for the estimated trend values of 2023, but we think it premature to project an indefinitely growing trend, and thus the estimate for 2019 represents a conservative estimate.

## 8 Final Remarks and Policy Implications

Our results are important for policy considerations, in particular, for monetary policy's effect on unemployment. As argued by Figura and Waller (2022), a steeper Beveridge curve could imply that tighter monetary policy would result in a larger decline in vacancies corresponding to only a mild increase in the unemployment rate.

In this paper, we attribute the Beveridge curve puzzle to the disproportional expansion of poaching vacancies. Our estimates combined with a theoretical model indicate that the slope of the Beveridge curve has indeed steepened from  $-1$  to at least  $-1.25$  and possibly  $-1.4$ . This coefficient implies that a decline in the vacancy rate from  $7\%$  to  $5\%$  should correspond to an increase in the unemployment rate from  $3.5$  to at most  $4.6$  percent, and possibly  $4.4$  percent, as opposed to  $4.9$  percent previously. Another consequence of the expansion of poaching vacancies is the outward movement of the Beveridge

curve which suggests that a coexistence of a 6% vacancy rate (rather than 4% vacancy rate) with a 4% unemployment rate may be the new normal. Consequently, a monetary tightening in the 2020s is likely to lead to a larger decline in job openings corresponding to a milder increase in the unemployment rate, consistent with a notion of a "soft landing."

The future is uncertain, however. The interpretation of the most recent behavior of the Beveridge curve depends on the reason for the expansion in poaching vacancies, and whether it is likely to continue. Among the possible explanations are both factors that reduced the costs associated with filling vacancies and factors that increased their benefits to firms. The first set of factors includes the effects of the expansion of online job search tools and increased use of AI (Acemoglu et al, 2022), the expansion of available temporary and remote work (Bloom et al, 2023), the expansion of the online gig economy (Stanton and Thomas, 2021). The second set of factors could include rising market concentration and markups (Autor et al, 2020, De Loecker et al, 2020) and the associated expansion of monopsony power of firms (Azar et al, 2019, Berger et al, 2022). If some of these factors are at play, the expansion of poaching vacancies could continue for as long as these trends continue. Therefore more changes in monetary policy could be absorbed by poaching vacancies, with little impact on vacancies designed for the unemployed and only a small increase in unemployment.

Alternatively, the expansion of poaching vacancies could be due to a reduction in mis-measurement: according to Davis et al. (2013), as of 2011, 42% of hires occurred at establishments that did not have any job openings. If those firms have gradually improved their reporting of vacancies that had not been reported previously, then the aggregate Beveridge curve has shifted outwards, but there are limits to such an expansion. In this case, the Beveridge curve will stabilize at a new level and slope.

The main lesson from our exercise, however, is that instead of looking at the classical Beveridge curve and interpreting its increasingly chaotic movements, we should shift our attention to the adjusted Beveridge curve, which is unlikely to change much, and will therefore remain a good indicator of the state of the labor market going forward.

Another important takeaway point is that economists and statistical agencies need to put resources into more and better measurement of the vacancy split, between vacancies designed for the unemployed and vacancies designed for poaching workers that already have a job. Surveys of firms conducted by statistical agencies could ask the firms a question that would shed light on this issue and enable direct measurement of the vacancy split. Such measurement would both enable the development of better theoretical models, and a better real-time assessment of the state of the labor market.

## 9 References

1. Acemoglu, Daron, Autor, David, Hazell, Jonathon, and Pascual Restrepo (2022) "Artificial Intelligence and Jobs: Evidence from Online Vacancies," *Journal of Labor Economics*, 40 (S1),

- pp. S293–S340. DOI: <https://doi.org/10.1086/718327>
2. Ahn, Hie Joo and Leland D. Crane (2020) "Dynamic Beveridge Curve Accounting," BOG DP No. 2020-027, DOI: <https://doi.org/10.17016/FEDS.2020.027>
  3. Ahn, Hie Joo, Hobijn, Bart and Ayşegül Şahin (2022) "The Dual U.S. Labor Market Uncovered," mimeo, DOI: <https://doi.org/10.20955/wp.2022.021>
  4. An, Sunbae and Frank Schorfheide (2007) "Bayesian analysis of DSGE models," *Econometric Reviews*, 26 (2-4), 113-172. DOI: <https://doi.org/10.1080/07474930701220071>
  5. Autor, David, Dorn, David, Katz, Lawrence F., Patterson, Christina, and John Van Reenen (2020) "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics*, 135 (2), pp. 645–709. DOI: <https://doi.org/10.1093/qje/qjaa004>
  6. Azar, José, Marinescu, Ioana, and Marshall Steinbaum (2019) "Measuring Labor Market Power Two Ways" *AEA Papers and Proceedings*, 109, pp. 317–321. DOI: <https://doi.org/10.1257/pandp.20191068>
  7. Berger, David W., Herkenhoff, Kyle F. and Simon Mongey (2022) "Labor Market Power" *American Economic Review* 2022, 112(4), pp. 1147–1193. DOI: <https://doi.org/10.1257/aer.20191521>
  8. Beveridge, William H. (1944) *Full Employment in a Free Society*. New York: W. W. Norton & Company, DOI: <https://doi.org/10.4324/9781315737348>
  9. Blanchard, Olivier, Domash, Alex and Lawrence H. Summers (2022) "The Fed is wrong: Lower inflation is unlikely without raising unemployment," *PIIE Realtime Economic Issues Watch*, August 2022, DOI: <http://dx.doi.org/10.2139/ssrn.4174601>
  10. Bloom, Nicholas, Davis, Steven J., Hansen, Stephen, Lambert, Peter John, Sadun, Raffaella, and Bledi Taska (2023) "Remote Work across Jobs, Companies, and Space" *NBER Working Paper No. 31007*. DOI: <http://doi.org/10.3386/w31007>
  11. Bonthuis, Boele, Jarvis, Valerie and Juuso Vanhala (2016) "Shifts in Euro Area Beveridge Curves and Their Determinants," *IZA Journal of Labor Policy*, 5 (20), 1-17, DOI: <https://doi.org/10.1186/s40173-016-0076-7>
  12. Cheremukhin, A., Restrepo-Echavarria P. and A. Tutino (2020) "'Targeted Search in Matching Markets." *Journal of Economic Theory*, vol 185, pp. 1-43. DOI: <https://doi.org/10.1016/j.jet.2019.104956>
  13. Chib, S., Jeliazkov, I. (2001). "Marginal likelihood from the metropolis-hastings output," *Journal of the American Statistical Association*, 96 (453), 270–281. DOI: <https://doi.org/10.1198/016214501750332848>

14. Crump, Richard K., Stefano Eusepi, Marc Giannoni, and Aysegul Sahin (2019): "A Unified Approach to Measuring  $u^*$ ," Brookings Papers on Economic Activity. DOI: <https://doi.org/10.3386/w25930>
15. Davis, Steven J., Faberman, R. Jason and John C. Haltiwanger (2013) "The Establishment-Level Behavior of Vacancies and Hiring," *The Quarterly Journal of Economics*, 128(2), 581-622, DOI: <https://doi.org/10.1093/qje/qjt002>
16. De Loecker, Jan, Eeckhout, Jan, and Gabriel Unger (2020) "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics*, 135 (2), pp. 561-644. DOI: <https://doi.org/10.1093/qje/qjz041>
17. Diamond, Peter A. and Ayşegül Şahin (2014) "Shifts in the Beveridge curve," Staff Reports 687, FRB NY, DOI: <http://dx.doi.org/10.2139/ssrn.2488141>
18. Dow, J.C.R. and L.A. Dicks-Mireaux (1958) "The Excess Demand for Labour. A Study of Conditions in Great Britain, 1946-56." *Oxford Economic Papers*, 10 (1), 1-33, DOI: <https://doi.org/10.1093/oxfordjournals.oep.a040791>
19. Elsby, Michael W. L., Michaels, Ryan and David Ratner (2015) "The Beveridge Curve: A Survey," *Journal of Economic Literature*, 53(3), 571-630, DOI: <https://doi.org/10.1257/jel.53.3.571>
20. Faberman, R. Jason, Mueller, Andreas I., Şahin, Ayşegül and Giorgio Topa (2022) "Job Search Behavior Among the Employed and Non-Employed," *Econometrica*, 90: 1743-1779, DOI: <https://doi.org/10.3982/ECTA18582>
21. Fallick, B., and C. A. Fleischman (2004) "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," BOG DP No. 2004-34.
22. Figura, Andrew, and Chris Waller (2022) "What does the Beveridge curve tell us about the likelihood of a soft landing?" BOG Fed Notes, July 2022, DOI: <https://doi.org/10.17016/2380-7172.3190>
23. Fujita, Shigeru and Moscarini, Giuseppe and Fabien Postel-Vinay (2020) "Measuring Employer to Employer Reallocation," NBER WP No. 27525, DOI: <https://doi.org/10.21799/frbp.wp.2021.22>
24. Geweke, J. (1999). "Using simulation methods for bayesian econometric models: inference, development and communication," *Econometric Reviews*, 18 (1), 1-126. DOI: <https://doi.org/10.1080/07474939908800428>
25. Hall, Robert E. and Marianna Kudlyak (2020) "Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics," FRB SF WP No. 2019-05, DOI: <https://doi.org/10.24148/wp2019-05>

26. Hobijn, Bart and Ayşegül Şahin (2012) "Beveridge Curve Shifts across Countries since the Great Recession," FRB SF WP No. 2012-24, DOI: <https://doi.org/10.24148/wp2012-24>
27. Lubik, Thomas A. (2021) "Revisiting the Beveridge Curve: Why has it shifted so Dramatically," FRB Richmond, Economic Brief No. 21-36, October.
28. Menzio, Guido and Shouyong Shi (2011) "Efficient Search on the Job and the Business Cycle," *Journal of Political Economy*, 119 (3), pp. 468-510. DOI: <https://doi.org/10.1086/660864>
29. Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Aabigail Wozniak (2016): "Understanding Declining Fluidity in the U.S. Labor Market," *Brookings Papers on Economic Activity*. DOI: <https://doi.org/10.1353/eca.2016.0015>
30. Pissarides, Christopher A. and Barbara Petrongolo (2001) "Looking into the Black Box: A Survey of the Matching Function," *Journal of Economic Literature*, 39(2), 390-431, DOI: <https://doi.org/10.1257/jel.39.2.390>
31. Rodgers, William M. and Alice L. Kassens (2022) "What Does the Beveridge Curve Tell Us about the Labor Market Recovery?" FRB SL, *On The Economy Blog*, July 2022.
32. Stanton, Christopher T. and Catherine Thomas (2021) "Who Benefits from Online Gig Economy Platforms?" NBER Working paper 29477. DOI: <https://doi.org/10.3386/w29477>