FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Phillips Meets Beveridge

Regis Barnichon
Federal Reserve Bank of San Francisco
CEPR

Adam Hale Shapiro
Federal Reserve Bank of San Francisco

July 2024

Working Paper 2024-22

https://doi.org/10.24148/wp2024-22

Suggested citation:

Barnichon, Regis, and Adam Hale Shapiro. 2024. "Phillips meets Beveridge." Federal Reserve Bank of San Francisco Working Paper 2024-22. https://doi.org/10.24148/wp2024-22

Phillips meets Beveridge*

Régis Barnichon[(a)] *and* *Adam Hale Shapiro*[(b)]

[(a)] Federal Reserve Bank of San Francisco and CEPR

[(b)] Federal Reserve Bank of San Francisco

7/22/24

Abstract

The Phillips curve plays a central role in the macroeconomics literature. However, there is little consensus on the forcing variable that drives inflation in the model, i.e., on the appropriate measure of "slack" in the economy. In this work, we systematically assess the ability of variables commonly used in the literature to (i) predict and (ii) explain inflation fluctuations over time and across U.S. metropolitan areas. In particular, we exploit a newly constructed panel dataset with job openings and vacancy filling cost proxies covering 1982-2022. We find that the vacancy-unemployment (V/U) ratio and vacancy filling cost proxies outperform other slack measures, in particular the unemployment rate. Beveridge curve shifts—notably, movements in matching efficiency—are responsible for the superior performance of the V/U ratio over unemployment.

JEL classification: E3, J63, J64.

Keywords: Phillips curve, Beveridge curve, matching efficiency

*We thank Katharine Abraham and Larry Ball, as well as participants at the NBER conference on "Inflation in the COVID Era and Beyond," for helpful comments and suggestions. The views expressed in this paper are the sole responsibility of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

1 Introduction

The main framework to explain inflation dynamics, the Phillips curve, links inflation to the amount of unused capacity (or slack) in the economy. The underlying intuition is that, as the economy heats up, demand tends to exceed capacity, causing upward pressure on prices and thus higher inflation. As inflation remained remarkably stable throughout successive business cycles over 1990-2020, many economists considered the Phillips curve to be "dormant," but the recent surge in inflation led to a revival of Phillips curve studies, notably on the ability of a Phillips type framework to account for the ups and downs in inflation in the post-COVID recovery (e.g., Ball et al., 2022; Benigno and Eggertsson, 2023; Blanchard and Bernanke, 2023).

Since the original paper linking the unemployment rate to (wage) inflation (Phillips, 1958), a number of theoretical studies have focused on deriving foundations for the Phillips curve —the aggregate supply (AS) relationship of the macroeconomy. Phelps (1967) and Friedman (1968) emphasized the concept of the unemployment *gap*; the deviation of unemployment from its natural level, a concept later microfounded by the New-Keynesian literature (e.g., Blanchard and Galí, 2010; Galí, 2015). However, unemployment is by far not the only measure of slack that has been proposed. Other popular slack candidates include average real marginal cost, the labor share, the output gap (see, e.g., Galí, 2015), the job-switching rate (Moscarini and Postel-Vinay, 2017, 2023), and most recently the vacancy–unemployment ratio (Barnichon and Shapiro, 2022; Ball et al., 2022).

There is currently little consensus on the most appropriate measure of slack, or more specifically on the most appropriate *forcing variable* in a Phillips curve framework. That is, which variable can best explain the movements in inflation caused by changes in aggregate demand? We tackle this question with a dual approach. First, we conduct an out-of-sample forecasting exercise, which is robust to overfitting issues inherent to in-sample analysis. We assess which slack measures best predict inflation at one-year horizons. From the interwar period to the upsurge in inflation of the 1960s, all the way up to the more recent post-COVID inflation surge, a set of variables consistently provide superior information about future inflation: the vacancy–unemployment (V/U) ratio and, more generally, proxies for vacancy filling costs—firms' cost of filling a job opening.

Second, we aim to assess whether the structural Phillips curve—the causal effect of slack on inflation—fits the data better using the V/U ratio instead of the traditional unemployment rate as a measure of slack. OLS estimates again confirm the superior performance of the V/U ratio over unemployment, though coefficient estimates are likely biased by endogeneity issues: unobserved inflation expectations, unobserved natural rates, unobserved supply shocks, and downward bias from countercyclical monetary policy (e.g., McLeay and Tenreyro, 2020; Barnichon and Mesters, 2020).

To address these endogeneity issues, we take a three-pronged approach. First, we estimate the model on a narrower measure of inflation, the San Francsico Fed's "cyclical core" inflation measure, which is plausibly less contaminated by supply disturbances. Second, we use the Romer and Romer (2004) monetary shocks as instrumental variables. Third, we exploit Hazell et al. (2022)'s insight that cross-sectional information can address (or at least lessen) endogeneity biases, and we use a newly assembled panel of V/U data at the Metropolitan Statistical Area (MSA) level over 1980-2022.The V/U ratio performs well across these three models, explaining inflation better than the unemployment rate.

The superior performance of the V/U ratio may seem surprising since vacancies and unemployment are highly correlated: the so-called Beveridge curve. In fact, one can conjecture that the V/U ratio has been ignored by the earlier Phillips curve literature for this very reason. At times, however, this correlation can deteriorate sharply due to shifts in the Beveridge curve. We show that these Beveridge curve shifts are responsible for the superior performance of the V/U ratio in explaining inflation dynamics. The post-COVID outburst in inflation is an example of such a Beveridge curve shift, and the V/U ratio explains the rise in inflation much better than the unemployment rate alone.

A simple accounting framework shows how Beveridge curve shifts are related to changes in matching efficiency—the efficiency with which the labor market matches job openings to job seekers. These shifts are relatively rare, explaining why a Phillips curve with the V/U ratio typically performs just as well as a traditional Phillips curve with unemployment. At times however, matching efficiency can decline markedly—in the 2008-2009 recession, for instance, or most strikingly in the aftermath of the COVID pandemic– and these drops are associated with higher inflation.

2 The Phillips curve forcing variable

Our starting point is the Phillips curve, which is a formal statement of the intuition that an expanding economy will result in a tight labor market where firms compete for workers, see rising labor costs, and thus raise prices. A standard formulation of the Phillips curve is the New Keynesian equation:

$$\pi_t = \gamma E_t \pi_{t+1} + \kappa x_t + \nu_t, \tag{1}$$

where x_t is the relevant measure of "slack,"or more specifically the Phillips curve *forcing variable*, and ν_t captures cost-push shocks. The Phillips curve is a central equation in macroeconomics. Despite its importance, however, there is is much uncertainty about which measure of slack is e most relevant and thus is the forcing variable that best explains inflation.

Economic slack

The most used forcing variables in the literature are proxies for tightness in the labor market, typically the unemployment rate or unemployment gap (Phillips, 1958). However, a potential drawback of the unemployment rate is that it ignores workers outside of the labor force.[1] To

[1]If nonparticipants return to the labor force during times of strong economic growth, they could reduce upward wage and price pressures by increasing the supply of workers available. In this case, the unemployment rate would overstate inflationary pressures; see Hobijn and Şahin (2021).

address this limitation, Hornstein et al. (2014) proposed an extended concept of unused labor. Their Non-Employment Index (NEI) includes potential job seekers outside of the labor force.

More recently, the vacancy–unemployment ratio (or V/U ratio for short) has been proposed as a proxy for labor market tightness (Barnichon and Shapiro, 2022; Ball et al., 2022)—an idea going back to Medoff and Abraham (1982). Intuitively, the ratio $\theta_t = \frac{V_t}{U_t}$ represents the number of job vacancies, or demand for labor, relative to the number of unemployed individuals, or supply of labor. Since the unemployment pool misses job seekers outside the labor force and misses on-the-job job seekers, Abraham et al. (2020) constructed a generalized V/U ratio—denoted by θ^*—that replaces unemployment with a measure of *effective* job searchers S_t that takes into account all possible job seekers. Specifically, $\theta^*_t = \frac{V_t}{S_t}$, where S_t is a weighted sum of unemployment, nonparticipants and employed job seekers where each job seeker type is weighted by its relative average job finding rate.

Marginal hiring costs

Despite the widespread use of these slack measures, the New Keynesian literature has made clear that the key determinant of inflation is not slack per se, but instead firms' real marginal costs. To address this concept, Gali and Gertler (1999) and Galí (2015) proposed using the share of output going to labor compensation—the labor share—as a proxy for firms' marginal costs. While the labor share is straightforward and easy to construct, it measures the *average* cost of labor, which need not coincide with the *marginal* cost of labor. Interestingly, the V/U ratio has also been proposed in this context, building on the intuition it should proxy for one key determinant of firms' marginal labor costs: the cost of finding and hiring an additional worker.

To see that point more formally, consider a standard model with search frictions (Pissarides, 2000). In that model, a key determinant of firms' real marginal cost is the vacancy filling cost—the cost of hiring a marginal worker—(e.g., Krause and Lubik, 2007; Krause et al., 2008),[2] which is given by

$$\chi_t = \frac{c}{q_t} \qquad \text{where} \qquad q_t \equiv \frac{m_t}{V_t} \tag{2}$$

where c is the cost of posting a vacancy, V_t the number of vacancies, and q_t the vacancy filling rate—the rate at which firms are able to hire to fill job openings—which is given by the flow of new matches at instant t (m_t) divided by the number of posted vacancies (V_t). Intuitively, the vacancy filling cost is the vacancy posting cost times the expected duration of that open vacancy, $1/q_t$.

In a standard search and matching model (Mortensen and Pissarides, 1994), the vacancy filling rate can be related to the V/U ratio by means of the matching function. The matching function relates the flow of new hires to the stocks of vacancies and unemployment. For a

[2]Technically, the real marginal cost is the sum of two terms: (i) a standard term (unit labor cost) that depends on the match surplus (productivity) and the surplus splitting rule, and (ii) a correction term capturing the difference between current vacancy filling costs and next period vacancy filling costs. If these costs follow approximately an AR(1), the correction term is summarized by the current vacancy filling cost χ_t; see Krause et al. (2008) and Galí (2015).

constant returns-to-scale matching function, the vacancy filling rate is given by $q_t = \frac{m(U_t,V_t)}{V_t} = q_t(\theta_t)$ where $\theta_t = \frac{V_t}{U_t}$ is the vacancy unemployment ratio. If we postulate that the matching function is Cobb-Douglas, we can write $m_t = m_{0t} U_t^{\sigma} V_t^{1-\sigma}$ with m_{0t} matching efficiency,[3] and simplify the vacancy filling cost as $\chi_t = c\frac{\theta_t^{\sigma}}{m_{0t}}$ or

$$\hat{\chi}_t = \sigma\hat{\theta}_t - \hat{m}_{0t} \ , \tag{3}$$

where "hats" denote (log) deviations from the steady state. Thus, if matching efficiency is constant ($\hat{m}_{0t} = 0$), the (log) V/U ratio $\hat{\theta}_t$ is a proxy for vacancy fillings costs, confirming the earlier intuition that the V/U ratio could be a relevant forcing variable for the Phillips curve.

The previous discussion ignores that firms can also hire directly from (i) the pool of employed workers who may search on the job or (ii) the pool of nonparticipants. With transitions from employment into employment and nonparticipation into employment, firms' vacancy filling rate becomes $q_t = \frac{m_t}{V_t} = \frac{p_t^{ue}U_t + p_t^{ee}E_t + p_t^{ne}N_t}{V_t}$ where p_t^{ue} is the Unemployment-to-Employment transition rate , p_t^{ee} is the employment-to-employment transition rate, and p_t^{ne} is the nonparticipation-to-employment transition rate. Under random matching,[4] the number of effective searchers S_t is given by

$$S_t = U_t + \frac{p_t^{ee}}{p_t^{ue}}E_t + \frac{p_t^{ne}}{p_t^{ue}}N_t \tag{4}$$

and the (generalized) vacancy filling cost becomes $\chi_t^* = c\frac{\theta_t^*}{m_0 t}$ with $\theta^* = \frac{V_t}{S_t}$ or

$$\hat{\chi}_t^* = \sigma\hat{\theta}_t^* - \hat{m}_{0t} \ . \tag{5}$$

The vacancy filling cost is similar to the V/U ratio, but in this instance to the generalized V/U ratio proposed by Abraham et al. (2020).

Our last marginal hiring cost proxy comes from Moscarini and Postel-Vinay (2017), who recently argued that, in a Postel-Vinay and Robin (2002) model, worker's job-switching rate can proxy for marginal costs. Different from the traditional random search and matching model of Mortensen and Pissarides (1994), in Postel-Vinay and Robin (2002) the share of the match surplus going to labor depends on employed worker's ability to elicit outside offers. In that case the ratio $\frac{p_t^{ee}}{p_t^{ue}}$ isolates the actual bargaining positions of employed workers. We will thus also evaluate the performance of the relative job switching rate $\hat{ee}_t = \hat{p}_t^{ee} - \hat{p}_t^{ue}$.

Other costs

While the cost of labor can drive up price pressures, it is not the only input of production for businesses and thus not the only factor determining marginal costs. Raw materials, machines, and other types of capital infrastructure also play an important role. Similar to labor costs,

[3]The Cobb-Douglas matching function is used in most macro models with search and matching frictions (Pissarides, 2000). The matching efficiency term m_{0t} can be seen as the residual of the Cobb-Douglas matching function. Matching efficiency is akin to the Hicks-neutral productivity term in an aggregate production function.

[4]Under random matching, the search intensity of a group (relative to the search intensity of the unemployed) is given by the transition rate into employment for that group relative to the unemployed's transition rate into employment.

increases in the cost of these inputs may force businesses to raise prices for their products to stay profitable. In this context, we consider an additional measure of slack based on industrial production: the Federal Reserve Board's measure of capacity utilization, which measures the fraction of resources used to produce goods in manufacturing, mining, and electric and gas utilities.

3 Forecasting inflation

In this section, we assess the forecasting performances of the forcing variables discussed above. To generate forecasts, we estimate local projections as detailed in Jordà (2005) of the form

$$\pi_{t+h} = \gamma_l \pi_{t-1} + \lambda x_t + \eta_{t+h}$$

where x_t is a forcing variable.

We run a horse race between candidate forcing variables to assess which is the most accurate in forecasting price inflation. We estimate the model using a 10-year rolling window and we create one-year-ahead ($h = 4$) forecasts following the last date in each rolling sample.[5] For each measure, we then calculate the forecast errors, which are the differences between their predicted values of inflation and actual inflation values. We measure overall forecasting performance using the mean of the squared values of these forecasting errors.

We consider the following variables as measures of slack: the raw unemployment rate (u^{raw}), the unemployment rate excluding temporary layoffs (u) in order to remove the COVID-specific spike of 2020, the Non-Employment Index (NEI), the log of the V/U ratio ($\hat{\theta}_t$), the level of the V/U ratio (θ_t), the log of the generalized V/U ratio of Abraham et al. (2020) ($\hat{\theta}^*_t$), vacancy filling cost proxies (described in more detail below), the relative job switching rate of Postel-Vinay and Robin (2002) ($\hat{e}e_t$), capacity utilization as estimated by the Board of Governors of the Federal Reserve, the unemployment gap and the log of the output gap as estimated from the Congressional Budget Office (CBO). The main series are depicted in the Appendix.

Data

The underlying data are publicly available except for the long series on worker transition rates and data on job openings. To construct transition rates between employment, unemployment and non-participation over 1968-2023, we follow Shimer (2012) and use Current Population Survey (CPS) micro data. For the the employment-to-employment transition rate, we follow Fujita et al. (2024) and obtain a series over 1995-2023. For job openings, we follow Barnichon (2010) to obtain a vacancy proxy series over 1951-2023. The series is constructed by splicing the Conference Board print newspaper help-wanted index with the Conference Board online help-wanted index following the methodology described in Barnichon (2010).[6] For the

[5] Using two-year-ahead forecasts ($h = 8$) gives similar conclusions.

[6] See Abraham (1987) for a discussion of the strengths and weaknesses of the Conference Board help-wanted index based on print newspapers.

generalized measure of job seekers (S_t) we use the series constructed by Abraham et al. (2020) over 1994-2023.

Last, to construct proxies for the recruiting filling cost χ_t and χ_t^*, we proceed as follows. For χ_t, we exploit the fact that, under a constant returns to scale matching function, we have $q_t = p_t^{ue}/\theta_t$, so that we can obtain a vacancy filling cost proxy from $\hat{\chi}_t = \hat{\theta}_t - \hat{p}_t^{ue}$ over 1967-2023. For χ_t^*, we proceed similarly by exploiting the fact that $q_t = \frac{m_t}{V_t} = \frac{S_t}{V_t} p_t^{ue}$ from (4) and get a generalized vacancy filling cost proxy from $\hat{\chi}_t^* = \hat{\theta}_t^* - \hat{p}_t^{ue}$ over 1994-2023. We consider two sample periods: (i) 1995-2023 where we could study and compare the performances of the largest number of forcing variables, and (ii) 1960-2023, which allows us to explore forecasting performance during the previous episode of high inflation—the 1960s and 1970s.

1995-2023

Figure 1a plots the mean-squared forecast errors (MSEs) of the different slack measures in predicting core personal consumption expenditures (core PCE) price inflation one year ahead over 2005–2023. The forecast errors are expressed in percentage terms relative to the baseline performance of the unemployment rate; thus, they can be interpreted as indicating how much better or worse they perform relative to the unemployment rate.

First, results show that the V/U ratio, and more generally all the vacancy filling cost proxies (in green), perform better than the other measures. The (log) generalized V/U ratio does slightly better than the raw V/U ratio; we cannot separate the performances of the level vs the log of the V/U ratio. Second, the vacancy filling cost proxies outperform the V/U ratio, and the best performing proxy is the generalized vacancy filling cost χ_t^*. The superior performances of χ_t^* relative to the V/U ratio indicate that two factors beyond V/U are important to understand inflation fluctuations: (i) time-varying matching efficiency ($\hat{m}_{0t} \neq 0$)—recall that $\chi_t = \sigma\hat{\theta}_t - \hat{m}_{0t}$—and (ii) time-varying hiring rates outside the unemployment pool, as the generalized vacancy filling cost χ_t^* outperforms the vacancy filling cost χ_t, which does not allow for hiring out of employment and nonparticipation. The superior performances of the vacancy filling cost proxies are all the more remarkable given the larger measurement error in the transition rates.[7]

Last, note that none of the traditional measures—the output gap, the unemployment gap, the non-employment index, or the labor share—outperform the V/U ratio. In addition, neither the output gap nor the unemployment gap outperform the V/U ratio, even though the gaps are constructed *ex-post* by the CBO, taking into account the later behavior of inflation and other indicators.

To better understand performances over time, Figure 1b plots the 10-year average rolling MSEs for the different forcing variables over 2005-2023. Two things to notice. First, the reasons for the superior performances of the vacancy filling cost proxies over the V/U ratios are to be found during the post-COVID inflation outburst. In particular, their superior performances

[7]Flow-based measures like χ_t or χ_t^* are more noisy than stock-based measures like the V/U ratio. See Figure A3.

Figure 1: Forecasting performances, 1995-2023

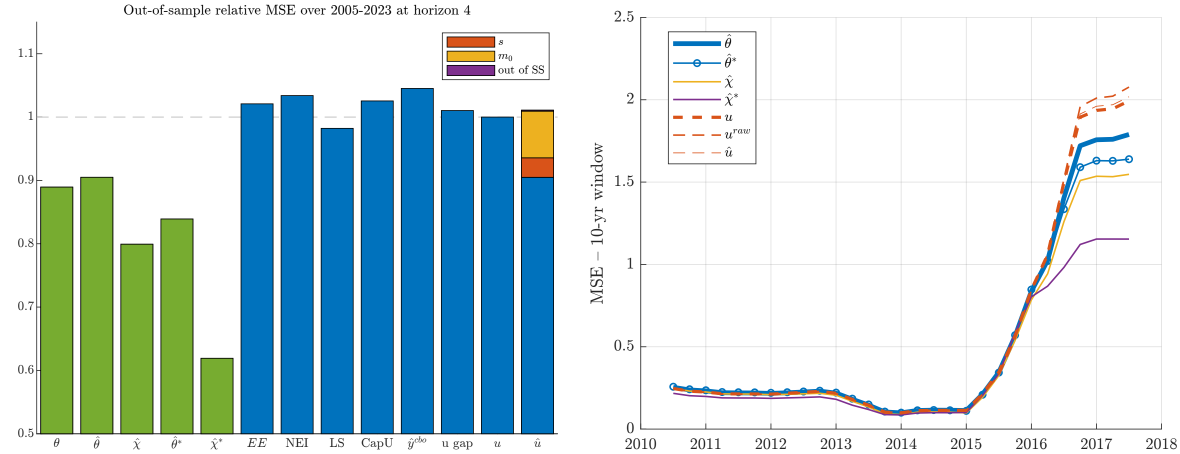

(a) Relative MSEs (b) 10-year Rolling MSEs

Notes: Inflation is measured as core PCE. Left panel: the mean-squared errors (MSE) are *relative* to the MSE of forecasts with the unemployment rate. $\hat{\theta}$ is the log V/U ratio, $\hat{\theta}^*$ is the generalized log V/U ratio, $\hat{u}$ is the log unemployment rate, $\hat{\chi}$ and $\hat{\chi}^*$ the vacancy filling cost proxies, EE the relative job switching rate, NEI is the Non-Employment Index, LS is the labor share, CapU is the Board of Governors capacity utilization rate, $\hat{y}^{cbo}$ and u gap are the output gap and unemployment gap estimated by the CBO, u^{raw} is the raw unemployment rate, u is the unemployment rate excluding temporary layoffs, and $\hat{u}$ is the log of u. The orange and red bars decompose the superior performances of $\hat{\theta}$ over $\hat{u}$ into the contribution of the Beveridge curve shifts due to movements in (i) the unemployment inflow rate s, (ii) matching efficiency (m_0) or (iii) out-of-steady-state dynamics ("out of SS").

may stem from a large decline in matching efficiency that raised firms' hiring costs since 2020, above what the behavior of the V/U ratio would suggest (see also Bagga et al., 2023). Second, the differences in performance are tiny during the stable inflation period of 2005-2019. It is only during the post-COVID recovery—during large inflation fluctuations—that the differences become noticeable. All forcing variable measures see large deterioration in forecasting performance, but the V/U ratios and the vacancy filling cost proxies ($\hat{\chi}_t$) do perform markedly better. This last observation indicates that large movements in inflation are necessary to discriminate between competing forcing variables. For these reasons, we next evaluate the performances of the V/U ratios and vacancy filling cost proxies during the large inflation movements of the 1960s and 1970s.

1960-2023

Figure 2 shows the same information as Figure 1 for a restricted set of slack measures over a longer sample period: 1960-2023, which allows us to include the buildup in inflation of the late '60s and the stagflation of the '70s. The available forcing variables are the log V/U ratio, the vacancy filling cost proxy $\hat{\chi}_t$, and the unemployment rate (in logs or level). Again, we find that the V/U ratio and the vacancy filling cost proxy outperform unemployment-based measures, notably by about 30 percent for the unemployment rate. Most importantly, this exercise confirms that the superior performances of 2022-2023 are not a unique occurrence: $\hat{\theta}_t$

Figure 2: Forecasting performances, 1968-2023

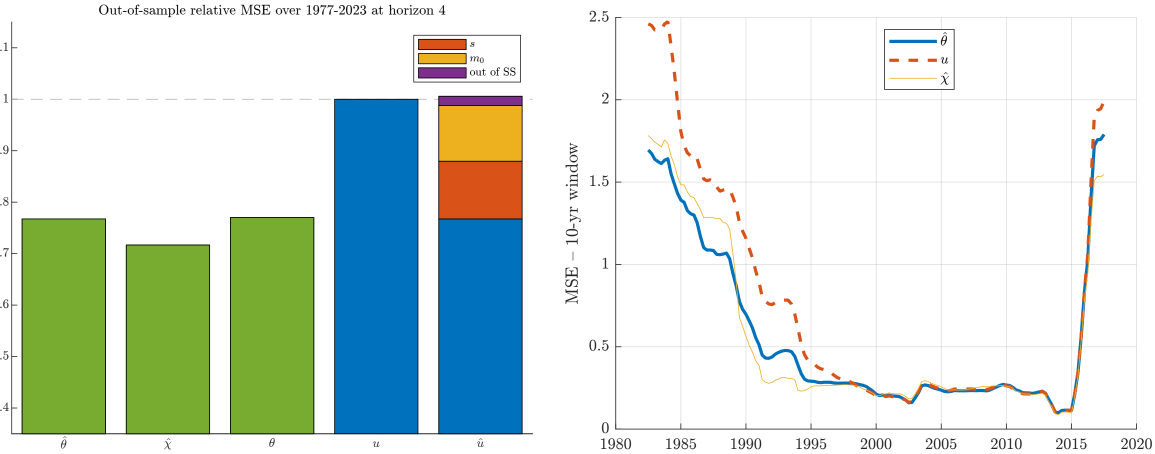

(a) Relative MSEs (b) 10-year Rolling MSEs

Notes: Inflation is measured from Core PCE. Left panel: the mean-squared errors (MSE) are *relative* to the MSE of forecasts with the unemployment rate. See Figure 1 for variable definitions.

(blue line) *systematically* outperforms u_t (red line) over 1960-2023.[8]

Longer samples

We find similar results using core CPI (Figure 3a) instead of core PCE over a slightly longer sample (1958-2023). In addition, we can extend the horse race between the V/U ratio and the unemployment rate to the period before World War II by incorporating vacancy data available from the NBER macro history database for "U.S. help-wanted advertising in newspapers, Metropolitan Life Insurance Company." Again, we find that the V/U ratio outperforms the unemployment rate in the interwar period (1919-1940); see Figure 3b.

Figure 3: Longer samples

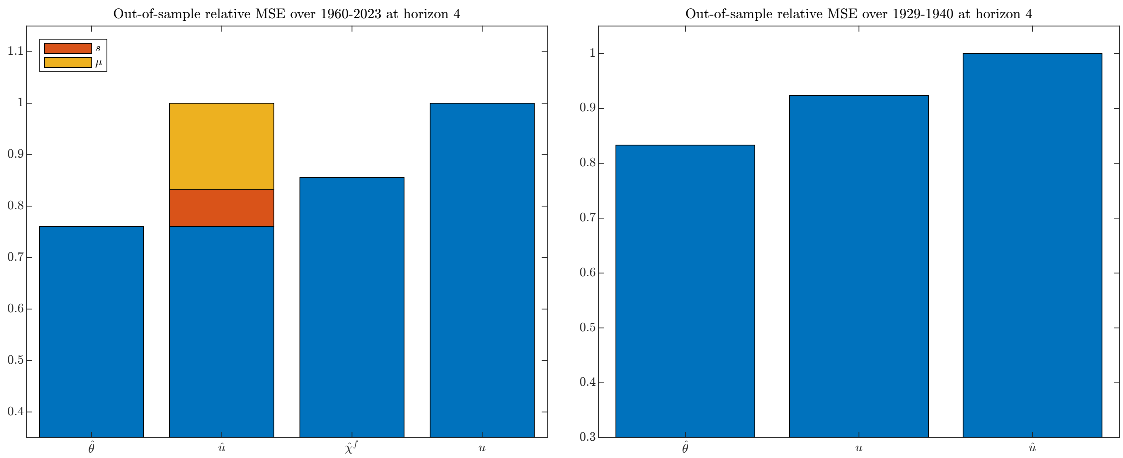

(a) 1958-2023 (b) 1919-1940

Notes: Inflation is measured as core CPI for panel (a) and headline CPI for interwar data in panel (b). See Figure 1 for variable definitions.

[8] Again, we note that it is difficult to separate competing forcing variables during periods of stable inflation (Figure 2b).

4 Estimating the structural Phillips curve

The discussion has so far focused on finding a measure that can best help forecast inflation, and we found that the V/U ratio or vacancy filling cost proxies were the most informative to predict inflation. A related but separate question is whether the superior prediction performance of the V/U ratio and vacancy filling cost capture a *structural* relation between hiring costs and inflation, or whether these variables simply proxy for other variables that are causing inflation.

This question is important for two reasons. First, it is hard to put much faith on superior forecasting performance alone without understanding the underlying reasons for such performance. If the V/U ratio only predicts better because it correlates with a variable that causes inflation over our evaluation sample, there is no guarantee that the superior performances continue to hold in other samples or in the future. This can be seen as an issue of external validity. In contrast, establishing that firms' hiring costs cause inflation is a more stringent test. Second, while a large literature has focused on consistently estimating a structural macro Phillips curve—the aggregate supply relationship of the economy—(Mavroeidis et al., 2014; Barnichon and Mesters, 2020), there is still much uncertainty about the appropriate forcing variable in such a Phillips curve. In this section, we therefore estimate and compare structural Phillips curves with different forcing variables: the V/U ratios, vacancy filling costs, and the unemployment rate.

4.1 Time series evidence

To estimate an aggregate Phillips curve, we use the representation of Hazell et al. (2022). Specifically, some manipulation of (1) gives

$$\pi_t = E_t\pi_\infty + \kappa E_t \sum_{j=0}^{\infty} \beta^j \left(\hat{x}_t - E_t\hat{x}_\infty\right) + \omega_t \ , \tag{6}$$

where $\hat{x}_t$ is a forcing variable (in deviation from the steady state), $E_t\hat{x}_\infty$ is its permanent component, and the residual captures all other transitory determinants of inflation beyond x_t.

If transitory fluctuations in the forcing variable follow an AR(1) with autocorrelation ρ, the expression simplifies to

$$\pi_t = \mathrm{E}_t\pi_\infty + \psi\left(\hat{x}_t - E_t\hat{x}_\infty\right) + \omega_t \tag{7}$$

where $\psi = \frac{\kappa}{1-\beta\rho}$.

We will thus estimate a Phillips curve of the form (7) with the regression (at quarterly frequency)

$$\pi_t = \alpha + \beta_x\hat{x}_{t-4} + \beta_\pi E_t\pi_\infty + \upsilon_t, \tag{8}$$

where inflation π_t is core PCE inflation, $E_t\pi_\infty$ is proxied with long-run inflation expectations taken from the Livingston survey, and the forcing variable $\hat{x}_t$ is either the unemployment rate u_t, the V/U ratio $\hat{\theta}_t$, or our vacancy filling cost proxies $\hat{\chi}_t$.

Table 1 reports the estimation results for the 1995-2023 sample period. All odd-numbered columns report "naive" ordinary least squares (OLS) estimates. Table 2 reports the same set of results for the 1960-2023 sample period. We z-scored the forcing variables (i..e, normalized them to have unit variance), so that the coefficients are directly comparable across columns—each coefficient capturing the effect of a one standard-deviation increase in the forcing variable on inflation. A larger coefficient thus indicates a larger explanatory power.

Confirming our out-of-sample prediction results, we can see that the V/U ratio outperforms the unemployment rate: the coefficient on $\hat{\theta}_t$ is more than 50 percent larger than the coefficient on u_t (column (3) vs column (1) in Table 1) and with a larger p-value. Similar results hold over 1960-2023 with a 22 percent larger coefficient on $\hat{\theta}_t$ (Table 2, columns (1) and (3)). In fact, over 1960-2023 the partial R^2—the R^2 of a regression where we first separated out the effect of π^e_∞—is twice as large using the V/U ratio than using unemployment alone. Further, the generalized V/U ratio of Abraham et al. (2020) does improve over the baseline V/U ratio, with a 30 percent higher coefficient than the baseline V/U ratio. The adjusted R^2 is also higher than with V/U—column (7) vs column (3). Last, and as in the out-of-sample forecasting exercise, the vacancy filling cost proxies do improve performance relative to their V/U ratio counterparts. The best performance overall is thus obtained with the generalized vacancy filling cost, though we caution that this is an in-sample result and could be prone to over fitting, especially given the nonstationary behavior of χ^*_t; see Figure A2 in Appendix.[9]

An important caveat of Tables 1 and 2 however is that these OLS coefficient estimates need not be informative about the structural Phillips curve (6), because OLS estimates could be biased by endogeneity issues. Indeed, the Phillips curve (6) postulates that inflation is determined by three main factors—expected future inflation, slack, and supply factors—all of which lead to endogeneity-related biases. Specifically, these biases reflect that (i) inflation expectations are measured with error, (ii) the long-run level of the forcing variable ($E_t\hat{x}_\infty$) is unobserved and (iii) supply shocks lead to confounding (see e.g., Barnichon and Mesters, 2020). An additional source of endogeneity bias is that of countercyclical policy: as the central bank works to mute the effects of aggregate shocks on inflation, the OLS estimate of the slope of the Phillips curve will be downward biased (McLeay and Tenreyro, 2020).

To address these endogeneity issues, we run three exercises. First, we estimate the Phillips curve on a narrower measure of inflation, the San Francisco Fed's "cyclical core PCE inflation" measure (Shapiro (2020)). This measure isolates those categories within the PCE price index that move systematically with the unemployment rate and are plausibly less contaminated by supply disturbances. This approach allows us to alleviate some of the endogeneity issues in the most recent sample period. Second, we use the Romer and Romer (2004) monetary shocks as instrumental variables in the Phillips curve regression, following Barnichon and Mesters (2020). While this approach will in principle address all endogeneity issues, the instrument is too weak in the post 1995 period—monetary shocks are small and rare during the Great Moderation (e.g., Ramey, 2016)—and can only be used over our longer sample. Third, we turn to MSA-level

[9]In the post-COVID recovery, the χ^*_t reached record highs, driven by a low relative employment-to-employment transition rate (which drives θ^*_t higher than θ_t) and a low matching efficiency.

data in order to estimate MSA-level Phillips curves, building on Hazell et al. (2022)'s insight that cross-sectional information makes it possible to address (or at least substantially lessen) these endogeneity biases.

Evidence using cyclical core PCE inflation

The San Francisco Fed cyclical core inflation measure is "trained" on data up to 2007, meaning each category (e.g., transportation services) is placed in the "cyclical" group based on its relationship with the unemployment rate between 1988 and 2007. For this reason, we estimate the Phillips curve on the cyclical inflation series between 2005 and 2023 so as to avoid any mechanical in-sample relationship. The results are shown in the even-numbered columns of Table 1. The coefficients are all larger in magnitude, and the fit of all models improves substantially. Again, $\hat{\theta}_t$ and $\hat{\theta}_t^*$ perform better than the unemployment rate, with the corresponding vacancy filling cost proxies performing even better.

Using monetary shocks as instrumental variables

The odd-numbered columns of Table 2 report the coefficients estimated using lags of monetary shocks as instrumental variables. The coefficients are bigger than using OLS—in line with a downward bias coming from supply shocks or systematic monetary policy—though the coefficients are now roughly of similar magnitudes: All forcing variables perform similarly. One reason could be that the instrumental variable (IV) estimator uses only a small share of the variation in the forcing variable—the fraction explained by monetary shocks—and there is no longer enough variation to discriminate among competing forcing variables. This can also be seen in the much larger standard errors. To discriminate among competing forcing variables and address the endogeneity issues, we will now exploit additional variation by estimating Phillips curves at the U.S. metropolitan level.

4.2 Evidence from U.S. Metropolitan Statistical Areas

Building on McLeay and Tenreyro (2020) and Hazell et al. (2022), we consider an MSA-level version of (7) with

$$\pi_{it} = \mathrm{E}_t \pi_{i\infty} + \psi \left(\hat{x}_{it} - E_t \hat{x}_{i\infty} \right) + \omega_{it} \ . \tag{9}$$

We exploit a new panel with information on labor market tightness over 17 MSAs between 1982 and 2022, estimating a panel regression of the form:

$$\pi_{i,t} = \psi \hat{x}_{i,t-4} + \delta_t + \alpha_{i0} + \alpha_{i1} t + \beta \mathbf{X}_{i,t-4} + \upsilon_{it}, \tag{10}$$

at the quarterly frequency where inflation, $\pi_{i,t}$, is measured as the four-quarter change in the core CPI in MSA i. Adding a cross-sectional dimension offers a number of advantages: (i) it allows for the inclusion of time fixed effects (δ_t) which control for time-specific factors common to all MSAs, such as unobserved inflation expectations (π_t^e), monetary policy, and global supply shocks; (ii) it includes MSA level fixed effects (α_{i0}), and MSA-specific time trends ($\alpha_{i1}t$) which control for unobserved natural MSA-level slack; and (iii) it considerably increases the effective sample size providing more variation in inflation and labor market slack. This is especially important to distinguish between competing labor market slack measures when differences are hard to detect. The vector $\mathbf{X}_{it}$ includes time-varying MSA-specific control variables, including lagged values of inflation and the relative price of goods and services.

Specification (10) alleviates many of the endogeneity issues discussed above. First, the time fixed effects control for movements in aggregate inflation expectations, movements in long-run marginal costs, aggregate supply shocks, and counter-cyclical monetary policy. Second, the inclusion of MSA fixed effects and MSA-specific linear trends allows for MSA-specific deviations of $\mathrm{E}_t\pi_{i,\infty}$ and $E_t\hat{x}_{i,\infty}$ from their aggregate counterparts as long as they follow a linear trend.[10]

Data construction

Our panel includes MSA-level data on unemployment, CPI inflation, and job openings between 1982 and 2022 for 17 MSAs. Unemployment data at the MSA level are available from the Bureau of Labor Statistics (BLS) Local Area Unemployment Statitics (LAUS), however, only back to 1990. To extend the sample back to 1982, we construct the MSA-level unemployment rate from CPS micro data, adjusting for the MSA redefinition in October 1985. Vacancy data are not readily available at the MSA level. However, three separate sources of information can be used to build a consistent time series for job openings at the MSA level over 1982-2022.

A first measure of vacancy postings is the Conference Board's Help-Wanted Index (HWI) available over 1951-2008. The HWI measures the number of help-wanted advertisements in 51 major newspapers. Since each newspaper advertises for the local job market, the Conference Board has also constructed an MSA-level HWI index over 1951-2008. Starting in the mid-1990s however, this "print" measure of vacancy postings became increasingly unrepresentative as advertising over the internet became more prevalent. A second measure of vacancy postings is online help-wanted advertising, which was published by the Conference Board spanning 2005-2010.

Building on Barnichon (2010), we combine these two series—"print" and "online" job advertising—to create a help-wanted index at the MSA level. A key variable in this exercise is the share of newspaper help-wanted advertising in total advertising. Since this print share is not directly observable, we model the development of online job advertising at each MSA level as the diffusion of a new technology —online job posting and job search—with a Mixed Information Source Model, which has been shown to successfully capture the diffusion of the internet in the U.S. population (e.g., Geroski, 2000). The appendix describes our data construction procedure in more details. Finally, our third source of vacancy data is from The Burning Glass Institute, which spans 2010 to 2022.

[10] This is an extension of McLeay and Tenreyro (2020) who posit that MSA-level deviations from aggregate inflation expectations are constant and can be controlled by region fixed effects. Our specification allows for time-varying $\mathrm{E}_t\pi_{it+\infty}$ and $E_t\hat{x}_{it+\infty}$ as long as they deviate "slowly" from their aggregate counterparts.

Results

Results of the MSA-level Phillips curve estimation are shown in Table 3, where we consider two forcing variables—the unemployment rate and the V/U ratio.[11] Estimates using the unemployment rate as the slack measure are shown in columns 1 and 2, while estimates using the log of the V/U ratio, $\hat{\theta}$, are shown in columns 3 and 4. Columns 5 and 6 report estimates with both the unemployment rate and the V/U ratio. We report models with no time or MSA fixed effects (columns 1,3, and 5) and including time and MSA fixed effects (columns 2, 4, and 6). Both forcing variables are normalized to a unit standard deviation for comparability. The inclusion of time and MSA fixed effects removes a great deal of upward bias on the unemployment rate and downward bias on $\hat{\theta}$. Both measures of slack are statistically significant, but again the coefficient on $\hat{\theta}$ is about 30 percent larger, with larger t-statistics, and the regression R^2 is higher with the V/U ratio as the forcing variable. Confirming our time series evidence, the MSA variation supports the V/U ratio as the better forcing variable over the unemployment rate. In fact, columns 5 and 6 show that the V/U ratio provides additional explanatory power over and above the unemployment rate. This indicates that the vacancy rate is providing additional information about inflation, which we expound on in the next section.

5 Phillips meets Beveridge

Based on our time series and MSA-level results, we conclude that the most successful specification for the Phillips curve is one with the V/U ratio ($\hat{\theta}_t$) or vacancy filling costs ($\hat{\chi}_t$) as the forcing variable. Notably, we find that $\hat{\theta}_t$ consistently outperforms the unemployment rate—the original forcing variable in the Phillips equation. This improvement can seem surprising in light of a well-known empirical regularity called the Beveridge curve: the existence of a tight relationship between vacancy postings and unemployment.

As illustrated in Figure 4, vacancy posting and unemployment rates comove negatively, and are highly correlated with a correlation of -.89 over the 1960-2023 period. An intriguing follow-up question is then the following: what additional information does the V/U ratio provide beyond the unemployment rate alone? To shed light on this issue, we dissect the theoretical underpinnings of the Beveridge curve—that is, the reasons underlying the high (but not perfect) correlation between unemployment and job openings. We will see that our results point to an important yet so far overlooked determinant of inflation: shifts in the Beveridge curve and, more specifically, changes in matching efficiency.

5.1 The Beveridge curve

To help understand the emergence of a Beveridge curve as well as the reasons behind its shifts, we consider a simple stock-flow accounting framework (e.g., Shimer, 2012) augmented with an aggregate matching function (e.g., Petrongolo and Pissarides, 2001).

[11] The sample sizes at the MSA level are too small to construct the worker transition rates needed (here p_t^{ue}) to recover the corresponding vacancy filling cost proxy.

Figure 4: The Beveridge curve

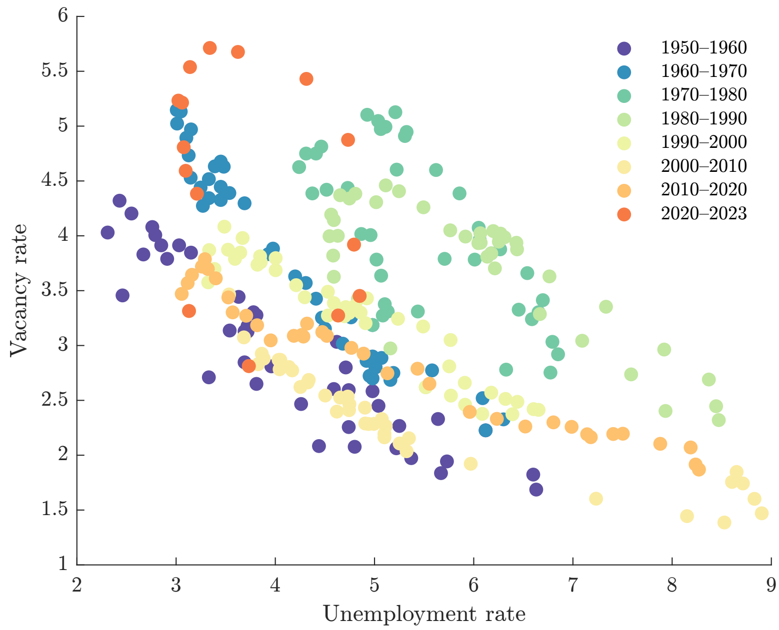

Steady-state unemployment

For ease of presentation, we consider the simpler case without on-the-job search and with constant labor force participation. In that case, the law of motion for unemployment is given by

$$u_{t+1} = s_t(1 - u_t) - f_t u_t + u_t \ , \tag{11}$$

where u_t is the unemployment rate, f_t the unemployment outflow probability, and s_t the unemployment inflow probability.

The steady-state unemployment rate is then defined as the "flow-consistent" unemployment rate—the unemployment rate that would ultimately prevail if the transition probabilities remained forever constant at their time-t values, f_t and s_t. In the U.S. labor market, the unemployment outflow probability is much larger than the unemployment inflow probability (by a factor of 10 or more), so that the steady-state unemployment can be approximated with

$$u_t^* \simeq \frac{s_t}{f_t} \tag{12}$$

and we can rewrite unemployment's law of motion as

$$u_{t+1} = f_t u_t^* + (1 - f_t) u_t \ . \tag{13}$$

The unemployment rate converges geometrically to its steady state at a rate $1 - f_t$, i.e., with a half-life of $\frac{-ln(2)}{ln(1-f_t)}$ or about two months for the US labor market with $f \approx 0.3$.[12] As argued by Shimer (2012), these large transition probabilities imply that the steady-state unemployment

[12]We measure f_t and s_t from short-term unemployment (see Shimer, 2012); see the Appendix for details.

rate provides a good approximation of the unemployment rate and $u_t \approx u_t^*$.

The shifting Beveridge curve

Using a Cobb-Douglas matching function, we can relate the flow of new hires to the stocks of vacancies and unemployment with $m_t = m_{0t} U_t^\sigma V_t^{1-\sigma}$. We obtain a very simple expression for the Beveridge curve with $u_t = \frac{s_t}{m_{0,t}\theta_t^{1-\sigma}}$ or

$$v_t = \delta_t u_t^{\frac{-\sigma}{1-\sigma}} \text{ where } \delta_t = \left(\frac{s_t}{m_{0t}}\right)^{\frac{1}{1-\sigma}} \tag{14}$$

where v_t is the number of vacancies relative to the size of the labor force. Expression (14) is a Beveridge curve: a log-log relationship between unemployment and vacancies, consistent with Figure 4.

Since out-of-steady state dynamics can play a role in times of large movements in unemployment, we generalize (14) and allow for out-of-steady state dynamics by taking a first-order approximation of unemployment around its steady state (Ahn and Crane, 2020) to obtain

$$\hat{v}_t = -\frac{\sigma}{1-\sigma}\hat{u}_t + \hat{\delta}_t \qquad \text{where} \qquad \hat{\delta}_t = \frac{1}{1-\sigma}\hat{s}_t - \frac{1}{1-\sigma}\hat{m}_{0t} - \frac{1}{f(1-\sigma)}\Delta\hat{u}_{t+1}, \tag{15}$$

where $\hat{\delta}_t$ captures shifts in the Beveridge curve and $\Delta\hat{u}_{t+1} = \hat{u}_{t+1} - \hat{u}_t$.[13]

Expression (15) highlights two important points. First, without shifts in the Beveridge curve, the V/U ratio and the (log) unemployment rate should provide the same information about future inflation. With $\hat{\delta}_t = 0$, we have $\hat{\theta}_t = -\frac{1}{1-\sigma}\hat{u}_t$, and the V/U ratio and (log) unemployment ($\hat{u}_t$) are perfectly collinear, and a horse race between $\hat{\theta}_t$ and $\hat{u}_t$ should be indeterminate.[14] In other words, a finding that the V/U ratio provides superior information about future inflation shows that Beveridge curve shifts are central to understanding inflation fluctuations—Phillips meets Beveridge.

Second, the Beveridge curve can shift for different reasons: (i) movements in the unemployment inflow rate ($\hat{s}_t$), (ii) movements in matching efficiency ($\hat{m}_{0t}$), and (iii) out-of-steady state dynamics $\Delta\hat{u}_{t+1} \neq 0$. Out-of-steady state dynamics correspond to the familiar "looping" of the Beveridge curve caused by the lagging response of unemployment to changes in vacancy postings (e.g., Blanchard and Diamond, 1989).

To measure the Beveridge curve shifts ($\hat{\delta}_t$), we run the regression $\hat{\theta}_t = \beta_u \hat{u}_t + e_t$ and take $\hat{\delta}_t$ as the regression residual. Similarly, to measure $\hat{m}_{0t}$—movements in matching efficiency—we run the regression

$$\hat{\theta}_t = \beta_u \hat{u}_t + \beta_s \hat{s}_t + \beta_o \Delta\hat{u}_{t+1} + e_t \ , \tag{16}$$

[13] Combine the matching function with unemployment's law of motion to get $v_t = \left(\frac{s_t(1-u_t-\Delta u_{t+1})}{m_{0t}u_t^\sigma}\right)^{\frac{1}{1-\sigma}}$ and take a first-order approximation with $(u_t, s_t, f_t, \Delta u_{t+1})$ around $(\bar{s}, \bar{f}, \bar{s}, \bar{f}, 0)$.

[14] In fact, it would likely favor the unemployment rate since the V/U ratio is more prone to measurement error given that job openings are only measured through a proxy from newspaper advertising.

and we take $\hat{m}_{0t}$ as the regression residual. See Figure A1 in the Appendix for a plot of our estimated matching efficiency series.

Figure 5 plots the time series for Beveridge curve shifts (in ppt of (log) vacancy rate) decomposed into the contribution of the unemployment inflow rate, out-of-steady-state dynamics and matching efficiency. While out-of-steady-state dynamics only play a small role, we can see that both the unemployment rate and matching efficiency lead to substantial shifts in the Beveridge curve. Most recently, these two forces partially compensated each other, with a lower job separation rate shifting the Beveridge curve inward and a deteriorating matching efficiency shifting the Beveridge curve outwards.

Figure 5: Beveridge curve shifts and matching efficiency

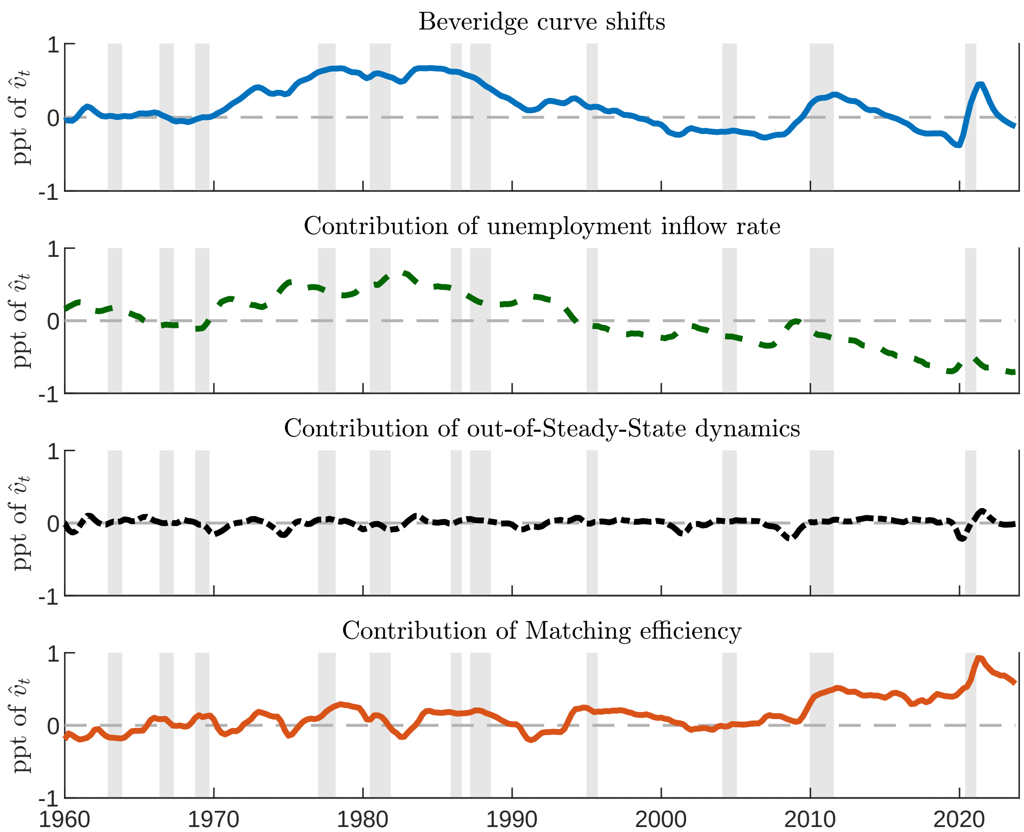

Notes: The Beveridge curve is estimated over 1951-2007. The blue line (Beveridge curve shifts) is the sum of the dashed green line (unemployment inflow rate), the dashed black line (out of steady state dynamics) and the red line (matching efficiency movements).

A number of factors can generate aggregate movements in matching efficiency: changes in workers' search intensity, changes in firms' recruiting intensity (Davis et al., 2013), changes in the composition of the unemployment pool (Barnichon and Figura (2015)), or changes in dispersion across labor markets (Barnichon and Figura (2015)) or mismatch (Şahin et al., 2014). While the pre-2007 cyclical pattern of matching efficiency has been attributed to changes in the composition of the unemployment pool —notably the share of long-term unemployed (see Barnichon and Figura, 2015)[15], matching efficiency has declined markedly since the end

[15]The long-term unemployed have a lower job finding rate than the short-term unemployed. In the early

of the financial crisis. The phenomenon worsened following COVID and the so-called *Great Resignation* (e.g., Barlevy et al., 2024). Overall however, much remains to be learned about the determinants of matching efficiency, and an important lesson from our study is that the behavior of matching efficiency is an important topic that extends beyond labor market studies: it has direct implications for our understanding of inflation, as we discuss next.

5.2 Beveridge curve shifts and inflation

To better understand how shifts in the Beveridge curve are important for inflation fluctuations, consider our baseline Phillips curve with V/U as the forcing variable:

$$\pi_t = \mathrm{E}_t\pi_\infty + \kappa\hat{\theta}_t + e_t \ . \tag{17}$$

From the Beveridge curve decomposition (16), we can split the effect of the V/U ratio on inflation into two components: (i) movements along the Beveridge curve, and (ii) shifts in the Beveridge curve. Indeed, from (15) with $\hat{\theta} = \hat{v}_t - \hat{u}_t$ and ignoring the small out-of-steady state dynamics for ease of exposition, we get

$$\hat{\theta}_t = \underbrace{-\frac{1}{1-\sigma}\hat{u}_t}_{\text{Mvts along BC}} + \underbrace{\frac{1}{1-\sigma}\hat{s}_t - \frac{1}{1-\sigma}\hat{m}_{0t}}_{\text{Shifts in BC}}, \tag{18}$$

so that the Phillips curve becomes

$$\pi_t = \mathrm{E}_t\pi_\infty \underbrace{-\frac{\kappa}{1-\sigma}\hat{u}_t}_{\text{Mvts along BC}} + \underbrace{\frac{\kappa}{1-\sigma}\hat{s}_t - \frac{\kappa}{1-\sigma}\hat{m}_{0t}}_{\text{Shifts in BC}} + e_t \ . \tag{19}$$

Expression (19) shows that the superior performance of the V/U ratio over unemployment comes from the shifts in the Beveridge curve.

Moreover, since $\hat{m}_{0t}$ is orthogonal to $\hat{s}_t$ by construction, we can separate the roles of the separation rate and matching efficiency in explaining inflation. Specifically, from (18) we can sequentially turn off a shifter of the Beveridge curve—the separation rate or matching efficiency—and study the performance of a counterfactual V/U ratio under either (i) a constant separation rate ($\hat{s}_t = 0$), or (ii) a constant matching efficiency ($\hat{m}_{0t} = 0$).[16] With these counter-factual V/U ratios in hand, we then replicate our two previous exercises —out-of-sample forecasting and in-sample Phillips curve estimation—, to isolate the contributions of each Beveridge curve shifters.

stages of recessions, bursts of layoffs tilt the pool of unemployed towards short-term unemployed and this raises matching efficiency: the aggregate job finding rate is higher than it would be given the level of the V/U ratio.

[16]Specifically, for (i) we evaluate the performance of a counterfactual V/U ratio $\tilde{\theta}_t^{(i)} \equiv (\hat{\beta}_u - 1)\hat{u}_t - e_t$, where $\hat{\beta}_u$ and e_t are estimated from (16). This counterfactual V/U ratio holds matching efficiency fixed and thus makes it possible to assess the contribution of the time-varying separation rate to the performances of $\hat{\theta}_t$ over $\hat{u}_t$. We proceed similarly to get the contribution of matching efficiency with the counterfactual V/U ratio $\tilde{\theta}_t^{(ii)} \equiv (\hat{\beta}_u - 1)\hat{u}_t - \beta_s\hat{s}_t$.

Figures 1 and 2 (orange and yellow bars) decompose the superior forecasting performances of $\hat{\theta}_t$ over $\hat{u}_t$ into the respective contributions of $\hat{s}_t$, $\hat{m}_{0t}$ and out-of-steady-state dynamics. In both cases, we can see that a large contributor to the superior forecasting performance of the V/U ratio comes from movements in matching efficiency.

Next, we estimate the more general Phillips curve specification implied by (19) and measure the importance of Beveridge curve shifts by running the regressions

$$\pi_t = \beta_\pi \mathrm{E}_t \pi_\infty - \beta_u \hat{u}_t + \underbrace{\beta_\delta \hat{\delta}_t}_{\text{Contribution of BC shifts}} + e_t$$

and

$$\pi_t = \beta_\pi \mathrm{E}_t \pi_\infty - \beta_u \hat{u}_t + \underbrace{\beta_s \hat{s}_t}_{\text{Contribution of } s_t} + \underbrace{\beta_{m_0} \hat{m}_{0t}}_{\text{Contribution of } m_{0t}} + e_t.$$

Table 4 confirms the importance of Beveridge curve shifts in column (3). The coefficient on $\hat{\delta}_t$ is significant: outward shifts in the Beveridge curve correlate strongly with rises in inflation. We can run a similar analysis using the MSA-level data. Specifically, to construct MSA-specific Beveridge curve shifts $\hat{\delta}_{it}$ we run MSA-level regressions of the log vacancy rate on the log unemployment rate including time and MSA fixed effects, along with MSA-specific trends. Table 5 shows results of this exercise on the MSA-level Phillips curve estimates. Column (1) includes the log of the unemployment rate alone, while column (3) includes MSA-level shifts in the Beveridge curve. The results show again that shifts in the MSA Beveridge curve ($\hat{\delta}_{it}$) are statistically significant, i.e., that MSA-level Beveridge curve shifts contain information about future inflation, over and above the inclusion of $\hat{u}$.

Last, column (4) of Table 4 shows that, consistent with (5.2), both the job separation rate and matching efficiency correlate strongly with inflation. While an exploration of the sources of the decline in matching efficiency is outside the scope of this paper, one lesson of our study is that the behavior of matching efficiency is an important topic that extends beyond labor market studies: it has direct implications for our understanding of inflation.

In fact, to get a sense of the contribution of Beveridge curve shifts to past inflation, Figure 6 repeats Figure 5 but where we converted the Beveridge curve shifts into units of additional inflation, using the estimate from Table 4, column (3). We can see that Beveridge curve shifts explain about 1 to 2 ppt of inflation fluctuations since 1960.

Figure 6: Beveridge curve shifts and matching efficiency

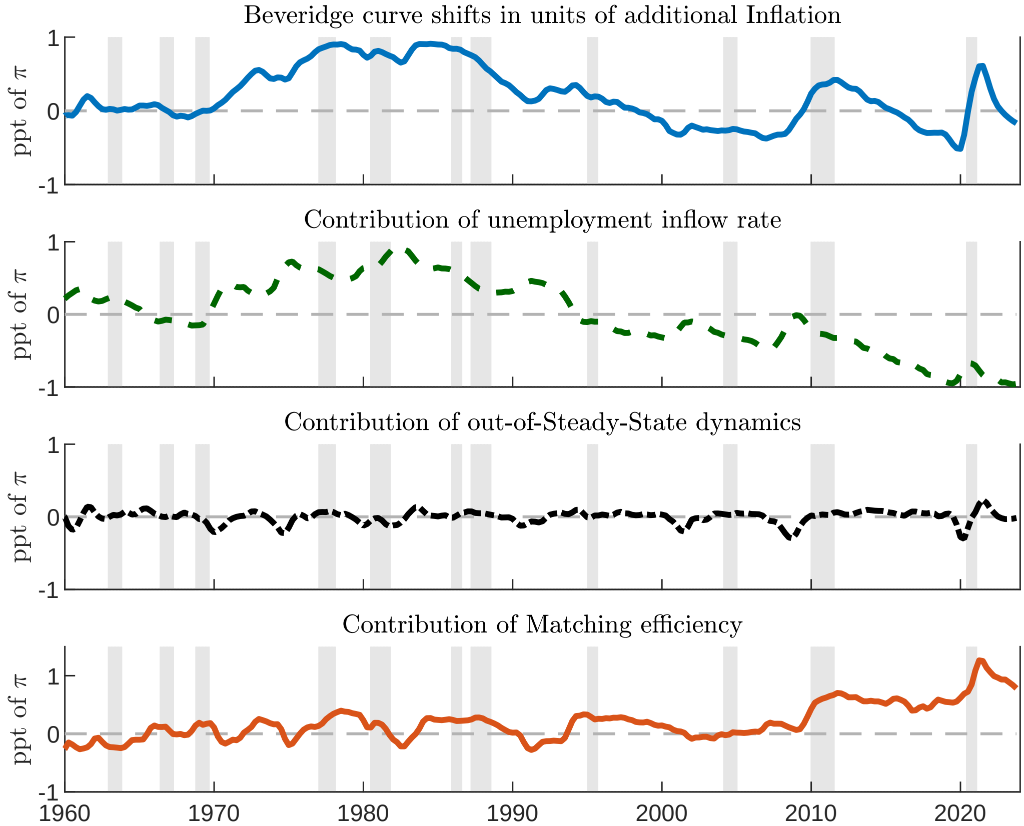

Notes: The Beveridge curve is estimated over 1951-2007. The blue line (Beveridge curve shifts) is the sum of the dashed green line (unemployment inflow rate), the dashed black line (out of steady state dynamics) and the red line (matching efficiency movements).

6 Nonlinear effects of slack on inflation

In light of the post-COVID surge in inflation, a number of recent works have argued that slack has a nonlinear effect on inflation; that the Phillips curve can steepen substantially in tight labor markets (Benigno and Eggertsson, 2023; Crust et al., 2023; Gitti, 2024). Using our Beveridge curve decomposition (18), we can explore the sources of that nonlinearity. Specifically, for our baseline Phillips curve with the log V/U ratio as the forcing variable, we have

$$\begin{aligned}\pi_t &= \mathrm{E}_t\pi_\infty + \beta_\theta(\theta_t)\hat{\theta}_t + \omega_t \\ &= \mathrm{E}_t\pi_\infty + \beta_u(\theta_t)\hat{u}_t + \beta_\delta(\theta_t)\hat{\delta}_t + \omega_t \end{aligned} \tag{20}$$

Clearly, if β_θ—the slope of the Phillips curve—depends on $\hat{\theta}_t$, then the Phillips curve is non-linear—a genuine nonlinearity. However, equation (20) suggests another possibility: that the Beveridge curve shifts (δ_t) systematically in tight labor markets; when the V/U ratio is high.[17] In that case, the Phillips curve can *appear* non-linear: in tight labor markets, systematic Beveridge curve shifts would move inflation (above and beyond $\hat{u}_t$), and give the impression of

[17]This could happen if the matching function is not exactly Cobb-Douglas: for instance, if matching efficiency declines systematically in tight labor markets.

a non-linear Phillips curve. To test between these different possibilities, we run the regression:

$$\pi_t = \beta_\pi \mathrm{E}_t \pi_\infty + \beta_u \hat{u}_t + \gamma_u \hat{u}_t \mathbb{1}_{\theta_t > \bar{\theta}} + \beta_\delta \hat{\delta}_t + \gamma_\delta \hat{\delta}_t \mathbb{1}_{\theta_t > \bar{\theta}} + \omega_t \ , \tag{21}$$

using as threshold variable $\bar{\theta}$ the median of θ_t.[18]

Table 6, column (2) confirms the presence of nonlinearities in the effect of the V/U ratio on inflation. Interestingly however, Table 6 column (3) shows that the nonlinearity appears to stem from systematic shifts in the Beveridge curve: the nonlinearity is entirely explained by $\hat{\delta}_t$. In tighter labor markets ($\theta_t > \bar{\theta}$), outward Beveridge curve shifts "raise" inflation but the converse is not true in slack labor markets.

Figure 7: Nonlinearities in the MSA-Level Phillips Curve

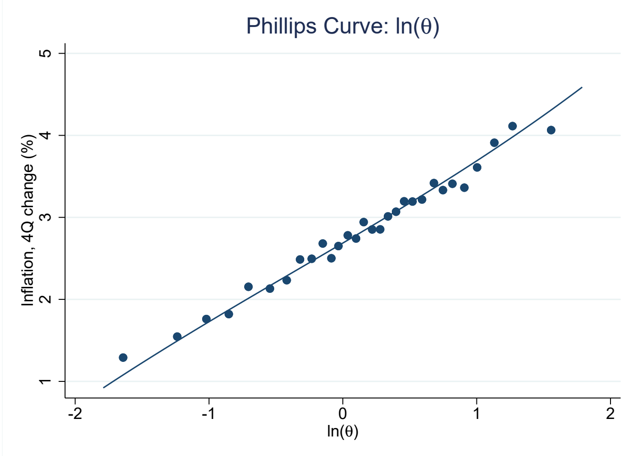

Notes: The left panel plots a binned scatter plot of the 4-quarter change in core inflation on the unemployment rate (4-quarter lagged value). The right panel shows the analogous plot on the log V/U ratio (4-quarter lagged). Controls include time and MSA fixed effects, MSA-specific trends, lagged core inflation and the lagged the lagged ratio of the goods and services price level.

We can use the MSA-level variation to further explore the presence of non-linearity. However, the evidence for nonlinearity is much weaker at the MSA level. In Table 7, we report tests for nonlinearities in $\hat{\theta}$ (column 2) as well as $\hat{\delta}$ (column 4).[19] The results show that the impact of $\hat{\theta}_t$ on inflation is quite linear, where the interaction term is positive but insignificant. Similarly, the effect of $\hat{\delta}_t$ shows minimal nonlinearity. This can be seen more clearly in Figure 7, which shows a binned scatter plot of the inflation rate (4-quarter change) and $\hat{\theta}$. There is little evidence of a nonlinear relationship.

7 Conclusion

In this work, we systematically assess the ability of popular variables at (i) predicting and (ii) explaining inflation fluctuations over time and across U.S. metropolitan areas. In particular, we exploit a newly constructed panel data set with job openings and vacancy filling cost proxies covering 1982-2022. We find that the vacancy-unemployment (V/U) ratio and

[18] Using $\bar{\theta} = E\theta_t + 1.6\sigma_\theta$ gives similar results.

[19] The threshold for the nonlinearity is 50th percentile by MSA, but results are generally similar when altering the thresholds as shown in Figure 7.

vacancy filling cost proxies outperform other slack measures, in particular the unemployment rate. Beveridge curve shifts—notably, movements in matching efficiency—are responsible for the superior performance of the V/U ratio over unemployment. While the determinants of movements in matching efficiency are still relatively poorly understood, one lesson of our study is that the behavior of matching efficiency is an important topic that extends beyond labor market studies: it has direct implications for our understanding of inflation.

References

Abraham, K. G. (1987). Help-wanted advertising, job vacancies, and unemployment. *Brookings Papers on Economic Activity*, 1987(1):207–248.

Abraham, K. G., Haltiwanger, J. C., and Rendell, L. E. (2020). How tight is the U.S. labor market? *Brookings Papers on Economic Activity*, 51(1):97–165.

Ahn, H. J. and Crane, L. D. (2020). Dynamic Beveridge curve accounting. *Federal Reserve Board of Governors, Finance and Economics Discussion Series.*

Bagga, S., Mann, L., Sahin, A., and Violante, G. L. (2023). Job amenity shocks and labor reallocation. *Working Paper.*

Ball, L., Leigh, D., and Mishra, P. (2022). Understanding U.S. inflation during the covid-19 era. *Brookings Papers on Economic Activity*, 1:54.

Barlevy, G., Faberman, R. J., Hobijn, B., and Şahin, A. (2024). The shifting reasons for Beveridge curve shifts. *Journal of Economic Perspectives*, 38(2):83–106.

Barnichon, R. (2010). Building a composite help-wanted index. *Economics Letters*, 109(3):175–178.

Barnichon, R. and Figura, A. (2015). Labor market heterogeneity and the aggregate matching function. *American Economic Journal: Macroeconomics*, 7(4):222–249.

Barnichon, R. and Mesters, G. (2020). Identifying modern macro equations using old shocks. *Quarterly Journal of Economics*, 135:2255–2298.

Barnichon, R. and Shapiro, A. H. (2022). What's the best measure of economic slack? *FRBSF Economic Letter*, 4:1–05.

Benigno, P. and Eggertsson, G. B. (2023). It's baaack: The surge in inflation in the 2020s and the return of the non-linear Phillips curve. Technical report, National Bureau of Economic Research.

Blanchard, O. J. and Bernanke, B. S. (2023). What caused the U.S. pandemic-era inflation? Technical report, National Bureau of Economic Research.

Blanchard, O. J. and Diamond, P. (1989). The Beveridge curve. *Brookings Papers on Economic Activity*, 1989(1):1–76.

Blanchard, O. J. and Galí, J. (2010). Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment. *American Economic Journal: Macroeconomics*, 2(2):1–30.

Crust, E., Lansing, K., and Petrosky-Nadeau, N. (2023). Reducing inflation along a nonlinear phillips curve. *FRBSF Economic Letter*, 17.

Davis, S. J., Faberman, R. J., and Haltiwanger, J. C. (2013). The establishment-level behavior of vacancies and hiring. *The Quarterly Journal of Economics*, 128(2):581–622.

Friedman, M. (1968). The role of monetary policy. *The American Economic Review*, 58(1).

Fujita, S., Moscarini, G., and Postel-Vinay, F. (2024). Measuring employer-to-employer reallocation. *American Economic Journal: Macroeconomics*, 16(3):1–51.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and its Applications*. Princeton University Press.

Gali, J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, 44:195–222.

Geroski, P. A. (2000). Models of technology diffusion. *Research Policy*, 29(4-5):603–625.

Gitti, G. (2024). Nonlinearities in the regional Phillips curve with labor market tightness. Technical report. working paper.

Hazell, J., Herreno, J., Nakamura, E., and Steinsson, J. (2022). The slope of the Phillips curve: evidence from U.S. states. *The Quarterly Journal of Economics*, 137(3):1299–1344.

Hobijn, B. and Şahin, A. (2021). Maximum employment and the participation cycle. Technical report, National Bureau of Economic Research.

Hornstein, A., Kudlyak, M., and Lange, F. (2014). Measuring resource utilization in the labor market. *FRB Economic Quarterly*, 100(1):1–21.

Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *The American Economic Review*, 95:161–182.

Krause, M. U., Lopez-Salido, D., and Lubik, T. A. (2008). Inflation dynamics with search frictions: A structural econometric analysis. *Journal of Monetary Economics*, 55(5):892–916.

Krause, M. U. and Lubik, T. A. (2007). The (ir)relevance of real wage rigidity in the New Keynesian model with search frictions. *Journal of Monetary Economics*, 54(3):706–727.

Mavroeidis, S., Plagborg-Møller, M., and Stock, J. H. (2014). Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve. *Journal of Economic Literature*, 52:124–188.

McLeay, M. and Tenreyro, S. (2020). Optimal inflation and the identification of the Phillips curve. *NBER Macroeconomics Annual*, 34(1):199–255.

Medoff, J. L. and Abraham, K. G. (1982). Unemployment, unsatisfied demand for labor, and compensation growth in the United States, 1956-1980, in M. Baily, ed. *Workers, Jobs, and Inflation*.

Mortensen, D. T. and Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3):397–415.

Moscarini, G. and Postel-Vinay, F. (2017). The relative power of employment-to-employment reallocation and unemployment exits in predicting wage growth. *American Economic Review*, 107(5):364–368.

Moscarini, G. and Postel-Vinay, F. (2023). The job ladder: Inflation vs. reallocation. Technical report, National Bureau of Economic Research.

Petrongolo, B. and Pissarides, C. A. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature*, 39(2):390–431.

Phelps, E. S. (1967). Phillips curves, expectations of inflation and optimal unemployment over time. *Economica*, 34:254–281.

Phillips, A. W. (1958). The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861-1957. *Economica*, 25:283–299.

Pissarides, C. A. (2000). *Equilibrium unemployment theory.* MIT Press.

Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium wage dispersion with worker and employer heterogeneity. *Econometrica*, 70(6):2295–2350.

Ramey, V. (2016). Macroeconomic Shocks and Their Propagation. In Taylor, J. B. and Uhlig, H., editors, *Handbook of Macroeconomics.* Elsevier, Amsterdam, North Holland.

Romer, C. D. and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, 94:1055–1084.

Şahin, A., Song, J., Topa, G., and Violante, G. L. (2014). Mismatch unemployment. *American Economic Review*, 104(11):3529–3564.

Shapiro, A. H. (2020). A simple framework to monitor inflation. Federal Reserve Bank of San Francisco. Working Paper 2020-29.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15(2):127–148.

Table 1: Philips Curve Estimates, 1995-2023

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
u_t	-0.20^{***} (0.09)	-0.77^{***} (0.11)	—	—	—	—	—	—	—	—
$\hat{\theta}_t$	—	—	0.36^{***} (0.09)	0.99^{***} (0.09)	—	—	—	—	—	—
$\hat{\chi}_t$	—	—	—	—	0.48^{***} (0.09)	1.15^{***} (0.08)	—	—	—	—
$\hat{\theta}_t^*$	—	—	—	—	—	—	0.47^{***} (0.09)	1.16^{***} (0.08)	—	—
$\hat{\chi}_t^*$	—	—	—	—	—	—	—	—	0.71^{***} (0.08)	1.26^{***} (0.09)
Inflation	core	cyclical	core	cyclical	core	cyclical	core	cyclical	core	cyclical
Sample	95-23	05-23	95-23	05-23	95-23	05-23	95-23	05-23	95-23	05-23
Adjusted R^2	0.03	0.32	0.11	0.53	0.20	0.65	0.21	0.68	0.40	0.66

Notes: The forcing variables were z-scored (demeaned and normalized to unit standard-deviation) for comparability across columns. See Figure 1 for variable definitions.

Table 2: Philips Curve Estimates, 1960-2023

	(1)	(2)	(3)	(4)	(5)	(6)
u_t	-0.25^{***} (0.06)	-1.12^{**} (0.57)	—	—	—	—
$\hat{\theta}_t$	—	—	0.33^{***} (0.06)	1.00^{**} (0.51)	—	—
$\hat{\chi}_t$	—	—	—	—	0.27^{***} (0.07)	1.05^{**} (0.63)
$E_t\pi_\infty$	1.04^{***} (0.04)	0.79^{*} (0.50)	0.95^{***} (0.07)	0.86^{*} (0.46)	0.95^{***} (0.07)	0.90^{*} (0.49)
Adjusted R^2	0.815		0.824		0.829	
Partial R^2	0.039		0.080		0.092	
IV	No	Yes	No	Yes	No	Yes

Notes: The forcing variables were z-scored (demeaned and normalized to unit standard-deviation) for comparability across columns. See Figure 1 for variable definitions.

Table 3: Philips Curve Estimates, MSA Level 1982-2022

	Dep. variable: Core Inflation (Δ4Q)					
	(1)	(2)	(3)	(4)	(5)	(6)
u	-0.0957***	-0.647***			-0.392***	-0.308**
	(0.0282)	(0.114)			(0.0644)	(0.123)
$\hat{\theta}$			0.280***	0.809***	0.584***	0.591***
			(0.0351)	(0.102)	(0.0625)	(0.128)
Observations	2431	2431	2431	2431	2431	2431
Adjusted R^2	0.329	0.689	0.358	0.695	0.378	0.698
Adj. Within R^2	0.329	0.199	0.358	0.215	0.378	0.223
MSA Fixed Effects	No	Yes	No	Yes	No	Yes
MSA Time Trends	No	Yes	No	Yes	No	Yes
Time Fixed Effects	No	Yes	No	Yes	No	Yes

All variables z-scored (demeaned and normalized to unit standard-deviation) for comparability across columns.
Controls included lagged inflation and the lagged ratio of the goods and services price level.
Standard errors clustered by MSA. See Figure 1 for variable definitions.

Table 4: Philips Curve Estimates: Testing for Shifts in Beveridge Curve, 1960-2023

	(1)	(2)	(3)	(4)
$\hat{u}_t$	-0.91^{***}	1.45^{*}	-0.42	-0.55^{**}
	(0.13)	(0.75)	(0.30)	(0.28)
$\hat{\theta}_t$	—	1.36^{***}	—	—
		(0.39)		
$\hat{\delta}_t$	—	—	-1.36^{***}	—
			(0.40)	
$\hat{s}_t$	—	—	—	1.03^{*}
				(0.54)
$\hat{m}_{0t}$	—	—	—	-1.43^{***}
				(0.40)
$E_t\pi_\infty$	1.01^{***}	0.88^{***}	0.88^{***}	0.89^{***}
	(0.07)	(0.07)	(0.07)	(0.07)
Adjusted R^2	0.817	0.826	0.826	0.826

Notes: See Figure 1 for variable definitions. $\hat{\delta}_t$ denotes Beveridge curve shifts, $\hat{s}_t$ the (log) unemployment inflow rate, and $\hat{m}_{0t}$ (log) matching efficiency.

Table 5: Philips Curve Estimates, MSA Level: Testing for Shifts in Beveridge Curve

	Dep. variable: Core Inflation (Δ4Q)		
	(1)	(2)	(3)
$\hat{u}$	-0.687***	-0.368**	-0.695***
	(0.0925)	(0.140)	(0.0856)
$\hat{\theta}$		0.506***	
		(0.159)	
$\hat{\delta}$			-0.122***
			(0.0384)
Observations	2431	2431	2431
Adjusted R^2	0.694	0.699	0.699
Adj. Within R^2	0.211	0.225	0.225
MSA Fixed Effects	Yes	Yes	Yes
MSA Time Trends	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes

All variables z-scored (demeaned and normalized to unit standard-deviation) for comparability across columns. Controls include lagged inflation and the lagged ratio of the goods and services price level. Standard errors clustered by MSA.

Table 6: Philips Curve Estimates: Testing for Curvature, 1960-2023

	(1)	(2)	(3)
$\hat{\theta}_t$	0.33^{***}	-0.10	—
	(0.06)	(0.11)	
$\hat{\theta}_t \mathbb{1}_{\theta>\bar{\theta}}$	—	0.56^{***}	—
		(0.10)	
$\hat{u}_t$	—	—	-0.29^{**}
			(0.13)
$\hat{u}_t \mathbb{1}_{\theta>\bar{\theta}}$	—	—	0.14
			(0.10)
$\hat{\delta}_t$	—	—	0.02
			(0.10)
$\hat{\delta}_t \mathbb{1}_{\theta>\bar{\theta}}$	—	—	-0.54^{***}
			(0.07)
$E_t \pi_\infty$	0.95^{***}	1.02^{***}	0.99^{***}
	(0.07)	(0.07)	(0.08)
Adjusted R^2	0.826	0.846	0.860

Note: Inflation is core PCE inflation. The threshold is $\bar{\theta} = E\theta$ in columns (2)-(3).

Table 7: Philips Curve Estimates, MSA Level: Testing for Curvature

	Dep. variable: Core Inflation (Δ4Q)			
	(1)	(2)	(3)	(4)
$\hat{\theta}$	0.809***	0.725***		
	(0.102)	(0.154)		
$\hat{\theta}$ x ($\theta > \bar{\theta}$)		0.203		
		(0.127)		
$\hat{u}$			-0.695***	-0.758***
			(0.0856)	(0.113)
$\hat{u}$ x ($\theta > \bar{\theta}$)				0.294
				(0.264)
$\hat{\delta}$			-0.122***	-0.0991*
			(0.0384)	(0.0541)
$\hat{\delta}$ x ($\theta > \bar{\theta}$)				-0.206
				(0.453)
Observations	2431	2431	2431	2431
Adjusted R^2	0.695	0.696	0.699	0.699
Adj. Within R^2	0.215	0.217	0.225	0.226
MSA Fixed Effects	Yes	Yes	Yes	Yes
MSA Time Trends	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes

All variables z-scored (demeaned and normalized to unit standard-deviation) for comparability across columns. Controls include lagged inflation, the lagged ratio of the goods and services price level, and threshold dummies. Thresholds are the 50th percentile by MSA. Standard errors clustered by MSA

Appendix

A Construction of the MSA job posting dataset

We exploit two data sources: (i) the MSA-level "print" help-wanted advertisement index (HWI) from the Conference Board over 1951-2011, and (ii) the number of "online" job postings in each MSA from the Conference Board over 2006-2011. As described in Barnichon (2010), the print HWI is affected by the diffusion of the internet after 1995 as firms switched from print newspaper job advertising to online job posting. To recover the total number of job postings from these two series, we model the evolution of the print share in job advertising as a smooth function. Specifically, the total number of advertisements (print and online) V_t is given by $V_t = V_t^p + V_t^o$ where V_t^p and V_t^o are the numbers of print and online job postings and

$$\omega_t = \frac{V_t^p}{V_t^p + V_t^o} \tag{22}$$

is the share of print help-wanted advertising in total advertising. To recover total job openings V_t from V_t^p, we posit that the trend in $\ln(V_t^p)$ after 1995 is driven by the diffusion of online advertising, and we estimate the trend from an HP-filter with $\lambda = 10^7$ at a monthly frequency in order to capture a slowly moving trend. As in Barnichon (2010), we model ω_t with a Mixed Information Source Model, which has been shown to successfully capture the diffusion of the internet in the U.S. population (e.g., Geroski, 2000):

$$\omega_t = 1 - \frac{1 - e^{-\gamma t}}{1 + \alpha e^{-\beta t}} \ , \tag{23}$$

which we estimate over the trend in V_t^p. We repeat this procedure separately for each MSA to obtain an MSA-specific print share ω_{it}.

B Matching efficiency

To measure $\hat{m}_{0t}$ —movements in matching efficiency— we run the regression

$$\hat{\theta}_t = \beta_u \hat{u}_t + \beta_s \hat{s}_t + \beta_o \Delta \hat{u}_{t+1} + e_t \tag{24}$$

where s_t is measured from short-term unemployment, and we take $\hat{m}_{0t}$ as the regression residual.

Figure A1 plots our estimated matching efficiency series. Notice two features. First, matching efficiency displays cyclical oscillations, generally rising in the early stages of recessions and deteriorating thereafter. This is mostly due to a composition effect, as the average unemployment duration of the unemployment pool varies with the cycle (see Barnichon and Figura, 2015). Second, matching efficiency displayed a large decline in the 2008-2009 recession, but did not fully recover before dropping again in the post-COVID recovery. One hypothesis is that the post-COVID decline in matching efficiency is linked to the Great Resignation — a general reconsideration of career choices and work-life balance (e.g., Bagga et al., 2023).

Figure A1: Matching efficiency

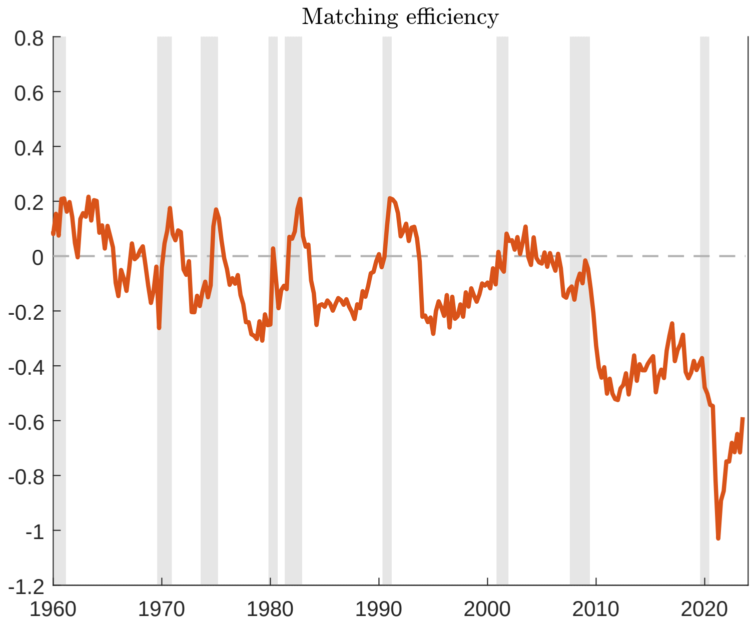

C Measuring inflow and outflow rates

To measure the monthly inflow and outflow rates at the national level, we use the difference equation

$$U_{t+1} = (1 - F_t)U_t + U_{t+1}^{<5wks}$$

where U_t and $U_t^{<5wks}$ denote respectively the total number of unemployed and the number of unemployed for less than 5 weeks (the newly unemployed during the month).

This unemployment outflow probability is then given by

$$F_t = 1 - \frac{U_{t+1} - U_{t+1}^{<5wks}}{U_t} \ ,$$

and the outflow rate is $f_t = -\ln(1 - F_t)$.

The unemployment inflow probability is obtained from

$$S_t(L_t - U_t) = U_{t+1}^{<5wks}$$

where L_t is the labor force size . The inflow rate is then $s_t = -\ln(1 - S_t)$.

C.1 Raw series for inflation and candidate forcing variables

Figure A2: Raw series 1994-2023

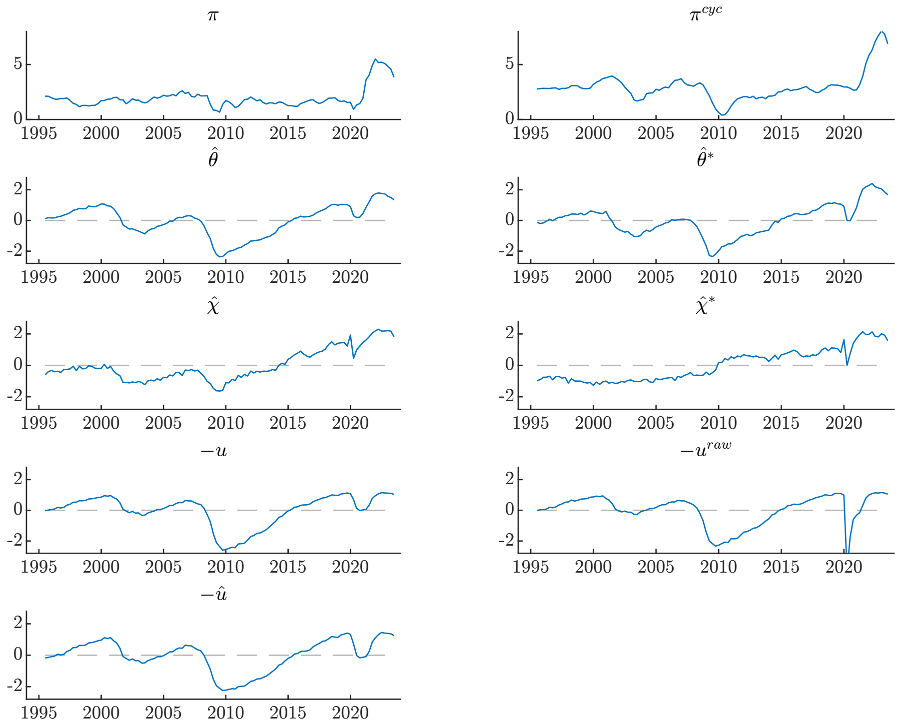

Figure A3: Raw series 1960-2023

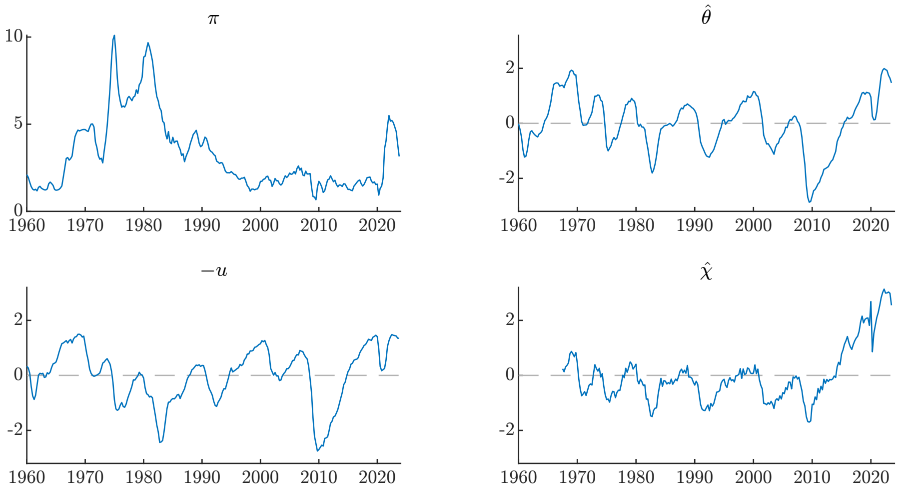

D Services inflation as the dependent variable

Table A1: Philips Curve Estimates, MSA Level 1982-2022, Services inflation

	Dep. variable: Services Inflation (Δ4Q)					
	(1)	(2)	(3)	(4)	(5)	(6)
u	-0.356***	-0.676***			0.0623	-0.317***
	(0.0453)	(0.122)			(0.0747)	(0.106)
$\hat{\theta}$			0.452***	0.851***	0.501***	0.626***
			(0.0516)	(0.152)	(0.0851)	(0.175)
Observations	2431	2431	2431	2431	2431	2431
Adjusted R^2	0.350	0.642	0.380	0.648	0.380	0.651
Adj. Within R^2	0.350	0.260	0.380	0.272	0.380	0.278
MSA Fixed Effects	No	Yes	No	Yes	No	Yes
MSA Time Trends	No	Yes	No	Yes	No	Yes
Time Fixed Effects	No	Yes	No	Yes	No	Yes

Standard errors clustered by MSA

Table A2: Philips Curve Estimates, MSA Level: Testing for Shifts in Beveridge Curve, Services inlfation

	Dep. variable: Services Inflation (Δ4Q)		
	(1)	(2)	(3)
$\hat{u}$	-0.804***	-0.557***	-0.810***
	(0.112)	(0.155)	(0.108)
$\hat{\theta}$		0.393*	
		(0.204)	
$\hat{\delta}$			-0.0945*
			(0.0495)
Observations	2431	2431	2431
Adjusted R^2	0.653	0.655	0.655
Adj. Within R^2	0.282	0.288	0.288
MSA Fixed Effects	Yes	Yes	Yes
MSA Time Trends	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes

Standard errors clustered by MSA