

NBER WORKING PAPER SERIES

REVISITING THE PHILLIPS AND BEVERIDGE CURVES:
INSIGHTS FROM THE 2020S INFLATION SURGE

Pierpaolo Benigno
Gauti B. Eggertsson

Working Paper 33095
<http://www.nber.org/papers/w33095>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2024

This paper was prepared for the "Federal Reserve Bank of Kansas City's 2024 Jackson Hole Economic Symposium." We especially thank Pascal Michaillat and Aysegul Sahin for educating us about labor market frictions, while absolving them of any errors, which are entirely our own. We would also like to thank our discussants, Guido Lorenzoni and Ivan Werning, as well as Hyeonsoo Lee, Adrien Foutelet, and Can Soylu for their outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Pierpaolo Benigno and Gauti B. Eggertsson. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Revisiting the Phillips and Beveridge Curves: Insights from the 2020s Inflation Surge
Pierpaolo Benigno and Gauti B. Eggertsson
NBER Working Paper No. 33095
October 2024
JEL No. E0, E30, E40, E42, E44, E47, E49, E5, E51, E58, J20, J39

ABSTRACT

This paper reexamines the Phillips and Beveridge curves to explain the inflation surge in the U.S. during the 2020s. We argue that the pre-surge consensus regarding both curves requires substantial revision. We propose that the Inverse-L (INV-L) New Keynesian Phillips Curve replace the standard New Keynesian Phillips Curve. The INV-L curve is piecewise-linear and more sensitive to labor market conditions when it crosses the *Beveridge threshold* — a point at which the labor market becomes excessively tight and enters a "labor shortage" regime. We introduce a modified Beveridge curve that features a near-vertical slope once the Beveridge threshold is passed, suggesting that in this region, adjustment in labor market tightness occurs almost exclusively through a drop in vacancies rather than an increase in unemployment. This feature matches the U.S. experience since the Federal Reserve's tightening cycle began in March 2022. We also observe a similar pattern in the data during five other inflation surges over the past 111 years where the Beveridge threshold was breached. We define a Beveridge-threshold (BT) unemployment rate. Once unemployment falls below this rate, policymakers must be alert to sharp inflationary pressures from demand or supply shocks. We explore several policy implications.

Pierpaolo Benigno
Department of Economics
University of Bern
3003 Bern
Switzerland
pierpaolo.benigno@unibe.ch

Gauti B. Eggertsson
Department of Economics
Brown University
64 Waterman Street
Providence, RI 02912
and NBER
gauti_eggertsson@brown.edu

1 Introduction

The inflation surge of the 2020s peaked at 6.2% for the core Consumer Price Index (CPI), while the overall CPI was almost double digits. The peak of the core CPI coincided with the start of the Federal Reserve's tightening cycle of Federal Funds Rates in Q1 2022. The rise in overall CPI continued into the following quarter, partly influenced by the Ukraine-Russia War, which began on February 24, 2022.

When the Federal Reserve began tightening policy in March 2022, the central policy question was: How costly would it be to bring inflation back to its 2% target? At the core of this debate are two fundamental macroeconomic relationships: the slope of the Phillips curve and the shape of the Beveridge curve.

The Phillips curve relates inflation to a measure of economic activity.¹ At that time, empirical evidence suggested that the Phillips curve was very flat. A leading example is the work by Hazzell, Herreno, Nakamura, and Steinsson (2022), which found that a 2.9-percentage-point increase in unemployment resulted in only a one-percentage-point decrease in inflation. This led pessimists to argue that, to achieve its inflation target, the Federal Reserve would need to accept a substantial increase in unemployment.

The Beveridge curve describes the relationship between the intensity with which firms are looking for workers (job vacancies or v) and how many workers are looking for a job (the unemployed or u). The Beveridge curve is typically plotted as shown in Figure 2. A common metric of labor market tightness is the ratio v/u , a concept that dates back to Beveridge's work in 1944.

If the Beveridge curve is flat, it indicates that cooling the labor market (reducing v/u) will inevitably result in significant unemployment. Drawing from several past episodes, authors such as Blanchard, Domash, and Summers (2022) adopted a pessimistic view in 2022, arguing that a substantial increase in unemployment would be necessary to cool the labor market and bring inflation back to target.

Others were more optimistic. Several economists, including the authors of this paper and most notably Federal Reserve Chairman Jerome Powell, argued for the possibility of a "soft landing" – reducing inflation without a substantial increase in unemployment. Benigno and Eggertsson (2023, hereafter BE) formalized this prediction, which current data increasingly supports. One contribution of this paper is to reconcile why prominent economists reached divergent conclusions despite using similar methodological approaches. Central to the resolution is that, unlike in our framework, existing models of the Beveridge curve are highly sensitive to a single parameter, the elasticity of the matching function with respect to unemployment, the value of which there is no consensus in the literature.²

¹An important aspect of the Phillips curve is expectations about future inflation. Long-term inflation expectations remained stable throughout the 2020s episode, as discussed in detail by Benigno and Eggertsson (2023).

²We discuss the debate between Blanchard, Domash, and Summers (2022) and Figura and Waller (2022) in Section 4. See also Crump et al. (2024) for discussion of the role of labor market and supply shock in shaping inflation dynamics during the inflation surge.

This paper revisits the soft landing prediction and clarifies how our modeling and empirical framework differ from existing approaches. Central to our analysis is a substantially modified Phillips curve, which we refer to as the Inverse-L (INV-L) New Keynesian Phillips Curve. This curve resembles the letter "L" turned backward, with each leg slanted, as shown in Figure 1. While our modification of the Beveridge curve is more modest, it carries significant implications. The most notable is that once a certain threshold is crossed, the Beveridge curve becomes nearly vertical.

We introduce a new term: the Beveridge threshold. This value marks a critical point in the labor market, measured by the ratio v/u , that triggers fundamental changes in both the Phillips curve and the Beveridge curve. Once the Beveridge threshold is crossed, the Phillips curve becomes steeper and more prone to generating inflation, while the Beveridge curve becomes nearly vertical, with most labor market adjustments occurring through changes in vacancies.

The Beveridge threshold can also be used to calculate an indicator of the unemployment rate that triggers stronger inflationary pressures, which we call the Beveridge threshold (BT) unemployment rate. This rate marks the point below which the Phillips curve steepens significantly.

The exact value of the Beveridge threshold is uncertain. Beveridge originally conjectured it to be 1, and this approximation works surprisingly well for U.S. data. However, the uncertainty around this threshold poses a significant challenge for policymakers.

The past 111 years of U.S. economic history include six major inflation surges: World War I, World War II, the Korean War, the Vietnam War, the Great Inflation of the 1970s, and the inflation surge of the 2020s. In all but one of these episodes (the Great Inflation of the 1970s), the labor market was extraordinarily tight, as indicated by v/u crossing the Beveridge threshold. One contribution of this paper is demonstrating that in these five episodes, the data suggest normalization primarily occurred through falling vacancies rather than rising unemployment.

Figure 1 illustrates three central lessons of the INV-L New Keynesian Phillips curve.

First, it explains why the inflation surge was unexpected by both markets and policymakers at the time (see panel (a)). In response to a substantial increase in demand, a flat Phillips curve, widely taken for granted prior to the inflation surge, would result in equilibrium at point E, with only a trivial inflationary effect.³ The INV-L Phillips curve explains the unpleasant surprise: beyond the Beveridge threshold, there is an unexpected burst of inflation for those who placed their faith in a flat Phillips curve.

Second, it illustrates how the costs of reducing inflation differ between the standard conventional flat New Keynesian (NK) Phillips curve and the INV-L Phillips curve. Consider, for example, the estimate

³This was a significant element in the formulation of the Federal Reserve's new policy framework announced in 2020, which emphasized the need to focus more on the employment part of the dual mandate, partly due to the judgment that the risk of inflation was low given the small slope of the Phillips curve. The idea was that there was little harm in letting the labor market run hot due to the low risk of inflation. In her presentation of the framework, for example, Governor Brainard argued that a key benefit of the 2020 policy framework was that it eliminated the 'longstanding presumption that accommodation should be reduced preemptively' as the labor market tightens 'in anticipation of high inflation that is unlikely to materialize.' See further discussion in Eggertsson and Kohn (2023)

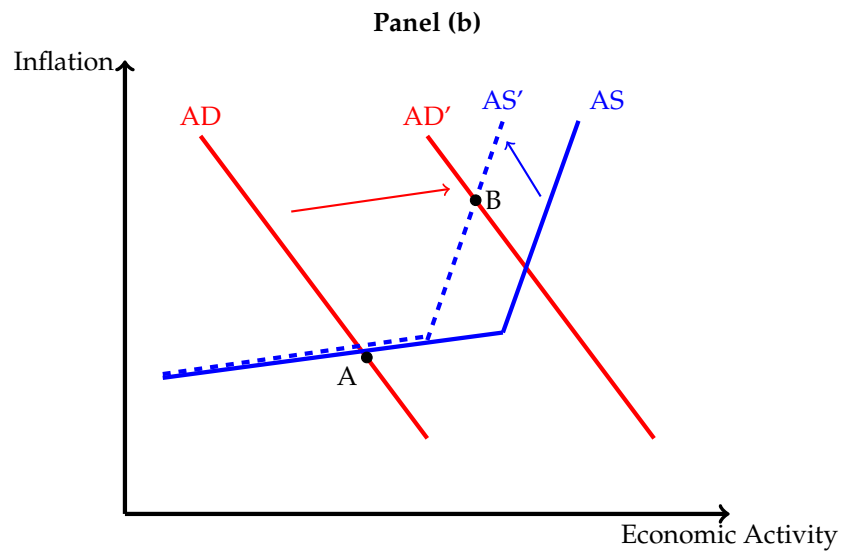
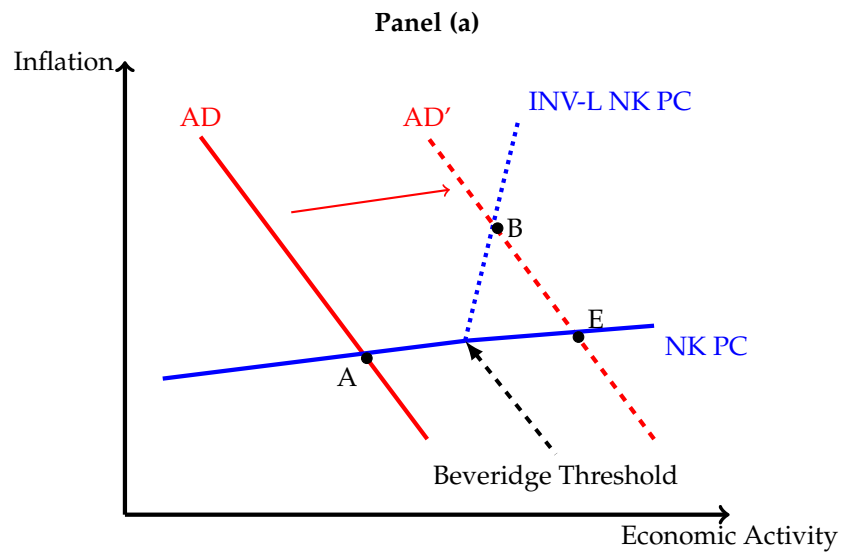


Figure 1: Panel (a): Demand shock impacting non-linear Phillips curve. Panel (b): Supply shock supercharged under a non-linear Phillips curve.

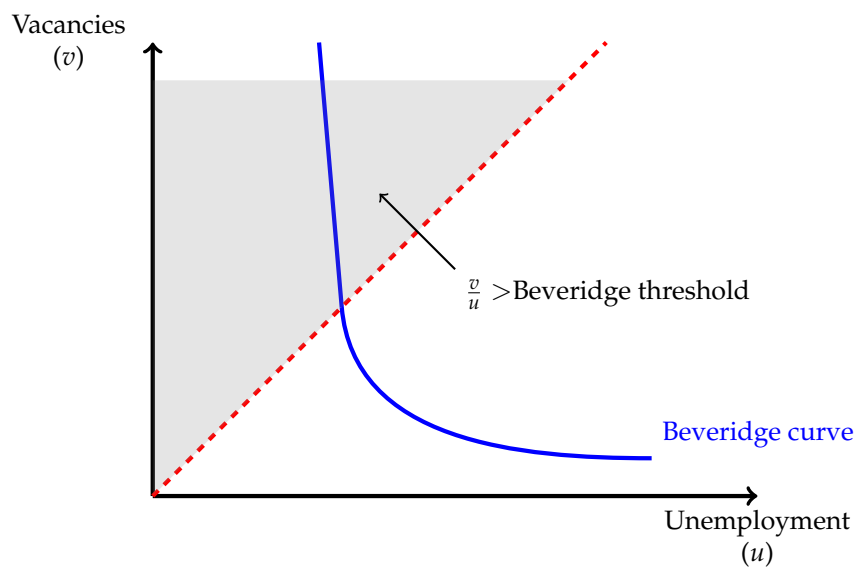


Figure 2: Beveridge curve during the 2020s inflationary surge. Grey area: region where $v/u > 1$.

by Hazell et al. (2022). Their analysis suggests that reducing inflation by 4.2 percent would require a 12.2 percentage-point increase in the unemployment rate. In contrast, the INV-L curve implies a much smaller loss. In the extreme case, if the curve were completely vertical beyond the Beveridge threshold, there would be no cost at all.

Third, panel (b) shows how supply shocks have out-sized inflationary impacts when demand intersects the Phillips curve beyond the Beveridge threshold. We present both theoretical and empirical evidence to support this, suggesting that supply shocks played a significant role in the surge.⁴

A more subtle implication of this point is that the debate over whether demand or supply drove the inflation surge is somewhat misleading in our framework. Panel (b) suggests that demand must push the economy beyond the Beveridge threshold for supply shocks to have a meaningful effect. In the text, we provide a simple decomposition based on BE, attributing roughly two-thirds of the surge to demand forces and one-third to supply. However, our empirical analysis indicates that supply shocks would not have had a statistically significant impact if the economy had not crossed the Beveridge threshold.

Figure 2 highlights our key insight about the Beveridge curve: once the economy crosses the Beveridge threshold (denoted by the dashed line and shaded region), the curve becomes nearly vertical, with adjustments occurring mainly through changes in vacancies rather than unemployment. This feature, central to our paper, aligns with recent data, as we will demonstrate. While some search-and-matching models can replicate this behavior, their results are typically highly sensitive to specific

⁴In panel (b) the interpretation of the x-axis, which we simply denote economic activity is output, as in Figure 19. In panel (a) the curves take the same shape regardless of if the x-axis refer to output or v/u .

parameter, for which empirical value there is little agreement. In contrast, this is a relatively robust feature of our model.

A key challenge in our framework is the uncertainty surrounding the value of the Beveridge threshold. While Beveridge's conjecture of 1 proves to be a surprisingly good approximation, there is nothing inherently special about this value; it can vary over time and with economic shocks. This uncertainty presents significant challenges for policymakers, especially given the lags in policy implementation. As of June 2024, the latest v/u was 1.2. Consequently, policymakers face two significant risks: if they are too slow to ease policy, they risk causing a "hard landing," characterized by high unemployment due to flatter Phillips and Beveridge curves. On the other hand, if they cut rates prematurely, they leave the economy vulnerable to inflationary supply shocks that could be exacerbated above the Beveridge threshold, or to unexpected demand shocks, which the analysis suggests have far larger impact on inflation once the equilibrium surpasses the Beveridge threshold. Our current assessment suggests the former risk outweighs the latter.

A useful property of the Beveridge curve we propose is that it allows us to construct a Beveridge-threshold (BT) unemployment rate, below which the Phillips curve steepens. The gap between the actual unemployment rate and the BT rate widened post-COVID, exceeding -1% at the peak of the inflation surge. As of this writing, the gap is narrowing and approaching zero, suggesting that further policy tightening may be unnecessary. The BT unemployment rate could become a useful metric for the Fed in evaluating its maximum employment mandate.⁵

2 The labor shortage, supply bottlenecks and the 2020s inflation surge

Two words dominated the economic discussion leading up to the inflation surge: "labor shortage" and "supply chain shortage." In this section, we focus on both terms. Of the two, we argue that the labor shortage is of primary importance, partly because it is a necessary condition for supply disruptions to have a significant impact on core inflation. We show that out of the six inflation surges in the U.S. during the last 111 years, five were associated with labor shortages.

2.1 Evidence of labor shortage

In this section we start by reviewing anecdotal evidence of labor shortage in the U.S. leading up to the inflation surge and review more formal measures we rely on later in the paper.

⁵Michaillat and Saez (2022) introduce the notion of an efficient unemployment rate, also identified by the Beveridge threshold and equivalent to the geometric average of the unemployment and vacancy rates. Our BT unemployment rate has a different foundation and quantitative implications as detailed in Section 5.



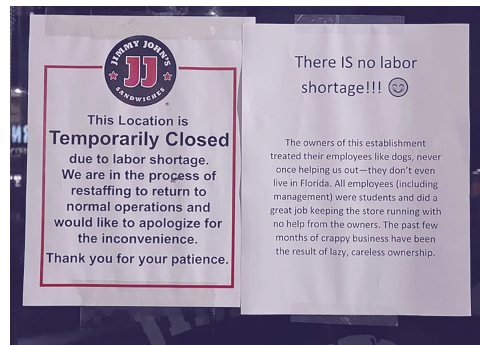
(a) Ohio labor shortage, June 2021



(b) Pennsylvania labor shortage, May 2021



(c) Virginia labor shortage, July 2022



(d) Florida labor shortage, January 2022

Figure 3: Labor shortage across different states

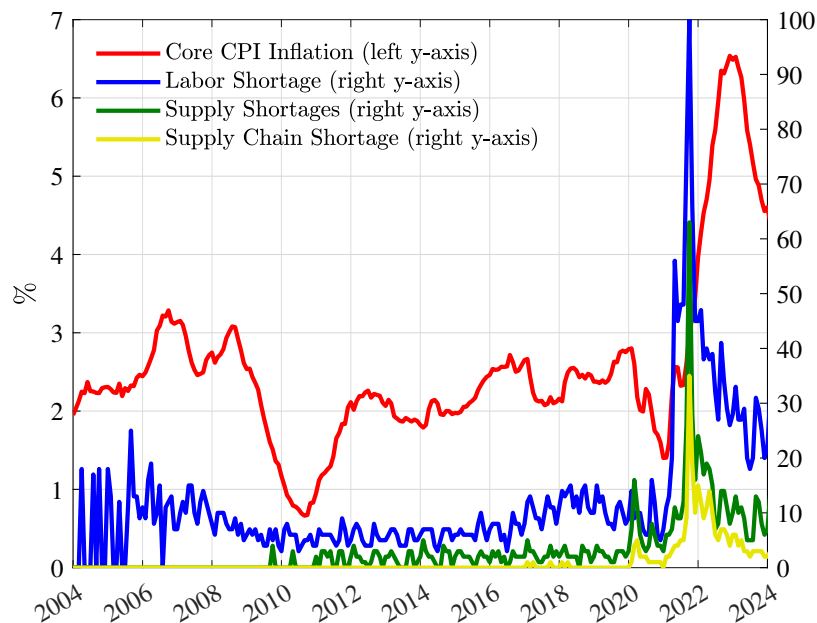


Figure 4: Core CPI inflation rate at annual rates (left y-axis), Google trend indexes on "labor shortage," "supply shortage," "supply chain shortage," (right y-axis). *Source: BLS and Google.*

Casual anecdotal evidence abounds of labor shortages during the inflationary surge of the 2020s. A family of five enters a restaurant in Providence in 2021 which is at 1/3 capacity. They are turned away on account of the restaurant being full. When the bewildered customers point out that 2/3 of the restaurant is empty, the answer is: that part is closed due to "labor shortage." Figure 3 illustrates anecdotal evidence of this kind, stories that were widespread throughout the U.S. during the inflation surge: Panel (a) shows a scene from Ohio where an establishment announced it is closed "due to no staff." Panel (b) shows that McDonald's in Pennsylvania began offering new employees a signing bonus of 500 dollars. Meanwhile, other McDonald's establishments tried to attract new employees by giving them the opportunity to "eat as much as they want" while working. Panel (c) shows a strategy employed in Virginia: It closed during certain days. An alternative common strategy was to close during specific hours of the day.⁶ Panel (d) shows a similar sign in Florida to the one in Ohio, announcing the establishment is closed due to a "labor shortage." To the right, a former worker puts an alternative announcement suggesting there is no such thing as a "labor shortage." According to the worker, all the company needs is to improve employment conditions.

The entry of "labor shortage" into public discussion is a relatively recent phenomenon. One way of seeing this is shown in Figure 4, which displays a Google trend index for "labor shortage" in the U.S., normalized to 100 when it peaked in October 2021, as the inflation surge was gaining momentum.⁷ The Google trend reached its maximum five months before core CPI peaked, in March 2022, and when the Federal Reserve tightened policy. Prior to the inflation surge, "labor shortage" barely registered in

⁶This strategy was adopted, for example, by the pharmacy CVS, which closed from 1-2 PM in many of its locations for a "lunch break."

⁷The numbers represent search interest relative to the highest point on the chart for the U.S. A value of 100 is the peak popularity for the term. A value of 50 means that the term is half as popular.

Google trends.

A more formal metric of labor market tightness is originally proposed by Beveridge (1944), the founding father of labor economics. He suggested that the tightness of the market should be considering the number of vacancies firms are seeking to fill (v) (demand for labor) relative to the number of unemployed workers looking for jobs (u) (supply of labor). Beveridge suggested that the labor market is tight when there are more firms looking to fill jobs than unemployed people seeking them, i.e., $\frac{v_t}{u_t} \equiv \theta_t > 1$. We call $v/u = 1$ the **Beveridge threshold** and define the region when it is crossed as a **labor shortage**. Michaillat and Saez (2022) provide theoretical rationale that the economy is efficient at the Beveridge threshold.

The bottom panels of Figures 5 and 6 show the CPI inflation rate since 1913. They illustrate the six major inflation surges in U.S. economic history over the past 111 years: World War I, World War II, the Korean War, the Vietnam War, the Great Inflation of the 1970s, and the COVID-19 inflation surge. The upper panel plots the v/u ratio marking the Beveridge threshold with a dashed line.⁸

We see a striking pattern when comparing the two panels in each figure. In the six major inflation surges observed in U.S. data since WWI, v/u surpasses the Beveridge threshold in all but one case: the Great Inflation of the 1970s. We analyze the correlation embedded in these relationships in next section.

In recent years, researchers have increasingly begun using the ratio of vacancies to unemployment (v/u) as an alternative to unemployment alone, which is the more conventional measure of labor market slack.⁹ Does v/u convey any relevant information beyond u alone? As we will discuss in Section 4, in the absence of shocks, the ratio v/u conveys the same information as u alone via the Beveridge curve. However, as we will see, COVID-19 caused measurable and significant shifts to the curve, among which is the fact that more vacancies are needed to create a match for a given level of unemployment.

Figure 7 shows that while u alone reasonably captures movements in v/u in the decade leading up to the inflation surge, this relationship breaks down after the COVID-19 pandemic and during the inflationary surge.¹⁰ Figure 31 in the Appendix shows other episodes where u deviates from v/u . This explains the finding in Furman and Powell (2022): v/u performs better than u alone even if one restricts oneself to predicting inflation using only data *prior* to the pandemic.

To see how the relationship broke down, consider the period when inflation breached the inflation target in the spring of 2021 (e.g., year-on-year core inflation was 2.96% in April 2021, the highest in 25 years). In May 2021, unemployment was at 5.75%, suggesting a slack labor market to many observers, since it was at 3.9% leading up to the pandemic. Meanwhile, several other measures, such

⁸These figures are adapted from Michaillat and Saez (2022), one of the original motivation of this analysis, even if they do not focus on inflation but instead emphasize the efficiency property of the Beveridge threshold.

⁹See for example Barnichon et al. (2021), Furman and Powell (2021), Ball et al. (2022), Domash and Summers (2022) and more recently Barnichon and Shapiro (2024).

¹⁰Our observation is built on a regression proposed by Kalantzis (2023) and is described in the subtitle of Figure 7. We have found that the Kalantzis regression is the best fitting one relative to various alternatives we have explored.

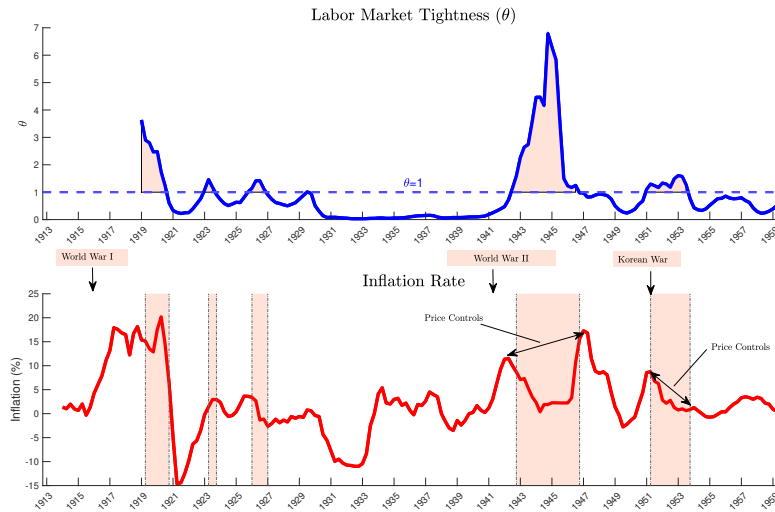


Figure 5: Top panel: (θ) vacancy-to-unemployed ratio. Bottom panel: CPI inflation rate at annual rates. Period 1913 Q1 – 1959 Q4. Source: Petrosky-Nadeau and Zhang (2021) and BLS.

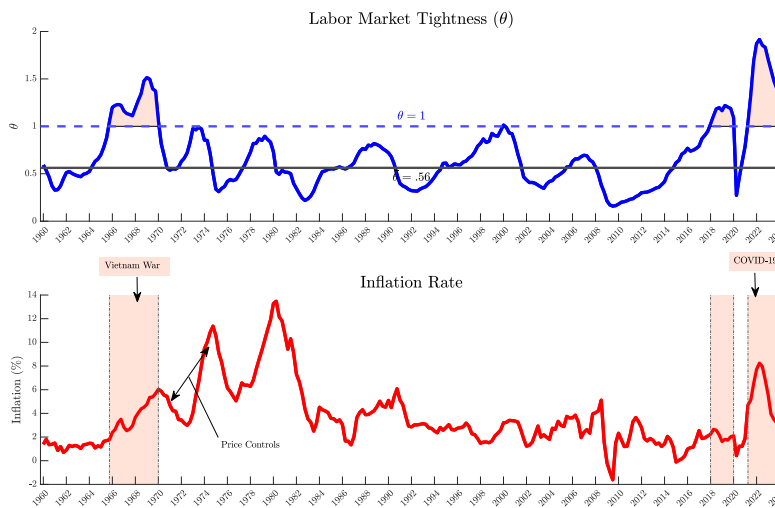


Figure 6: Top panel: (θ) vacancy-to-unemployed ratio. Bottom panel: CPI inflation rate at annual rates. Period 1960 Q1 – 2024 Q2. Source: Petrosky-Nadeau and Zhang (2021), Barnichon (2010) and BLS.

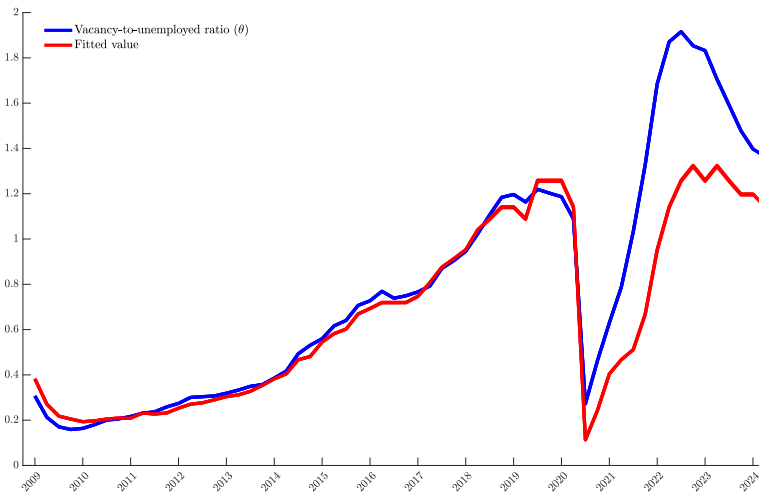


Figure 7: Vacancy-to-unemployed ratio and its fitted value using a regression proposed by Kalantzis (2023): $\ln \theta_t = a + b \ln(u_t / (1 - u_t)) + \varepsilon_t$ on the sample 2001 Q1 – 2024 Q2. The Figure shows the time-series for the sample 2009 Q1 – 2024 Q2. *Source: JOLTS and BLS.*

as the prime-age employment-to-population ratio frequently cited by Fed officials at the time, also pointed to a weak labor market.¹¹ At the same time, the vacancy-to-unemployed ratio surpassed the Beveridge threshold, signaling the beginning of a tight labor market.

Using the unemployment rate alone fails to capture both the demand and supply sides of the labor market, as emphasized by Beveridge, who argued that both sides were better proxied by v/u , with v representing demand and u supply. The Beveridge metric during the inflation surge of 2020s suggested a significant imbalance between the demand and supply of labor. Despite a large number of workers looking for jobs, there was an even larger number of firms trying to fill them, as suggested by our anecdotal evidence at the beginning of this section. Indeed, v/u reached its highest level since WWII during this episode, as shown in Figures 5 and 6.

2.2 Evidence of supply shortages

Aside from the popular discussion of labor shortage, the topics of "supply-side bottlenecks" and "supply chain disruptions" also dominated the discussion leading up to the inflation surge.

As in the case of the labor shortage, there is a considerable amount of anecdotal evidence in this vein. Figure 8 illustrates a notable event from the spring of 2021, when a single freight vessel blocked the Suez Canal, causing delays for up to 369 vessels. Similarly, U.S. ports experienced severe congestion as the economy recovered, with a peak of 150 vessels idling while awaiting their turn to offload cargo. A well-known instance involved about 100 ships drifting idle off the coast of Los Angeles, a critical hub where approximately 40 percent of containerized U.S. imports are processed through the ports

¹¹See Eggertsson and Kohn (2023) for a discussion.



(a) March 2021: blockage of the Suez Canal by a vessel.



(b) September 2021: Over 100 idle cargo ships waiting to offload outside of Los Angeles

Figure 8: Anecdotal evidence of global shipping disruptions

of Los Angeles and Long Beach. Part of the reason for these supply disruptions, aside from pure accidents like the Suez incident, was that the recovery from COVID-19 was uneven. Demand for goods outpaced the demand for services, a topic we return to later.

Again it is useful to consider Google Trends to formalize these anecdotal evidences. We see in Figure 4 that different measures of supply shortages spiked prior to the 2020s inflation surge.

A major challenge in measuring supply shortages in the literature following the inflation surge is the lack of a uniform method. Myriad measures have been proposed, see among others the Global Supply Chain Pressure Index of the Federal Reserve Bank of New York discussed in Benigno et al. (2022). These metrics are sensitive to post hoc pattern identification. While there is undoubtedly much to be learned from these measures, particularly as more data becomes available, our preferred approach takes a different direction. While not perfect, we prefer to use exactly the same measures that were commonly used to capture supply disturbances in the literature on the Phillips curve *prior* to the inflation surge of the 2020s, thus avoiding any concerns of ex post fitting.

The three most common measures are shown in Figure 9: headline shocks to overall CPI or PCE, and the difference between the growth of the import price deflator and the GDP deflator (see discussion in Ball et al. (2022)).

The most striking feature of these measures, relative to the prominence of supply disruptions in public discussion, is that they do not seem particularly volatile during the inflation surge. As we will see in the next section, however, once combined with labor shortage, they can have a substantial effect.

We view it as quite plausible that this measure of supply shocks leaves out some important components specific to COVID-19. In the empirical analysis we report, there is a substantial unexplained spike in inflation in Q2 of 2021 (see Figure 12 in the next section). A natural explanation for this is that our measure of supply shocks does not fully capture the special nature of supply-side bottlenecks during that quarter. For example, it was at the very end of Q1 2021 that the Suez Canal was

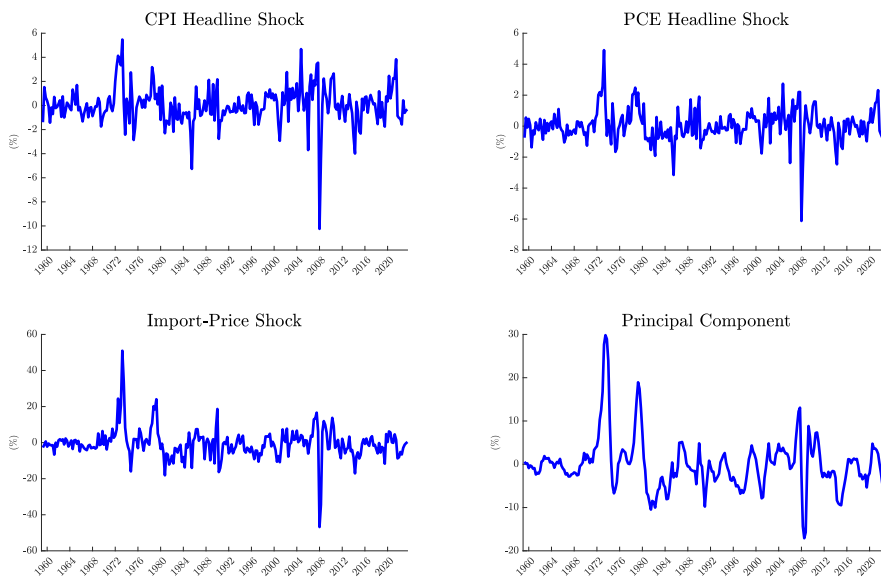


Figure 9: Measures of supply shock and their principal component. Panel a: CPI Headline Shock. Panel b: PCE Headline Shock. Panel c: Import-Price Shock. d: Principal Component of the three shocks. *Source: BLS, Authors' computation.*

clogged.¹²

3 The return of the non-linear Phillips curve and the role of super-charged supply shocks

The last section suggested a relationship between inflation and labor market tightness, showing that in five out of six inflation surges in the U.S. during the past 111 years, the labor market was extraordinarily tight. Here, we show that once the economy crosses the Beveridge threshold and enters a period of labor shortage, i.e., $v/u > 1$, there is strong evidence of the Phillips curve becoming substantially steeper. This has two major implications. First, demand shocks will exert a substantially larger impact on inflation. Second, negative supply shocks have a larger impact on inflation.

3.1 Basic correlations

Sometimes, a picture says more than a thousand words. Figure 10, reproduced from BE, shows a scatter plot of aggregate quarterly data with inflation on the y-axis and the logarithm of θ on the x-axis. Labor shortage is characterized by crossing the Beveridge threshold so that $\theta > 1$, or equivalently $\log(\theta) > 0$. We split the period into four sub-periods, with panels (a) and (d) both showing instances

¹²The incident occurred on March 23, 2021, and the canal was cleared for traffic on March 29, 2021.

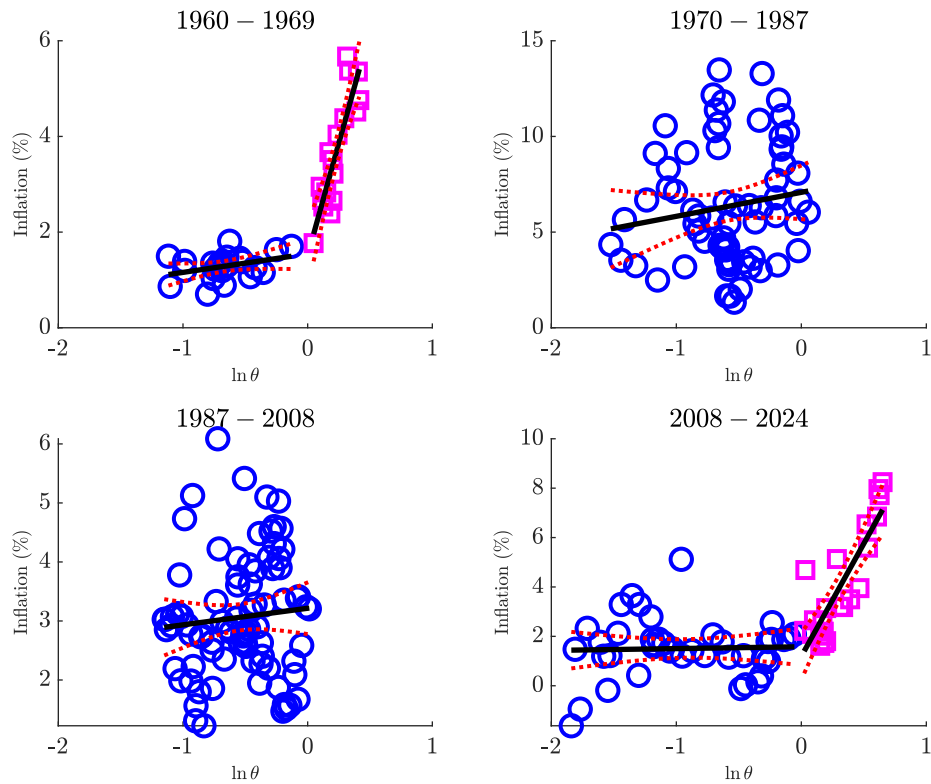


Figure 10: Inflation: CPI inflation rate at annual rates. $\ln \theta$: log of the vacancy-to-unemployed ratio. Sample 1960 Q1 – 2024Q2. Source: BLS and Barnichon (2010).

where $\theta > 1$, corresponding, respectively, to the late 1960s and the 2020s. The periods of labor shortage are shown with purple squares, while what we define as the "regular labor market" is shown by blue circles. As the Beveridge threshold is crossed, just looking at the raw data seems to hint at a substantial difference in the relationship between v/u and inflation.

An obvious limitation of time series data is that there is a limited number of episodes of labor shortages. Gitti (2023) considers data on inflation in 21 Metropolitan Statistical Areas (MSAs) in the United States from December 2000 to April 2023 and constructs a new time series for v/u corresponding to each region. The raw data that her overall findings are based on are shown in Figure 11. The same pattern emerges as in the aggregate time series: there appears to be a clear break as one crosses the Beveridge threshold around $\log(\theta) = 0$.

3.2 Traditional statistical analysis of Phillips curves allowing for non-linearity

Theoretically, a scatter plot represents equilibrium points where aggregate demand and supply intersect. One cannot claim that it provides any reliable information about either demand or supply, such as a Phillips curve.

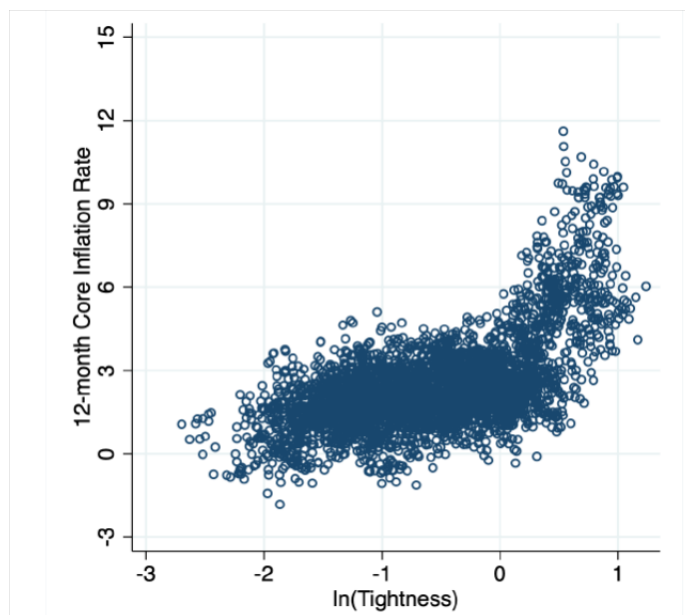


Figure 11: Inflation: CPI inflation rate at annual rates. In θ : log of the vacancy-to-unemployed ratio for 21 Metropolitan Statistical Areas in the U.S. from 2000-2023. *Source: Gitti (2023).*

And yet, it is hard to look at these figures without suspecting that during periods of labor shortage, there are strong hints of higher inflation pressures. Here, we report results from BE that show this suspicion can be supported with conventional regression analysis, controlling for confounding influences, particularly supply shocks, as suggested by McLeay and Tenreyro (2019) as one way of solving the inherent identification problem when estimating Phillips curves.

BE replicate conventional regressions common in the literature, where core inflation is regressed on $\ln \theta$ and a number of control variables, including proxies for supply shocks and inflation expectations. Their innovation relative to the existing literature is simple. Existing literature, such as Ball et al. (2022), considers empirical specifications for non-linearities at all times by including second-order terms. Instead, BE introduce a dummy variable for the periods when $\theta > \theta^*$, where θ^* is an estimate of the break point in the slope of a piece-wise linear Phillips curve. The dummy variable is introduced for both the θ variable and the supply shock, and they find that it is statistically and economically significant in both cases. Moreover, by examining the likelihood function of their empirical specification, they find that $\theta^* \approx 1$ provides a reasonably good fit, thus supporting the Beveridge hypothesis that crossing the unitary threshold leads to an excessively tight labor market. Yet, it is worth emphasizing that this threshold is not precisely estimated. As discussed in the introduction, this presents significant challenges for policymakers.

The motivation for the specification, shown in Appendix A along with baseline results, is both empirical and theoretical. Inspection of Figures 10 and 11 suggests that a piece-wise linear function may provide a reasonable approximation of the data. Moreover, BE show theoretically that a characterization of this kind arises quite naturally, if one adapts a model of wage setting in the spirit of Phillips

(1958) as we will discuss in Section 3.4.

The main conclusions of the benchmark empirical result in BE are summarized below. We refer the reader to the original paper for precise definitions of variables and a number of robustness checks.

The first main conclusion is that the slope of the Phillips curve, defined in terms of the relation between core inflation and θ , increases by a factor of ten once the Beveridge threshold is crossed. It is very flat in the absence of a labor shortage (0.52) compared to when there is a labor shortage (5.35), using the last part of the sample from 2009 onwards.¹³ The second major conclusion is that labor shortages supercharge supply shocks. While supply shocks do not have a statistically significant effect on core inflation under normal circumstances according to the benchmark estimates, they do so during a labor shortage, as we will show below.

Figure 12 shows the evolution of core CPI inflation accounted for by the various right-hand side variables of the regression in BE.¹⁴

The first major point is that the majority of the run-up is explained through the combination of the hatched gold bars (the contribution of θ interacted with the dummy variable for labor shortage) and the hatched red bars (the contribution of the supply shocks interacted with the dummy variable for labor shortage). At the beginning of the surge, in the second quarter of 2021, supply shocks explain a larger component. However, starting from the following quarter, labor tightness becomes predominant in explaining the run-up, even though supply shocks still play a non-trivial role.

A second observation is that a significant part of the initial decline in inflation is due to supply shocks exerting downward pressure on inflation, playing an even bigger role than they did in the run-up. However, the component capturing labor market tightness remains an important headwind and explains why inflation has not fallen promptly to the two-percent target. In the second quarter of 2024, when the estimation period ends, the v/u was at 1.2 as of June 2024, down from its peak of 2 in March 2022 once the Federal Reserve started raising rates. Thus, the U.S. economy has yet to return to the Beveridge threshold, and accordingly, inflation remains more sensitive to both supply shocks and labor market tightness.

The gray bars in Figure 12 represent the residual that the econometric specification is unable to account for. Interestingly, it plays the largest role in the first quarter of the run-up in inflation, i.e., Q2 of 2021. A plausible explanation for this, as we have already discussed, is that our measure of supply shock is based on traditional frameworks developed by the literature prior to the inflation surge. As a result, it may not capture the unique features that characterized the inflation of the 2020s, such as supply chain disruptions, which were likely of particular importance early in the period.¹⁵

¹³See column 4 in Table 1 in the Appendix.

¹⁴Figure 32 in the Appendix show the decomposition starting from 2017.

¹⁵The most common traditional measure is the difference between headline and core inflation. Yet it does not account for the highly unusual supply bottlenecks that were especially prominent at the onset of the inflation surge, as indicated by the Global Supply Chain Pressure Index of the Federal Reserve Bank of New York. See Benigno et al. (2022).

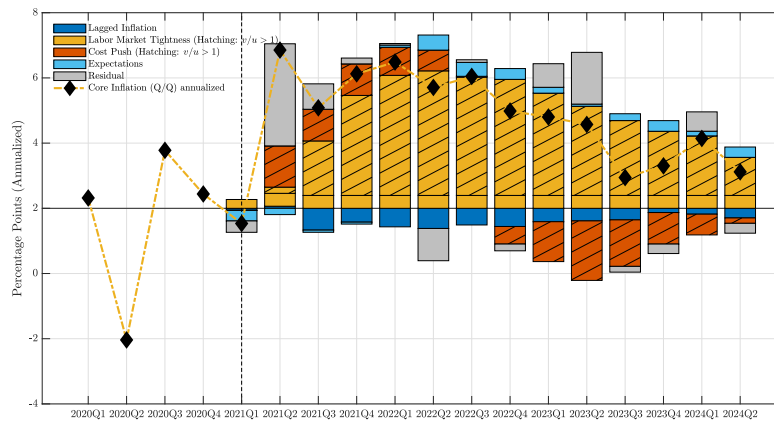


Figure 12: Decomposition of the baseline regression in BE between the contributions of the various regressors: lag inflation, $\ln \theta$, supply shocks, inflation expectations. For the variable $\ln \theta$, hatching corresponds to the contribution of the variable for the portion of θ that exceeds the unitary value. For the supply shock, hatching corresponds to the contributions of the variable when $\theta > 1$. Core inflation and all the components are plotted at annualized quarterly rates. *Source: Benigno and Eggertsson, 2023.*

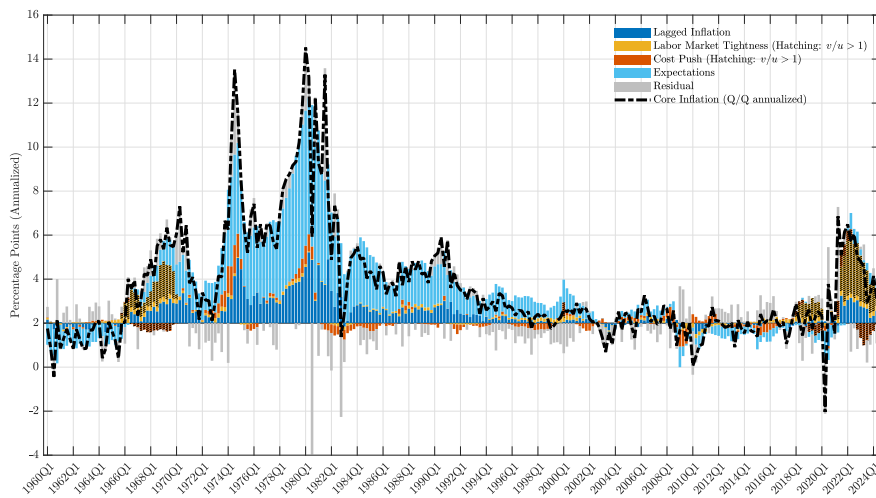


Figure 13: Decomposition of the baseline regression in BE between the contributions of the various regressors: lag inflation, $\ln \theta$, supply shocks, inflation expectations. For the variable $\ln \theta$, hatching corresponds to the contribution of the variable for the portion of θ that exceeds the unitary value. For the supply shock, hatching corresponds to the contributions of the variable when $\theta > 1$. Core inflation and all the components are plotted at annualized quarterly rates. *Source: Benigno and Eggertsson, 2023.*

3.3 The Fed credibility and the inflation surges of 1960s vs 2020s

A notable feature of Figure 12 is that inflation expectations play virtually no role in the inflation dynamic. This is consistent with the observation that long-term inflation expectations were stable during this period.

Conventional wisdom, however, attributes the unanchoring of inflation expectations as an important factor in explaining the Great Inflation of the 1970s. Of the six major inflation surges of the past 111 years discussed in the last section, this is the only one where v/u remained below the Beveridge threshold. Figure 13 performs the same type of decomposition but for the entire period 1960 Q1 –2024 Q2.

A central result is that inflation expectations play a large role in explaining the Great Inflation of the 1970s, consistent with conventional wisdom. Their contribution also remained persistently high as the Federal Reserve brought down inflation in the 1980s. This is one of the fundamental insight of the literature explaining the depth and duration of the Volcker recession which brought inflation back under control (see Erceg and Levin (2003) and Goodfriend and King (2005) for two classic references).

For our purposes, it's crucial to distinguish between the inflation surges of the 1960s and the 2020s. While both were triggered by labor market tightening beyond the Beveridge threshold ($v/u = 1$), there is an important difference. As shown in Figure 13, inflation expectations in the 1960s *began to become unanchored* towards the end of the surge, which did not happen in the 2020s. This unanchoring in the 1960s may have contributed to the Great Inflation of the 1970s, an episode often overlooked by economists. The 1960s surge likely made inflation expectations more vulnerable to rapid increases following the oil price shocks of the 1970s.

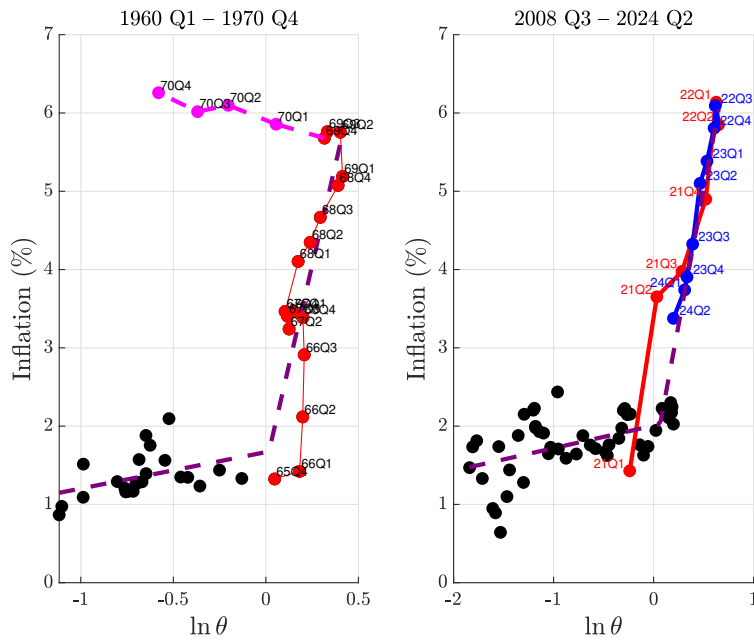


Figure 14: Scatter plots of core CPI inflation, quarterly rates (annualized), (y-axis), and vacancy-to-unemployed ratio, $\ln v/u$, (x-axis). Left panel: period 1960 Q1 – 1970 Q4. Right panel: period 2008 Q3 – 2024 Q2. Source: BLS and Barnichon (2010), Authors' computation.

To develop this argument further, consider the scatter plots from Q1 1960 to Q4 1970 and from Q3 2008 to Q2 2024 as shown in Figure 14. First, examine the run-up in inflation during both episodes. We see that inflation travels along the steep part of the Phillips curve once $\log(\theta) > 0$. That is where the similarities end, however. Once θ starts dropping during the inflation surge of the 2020s, inflation travels down the same curve. In sharp contrast, during the 1960s surge, as θ started dropping in the early 1970s, inflation kept increasing.

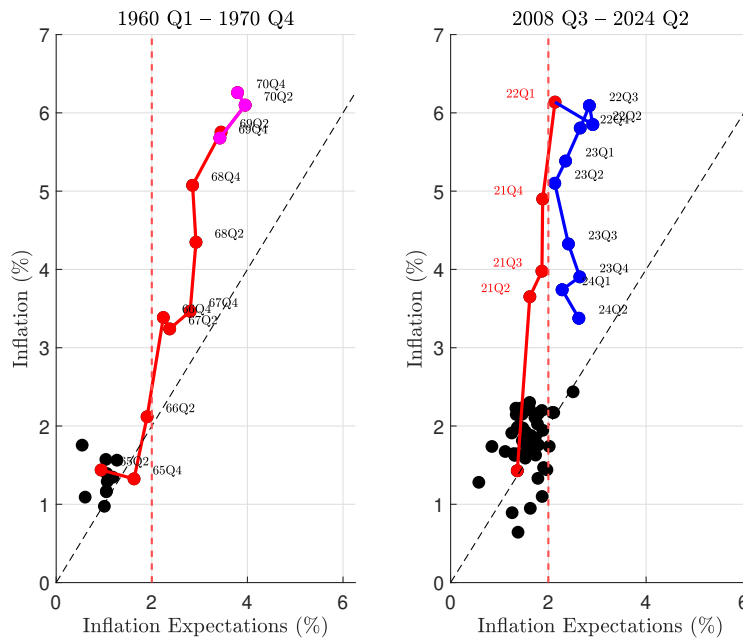


Figure 15: Scatter plots of core inflation, quarterly rates (annualized), (y-axis), and inflation expectations, (x-axis). Left panel: period 1960 Q1 – 1970 Q4, bi-annual frequency; Livingston survey inflation expectations. Right panel: period 2008 Q3 – 2024 Q2, quarterly frequency; 2-year Cleveland Fed inflation expectations. Source: BLS, Federal Reserve Bank of Cleveland, Federal Reserve Bank of Philadelphia, Authors' computation.

This is explained by a rise in inflation expectations, as a careful examination of Figure 13 reveals. To develop this argument in a more transparent way, Figure 15 shows inflation expectations on the x-axis and actual inflation on the y-axis. If expectations move one-to-one with realized inflation, we should expect the data to line up along the 45-degree line shown with a dashed black line. If they are insensitive to realized inflation, however, they should concentrate around the vertical red line denoting a 2 percent inflation target. On both figures, the black dots denote the data before the Beveridge threshold is crossed. The red line in both graphs shows the rise in inflation associated with increased tightness in the labor market as v/u crosses 1.

A first observation is that the red line in the 1960s is much closer to the 45-degree line than during the 2020s, where it does not deviate substantially from it. Of even greater interest is to consider once the labor market starts weakening again. This is shown with purple dots in the 1960s: As labor tightness eases, inflation continues to increase in tandem with inflation expectations. In sharp contrast, consider now the recent episode shown in the right-side panel of Figure 15. The blue dots denote the period after March 2022, when the Federal Reserve started tightening interest rates, at which point labor market tightness eases. In response to this, actual inflation falls, in sharp contrast to the end of the inflation surge of the 1960s. Moreover, we see that inflation expectations stay much closer to the 2 percent target line throughout the 2020s inflation surge, relative to the 1960s inflation surge.

3.4 The relationships between inflation and v/u , and the INV-L New Keynesian Phillips Curve

So far, we have limited ourselves to describing the data without explaining the theoretical reasoning that motivates our thinking. Here, we describe how one can, on the basis of our previous work in BE, arrive at a straightforward theoretical relationship that rationalizes a piece-wise linear relationship, grounded in the empirical evidence we have already presented. The resulting relationship is what we term the Inverse-L New Keynesian Phillips curve. Relative to the standard New Keynesian Phillips curve, there are two major differences. The first is that we replace traditional measures of slack, such as the output gap, with $\theta = v/u$. The second is that the model can be approximated as being piece-wise linear; once a certain threshold is breached and $\theta > \theta^*$, the slope of the Phillips curve changes.

We arrive at this relationship in two steps, which are described in detail in BE. In the first step, a New Keynesian (NK) Phillips curve is derived where inflation depends on marginal costs, but our interpretation of how this marginal cost is computed differs from the standard approach, namely by emphasizing the wages of new hires rather than the average wage rate. In the second step, we postulate a wage-setting behavior as in Phillips (1958), which is non-linear. It is thus the labor market that is the fundamental source of the non-linearity of the Phillips curve. In the second step, wages of new hires can be expressed in terms of labor market tightness, i.e., θ .

Following the aforementioned ordering, in the first step, BE obtains:

$$\pi_t = \kappa_w \underbrace{\hat{w}_t^{new}}_{\text{Marginal Cost of Labor}} + \kappa_v \underbrace{\hat{v}_t}_{\text{Cost Push Shocks}} + \beta E_t \pi_{t+1} \quad (1)$$

Where w_t^{new} represents the real wages of new hires, \hat{v}_t are supply/cost push shocks, π_t is inflation, E_t is an expectation operator, and $\kappa_v > 0$, $\kappa_q > 0$, and $0 < \beta < 1$ are coefficients. The key difference, relative to the standard Phillips curve, is that marginal costs are proxied by newly negotiated wages compared to some measure of average wages. This innovation also clarifies the difference between our results and an influential paper by Bernanke and Blanchard (2023). They attribute a more trivial role to labor market tightness (especially toward the beginning of the surge). Bernanke and Blanchard (2023) follow the New Keynesian tradition and measure marginal cost with unit labor cost, proxied by average hourly labor cost.¹⁶

Here, instead, we distinguish between the salaries of *existing* workers and *new* workers. We argue that it is the wages of new workers that are relevant to marginal cost if a firm needs to increase production. The wages of new and existing workers may be very close under normal circumstances. BE, for example, assume for simplicity that they are exactly the same in the absence of labor shortage. When there is a labor market shortage, however, they may become very different, a difference that

¹⁶Moreover, Bernanke and Blanchard (2023) assume flexible prices and constant returns to labor, which implies that price inflation is equal to wage inflation and variations in supply shocks. This tends to put a large weight on supply shocks by necessity, since average wage inflation was low towards the beginning of the surge.

can be exaggerated by inflation to the extent that existing workers have negotiated their salaries in nominal terms.

The difference between the salaries of new workers relative to existing workers became particularly stark during the inflation surge. One way of seeing this is to compare the wage growth of those workers accepting new jobs to the wages of existing workers.

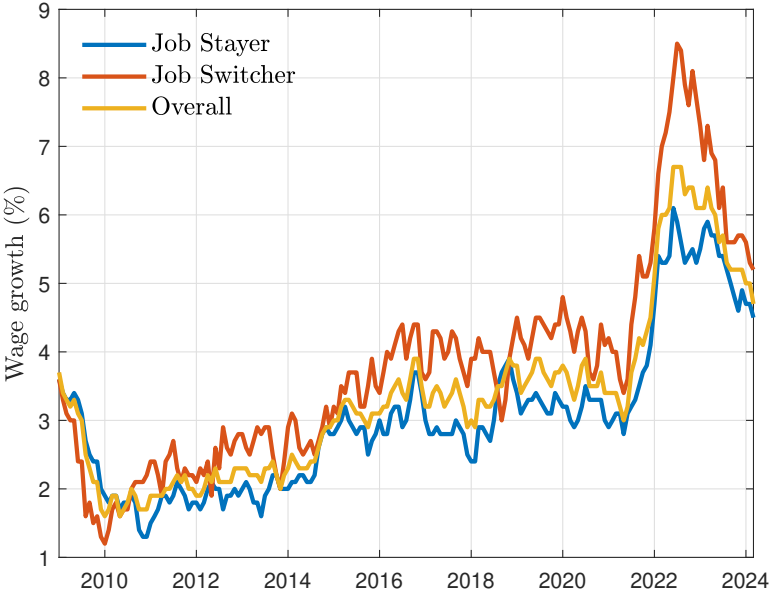


Figure 16: Wage growth (%), overall, and decomposition between job switchers and job stayers. *Wage Growth Tracker, Federal Reserve Bank of Atlanta.*

Figure 16 show that workers switching from one job to another, who correspond to approximately 40 percent of all new hires,¹⁷ saw a much more substantial increase in their nominal wage growth than existing workers. One aspect of the data on job-to-job switchers is that there may be compositional effects; perhaps one reason for the wage hike of a person switching jobs is that they found a more productive way of exploiting their talents. For this reason, some researchers have focused on wage growth for particular vacancies, which do not pose the same challenge. A recent estimate of wage growth taking this approach is shown by Crump et al. (2023) using Burning Glass Data and the methodology developed by Cattaneo et al. (2024).

¹⁷See Sedlacek (2016)

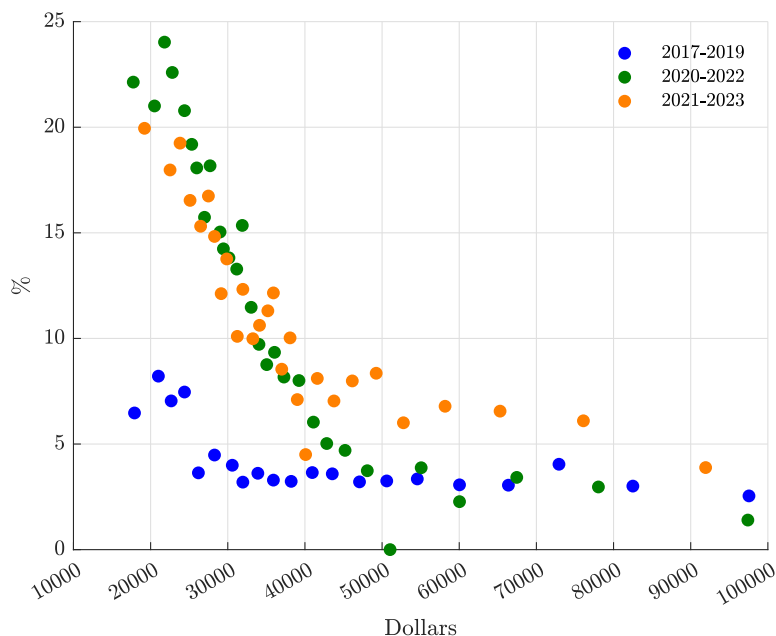


Figure 17: 2-year posted wage growth before and after the pandemic conditional on initial wage level, nonparametric estimates of the conditional median function. *Source: Crump et al. (2024).*

Figure 17 compares the two-year growth rate of nominal wages for particular vacancies in the period 2017-2019 relative to 2020-2022 and 2021-2023. As shown in this figure, the wage growth in 2017-2019 is far below that of the latter two-year periods, which is especially true at the lower end of the distribution. In sum, the evidence in Figures 16 and 17 suggests that it makes a quantitatively important difference if the wages of newly hired workers represent the true marginal costs for firms during a labor shortage, rather than average wages.

Starting from equation (1), BE show in a second step that the non-linear INV-L NK Phillips curve arises due to the way in which wages are set, bringing us back to the original Phillips curve. It is often overlooked that Phillips' (1958) article differs in two fundamental ways from existing formulations that carry his name. First, the y-axis represents nominal wage growth. Second, and this was one of his central point, it is highly non-linear.

Figure 18 shows the original Phillips curve, which is very flat at high unemployment rates and becomes very steep once the labor market is tight. The reason Phillips gives for the shape of his curve is that workers are very reluctant to accept wages below *the prevailing wage*, so wages fall only slowly even if unemployment is high. The converse is not true, however. Workers are happy to accept *new wages* that are higher than the prevailing ones, so in a tight labor market "we should expect employers to bid wage rates up quite rapidly." BE embed this idea into the Phillips curve in Figure 1.¹⁸ Importantly, BE's theoretical framework allows them to express new wages in terms of v/u . This

¹⁸See also Forbes, Gagnon, and Collins (2021) for a discussion of how downward nominal rigidities can bend the Phillips curve and for providing worldwide empirical evidence.

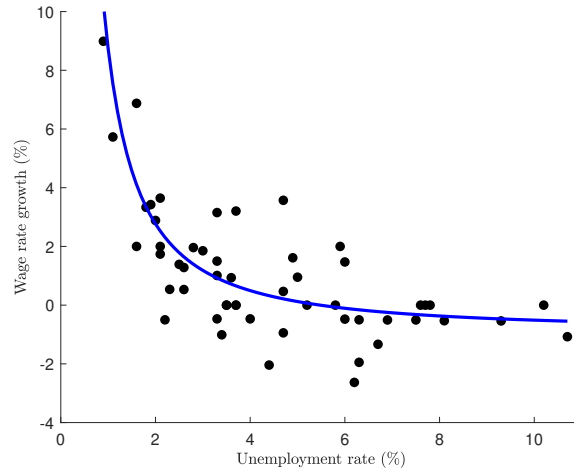


Figure 18: Phillips curve (1861 – 1913, United Kingdom, Phillips, 1958): Wage rate growth (Source: Brown and Hopkins, 1950) vs. unemployment rate (Source: Feinstein, 1972, adjusted by Board of Trade).

is significant since data on vacancies and unemployment are readily available, while good and representative data on new wages is not as easy to come by.¹⁹ Below, we present a simplified version relative to BE:

$$\pi_t = \begin{cases} \tilde{\kappa}^{tight} \hat{\theta}_t + \tilde{\kappa}_v^{tight} \hat{v}_t + \beta E_t \pi_{t+1} & \hat{\theta}_t > \hat{\theta}_t^* \\ \tilde{\kappa} \hat{\theta}_t + \tilde{\kappa}_v \hat{v}_t + \tilde{\kappa}_\beta E_t \pi_{t+1} & \hat{\theta}_t \leq \hat{\theta}_t^* \end{cases} \quad (2)$$

in which $\tilde{\kappa}^{tight} > \tilde{\kappa} > 0$, $\tilde{\kappa}_v^{tight} > \tilde{\kappa}_v > 0$, whereas the relationship between $\tilde{\kappa}_\beta$ and β is ambiguous depending on the specifics of the calibration.

It is worth highlighting one aspect of the result we mentioned above, even if we do not develop it further in this analysis. We have taken for granted that the Beveridge threshold is fixed at 1, as suggested by Beveridge himself. BE analysis suggests, however, that it can take on values different from 1, depending on institutional details, and moreover, it may vary over time. The exact region of labor market tightness when it becomes more and more inflationary is a topic that is largely unexplored.

3.5 Using the INV-L Phillips curve to understand the inflation surge of 2020s

We can use the simple framework we developed to gain insights into the inflation surge of the 2020s and give the reader a quantitative feel for the illustrative figures in the introduction. The Phillips curve derived in the last subsection maps directly onto our empirical estimates. Figure 19 casts the analysis in an Aggregate Demand (AD) and Aggregate Supply (AS) framework, rewriting the Phillips curve in terms of output instead of v/u – a more familiar exposition for most readers. The numerical

¹⁹Moreover, allowing for decreasing returns in labor, equation (1) would also show a term referring to the employment rate, which is a direct function of v/u .

example uses the empirical estimates of the regression we have discussed (see further details in BE and Table 1 in the Appendix, column 4).

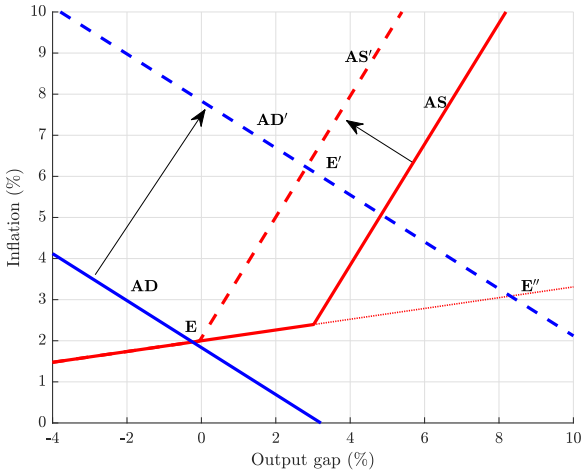


Figure 19: The 2020s Inflationary Surge: Inflation and output determination using an AS – AD framework in response to an increase in demand and a supply shock. *Source: Benigno and Eggertsson (2023).*

Consider first the increase in demand alone. If the slope of the Phillips curve is unchanged, then the demand would only have increased inflation to about 3 percent. In the numerical example, the increase in demand and negative supply shocks are chosen to match the relative contribution of labor market tightness and supply shocks as shown in Figure 12 for the second quarter of 2022. When the labor market is slack, on the flatter part of the curve, the supply shocks do not have a statistically significant effect according to our estimate.

Consider now the consequence of a nonlinear Phillips curve. Not only does it amplify the effect of the demand shock, but it also amplifies the supply shocks’ effect so that about three-quarters of the rise in core CPI from 2 to 6.2 percent is accounted for by the shift in AD, while one-quarter is explained by supply shocks.

An interesting property of the numerical example is that if the supply shocks revert back to pre-surge levels, then inflation can be brought down to about 2.5 percent without any drop in output with suitable monetary tightening: a soft landing. We think this provides a reasonable description of the current policy environment, a claim buttressed by the regression decomposition in Figure 12.

A word of caution is appropriate. As of writing, the labor market is still in the labor shortage region with $v/u = 1.2$ as of June 2024. This implies that if the deflationary supply pressures observed from Q1 2022 to the present reverse themselves, the Fed might be more cautious about cutting rates before crossing the Beveridge threshold. Yet, given the uncertainty surrounding the Beveridge threshold, as will become clearer in the next section, policymakers also risk over-tightening and entering a zone with far less favorable inflation-output trade-offs, where further reductions in inflation would come at the cost of substantial output contraction and higher unemployment.

One of the key arguments against the possibility of a soft landing, which now seems within reach, was based on the Beveridge curve, suggesting that cooling the economy would inevitably lead to a significant rise in unemployment. However, as with output, this concern has proven overly pessimistic, provided the Fed's contractionary policies have not been pushed too far. The reasons for this will be discussed in the following sections.

4 Revisiting the Beveridge curve

In the last section, we linked inflation to economic slack in a non-linear way, suggesting that the ratio of job vacancies to unemployed workers, v/u , is a useful measure of slack. A key takeaway from this is that inflation can be reduced without a significant drop in output. Instead, it is labor market tightness, v/u , that decreases. But how does this adjustment occur? Through a reduction in vacancies or an increase in unemployment?

The empirical relationship between v and u is known as the Beveridge curve. If this relationship is stable, then either v or u can serve as a sufficient statistic for the other. In this section, we show that during a labor shortage, a reduction in v/u not only leads to a small drop in output but also occurs primarily through a decrease in vacancies, with little change in unemployment. We first establish this empirically by analyzing data from the 2020s inflation surge, followed by data from the four other inflation surges over the past 111 years when v/u crossed the Beveridge threshold of 1.

We then provide a theoretical rationale for this result and compare it with prominent theoretical frameworks used to predict the likelihood of a hard or soft landing when the Federal Reserve began tightening in March 2022.

Unlike the Phillips curve, which is named after Phillips and was first drawn by him, the Beveridge curve was not drawn by its namesake. As in the case of the Phillips curve, the first study plotting the Beveridge Curve was based on data from the United Kingdom, using a subsample of Phillips' analysis, and published in the same year by Dow and Dicks-Mireaux (1958).

Dow and Dicks-Mireaux (1958) begin their work with the premise that unemployment and unfilled vacancies are frequently used to characterize labor demand. They state that "an increase in the pressure of demand will then always increase the level of unfilled vacancies reported and reduce unemployment" (Dow and Dicks-Mireaux, 1958, p. 4). This suggests an inverse relationship between vacancies and unemployment, which is the Beveridge curve. Interestingly, they further argue that "if employers give a correct statement of their vacancies, the line dividing areas of high and low demand would be at exactly 45 degrees through the origin" (Dow and Dicks-Mireaux, 1958, p. 4). This is identical to what we have termed the Beveridge threshold.

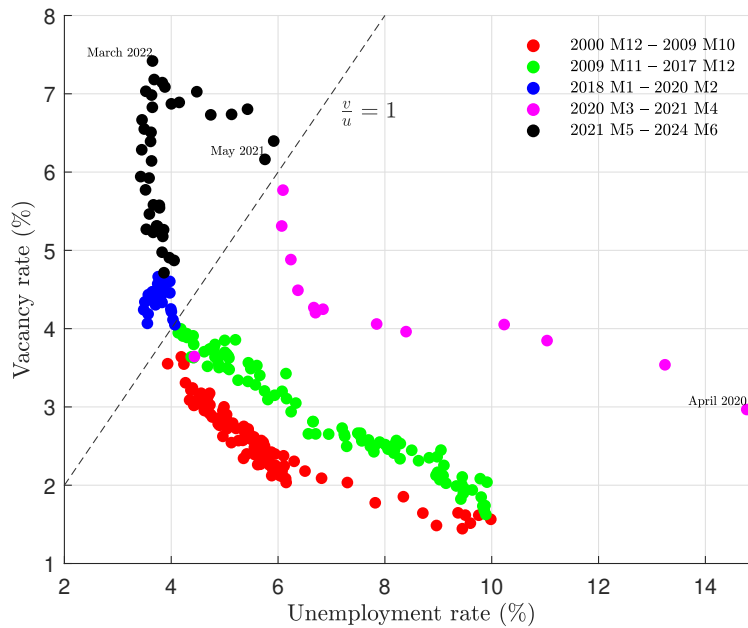


Figure 20: United States: scatter plots of job vacancy rate (v) and unemployment rate (u), monthly frequency, sample December 2000 – June 2024. *Source: BLS, JOLTS.*

4.1 The U.S. Beveridge curve in the 21st century

Figure 20 plots the job vacancy rate, calculated as the number of vacancies as a fraction of the labor force, and the unemployment rate, calculated as the number of unemployed individuals as a fraction of the labor force. The Beveridge threshold, $v/u = 1$, is shown with a dashed line. The data are monthly, sourced from the Job Openings and Labor Turnover Survey (JOLTS) of the U.S. Bureau of Labor Statistics (BLS), and cover the period from the survey's start in December 2000 to the most recent observation.

We can identify five distinct periods, or "regimes," each exhibiting unique dynamics. The first period, highlighted in red, extends from the start of the sample to the peak of unemployment during the financial crisis in October 2009. This period shows a stable pattern, resembling a curve that negatively correlates the vacancy rate with the unemployment rate.

Following the trough in October 2009, there was a continuous improvement in labor market conditions up to February 2020. This period also displayed a stable pattern but on a curve that was outwardly shifted compared to the previous period. Within this time frame, we have highlighted the points from January 2018 to February 2020 in blue, where the vacancy rate exceeded the unemployment rate and crossed the Beveridge threshold, just before COVID-19 and well before the inflation surge.²⁰

²⁰Figure 32 in the Appendix shows that, during this period, negative supply shocks mitigated the inflationary impact of labor market tightness.

The blue period is followed by an abrupt shift due to the eruption of the COVID-19 pandemic, captured by the purple point on the far right of the graph, where the unemployment rate reached 15%. This fourth period is characterized by an irregular recovery in the labor market.

As emphasized in the last section, the final part of this recovery is marked by a tight labor market, with the job vacancy rate rising above the unemployment rate. This period is depicted with black dots, starting from May 2021. During this time, there is a steady decline in the unemployment rate and an increase in vacancies until the Federal Reserve begins raising rates in March 2022. Notably, the most recent points of decline align closely with the pre-COVID observations depicted in blue.

Should the sharp reduction in vacancies with little or no change in unemployment be surprising? One way to address this question is from a purely theoretical perspective, which we will explore shortly. However, it seems natural first to consider what can be learned from the other four inflationary surges of the past 111 years when the Beveridge threshold was crossed. Once the Beveridge threshold was breached, did the labor market adjust through an increase in unemployment or a reduction in vacancies? As we will see, a drop in labor market tightness driven by a reduction in vacancies, with little change in unemployment, is a pattern shared across these episodes, with the possible exception of the aftermath of WWI.

4.2 Four inflation surges of the 20th century and the adjustment of v/u

There are four major inflation surges in the 20th century associated with labor shortages ($v/u > 1$): WWI, WWII, the Korean War, and the Vietnam War. Our question is: Is it possible to cool down the labor market, $\theta = v/u$, without increasing unemployment? To answer this, it makes sense to temporarily set aside the traditional representation of the Beveridge curve and instead plot unemployment on the x-axis and θ on the y-axis in place of v . We focus on periods when $\theta > 1$ and seek to understand whether bringing it back to 1 requires a significant increase in unemployment. Our general finding is that throughout these episodes, the adjustment largely occurs via a fall in vacancies.

Figure 21 plots the data on a diagram with θ on the y-axis and the unemployment rate on the x-axis.²¹ The figure uses solid lines to represent periods when θ is rising and a dashed line to indicate periods when it is declining.

The left panel documents the two most extreme episodes: WWI and WWII. The increase in θ was sharpest during WWII, reaching a peak of 7. Moving from 7 back to 1 increased unemployment from 1.2% to 4%. Meanwhile, the vacancy rate dropped from 8.4% at its peak to 4%, aligning with the Beveridge threshold of 1. This suggests that about 60 percent of the adjustment occurred through a reduction in vacancies. However, this adjustment was clearly influenced by the fact that unemployment started at very low levels, presumably well below its natural rate.

²¹Data are presented on a quarterly basis, i.e., as averages of the monthly rates.

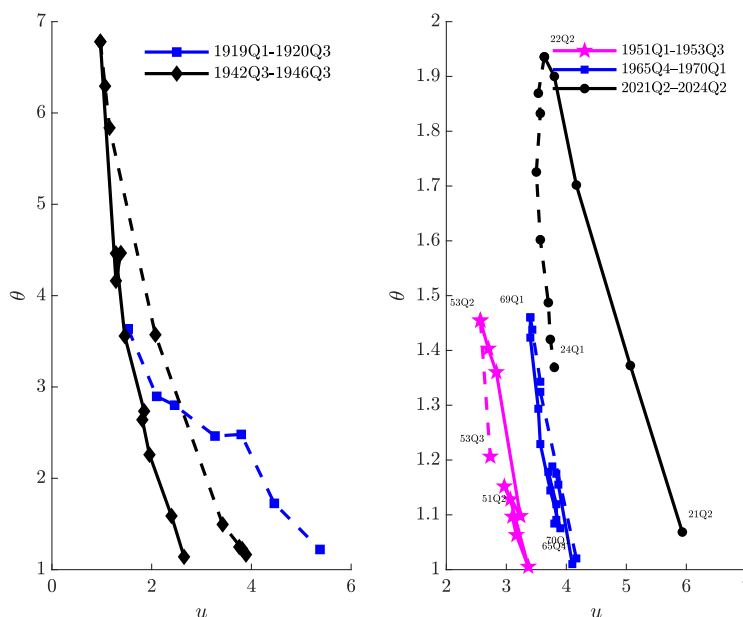


Figure 21: United States: scatter plots of vacancy-to-unemployed ratio, θ , and unemployment rate, u , at quarterly frequency. Periods in which θ is above the unitary value. Left panel: 1919 Q1 – 1920 Q3, 1942 Q3 – 1946 Q3. Right panel: Periods 1951 Q1 – 1953 Q3, 1965 Q4 – 1970 Q1, and 2021 Q2 – 2024 Q2. Source: BLS, Barnichon (2010), Petrosky-Nadeau and Zhang (2021).

The inflationary surge in the 2020s, shown in the right-side panel, began with a relatively higher unemployment rate, close to 6%. As the labor market cooled, nearly all of the adjustment occurred through a reduction in vacancies, likely because the economy was near the natural rate of employment. The same general pattern holds for the Korean War and the Vietnam War: θ dropped back to the Beveridge threshold with relatively modest changes in unemployment.²²

We do not have data on the run-up in vacancies during World War I, as our analysis is based on a series constructed by Petrosky-Nadeau and Zhang (2021), which only goes back to 1919. Generally, the drop in labor market tightness after World War I led to a greater increase in unemployment than in other episodes. This was largely driven by extremely contractionary monetary policy, which briefly resulted in a deflation rate of -15% – a record in U.S. history, even surpassing that of the Great Depression. The sharply contractionary monetary policy at the time is largely attributed to the constraints imposed by the gold standard (see, e.g., the discussion in Friedman and Schwartz (1963)).

What emerges from these figures is that at low unemployment rates, unemployment becomes relatively insensitive to variations in labor demand, as originally suggested by Dows and Dicks-Mireaux (1958). This observation seems to capture the concept of a minimum unemployment rate or maximum employment.

²²The 2018-2020 episode of labor market tightness is shown in Figure 33 in the Appendix displaying similar characteristics.

4.3 A minimalistic model of a Beveridge Curve

The analysis presented so far in this section has been purely descriptive, based on data. Here, we summarize a parsimonious search and matching model that underlies the analysis of BE. This model effectively describes the dynamics of labor tightness during periods of labor shortages. In Section 5, we generalize the model to better capture periods of slack in the labor market. We then contrast this model with two influential analyses that were hotly debated as the Federal Reserve began tightening policy in 2022: the work of Blanchard, Domash, and Summers (2022, hereafter BDS) and Figura and Waller (2022, hereafter FW).

The labor market literature is rich in theories of the Beveridge curve. A comprehensive survey of this literature is beyond the scope of this paper.²³ We propose an alternative view of the Beveridge curve based on BE. Our objective is to link inflation determination to measures of slack, specifically v/u , in the most parsimonious way possible. Accordingly, we make several stark simplifying assumptions.

The model has three key aspects. First, at the beginning of each period, the household makes a labor force decision, F_t . Second, at the same time, an exogenous fraction of the labor force participants $(1 - z_t)$ are "attached" to existing firms, while a fraction z_t are "unattached" and search for employment (for now, this fraction is exogenous; see Section 5 for an extension). Third, the level of employment/unemployment is determined at the end of the period via a standard matching function from the search and matching literature, after which production takes place.

An important aspect of our interpretation of the model is that the unattached workers, z_t , include previously unemployed individuals, those laid off, people entering the labor force, and those transitioning from one job to another. Only about 20 percent of new jobs are filled by unemployed workers, while about 40 percent are filled by people moving from one job to another, with the remainder primarily consisting of those who were previously categorized as non-participants in the labor force.

At the end of each period, unattached workers are either employed or unemployed. Therefore, we have:

$$z_t F_t = H_t + U_t, \tag{3}$$

where H_t is the number of hires in the period and U_t is the number of unemployed people at the end of the period. Following the labor-market search literature, we assume that the number of hires is a Cobb-Douglas function of unemployed workers and job vacancies:²⁴

²³For surveys covering the literature on Beveridge curves, see, for example, Petrongolo and Pissarides (2001), Elsby et al. (2015), and Barlevy et al. (2023).

²⁴Since H_t represents gross hires, one could argue that the matching function is misspecified if the composition of gross hires changes over time due to variations in flows from different segments of the hiring pool. See, e.g., BDS for a discussion, where they argue that in 2022 the proportions of the different flows were similar to what they were before the pandemic, suggesting that U_t is a reasonable proxy for the flows of other groups into new jobs.

$$H_t = m_t \left(U_t^\eta V_t^{1-\eta} \right), \quad (4)$$

where m_t represents matching efficiency and η is a parameter satisfying $0 < \eta < 1$. Using lower-case letters to denote unemployment and vacancy rates, i.e., $u_t \equiv \frac{U_t}{F_t}$ and $v_t \equiv \frac{V_t}{F_t}$, we can combine equations (3) and (4) by substituting for $\frac{H_t}{F_t}$ to obtain the following Beveridge curve:

$$z_t - u_t = m_t v_t^{1-\eta} u_t^\eta, \quad (5)$$

which will turn out to be a useful expression when contrasting with other works and can be more neatly expressed as:

$$v_t = \left(\frac{z_t - u_t}{m_t u_t^\eta} \right)^{\frac{1}{1-\eta}}. \quad (6)$$

This Beveridge curve holds at any time t , unlike in many labor market models where the Beveridge curve represents a steady-state relationship. The key to deriving a Beveridge curve that applies at all times, without any transition dynamics, is the assumption that the fraction of workers attached to existing firms, $(1 - z_t)$, is chosen at random and independently of the workers' history. A similar modeling approach is used in the familiar Calvo model of price setting.

The Beveridge curve describes a negative relationship between the vacancy rate and the unemployment rate and is plotted in Figure 22. The figure shows data ranging from January 2020, before the COVID-19 pandemic, to the present day. We plot a family of curves because they can be shifted by z_t and m_t . Both can be computed from the data and are shown in Figure 23. The value of z_t can be backed out directly,²⁵ while we need to make an assumption about η to extract the time series for m_t .²⁶

By construction, each curve passes through every data point. The green curve, for example, corresponds to the Beveridge curve in January 2020, on the eve of COVID-19. The largest shift in the curve occurs in April 2020, when z_t spikes and matching efficiency m_t drops, with unemployment reaching its peak and the vacancy-to-unemployment ratio hitting its minimum.

²⁵We can back out the time series for z_t using equation (3), as both H_t and U_t can be empirically measured using JOLTS data, and F_t is also measured by the BLS.

²⁶We can rewrite (4) as

$$\frac{H_t}{U_t} = m_t \theta^{1-\eta},$$

and regress the log of the left-hand side on $\log(\theta)$ to obtain the estimate $\eta = 0.3155$ in the sample from January 2009 to December 2019, which is close to the value assumed by BDS of 0.4. Using this value, we obtain a time series for m_t using as input the time series for H_t , U_t , and V_t , i.e.,

$$m_t = \frac{H_t}{U_t^\eta V_t^{1-\eta}}.$$

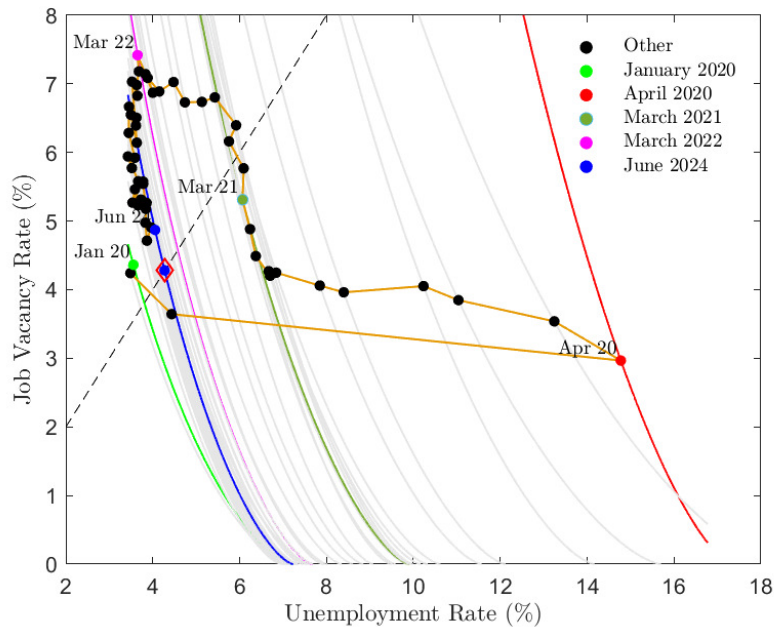


Figure 22: Scatter plot of job vacancy rate versus unemployment rate, sample January 2020 – June 2024, and Benigno-Eggertsson Beveridge curve (6). *Source: JOLTS, BLS. Authors' computation.*

As the labor market began to recover after the initial COVID-19 shock, the Beveridge curve shifted back to the left. Around November 2020, it stabilized for some time, as seen from the points climbing up the green curve, with unemployment decreasing and vacancies surging significantly. During this period, θ rose from 0.6 to 1.1 by May 2021. Afterward, labor market conditions continued to improve, with matching efficiency increasing and the fraction of unattached workers decreasing.

Notably, both matching efficiency and the fraction of unattached workers stabilized at the end of 2022, just before the Federal Reserve began raising rates. This suggests that the Beveridge curve has largely remained stable since then, with only a minor drift to the left.

The blue Beveridge curve in Figure 22 is defined as the curve that passes through the most recent data point, from June 2024. As is evident from this curve, the cooling of the labor market has closely followed this curve since shortly after the Federal Reserve began tightening policy. The BE Beveridge curve thus predicted that the labor market would stabilize largely through a reduction in vacancies, at the cost of a relatively modest increase in unemployment, as observed in the data.

While the exogenous series z_t has largely returned to pre-pandemic levels, matching efficiency remains notably below its pre-COVID level (see Figure 23). On the eve of COVID-19 in January 2020, matching efficiency stood at 0.88. It dropped significantly at the start of the pandemic, as filling vacancies became more difficult, then briefly recovered, only to decline again. Although it began to modestly improve during the inflation surge, matching efficiency still remains below pre-COVID levels, ending our sample at 0.75.



Figure 23: Unattached workers (z), matching efficiency (m), sample December 2000 – June 2024. Source: JOLTS, BLS. Authors' computation.

One explanation for the decline in matching efficiency, which suggests that more vacancies are needed to generate a successful hire at any given level of unemployment, is a structural change in aggregate spending. It likely takes longer for a worker to find a good job match when moving across sectors than when switching jobs within the same sector. Eggertsson and Kohn (2023) show that the COVID-19 recovery was highly uneven, with spending in the goods market greatly outpacing the recovery in services. The spending mix has still not returned to its pre-COVID levels. The possibility that the COVID-19 crisis could trigger a persistent or permanent sectoral shift was anticipated by Guerrieri et al. (2021), who study optimal monetary policy in the context of structural reallocation.²⁷

One of the central lessons of this episode is that viewing unemployment in isolation provides limited information about the tightness of the labor market. Recall that before the COVID-19 crisis, unemployment was 3.9%. This created the perception that any unemployment above that level indicated a slack labor market. However, Equation (6) clarifies that a tight labor market, such as when $\theta > 1$, can be consistent with a relatively high unemployment rate of 5.75%, as observed during the run-up of inflation in 2021. The reason for this is that there were so many firms urgently looking to hire, resulting in a high vacancy rate.²⁸

Looking ahead, if the Beveridge curve remains stable as indicated by the blue line, then the unemployment rate will be 4.28%, as shown by the diamond, once the labor market reaches balance at the

²⁷Their main finding is that structural reallocation can, in principle, justify inflation being above or below its target level for some time.

²⁸In Section 5.1, we refine the conditions under which the labor market is considered tight with respect to the unemployment rate by introducing the concept of the Beveridge-threshold unemployment rate, which indicates the value below which the labor market is tight.

Beveridge threshold.²⁹

A careful reader will observe that while the Beveridge curve we have outlined here does a reasonable job of describing the data in a tight labor market, a look at Figure 20 reveals that its success is limited by its inability to replicate the two relatively stable regimes shown by the green and red dots. We extend the model in Section 5 to address this by making z_t endogenous. Before we get there, however, it is useful to put the empirical success of this simple Beveridge curve in context by comparing it with more standard ones. To focus the discussion, we consider two papers aimed at predicting the effect of the Federal Reserve's tightening of policy in 2022 on unemployment.

4.4 The Beveridge curve and the policy debate when the Fed tightened policy in 2022

There was a lively debate after the Federal Reserve began tightening policy in March 2022 about how costly it would be in terms of rising unemployment, often framed as whether there would be a "soft" or "hard" landing following the inflation surge. We first consider the analysis by BDS, who predicted a hard landing in July 2022 based on April 2022 data, and then turn to the analysis by FW. Our general finding is that their analyses are not as different as they might initially appear. The key difference lies in their assumptions about the elasticity of the matching function. There's little consensus on this elasticity, and both use values well within the typical range found in the literature. As we will see, the BE model we reviewed in the last section is not as sensitive to different assumptions about this elasticity.

The basis of the BDS curve is the standard matching function (4) that we also assumed in our analysis, which can be written in rates as:

$$h_t = m_t v_t^{1-\eta} u_t^\eta \quad (7)$$

or more neatly as:

$$v_t = \left(\frac{h_t}{m_t u_t^\eta} \right)^{\frac{1}{1-\eta}}. \quad (8)$$

This represents a negative relationship between the vacancy rate and the unemployment rate. The key assumption is that $h_t \equiv H_t/F_t$, the gross number of hires with respect to the labor force, which they label reallocation intensity, is an exogenous variable, shown in the top panel of Figure 34.³⁰

²⁹Section 5.1 improves this prediction, suggesting a landing at 4.42%.

³⁰They further assume that gross separations and gross hires are equalized at each point in time in their model.

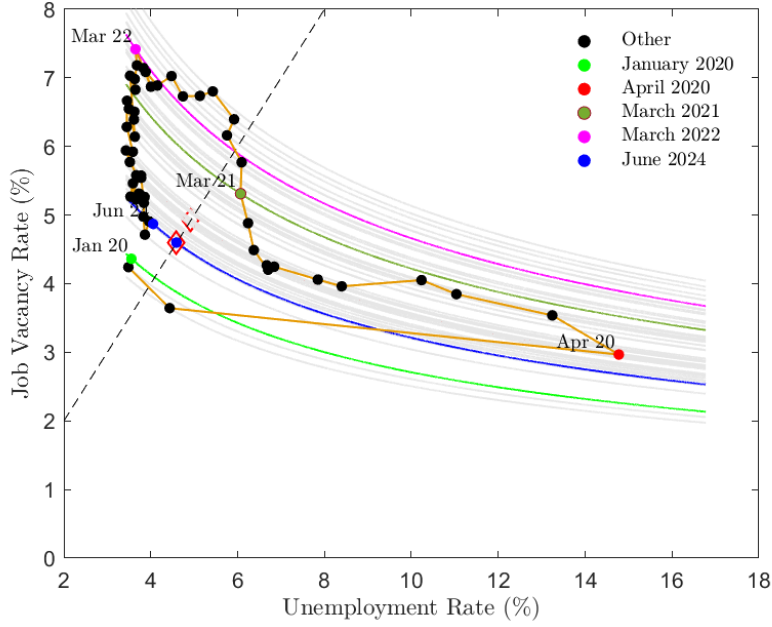


Figure 24: Scatter plot of job vacancy rate versus unemployment rate, sample January 2020 – June 2024, and Blanchard-Domash-Summers Beveridge curve (8). *Source: JOLTS, BLS. Authors' computation.*

Comparing the BDS Beveridge curve in (7) and the BE curve in (5), we see that the *only* difference is that in BE, h_t is endogenous with $h_t = z_t - u_t$.

Figure 24 depicts our replication of their results, plotting the BDS Beveridge curve for a given h . As shown by the pink line, drawn when the Federal Reserve began raising rates, the BDS Beveridge curve implies a substantial increase in unemployment to 5.8% to reach the Beveridge threshold of $v/u = 1$. In contrast, the BE Beveridge curve implies a value of 4.5%, which is closer to the rate at which the U.S. economy is currently landing.

Comparing the Beveridge curves in Figure 22 and Figure 24, we see that BDS requires a continuous shift in the exogenous variable h since March 2022 to rationalize the data. In contrast, the data in the BE Beveridge curve moves along the same curve, denoted in blue.

FW consider a textbook dynamic of the unemployment rate in labor-market search models, as presented in Pissarides (2000). Abstracting from variations in the labor force, which is normalized to 1, the unemployment rate at time $t + 1$ is

$$u_{t+1} = u_t + s_t(1 - u_t) - f_t u_t. \tag{9}$$

Variations in the unemployment rate are driven by workers separating from job relationships, represented by the term $s_t(1 - u_t)$, where the separation rate is defined as $s_t \equiv S_t/N_t$, with S_t as the

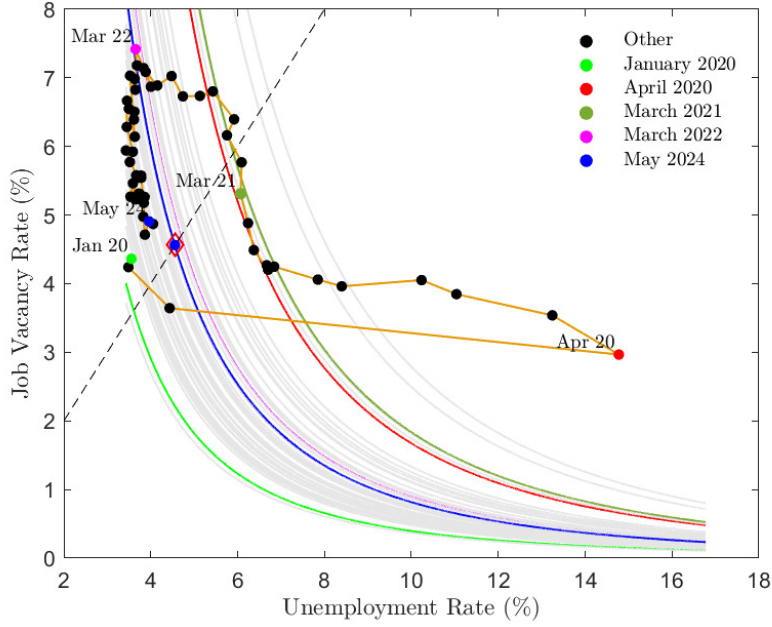


Figure 25: Scatter plot of job vacancy rate versus unemployment rate, sample January 2020 – June 2024, and a standard Beveridge curve (11) as in Figura and Waller (2022). *Source JOLTS, BLS. Authors' computation.*

number of separations and N_t as the number of employed workers. The number of unemployed workers entering into job relationships, given by $f_t u_t$, depends on the job-finding rate, defined as $f_t \equiv H_t / U_t$.

The job-finding rate can be related to the vacancy-to-unemployment ratio through the same matching function (4) as:

$$f_t = \frac{H_t}{U_t} = m_t \left(\frac{v_t}{u_t} \right)^{1-\eta}. \quad (10)$$

Using equation (10) to substitute for f_t in equation (9), and assuming that unemployment is at steady state, so that $u_{t+1} = u_t$, we obtain the benchmark Beveridge curve commonly found in the literature:

$$v_t = \left(\frac{(1 - u_t) s_t}{m_t u_t^\eta} \right)^{\frac{1}{1-\eta}}, \quad (11)$$

where m_t and s_t are allowed to vary over time, and we assume that the steady-state relationship holds period by period, as is common in the literature.

Figure 25 is our replication of the FW results. Their main input in the calculation is the value of the coefficient η , which we estimate to be 0.663 using their model, close to their assumed value of 0.7.³¹ Unlike BE and BDS, the curves do not pass directly through the data points; for example, the blue curve for the last observation in May 2024 is well to the right of the data point.³²

As the figure indicates, FW perform considerably better than BDS in tracking the drop in vacancies. However, their Beveridge curve does not align as cleanly with the actual observations as in BE, with the data points slightly to the left of the blue curve.³³ Yet, the overall fit conveys the same message: the adjustment is taking place almost exclusively through a drop in vacancies.

What accounts for the different findings? To address this question, it is useful to first summarize the differences in the three approaches, which are mainly: (i) FW assume s_t is exogenous, while BDS assume h_t is exogenous, and BE assume z_t is exogenous; (ii) the steady-state analysis of FW implies that their Beveridge curve does not necessarily pass through the data points, unlike BDS and BE; (iii) FW estimate m_t and s_t differently and use a different dataset, while BE and BDS use JOLTS data.

A key input in all these exercises is the elasticity of the matching function with respect to unemployment, η . Most importantly, the different methods used to estimate matching elasticities significantly impact the results. We assume the same value for η in both BE and our replication of BDS, with $\eta = 0.3155$, while for the replication of FW, we use $\eta = 0.663$.³⁴

Both of these estimates, however, fall well within the range established by various authors. Sahin et al. (2014), for example, cite various studies that provide estimates ranging from 0.28 to 0.72 and, after reviewing several studies, settle on a value of 0.5. Petrongolo and Pissarides (2001) report a similar range in their survey of the literature, highlighting the dependence on estimation methods. Therefore, before proceeding further, it is worth considering the results of BE, FW, and BDS using the same elasticity value, which is common in the literature.

Figure 26 shows the Beveridge curve implied by the three approaches, assuming a common value of

³¹The job finding rate f in equation (10) is approximated following Shimer (2005) as:

$$f_t = 1 - \frac{U_{t+1} - U_{t+1}^s}{U_t}$$

where U_t is the number of unemployed individuals and U_{t+1}^s is the number of short-term unemployed, i.e., workers unemployed for less than five weeks in month $t + 1$, as measured by the BLS. Thus, f_t is the probability that a worker unemployed in month t finds a job by $t + 1$. Figure 34 shows the job finding rate for the sample from December 2020 to June 2024. To obtain an estimate of η , we regress $\ln f_t$ on a constant and $\ln(v_t/u_t)$ in the sample from January 2009 to December 2019, obtaining a point estimate for η equal to 0.663. Note that this differs substantially from our estimate of 0.315. We then retrieve m_t using (10), which is shown in Figure 34 and differs from the one used in our analysis, as shown in Figure 23. To obtain the separation rate, following Ahn and Crane (2020), we use equation (9) to retrieve the sequence s_t , given the sequences of the unemployment rate and finding rate. This is shown in the bottom panel of Figure 34, with numbers in line with the literature.

³²Due to the construction of the finding rate, one observation is lost, so the last curve is plotted for May 2024 instead of June 2024.

³³Figure 35 in the Appendix neatly compares the three models using the same graph.

³⁴The reason for the different estimates is as follows: BDS, like us, consider gross hires, which include job-to-job transitions, whereas FW consider net hires, focusing only on transitions from the unemployment pool to employment. The latter approach is more suitable for using the matching function (4), while our approach would require more careful consideration of different matching functions for unemployed workers versus those transitioning from job-to-job. However, FW focus on new hires coming from the pool of unemployed, which constitutes only 20 percent of new hires, with the remainder approximately evenly divided between those joining the labor force and those making job-to-job transitions.

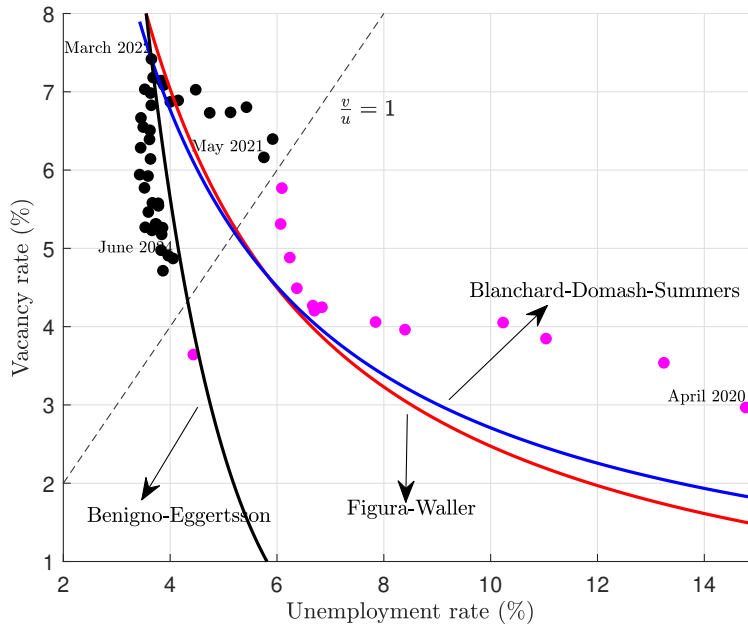


Figure 26: Comparisons between Beveridge curves of Benigno and Eggertsson (2023), Figura and Waller (2022) and Blanchard, Domash and Summers (2022). Scatter plot of job vacancy rate versus unemployment rate, sample January 2020 – June 2024. *Source: JOLTS, BLS. Authors' computation.*

$\eta = 0.5$.³⁵ As the figure reveals, the results from FW and BDS are essentially the same, suggesting a landing somewhere between the hard and soft predictions of the two.

Furthermore, as shown in Figure 36, if we swap the initially assumed elasticities of the two models, assigning $\eta = 0.3155$ to FW and $\eta = 0.663$ to BE and BDS, the results reverse: BDS predicts a soft landing, while FW predicts a hard landing. Interestingly, however, the BE model is relatively insensitive to this parameter and consistently predicts a soft landing.

Yet, this insensitivity may also be considered a weakness, as we have already noted. While the BE Beveridge curve is well-suited to describing a tight labor market under conditions of full employment, it is unable to account for the Beveridge curve when the labor market is loose. This is what we now turn to.

5 A generalized Beveridge curve and the BT unemployment rate

The appeal of the Beveridge curve we proposed in equation (6) lies in its ability to accurately predict behavior in a tight labor market. The curve captures the fact that when v/u crosses the Beveridge threshold, most of the adjustment occurs through a reduction in vacancies. However, when explaining periods in which the labor market is below the Beveridge threshold, equation (6) performs less

³⁵The time series for m_t is appropriately recomputed when considering a different value for the matching elasticity in the three models.

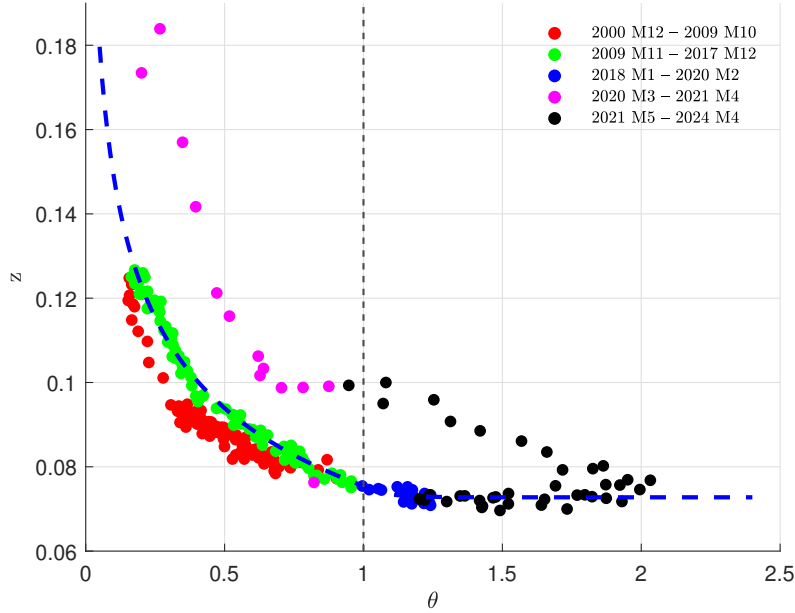


Figure 27: Scatter plot of fraction of unattached workers, z , as in Figure 23, and θ , job vacancy-to-unemployed ratio. Sample: December 2020 – June 2024. *Source: JOLTS, BLS. Authors' computation.*

well and requires constant adjustments to the exogenous variables to match the data. Our interpretation of this result is that while our assumption of an 'exogenous' fraction of 'unattached workers' is a reasonable approximation in a tight labor market, it is less effective when $v/u < 1$. Below, we outline a simple approach to address this issue.

In Figure 27, we present a scatter plot of θ , the vacancy-to-unemployed ratio, and the variable z , whose time series is shown in Figure 23. We use different colors to distinguish five periods in the sample. The red points, representing the pre-financial crisis period (December 2000 – September 2009), and the green points, representing the post-financial crisis period up to the end of 2017, are aligned on a curve specific to each period.³⁶ In contrast, the data where θ is greater than unity align well on a nearly horizontal line for the post-COVID period, particularly when the sample is further restricted starting from March 2022. This suggests a relationship where the number of unattached workers, or conversely, the number of attached workers, is insensitive to labor market conditions captured by θ when it is well above unity. However, in a loose labor market, a lower θ is associated with a higher fraction of unattached workers, likely due to higher unemployment. Based on this insight, we hypothesize a relationship of the following form:

$$z(\theta_t) = \begin{cases} (\bar{z}_t - a_h) + a_h \theta_t^{-b_h} & \theta_t \geq 1 \\ (\bar{z}_t - a_l) + a_l \theta_t^{-b_l} & \theta_t < 1 \end{cases} \quad (12)$$

for some non-negative parameters a_h , a_l , b_h and b_l and a variable \bar{z}_t that might shift over time. The

³⁶The blue points, denoting the periods from January 2018 to February 2020, also align well with the green points.

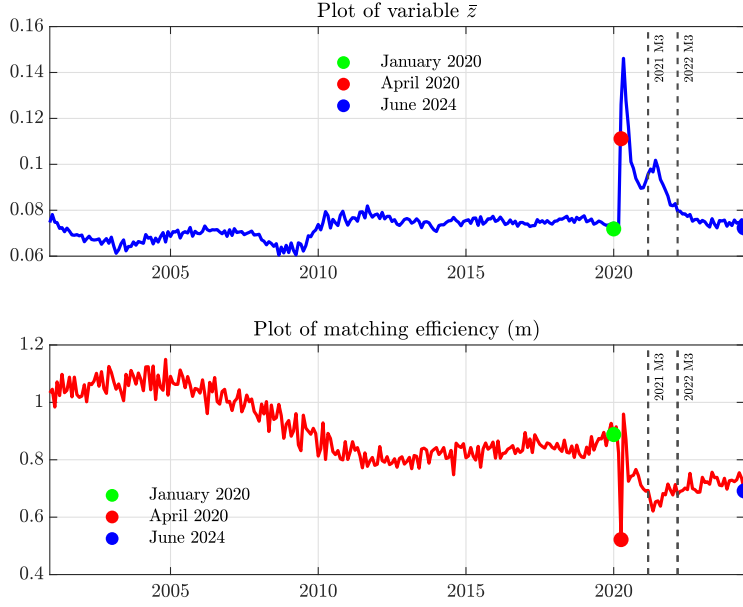


Figure 28: Plot of \bar{z} , as in equation (12), and of matching efficiency m . Sample: December 2020 – June 2024. Source: JOLTS, BLS. Authors' computation.

function is continuous at $\theta_t = 1$; moreover $b_h > b_l$ and $a_h < a_l$ capture the features of the data underlined in Figure 27. We estimate the parameters of the functional form (12) on the sample from September 2009 to January 2020, including the green and blue points, and that from March 2022 to June 2024, which looks more vertically aligned. We fix \bar{z}_t at the value that our original variable z_t had in April 2018, i.e., at 0.0754, when θ was close to the unitary value at 0.9956. The fit of the estimation is shown in Figure 27, demonstrating how well the proposed dashed-blue curve aligns with the green and blue points, as well as the black points, which represent data since March 2022. The curve also becomes asymptotic already at values of θ greater than 1.15.³⁷ We can then use the estimates of a_h , b_h , a_l and b_l , and the time series for z_t and θ_t to back up the primitive disturbance \bar{z}_t . This is plotted in the top panel of Figure 28. It is interesting to note the stability of \bar{z}_t before COVID-19, with an upward shift after the 2007-2009 financial crisis. Then, the extraordinary effect of the COVID-19 pandemic on the U.S. labor market is evident, with a sudden and unprecedented spike in \bar{z}_t and its return to pre-COVID values during the last couple of years.

We can now use this analysis to improve the Beveridge curve by substituting the function $z(\theta_t)$ for z_t in equation (6), obtaining it implicitly defined in:

$$0 = \begin{cases} m_t u_t^\eta v_t^{1-\eta} - \bar{z}_t - a_h (u_t^{b_h} v_t^{-b_h} - 1) + u_t & \theta_t \geq 1 \\ m_t u_t^\eta v_t^{1-\eta} - \bar{z}_t - a_l (u_t^{b_l} v_t^{-b_l} - 1) + u_t & \theta_t < 1 \end{cases} . \quad (13)$$

³⁷The estimates are the following: $a_h = 0.0027$, $b_h = 14.2464$, $a_l = 0.1140$, $b_l = 0.2168$.

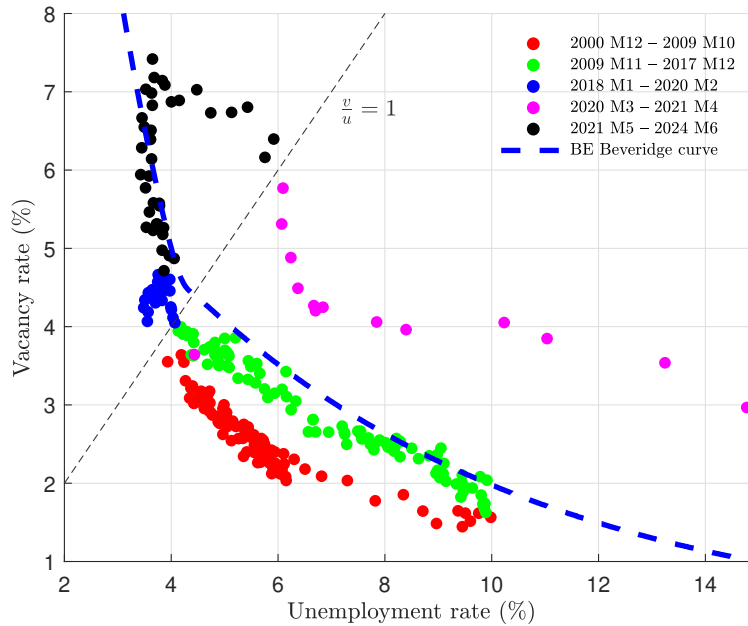


Figure 29: United States: scatter plot of job vacancy rate v and unemployment rate u ; plot of the Beveridge curve in equations (13). Sample: December 2020 – June 2024. *Source: JOLTS, BLS. Authors' computation.*

Figure 29 shows the resulting Beveridge curve. By allowing for an endogenous evolution of z_t , we see that it can match both the steepness of the curve during labor shortages and its slope below the Beveridge threshold.

Key to this derivation is the introduction of the function in equation (12). The interpretation of this function is straightforward. It appears to approximate the data relatively well by assuming that the number of unattached workers searching for jobs was exogenous under conditions of extreme labor tightness. As labor market conditions worsen (i.e., as θ decreases), the first-order effect would be an increase in z_t due to a larger number of unemployed individuals. With job prospects uncertain, it is also plausible that workers become more insecure about their current jobs and begin looking for alternatives, and that members of households who did not work before might try to re-enter the job market if a household member becomes unemployed. These are all mechanisms that are ripe for future research.³⁸

5.1 Beveridge-Threshold (BT) unemployment rate

In this section, we derive an index that we believe could be of practical interest to policymakers. It answers the following question: At what level of unemployment should the policymaker be concerned

³⁸For simplicity, we assumed that the $z(\theta_t)$ functional form had different parameters below and above the Beveridge threshold. However, this is not necessary. We have also considered more general invariant functional forms. One example that matches the data relatively well is a generalized Sigmoid function. We leave further exploration of this issue, as well as detailed microfoundations, to future research.

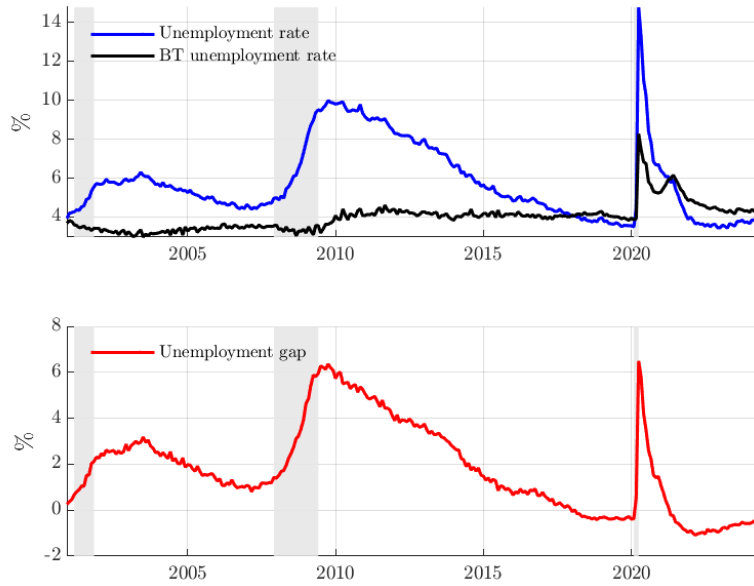


Figure 30: Top panel: unemployment rate, u , and Beveridge-Threshold (BT) unemployment rate, equation (14). Bottom panel: unemployment gap, defined as the difference between the unemployment rate and the BT unemployment rate. Sample: December 2000 – June 2024. *Source: JOLTS, BLS. Authors' computation.*

that inflation becomes excessively sensitive to demand and supply shocks?

By construction, the generalized Beveridge curve is continuous at the point $\theta_t = 1$, but it has a kink at that point. At the kink, the vacancy and unemployment rate are given by:

$$v(\theta_t = 1) = u(\theta_t = 1) = \frac{\bar{z}_t}{1 + m_t}. \quad (14)$$

The curve now has as primitive shifters the matching efficiency m_t and the variable \bar{z}_t , both plotted in Figure 28. An increase in \bar{z}_t or a fall in m_t both lead to an increase in the unemployment rate at which the economy lands when θ reaches unity. Given the latest values of \bar{z}_t and m_t in June 2024, the prediction is that unemployment will land at 4.42%. With the higher value of matching efficiency pre-COVID, it would have been at 3.96%. Understanding the shifters for the kink point is also important for our Phillips curve analysis, indicating the unemployment rate at which the curve becomes steep and when inflationary pressure from movements in demand or supply shocks could become more significant. We label this unemployment rate, the Beveridge-Threshold (BT) unemployment rate, which can be backed up using the sequences \bar{z}_t and m_t .

Figure 30 plots in the top panel the unemployment rate in comparison with the BT counterpart, while the bottom panel plots the unemployment gap given by their difference. It is interesting to observe

that the BT unemployment rate was consistently below 4% before the Great Financial Crisis, even reaching 3%. With the financial crisis, it increased above 4%, never crossing above 4.5%, and returned just below 4% before the COVID-19 pandemic. It is with the pandemic that this concept of the unemployment rate reached 8%, then gradually decreased to 4.42% at its latest June 2024 observation.

There are only two episodes during the sample in which the unemployment rate went below the Beveridge threshold, and they correspond to the periods in which the ratio v/u went higher than the unitary value. However, during the first episode in 2018 lasting until the pandemic, the gap did not go lower than -0.5%. Instead, the post-pandemic episode started in May 2021, with the gap suddenly widening to a maximum of -1.08% exactly in March 2022 at the peak of the inflationary surge. At that time, the unemployment rate was at 3.65%, while the Beveridge-threshold rate lagged behind at 4.73%. This gap is another way to read the pressure coming from the labor market to inflation, combined with the steep part of the Phillips curve. In contrast, the positive unemployment gap experienced for all the other parts of the sample, and evidently during the financial crisis and COVID pandemic, did not substantiate into deflation or very low inflation because of the flatness of the Phillips curve.

More work is surely needed to appropriately interpret the Beveridge-threshold unemployment rate we have been proposing here. However, it is tempting to relate it to the concept of the natural unemployment rate, capturing either maximum or potential employment levels.³⁹ The fact that it is derived in a more independent way from the unemployment rate makes it attractive and less prone to underestimation of the unemployment gap that affects most measures based on filtering the unemployment rate itself. The fact that it is rarely touched by the U.S. economy also aligns with other analyses, such as Gagnon and Sarsenbayev (2022), who have emphasized that advanced economies have mostly been running at an excessive unemployment rate, based on a conceptual framework with downward nominal wage rigidities, a key feature of our theoretical apparatus as well.⁴⁰

However, regardless of whether it is close to the natural rate of unemployment, it may still represent a useful instrument to advise policymakers on where inflationary pressures could start to become a significant concern. As of our latest data from June 2024, the gap is just 0.37% and is likely to have shrunk further in July after the latest unemployment rate reading of 4.3%. This is good news regarding the tightening actions undertaken by the Federal Reserve since the start of the hiking cycle.

Conversely, a positive gap is not necessarily a signal of an impending recession, although its sudden increase has been related to recessions in the U.S. economy, as evidenced by the grey areas.

This brings us to ask importantly what will happen on the other side of the gap to evaluate the risk of overshooting too much the Beveridge-unemployment threshold. To answer this issue, we come back to the Beveridge curve (13).

³⁹Our concept of the BT unemployment rate is closer to the efficient unemployment rate proposed by Michailat and Saez (2022) because both are identified by the Beveridge threshold, but it has a different theoretical foundation and quantitative implications.

⁴⁰Our BT unemployment rate is therefore different from the secular trend unemployment rate estimated by Crump et al. (2024), which is more symmetric with respect to the actual rate.

5.2 Landing beyond the Beveridge threshold

Using the latest values of \bar{z}_t and m_t , we can plot the Beveridge Curve shown in Figure 29, along with scatter plots of the vacancy and unemployment rates since the start of the JOLTS data in December 2000. The figure illustrates the flatter portion of the curve when the labor market is loose, i.e., $\theta < 1$, highlighting the potential costs in terms of higher unemployment in a weak labor market. The most recent reading of θ in June 2024 is 1.20, corresponding to an unemployment rate of 4.06% and a vacancy rate of 4.87%. Reducing the vacancy rate to 4.42%, thereby bringing θ to 1, would increase unemployment by only 0.36%. However, if the vacancy rate were to drop further, say to 4.04%, unemployment would rise to 5.05%, with θ falling to 0.8, as shown in the table below.

θ	u (%)	v (%)
1.20	4.06	4.87
1.10	4.19	4.61
1.00	4.42	4.42
0.95	4.56	4.33
0.90	4.71	4.24
0.80	5.05	4.04

Table 1: Corresponding u and v values for different θ values

This highlights the different slope of the Beveridge curve at the kink point and underscores the potential costs in terms of the unemployment rate if monetary policy contraction goes beyond what is necessary to bring the labor market to ‘neutral’ conditions. As of July 2024, the most recent unemployment rate is 4.3%, although vacancy data are not yet available. Assuming the stability of the latest Beveridge curve shown in Figure 29, we can infer that the vacancy rate should settle around 4.5%, resulting in $\theta = 1.046$, which is very close to the Beveridge threshold. Any further reduction in θ would come with higher unemployment costs: 4.42% at $\theta = 1$, 4.56% at $\theta = 0.95$, and 4.71% at $\theta = 0.90$.

As shown in Figure 29, the dashed-blue curve plotted with the latest data on \bar{z} and m aligns well with the black points since mid-2022 and a few of the green points, mostly following the 2007-2009 financial crisis. However, the curve also shifts with changes in \bar{z} and m . To shift the curve downward and better fit the green and blue points, matching efficiency would need to improve to pre-COVID levels. Achieving a good fit with all the red points would require matching efficiency similar to pre-financial crisis levels, while the purple points are largely explained by the unprecedented increase in \bar{z} during COVID.

6 Conclusion: policy implications

The bad news for policymakers emerging from our analysis of the Phillips curve is that at very low unemployment rates, they may encounter a steeper Phillips curve, which has the added disadvantage of amplifying the impact of supply shocks. The good news, however, is that as long as inflation expectations remain stable, the cost of reducing inflation in terms of increased unemployment is relatively low. The contrast between the inflation surge of the 1960s, which triggered a persistent change in inflation expectations, and the Federal Reserve's success today in maintaining stable inflation expectations is striking.

We conclude with two warnings. The first is that once the labor market crosses back over the Beveridge threshold so that $v/u < 1$, further reductions in inflation are likely to be more costly due to the flatness of the Phillips curve and the less steep Beveridge curve (and if the actual Beveridge threshold is higher, these costs will emerge even sooner). The second warning is that the Beveridge threshold may not yet have been reached. Much of the current reduction in inflation is in part due to the easing of supply shocks. With v/u still greater than 1, this suggests that adverse supply shocks could have significant effects on inflation.

References

- [1] Ahn, Hie Joo, and Leland Crane. 2020. Dynamic Beveridge Curve Accounting. Finance and Economics Discussion Series 2020-027, Board of Governors of the Federal Reserve System.
- [2] Ball, Laurence, Daniel Leigh and Prachi Mishra. 2022. Understanding U.S. Inflation During the COVID Era. Brookings Papers on Economic Activity.
- [3] Barlevy, Gadi, Jason R. Faberman, Bart Hobijn, Aysegul Sahin. 2023. The Shifting Reasons for Beveridge-Curve Shifts. NBER Working Paper 31783.
- [4] Barnichon, Regis. 2010. Building a Composite Help-Wanted Index. *Economics Letters* 109(3): 175-178.
- [5] Barnichon, Regis and Adam Hale Shapiro. 2022. What's the Best Measure of Economic Slack? FRBSF Economic Letter 2022-04.
- [6] Barnichon, Regis, and Adam Hale Shapiro. 2024. Phillips meets Beveridge. Federal Reserve Bank of San Francisco Working Paper 2024-22.
- [7] Benigno, Gianluca, Di Giovanni, Julen, Groen, Jan J., and Adam Noble. 2022. The GSCPI: A new barometer of global supply chain pressures. FRB of New York Staff Report, (1017).
- [8] Benigno, Pierpaolo, and Gauthi B. Eggertsson. 2023. It's baaack: The Surge in Inflation in the 2020s and the Return of the Non-linear Phillips Curve. NBER Working Paper, No. 31197.
- [9] Benigno, Pierpaolo, and Gauthi B. Eggertsson. 2024. The Slanted-L Phillips Curve. *AEA Papers and Proceedings*. Vol. 114.
- [10] Beveridge, William H. 1944. *Full Employment in a Free Society*. London, George Allen & Unwin, Ltd.
- [11] Blanchard, Olivier J., and Ben S. Bernanke. 2023. What Caused the US Pandemic-Era Inflation?. NBER Working Paper No. 31417.
- [12] Blanchard, Olivier, Domash, Alex and Lawrence Summers. 2022. Bad News for the Fed from the Beveridge Space. Peterson Institute for International Economics Policy, Policy Brief, 22-7.
- [13] Cattaneo, Mathias D., Crump, Richard K., Farrell, Max and Yingjie Feng. (2024). Nonlinear Binscatter Methods. ArXiv preprint arXiv:2407.15276.
- [14] Crump, Richard K., Stefano Eusepi, Marc Giannoni and Aysegul Sahin. 2024. The Unemployment-Inflation Trade-Off Revisited: The Phillips Curve in COVID Times. *Journal of Monetary Economics* 145.
- [15] Crump, Richard, Eusepi, Stefano Eusepi and Aysegul Sahin. 2024. Expectations and the final Mile of Disinflation. Liberty Street Blog, March 5th.
- [16] Domash, Alex, and Lawrence H. Summers. How Tight Are U.S. Labor Markets? NBER Working Paper No. 29739.

- [17] Dow, J. Christopher R., and Leslie Arthur Dicks-Mireaux. 1958. The Excess Demand For Labour A Study Of Conditions In Great Britain, 1946â56. *Oxford Economic Papers* 10.1: 1-33.
- [18] Eggertsson, Gauti B., and Donald Kohn. 2023. The Inflation Surge of the 2020s: The Role of Monetary Policy. Presentation at Hutchins Center, Brookings Institution.
- [19] Elsby, Michael, Ryan Michaels, and David Ratner. 2015. The Beveridge Curve: A Survey. *Journal of Economic Literature* 53: 571-630.
- [20] Erceg, Chris and Andrew Levin. 2003. Imperfect Credibility and Inflation Persistence. *Journal of monetary economics* 50(4): 915-944.
- [21] Figura, Andrew and Chris Waller. 2022. What does the Beveridge Curve Tell Us about the Likelihood of a Soft Landing? *FEDS Notes* July 29, 2022.
- [22] Forbes, Kristin J., Joseph E. Gagnon, and Christopher G. Collins. 2021. Low Inflation Bends the Phillips Curve around the World: Extended Results. Peterson Institute for International Economics, Working Paper No. 21-5.
- [23] Furman, Jason and Wilson Powell III. 2021. What's the Best Measure of Labor Market Tightness? Peterson Institute for International Economics, November 22, 2021.
- [24] Friedman, Milton and Anna J. Schwartz. 1963. *A Monetary History of the United States, 1867-1960*. Princeton University Press.
- [25] Gagnon, Joseph E. and Christopher G. Collins. 2022. 25 Years of Excess Unemployment in Advanced Economics. Peterson Institute for International Economics, Working Paper No. 22-17.
- [26] Goodfriend, Marvin and Robert G. King. 2005. The incredible Volcker disinflation. *Journal of Monetary Economics* 52(5): 981-1015.
- [27] Gitti, Giulia. 2023. Nonlinearities in the Regional Phillips Curve with Labor Market Tightness. Unpublished manuscript, Brown University.
- [28] Guerrieri, Veronica, Lorenzoni, Guido, Straub, Ludwig and Ivan Werning, 2021. Monetary Policy in Times of Structural Reallocation. Jackson Hole Economic Symposium.
- [29] Hazell, Jonathon, Juan Herreno, Emi Nakamura and Jon Steinsson. 2022. The Slope of the Phillips Curve: Evidence from U.S. States. *The Quarterly Journal of Economics* 137(3): 1299-1344.
- [30] Kalantzis, Yannick (@kalantzis). 2023. Why Does Everybody Now Talk About V/U and Hardly Anyone Mentioned it Before. Published online on June 28, 2023. Accessed 08/02/24: <https://x.com/kalantzis/status/1674097958298451980>
- [31] McLeay, Michael and Silvana Tenreyro. 2019. Optimal Inflation and the Identification of the Phillips Curve. *NBER Macroeconomics Annual* 34, 199-255.
- [32] Michaillat, Pascal, and Emmanuel Saez. 2022. $u^* = \sqrt{u\bar{v}}$. NBER Working Paper No. 30211.

- [33] Petrongolo, Barbara and Chris Pissarides. 2001. Looking into the Black Box: A Survey of the Matching Function. *Journal of Economic literature* 39(2): 390-431.
- [34] Petrosky-Nadeau, Nicolas and Lu Zhang. 2021. Unemployment Crises. *Journal of Monetary Economics* 117: 335-353.
- [35] Phillips, Alban W. 1958. The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957. *Economica* 25(100): 283–299.
- [36] Pissarides, Christopher A. 2000. *Equilibrium unemployment theory*. MIT press.
- [37] Sedlacek, Petr. 2016. The Aggregate Matching Function and Job Search from Employment and Out of the Labor Force. *Review of Economic Dynamics* 21: 16-28.
- [38] Sahin, Aysegul, Song, Joseph, Topa, Giorgio and Giovanni L. Violante, 2014. Mismatch Unemployment. *American Economic Review* 104(11): 3529-3564.

Appendix A

Here we reproduce benchmark regression from BE. It is an ordinary least squares regression:

$$\pi_t = \beta_c + \beta_\pi \pi_{t-1} + (\beta_\theta + \beta_{\theta_d} D_t) \ln \theta_t + (\beta_v + \beta_{v_d} D_t) v_t + \beta_{\pi^e} \pi_t^e + \varepsilon_t, \quad (\text{A.15})$$

where $\beta_c, \beta_\pi, \beta_\theta, \beta_{\theta_d}, \beta_v, \beta_{v_d}, \beta_{\pi^e}$ are parameters, and ε_t is a zero-mean normally-distributed error. D_t is a dummy variable that takes value one if $\theta_t \geq 1$. $\pi_t \equiv \ln P_t / \ln P_{t-1}$ is inflation, π_{t-1} is its one-quarter lag, $\ln \theta_t$ is the logarithm of the vacancy-to-unemployed ratio, v_t is a supply shock, and π_t^e is inflation expectations.

Table 1: Phillips Curve Estimates

	(1)	(2)	(3)	(4)
	1960-2024	2008-2024	1960-2024	2008-2024
<i>Inflation lag</i>	0.3707*** (0.0949)	0.2668 (0.2503)	0.2572*** (0.0933)	-0.1377 (0.1951)
$\ln \theta$	0.6748*** (0.1779)	0.7267* (0.3708)	0.2367 (0.1993)	0.5227 (0.3188)
$\theta \geq 1$			3.7165*** (0.8248)	5.3565*** (0.8936)
<i>v shock</i>	0.0377** (0.0192)	0.0177 (0.0393)	0.0446** (0.0204)	-0.0093 (0.023)
$\theta \geq 1$			0.1015 (0.0993)	0.275** (0.1212)
<i>Inflation expectations</i>	0.6596*** (0.1064)	0.8263 (0.6225)	0.8072*** (0.1016)	0.5091 (0.5048)
<i>Constant</i>	0.5559*** (0.1538)	0.9406** (0.4176)	0.1977 (0.1662)	0.3954 (0.3822)
R^2 adjusted	0.8139	0.5137	0.8264	0.6603
Observations	258	64	258	64

· ***, **, * denote statistical significance at the 1, 5, and 10 percent level, respectively.

· Newey-West standard errors.

· (1) and (3): sample 1960 Q1 – 2024 Q2

· (2) and (4): sample 2008 Q3 – 2024 Q2

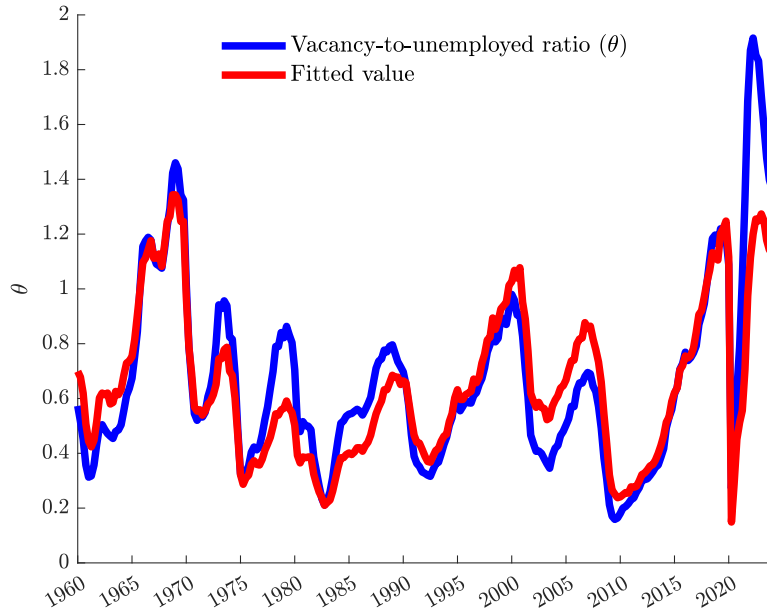


Figure 31: v/u relative to fitted value based upon the regression reported in Figure 7

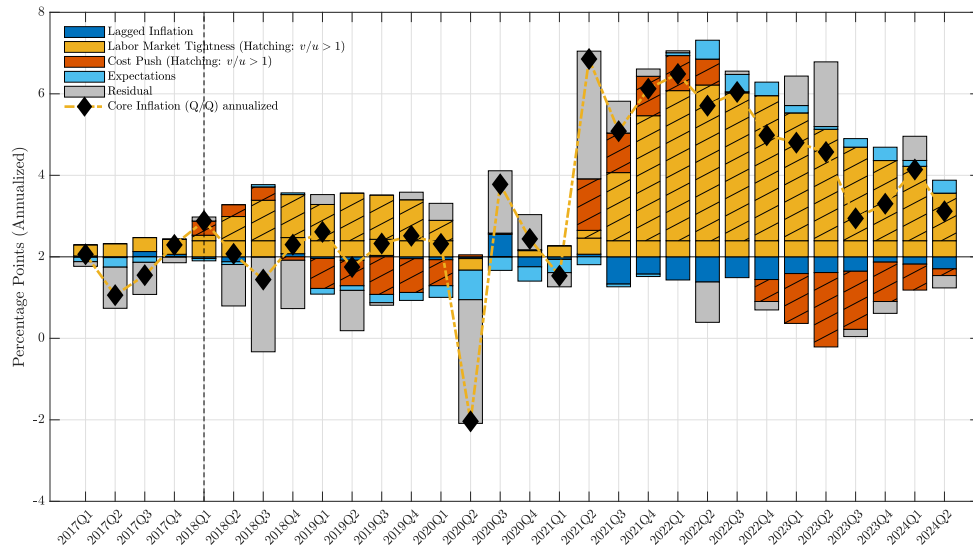


Figure 32: Decomposition of the baseline regression in BE between the contributions of the various regressors: lag inflation, $\ln \theta$, supply shocks, inflation expectations. For the variable $\ln \theta$, hatching corresponds to the contribution of the variable for the portion of θ that exceeds the unitary value. For the supply shock, hatching corresponds to the contributions of the variable when $\theta > 1$. Core inflation and all the components are plotted at annualized quarterly rates. *Source: Benigno and Eggertsson, 2023.*

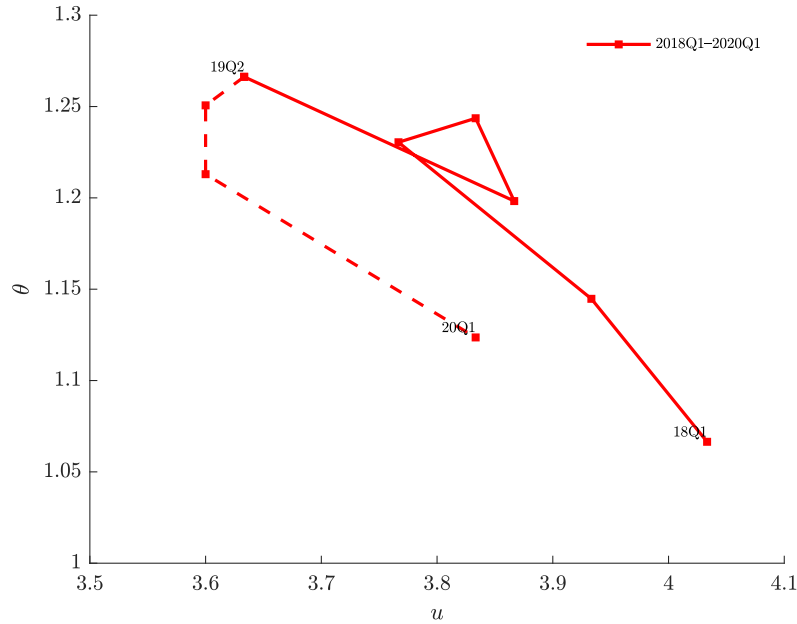


Figure 33: United States: scatter plots of vacancy-to-unemployed ratio, θ , and unemployment rate, u , at quarterly frequency. Period 2018 Q1 – 2020 Q1. Source: BLS.

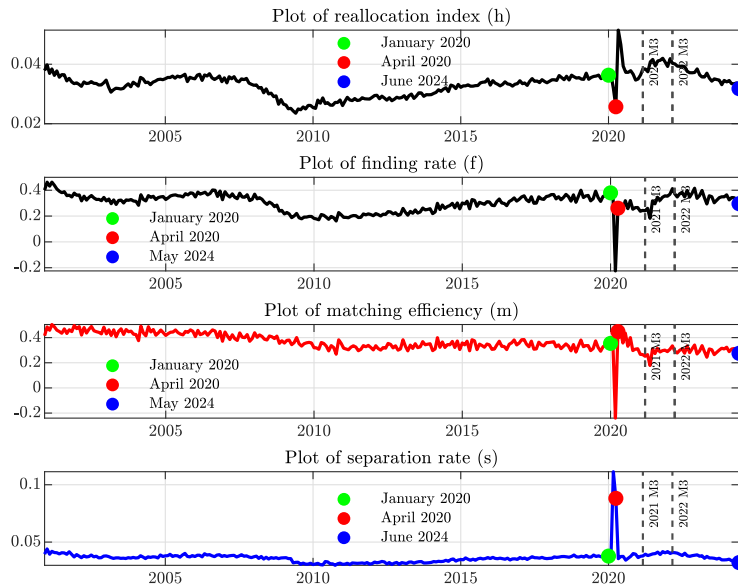


Figure 34: Blanchard, Domash and Summers (2022): reallocation index (h). Figura and Waller (2022): Job finding rate (f), matching efficiency (m), separation rate (s), sample December 2000 – June 2024. Source: JOLTS, BLS. Authors' computation.

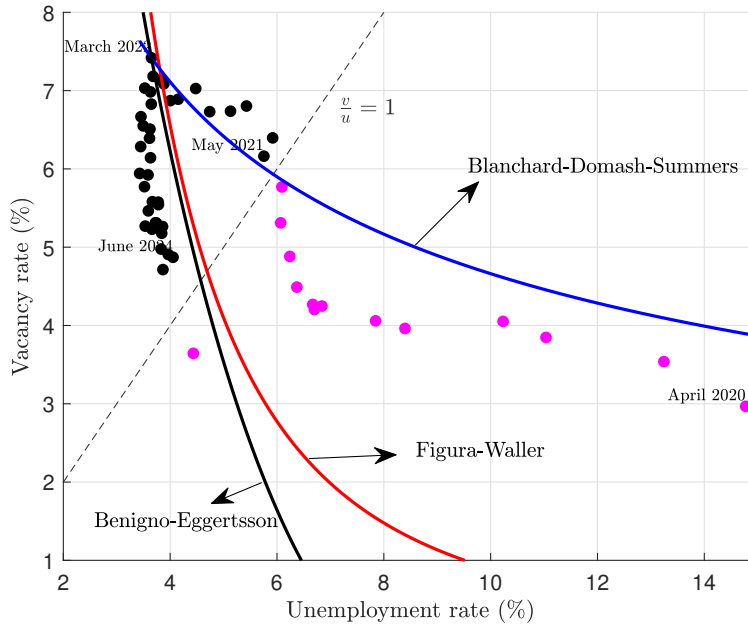


Figure 35: Comparisons between Beveridge curves. Scatter plot of job vacancy rate versus unemployment rate, sample January 2020 – June 2024. *Source: JOLTS, BLS. Authors' computation.*

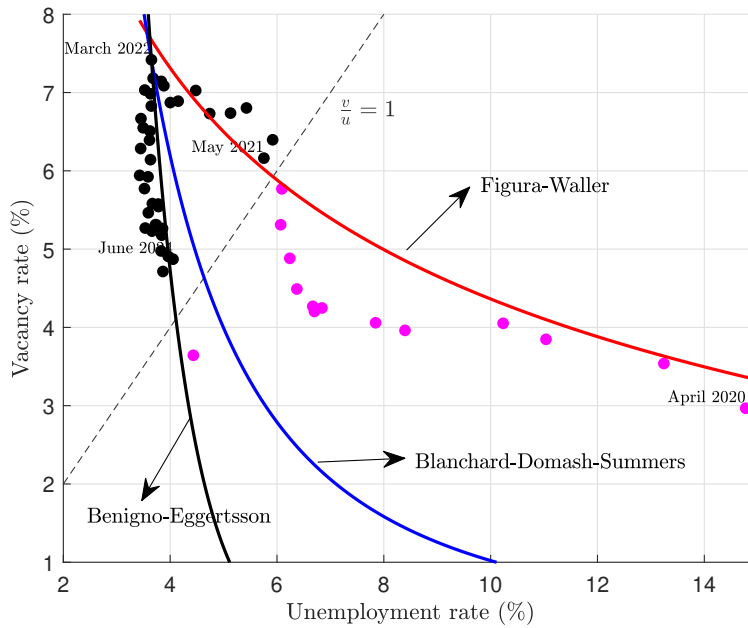


Figure 36: Swapping elasticities: comparisons between Beveridge curves. Scatter plot of job vacancy rate versus unemployment rate, sample January 2020 – June 2024. *Source: JOLTS, BLS. Authors' computation.*