Decomposing Shifts in the Beveridge Curve: Implications for Labor Market Dynamics and Inflation

By

Katharine Abraham, John Haltiwanger and Lea Rendell[1]

Extended Abstract

Shifts in the empirical Beveridge Curve are closely monitored by labor market analysts. Movements involving simultaneous increases or decreases in both unemployment and vacancies are often interpreted as changes in matching efficiency, potentially driven by shifts in skill or geographic mismatches in the labor market. Recent work by Barnichon and Shapiro (2024) emphasizes the importance of such shifts for understanding inflation dynamics. They find that the vacancy-to-unemployment ratio—a traditional measure of labor market tightness—outperforms the unemployment gap in Phillips Curve estimations. Their results suggest that shifts in the Beveridge Curve—and in particular shifts in matching efficiency—play a key role in this improved predictive performance.

Abraham, Haltiwanger, and Rendell (2020) (hereafter AHR) propose a generalized matching function that incorporates effective searchers and effective vacancies, extending the standard approach based on unemployment and vacancies. Their model accounts for several sources of misspecification in the traditional matching function. First, many hires originate not from the unemployed, but from job-to-job transitions or individuals classified as not in the labor force. Thus, the effective pool of job seekers encompasses the entire working-age population, with varying search intensities across different groups—especially high among the recently

[1]Abraham and Haltiwanger are faculty members at the University of Maryland and Research Associates of the NBER; Rendell is a faculty member at Kenyon College. Both Abraham and Haltiwanger have current grants from the Sloan Foundation and were co-PIs on a recent grant from the National Science Foundation. Abraham also is currently receiving research support from the W.E. Upjohn Institute for Employment Research.

unemployed but non-zero for all groups. Second, recruiting intensity differs systematically across vacancies, with fast-growing firms, for instance, exerting greater effort to fill positions. By accounting for these factors, AHR demonstrate that their generalized matching function better captures actual job-filling and job-finding dynamics. They also note—but do not fully explore—the potential implications for interpreting shifts in the standard Beveridge Curve.

In this paper, we leverage the generalized matching function to analyze observed shifts in the Beveridge Curve. We provide an exact decomposition that separates shifts due to true changes in matching efficiency from those stemming from the aforementioned sources of misspecification. Our analysis spans 1990:Q1 to 2025:Q4, covering key episodes such as the Great Financial Crisis, the COVID-19 pandemic, and the subsequent recovery.

Beyond examining Beveridge Curve shifts, we also revisit Barnichon and Shapiro's (2024) findings on inflation dynamics. While they test the AHR generalized tightness measure and find it superior to the standard tightness measure for their Phillips Curve analysis (which itself outperforms the unemployment gap), their analysis of Beveridge Curve shifts focuses on the standard, not the generalized, version. Our decomposition allows us to distinguish between true changes in matching efficiency versus shifts in the composition of effective searchers and effective vacancies as drivers of these shifts and thus of observed inflation dynamics. This framework also enables us to investigate whether particular groups of job seekers play an outsized role in these dynamics.

A Motivating Decomposition

Consider measured matching efficiency from a standard matching function:

$$\log(\mu_t^s) = \log(H_t) - \alpha^s \log(V_t) - (1-\alpha^s)\log(U_t) \tag{1}$$

where μ_t^s is standard matching function efficiency, H_t is hires, V_t is vacancies, U_t is unemployment, and α^s is the elasticity of hires with respect to vacancies, with $1-\alpha^s$ the elasticity of hires with respect to unemployment assuming a constant returns to scale matching function.

Matching efficiency for the generalized matching function is given by:

$$\log(\mu_t^g) = \log(H_t) - \alpha^g \log(EV_t) - (1-\alpha^g)\log(ES_t) \qquad (2)$$

where EV_t is effective vacancies, ES_t is effective searchers, α^g is the elasticity of the generalized matching function with respect to EV_t and (1- α^g) is the elasticity of the generalized matching function with respect to ES_t, again assuming constant returns to scale in matching. Effective vacancies and effective searchers are given by $EV_t = \sum_j \rho_t^{v_j} V_{jt}$ and $ES_t = \sum_i \rho_t^{S_i} S_{it}$ where V_{jt} represents the number of job openings from firms of type *j* at time *t*, S_{it} represents the number of job searchers of type *i* at time *t*, $\rho_t^{v_j}$ represents the intensity of employer recruiting effort for firms of type *j* at time *t*, and $\rho_t^{s_i}$ represents the intensity of job search on the part of searchers of type *i* at time *t*.

Combining these two approaches yields:

$$\log(\mu_t^s) = \log(\mu_t^g) + \alpha^g \log(EV_t) - \alpha^s \log(V_t) + (1-\alpha^g)\log(ES_t) - (1-\alpha^s)\log(U_t) \qquad (3)$$

Equation (3) highlights that shifts in matching efficiency from the standard matching function can reflect shifts in the matching efficiency of the generalized matching function but also shifts in effective vacancies relative to simple vacancy counts and effective searchers relative to the count of the unemployed. The results in AHR already provide *prima facie* evidence that this decomposition is likely to be informative. More specifically, their generalized matching function

tracks actual job filling and finding rates more closely than the standard matching function. Put differently, the latter requires greater implied changes in matching efficiency to match observed job filling and job finding rates.

Implementation Issues

AHR generate a measure of effective searchers based upon twenty-two distinct groups of the working age population. These groups include the employed distinguishing between those who are involuntarily part time and the remainder; thirteen groups among the unemployed based on reason for and duration of unemployment; and seven groups among those out of the labor force, including three groups among those who say they want a job but have not actively searched in the current month, along with people who say they are in school, retired, disabled, or out of the labor force for some other reason. This will permit us to isolate which groups are particularly relevant for shifts in measured matching efficiency.

One implementation issue in measuring the relative search intensities of the twenty-two groups is how to handle changes in these relative intensities over time. In AHR, for twenty-one of the groups, relative search intensities were set equal to the job finding rates in 2006 relative to the 2006 job finding rate for a reference group; the reference group was assigned a relative search intensity of one. Over time, however, relative search intensities may change. We will explore different approaches to accounting for trend changes in these relative search intensities. One possible approach could allow for linear or quadratic trends in the relative intensities; another could capture trend variation by interpolating between the values of the relative intensities at cyclical peaks.

Another challenging implementation issue is how to account for variation in recruiting intensity, especially for disaggregated groups. Davis, Faberman, and Haltiwanger (2013)

(hereafter DFH) provide one method for generating aggregate variation in recruiting intensity over time that builds on the micro-level relationship they establish between job filling rates and the number of jobs an employer is seeking to fill. AHR used the national aggregate time series produced using this method. We could apply the same approach to published JOLTS data by industry and/or region. Other studies have used survey information on proxies for recruiting intensity. Most closely related to our work, Forsythe and Weinstein (2024) made use of data for large employers surveyed over the period from 2006 to 2016 about their college recruiting efforts. Consistent with DFH, they find that employers recruit more intensively when they have a larger number of openings to fill and that more intensive recruiting efforts are associated with a higher job filling rate. Additional information on employer recruiting intensity is available from a short survey conducted by the Richmond Fed in 2022 (see https://www.richmondfed.org/publications/research/economic_brief/2022/eb_22-36). While these latter sources can't be used directly to produce a monthly recruiting intensity time series for our 1990-2025 sample period, they may provide insights that are useful for cross checking and validating alternative approaches to capturing recruiting intensity.

References

Abraham, Katharine, John Haltiwanger, and Lea Rendell. 2020. "How Tight is the US Labor Market?," *Brookings Paper on Economic Activity, Spring 2020*, 97-165.

Barnichon, Regis and Adam Shapiro. 2024. "Phillips meets Beveridge," *Journal of Monetary Economics*, 148(S), 103660.

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

Forsythe, Elizabeth, and Russell Weinstein. 2024. "Recruiting Intensity, Hires, and Vacancies: Evidence from Firm-Level Data," Working Paper, University of Illinois.