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

5/6/24

Abstract

Despite the popularity of the Phillips curve, there is little consensus on the forcing variable driving inflation, i.e., on the appropriate measure of "slack" in the economy. In this work, we systematically assess the ability of popular variables at (i) predicting and (ii) explaining inflation fluctuations over time and across US 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, vacancy posting

*Preliminary. 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 Phillips paper (Phillips, 1958) linking the unemployment rate to (wage) inflation, 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) first emphasized the relationship between inflation and 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. Instead, average real marginal cost, the labor share or the output gap (see e.g., Galí, 2015), and more recently the job-switching rate (Moscarini and Postel-Vinay, 2017, 2023) and the vacancy–unemployment ratio (Barnichon and Shapiro, 2022; Ball et al., 2022) have been proposed as slack candidates.

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 over-fitting issues inherent to in-sample analysis. We assess which slack measures best predicts inflation at one-year horizons. From the post-Covid period all the way back to the upsurge in inflation of the 1960s as well as the interwar period, 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 performances of the V/U ratio over unemployment, though coefficient estimates are likely biased by endogeneity issues: unobserved inflation expectations, unobserved natural rates, confounding from supply shocks, and downward bias from counter-cyclical monetary policy (e.g., McLeay and Tenreyro, 2020; Barnichon and Mesters, 2020).

To address these endogeneity issues, 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 MSA level over 1980-2022. Again, the V/U ratio performs very well, explaining inflation better than the unemployment rate.

The superior performance of the V/U ratio may seem surprising since vacancy 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—because V/U and unemployment are so highly correlated. At times, however, this correlation can deteriorate sharply due to shifts in the Beveridge curve shifts. 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 the most relevant measure of slack, that is about the forcing variable that best explains inflation.

Economic slack

The most popular forcing variables are proxies for tightness in the labor market, typically the unemployment rate or unemployment gap (Phillips, 1958). A potential drawback of the unemployment rate however is that it ignores workers outside of the labor force.[1] To 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, which accounts for potential cyclical fluctuations in the size of the labor force.

More recently, a new proxy for labor market tightness has recently been proposed (e.g.,

[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 Hobijn and Şahin (2021).

Barnichon and Shapiro, 2022; Ball et al., 2022): the vacancy–unemployment ratio (or V/U ratio for short). Intuitively, the ratio represents the number of job vacancies, or demand for labor, relative to the number of unemployed individuals, or supply of labor. As with the unemployment rate however, the V/U ratio misses job seekers outside the labor force, and it also misses on-the-job job seekers, that is employed workers who search for another job. To take into account all possible job seekers, Abraham et al. (2020) constructed a generalized V/U ratio that replaces unemployment with a measure of *effective* job searchers.

Marginal hiring costs

Despite the popularity 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 (Gali and Gertler, 1999; Galí, 2015). Instead, the literature proposed using the share of output going to labor compensation —the labor share— as a proxy for firm's 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.

Instead, the vacancy–unemployment ratio (Barnichon and Shapiro, 2022; Ball et al., 2022) and the job-switching rate (Moscarini and Postel-Vinay, 2017) have recently been proposed, building on the intuition that they should proxy for firms' marginal costs.

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 cost of hiring a marginal worker —the vacancy filling cost— (e.g., Krause and Lubik, 2007; Krause et al., 2008), 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 is the vacancy filling rate —the rate at which firms fill vacancies— which is given by the flow of new matches at instant t (m_t) divided 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, and that expected duration is $1/q_t$.

To proxy for χ_t, one approach could be to exploit data on firms' vacancy filling rate. While these data were not available until recently, we can exploit the properties of the constant returns to scale matching functions to proxy for q_t. Indeed, in a standard search and matching model (Mortensen and Pissarides, 1994), the vacancy filling rate can be related to the vacancy–unemployment ratio by means of the matching function. The matching function relates the flow of new hires to the stocks of vacancies and unemployment, and for a 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, while the unemployed job finding rate is given by $p_t^{ue} = \frac{m(U_t,V_t)}{U_t} = q_t(\theta_t)\theta_t$. Combining with the expression for χ_t, we get our first hiring cost proxy $\chi_t^{ue} = c\frac{\theta_t}{p_t^{ue}}$, or expressed in log deviations from steady-state

$$\hat{\chi}_t^{ue} = \hat{\theta}_t - \hat{p}_t^{ue} \ , \tag{3}$$

where "hats" denote (log) deviations from steady-state. In other words, we can recover firms' vacancy filling rate from the V/U ratio and the unemployment outflow rate. This property is attractive, because both the V/U ratio and the unemployment outflow rate are available over a long time sample period, as we will see.

Further, 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,[2] and simplify the vacancy filling cost as $\chi_t = c\frac{\theta_t^{\sigma}}{m_{0t}}$ or

$$\hat{\chi}_t^{\theta} = \sigma\hat{\theta}_t - \hat{m}_{0t}. \tag{4}$$

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 benefit of the V/U ratio is that it is available over a long sample period, in fact all the way back to 1919. The cost is that it relies on the assumption of constant matching efficiency. If matching efficiency is not constant, (2) makes clear that the V/U ratio will not perfectly capture variations in firms' recruiting costs.

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 $q_t = \frac{m_t}{V_t} = \frac{p_t^{ue}U_t + p_t^{ee}E_t}{V_t} + p_t^{ne}N_t$where p_t^{ee} is the Employment-to-Employment transition rate and p_t^{ne} is the Nonparticipation-to-Employment transition rate. We get the (generalized) vacancy filling cost proxy[3]

$$\hat{\chi}_t^{ee} = \hat{\theta}_t - \hat{p}_t^{ue} - \hat{\gamma}_t \qquad \text{where} \qquad \gamma_t = 1 + \frac{p_t^{ee}}{p_t^{ue}}\frac{1-u_t}{u_t} + \frac{p_t^{ne}}{p_t^{ue}}\frac{1-l_t}{l_t u_t}\ , \tag{5}$$

where u_t and u_t are the unemployment and job openings numbers expressed as fractions of the labor force and l_t is the labor force participation rate. Comparing with (3) —vacancy filling costs under hiring from unemployment alone—, we can see that the term γ_t is a correction factor that allows for changes in the relative importance of hiring from Employment ($\frac{p_t^{ee}}{p_t^{ue}}$) or Nonparticipation ($\frac{p_t^{ne}}{p_t^{ue}}$).

As with expression (3), the generalized vacancy filling cost is attractive because it can be measured over a relatively long period of time by exploiting CPS micro data, as we will see next.

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,

[2]The Cobb-Douglas matching function is used in most macro models with search and search and matching frictions (Pissarides, 2000). The matching efficiency term $m_{0,t}$ 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.

[3]Start from $q_t = \frac{m_t}{V_t} = \frac{p_t^{ue}U_t + p_t^{ee}E_t + p_t^{ne}N_t}{V_t}$, combine with $\chi_t^{ee} = \frac{c}{q_t}$ and log-linearize around the steady-state.

increases in the cost of these inputs may force businesses to raise prices for their products to stay profitable. In this context, we consider two additional measures of slack based on industrial production. One is 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. The second measure is the industrial production (IP) share, or the share of output in the economy attributable to industrial production, as proposed in Shapiro (2008). These two measures capture the intuition that material inputs, such as primary metals, wood, and machinery, become relatively more costly during an economic expansion.

3 Forecasting inflation

In this section, we assess the forecasting performances of the forcing variables discussed above. To generate forecasts, we estimate local projections (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.[4] 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.

As measures of slack we consider the following variables: 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 V/U ratio ($\hat{\theta}_t$), the generalized V/U ratio ($\hat{\theta}^*_t$) of Abraham et al. (2020), the vacancy filling cost ($\hat{\chi}^{ue}_t$), the generalized vacancy filling cost ($\hat{\chi}^{ee}_t$), capacity utilization, the IP share, the unemployment gap and the log output gap as estimated from the CBO. The main series are depicted in the Appendix.

To measure the variables necessary to construct the vacancy filling costs proxies, we can exploit CPS data to construct $\hat{\chi}^{ue}_t = \hat{\theta}_t - \hat{p}^{ue}_t$ back to 1967, or a similar proxy $\hat{\chi}^{f}_t = \hat{\theta}_t - \hat{f}_t$, where f_t is the unemployment outflow rate (Shimer, 2012). To construct a proxy for $\hat{\gamma}_t$ in $\hat{\chi}^{ee}_t = \hat{\chi}^{ue}_t - \hat{\gamma}_t$, we again use CPS data measure the Nonparticipation-to-Employment transition rate over 1967-2023, and we use the Employment-to-Employment over 1995-2023 transition rate from Fujita et al. (2020).

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 60s and 70s—. The Appendix also reports results from the Interwar period (1919-1940).

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

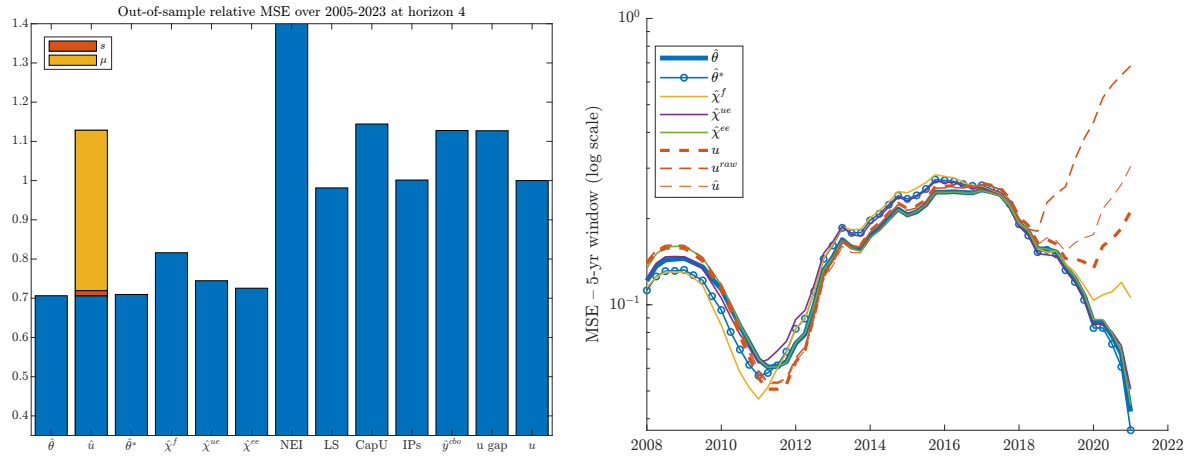
1995-2023

Figure 1a plots the mean-squared forecast errors 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 (green bar); thus, they can be interpreted as indicating how much better or worse they perform than the unemployment rate.

We can see that the V/U ratio, or more generally all the vacancy filling costs proxies, performs much better than the other measures. In particular, the V/U ratio outperforms the traditional unemployment rate by about 30%, a substantial improvement. The other vacancy filling cost proxies outperform the unemployment rate by similar amounts. Perhaps surprisingly, our most general vacancy filling cost proxy $\hat{\chi}_t^{ee}$, which allows for hiring from unemployment, employment and nonparticipation does not materially improve performance over the simpler V/U ratio. While this could point towards constant matching efficiency ($\hat{m}_{0t} = 0$) and a constant hiring rate outside the unemployment pool ($\hat{\gamma}_t = 0$), this could also be the result of larger measurement error in the transition rates.[5]

In contrast, none of the traditional measures —the output gap, the unemployment gap, the non-employment index or the labor share— outperform unemployment. Interestingly, 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.

Figure 1: Forecasting performances, 1995-2024



(a) Relative MSEs (b) 5-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. "$\hat{\theta}$" is the log V/U ratio, "$\hat{\theta}^*$" is the generalized log V/U ratio, "$\hat{u}$" is the log unemployment rate, the "$\hat{\chi}$s" are the three vacancy filling cost proxies, "NEI" is the Non-Employment Index, "LS" is the labor share, "CapU" is the Board of Governors capacity utilization rate", "IPs" is the (detrended) share of industrial production in GDP, "$\hat{y}^{cbo}$" and "u gap" are the output gap and unemployment gap estimated by the CBO, and "u" is the unemployment rate. The orange and red bars decompose the superior performances of $\hat{\theta}$ over $\hat{u}$ into the contribution of the Beveridge curve shifts, see Section the

To better understand performances over time, Figure 1b plots the 5-year average rolling

[5]Flow-based measures like χ_t are more noisy than stock-based measures like the V/U ratio. See Figure 9.

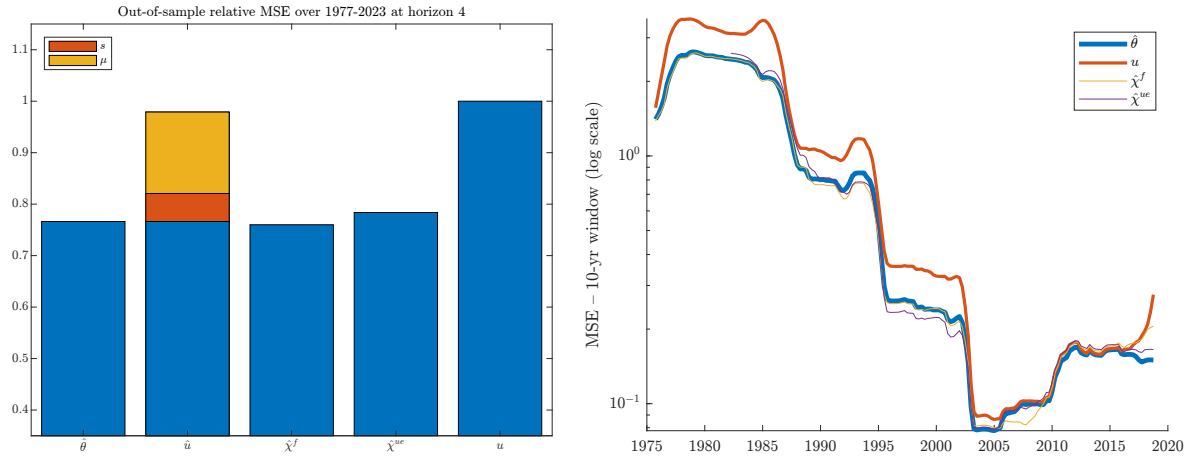
MSEs for the different models over 2005-2023. Two things to notice. First, the V/U ratio and vacancy filling cost proxies generally outperform the unemployment rate measures over the entire sample period, though the differences are small during the stable inflation period of 2005-2019. Second, the differences become large during the post-COVID recovery where all forcing variables measures see large deterioration in forecasting performance, *except* for the V/U ratios and the vacancy filling cost proxies ($\hat{\chi}_t$). In fact, the generalized V/U ratio of Abraham et al. (2020) performs the best in forecasting inflation during the Covid recovery (though its performances relative to the V/U ratio are less good over 2012-2019).

These observations indicate that large movements in inflation are necessary to discriminate between competing forcing variables. For these reasons, we will now evaluate the performances of the V/U ratios and vacancy filling costs proxies during the large inflation movements of the 60s and 70s.

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 build up 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 proxies $\hat{\chi}^{ue}$ and $\hat{\chi}^{ee}$ that includes job seekers from outside the labor force.[6]

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.

Again, we find that the V/U ratio and the vacancy filling cost proxies outperform other measures, notably the unemployment rate by about 30 percent. Most importantly, this exercise confirms that the superior performances of 2022-2023 is not a unique occurrence: $\hat{\theta}_t$ (blue line) *systematically* outperforms u_t (red line) over 1970-2023.[7]

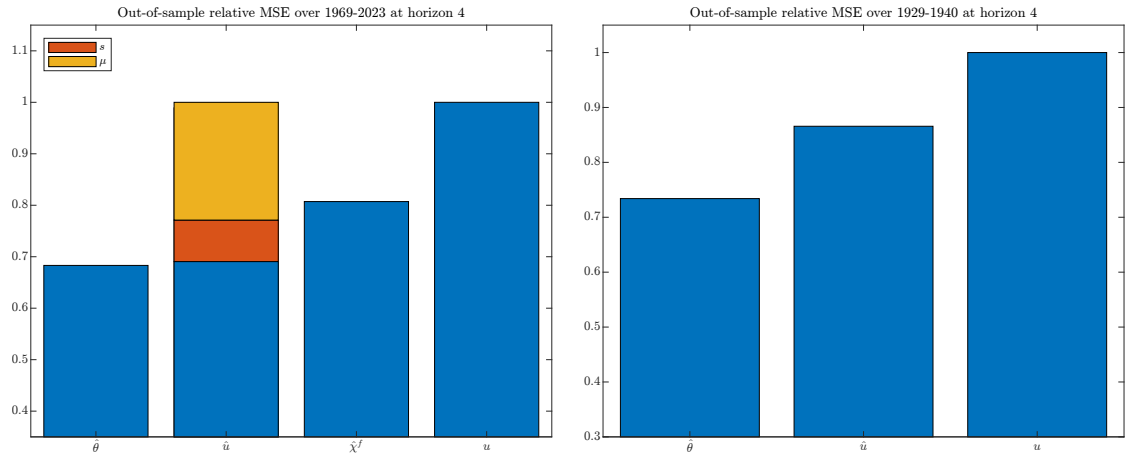
[6]Since Employment-to-Employment transitions are not available before 1995, in this section we use omit these transitions from $\hat{\chi}^{ee}$.

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

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 our V/U ratio—unemployment horserace before WII, since vacancy data ("Advertising in Newspapers, Metropolitan Life Insurance Company") are available from the NBER macro history database. Again, we find that the V/U ratio outperforms unemployment in the interwar period (1919-1940), see Figure 3b.

Figure 3: Longer samples



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

Notes: Inflation is measured from Core CPI for panel (a) and Healdine CPI for Interwar data in panel (b).

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 and vacancy filling costs proxies were the most informative to predict inflation. A related, but separate, question is whether the superior prediction performance of the V/U ratio or the vacancy filling cost captures 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 as 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 structural macro Phillips curve —the (AS) relation 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 thus estimate and compare structural Phillips curves with different forcing variables: the V/U ratio, 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 steady-state), $E_t\hat{x}_\infty$ its permanent component $\hat{x}_t$, 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 $\psi = \frac{\kappa}{1-\beta\rho}$.

At this stage, to make progress and estimate ψ with time series information, we need an identifying assumption to identify $E_t\hat{x}_\infty$, and we posit that the long-run component of hiring cost depends on slow moving structural parameters of the economy, and that $E_t\hat{x}_\infty$ is a slow-moving component that is independent of $\hat{x}_t$, so that we have $E\hat{x}_t E_t\hat{x}_\infty = 0$. We will relax this assumption later, when we exploit MSA-level variation.

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

$$\pi_t = \psi\hat{x}_{t-4} + \beta E_t\pi_\infty + \upsilon_t, \tag{8}$$

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" 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 50 percent larger than the coefficient on u_t over 1995-2023 —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 partialled out the effect of π^e_∞— is twice as large using the V/U ratio than using unemployment alone. As in the forecasting exercise, the vacancy cost filling proxy does appear to perform better than unemployment, but the evidence is not as conclusive, likely because of higher measurement error in $\hat{\chi}_t$, see Figure 9. Last, the generalized V/U ratio of Abraham et al. (2020) does perform best overall, with a 15 percent higher coefficient than the baseline V/U ratio. The adjusted R2 is

also higher than with V/U —column (7) vs column (3)—.

An important caveat of Tables 2 and 1 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: (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 counter-cyclical policy: as the central bank works to mute the effects of aggregate shocks on inflation, 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 Federal Reserve Bank of San Francisco's (FRBSF) "cyclical core PCE inflation" measure (Shapiro (2020)). This measure isolates those categories within the PCE price index that move systematically with the unemployment rate, which are plausibly less contaminated by supply disturbances. This approach allows us to alleviate some of the endogeneity issues on the most recent sample period and thus include the most promising forcing variable —the generalized V/U ratio—. 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 post 1985 —monetary shocks are small and rare during the Great Moderation (e.g., Ramey, 2016)— and can only be used over the longer 1960-2023 sample. Third, we turn to MSA-level data in order to estimate MSA-level Phillips curves, building on Hazell et al. (2022)'s insight that cross-sectional information allows to address (or at least substantially lessen) these endogeneity biases.

Evidence using cyclical core PCE inflation

The FRBSF cyclical core inflation measure is "trained" on data up to up 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 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 fit of all models improves substantially, using the cyclical inflation measure. Furthermore, the V/U ratio and in particular, $\hat{\theta}_t^*$, perform better with this alternative inflation measure than the unemployment rate.

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 not roughly of similar magnitudes: All forcing variables perform similarly. One reason could be that the 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 US metropolitan level.

4.2 Evidence from US Metropolitan Statistical Areas

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

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

To address endogeneity issues, we posit that the MSA-specific deviations of long-run inflation expectations and long-run component of x_t from aggregate values ($\mathrm{E}_t\pi_{i,\infty} - \mathrm{E}_t\pi_{t+\infty}$ and $E_t\hat{x}_{i,\infty} - E_t\hat{x}_{\infty}$) are slow moving and follow a polynomial trend, while movements in aggregate inflation expectations and long-run marginal costs are controlled by time fixed effects.[8] Provided that there are no MSA-specific supply shocks, the time fixed effects will also control for aggregate supply shocks.

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 level, 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.

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

[8] 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.

Bureau of Labor Statistic's (BLS) Local Are 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 posting 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, an MSA-level HWI index has also been constructed by the Conference Board over 1951-2008. Starting in the mid-1990s however, this "print" measure of vacancy posting became increasingly unrepresentative as advertising over the internet became more prevalent.

A second measure of vacancy positing 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 an 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 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 US population (e.g., Geroski, 2000). The model is then estimated over the subsample when both vacancy series overlap. Finally, our third source of vacancy data is from The Burning Glass Institute, which spans 2010 to 2022.

Results

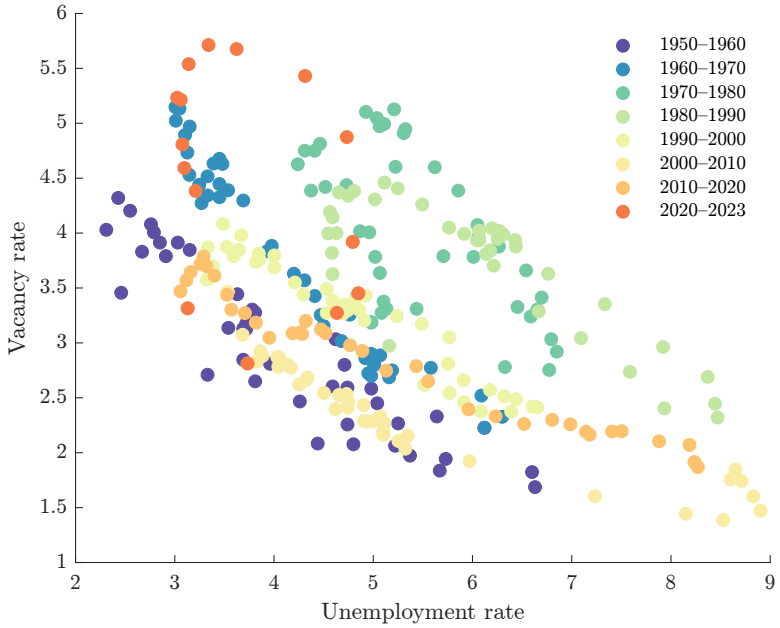
Results of the MSA-level Phillips curve estimation are shown in Table 3, where consider two forcing variables —the unemployment rate and the V/U ratio—. Estimates using the unemployment rate as the slack measure are shown in columns 1 and 2, while estimates using the log of the V/U ratio ($\hat{\theta}$) 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 $\hat{\theta}$. This indicates that the omitted-variable bias acts to attenuate the size of the coefficients (positive bias in the case of the unemployment rate and negative in the case of the vacancy-to-unemployment ratio). 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 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 forcing variable. Notably, we find that $\hat{\theta}_t$ or $\hat{\chi}_t$ substantially outperform 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 relation between vacancy posting and unemployment.

Figure 4: The Beveridge curve



As illustrated in Figure 4, vacancy posting and unemployment rates are negatively correlated, with periods of high unemployment corresponding to periods of low levels of job openings and vice-versa. In that context, it may seem surprising that a Phillips curve with the unemployment rate as forcing variable does not perform as well as a Phillips curve with the V/U ratio, since the two variables are higher correlated with a correlation of $-.89$ over 1960-2023. An intriguing follow-up question is then the following: what additional information does the V/U ratio bring above and beyond the unemployment rate alone?

To shed some light on this issue, we can go back to the theoretical underpinnings of the Beveridge curve, i.e., to 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).

Steady-state unemployment

Let U_t, E_t, and I_t denote the number of unemployed, employed and inactive (out of the labor force) individuals, respectively, at instant $t \in \mathbb{R}_+$. Letting p_t^{AB} denote the hazard rate of transiting from state $A \in \{E, U, I\}$ to state $B \in \{E, U, I\}$, unemployment, employment and inactivity (i.e., out of the labor force) will satisfy the system of differential equations

$$\begin{cases} \dot{U}_t = p_t^{EU} E_t + p_t^{IU} I_t - (p_t^{UE} + p_t^{UI}) U_t \\ \dot{E}_t = p_t^{UE} U_t + p_t^{IE} I_t - (p_t^{EU} + p_t^{EI}) E_t \\ \dot{I}_t = p_t^{EI} E_t + p_t^{UI} U_t - (p_t^{IE} + p_t^{IU}) I_t \end{cases} \tag{11}$$

As first argued by Shimer (2012), the magnitudes of the hazard rates is such that the half-life of a deviation of unemployment from its steady state value is about a month. As a result, at a quarterly frequency, the unemployment rate $u_t = \frac{U_t}{LF_t}$ is very well approximated by its steady-state value u_t^{ss} so that

$$u_t \simeq \frac{s_t}{s_t + f_t} \equiv u_t^{ss} \tag{12}$$

with s_t and f_t defined by

$$\begin{cases} s_t = p_t^{EU} + \frac{p_t^{EI} p_t^{IU}}{1 - p_t^{II}} \\ f_t = p_t^{UE} + \frac{p_t^{UI} p_t^{IE}}{1 - p_t^{II}}. \end{cases}$$

Expression (12) generalizes the simpler two-states case without movements in-and-out of the labor force where U_t satisfies $\dot{U}_t = p_t^{EU} E_t - p_t^{UE} U_t$ and $u_t^{ss} = \frac{p_t^{EU}}{p_t^{EU} + p_t^{UE}}$. With movements in-and-out of the labor force, workers can transition between U and E either directly (U-E) or in two steps by first leaving the labor force (U-I) and then by finding a job directly from inactivity (I-U). As a result, f_t, the "U-E transition probability" that matters for steady-state unemployment rate is a weighted average of p_t^{UE} and $p_t^{UI} p_t^{IE}$, with weights of 1 and $\frac{1}{1 - p_t^{II}}$, the average time that a worker going U→I→E spends transitioning through state I.[9] s_t has a similar expression.

In practice, the unemployment outflow rate is much larger than the unemployment inflow rate (by a factor of 10 or more), so that the steady-state unemployment can be approximated with

$$u_t \simeq \frac{s_t}{f_t}. \tag{13}$$

[9]Figure 7 in the appendix shows the behavior of p_t^{UE} and $\frac{p_t^{UI} p_t^{IE}}{1 - p_t^{II}}$ —the determinants of f_t—. The two series are highly correlated with a raw correlation of 0.75 and a correlation of 0.90 after detrending with a quadratic polynomial.

The matching function

Using a Cobb-Douglas matching function $m_t = m_{0t} U_t^{\sigma} V_t^{1-\sigma}$, we can relate the flow of new hires to the stocks of vacancies and unemployment. and express the unemployment exit f_t —the ratio of new hires to the stock of unemployed— as

$$f_t = \frac{m_t}{U_t} = m_{0,t} \theta_t^{1-\sigma}.$$

The top of figure 5 plots the actual unemployment outflow rate over 1948-2023 along with its fitted value. Abstracted from the recent episode, the matching function does a very good job at capturing fluctuations in the outflow rate.

The shifting Beveridge curve

The matching function is the cornerstone of the Beveridge curve. Combining with the steady-state approximation, we get $u_t = \frac{s_t}{m_{0,t}\theta_t^{1-\sigma}}$ or

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

Expression (14) is the Beveridge curve, and with a Cobb-Douglas matching function, the Beveridge curve is a log-log relationship between unemployment and vacancies, consistent with Figure 4.

Log-linearizing (14), we get

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

where $\hat{\mu}_t$ captures shifts in the Beveridge curve.

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 $\hat{\mu}_t = 0$, 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.[10] In other words, a finding that the V/U ratio provides superior information about future inflation shows that Beveridge curve shifts are central to understand inflation fluctuations —Phillips meets Beveridge—.

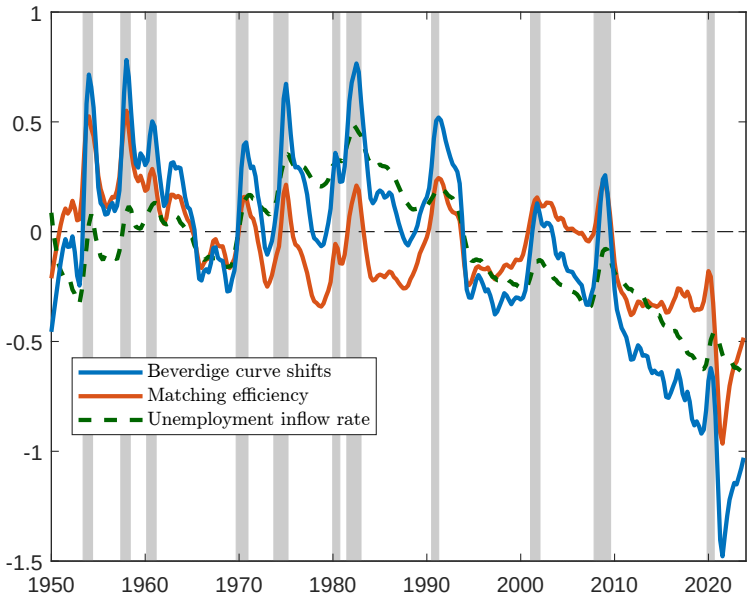
Second, the Beveridge curve can shift for different reasons. In this framework, the Beveridge curve can shift with: (i) movements in the unemployment inflow rate ($\hat{s}_t$) and (ii) movements in matching efficiency ($\hat{m}_{0t}$). To measure $\hat{\mu}_t$ —shifts in the Beveridge curve—, we run the regression $\hat{\theta}_t = \beta_u \hat{u}_t + e_t$ and take $\hat{\mu}_t$ as the regression residual. Similarly, to measure $\hat{m}_{0t}$ —movements in matching efficiency— we run the regression $\hat{\theta}_t = \beta_u \hat{u}_t + \beta_s \hat{s}_t + e_t$ where s_t is measured from short-term unemployment (see Shimer, 2012), and we take $\hat{m}_{0t}$ as the regression

[10] 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.

residual.

Figure 5 plots the time series for Beveridge curve shifts since 1951, decomposed into the contribution of the job separation rate and matching efficiency, which we define as the residual of a regression of $\hat{v}_t$ on $\hat{u}_t$ and $\hat{s}_t$. We can see that matching efficiency displays cyclical fluctuations, increasing in the early stages of recessions and worsening in the early stages of the recovery.

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 red line (matching efficiency movements) and dashed green line (unemployment inflow rate).

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).[11]

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)—[12], matching efficiency has declined markedly since the end of the financial crisis. The phenomenon worsened following COVID and the so-called *Great Resignation* (e.g., Barlevy et al., 2023).

[11]Two additional factors that can shift the Beveridge curve are (i) out-of-steady state dynamics and (ii) on-the-job search. First, the residual term captures the out-of-steady-state transition dynamics. Though small, out-of-steady-state dynamics can explain the slight time lag between the unemployment rate and the vacancy rate. Second, with on-the-job search, the Beveridge curve residual could also capture variation in employed search intensity over time (see e.g., Bagga et al., 2023). Since we measure matching efficiency as a residual of a Beveridge curve regression, we can think of our $_{0t}$ measure as capturing all these possible mechanisms.

[12]Long-term unemployed have a lower job finding rate than the short-term unemployed. In the early 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.

5.2 Beveridge curve shifts and inflation

To better understand how shifts in the Beveridge curve are important for inflation fluctuations, we replicate our two previous exercises —out-of-sample forecasting and Phillips curve estimation—, and we split the (log) V/U ratio into two components: (i) movements along the Beveridge curve, and (ii) shifts in the Beveridge curve.

Combining (15) with $\hat{\theta} = \hat{v}_t - \hat{u}_t$, we get

$$\hat{\theta}_t = \underbrace{-\frac{1}{1-\sigma}\hat{u}_t}_{\text{Mvts along BC}} + \underbrace{\frac{1}{1-\sigma}\hat{s}_t - \frac{1}{1-\sigma}\hat{m}_{0t}}_{\text{Shifts in BC}} \tag{16}$$

so that we can decompose the performance of the log V/U ratio into the contribution of log unemployment (movements along the curve) and the independent contribution of shifts in the Beveridge curve ($\hat{\mu}_t$).

Figures 1 and 2 (orange and red bars) use (16) to decompose the superior forecasting performances of $\hat{\theta}_t$ over $\hat{u}_t$ into the respective contributions of matching efficiency and the unemployment inflow rate over 2005-2023 and 1970-2023, building on decomposition (15). In both cases, we can see that most of the superior forecasting performance of labor market tightness over unemployment comes from movements in matching efficiency.

Next, we can split the forcing variable $\hat{\theta}_t$ of our Phillips curve regressions into the separate contributions of movements along the Beveridge curve and shifts in the Beveridge curve. Specifically, we run the regressions

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

Table 4 confirms the importance of Beveridge curve shifts in column (3). The coefficient on $\hat{\mu}_t$ is significant: outward shifts in the Beveridge curve correlate strongly with rises in inflation. Table 4, column (4) further shows that both the job separation rate and matching efficiency correlate strongly with inflation. Last, column (5) shows that matching efficiency contains additional information —above and beyond the V/U ratio— about future inflation. This finding is consistent with vacancy filling cost being the relevant forcing variable in the Phillips curve. Indeed, recall that we had $\hat{\xi}_t = \sigma\theta_t - \hat{m}_{0t}$. In other words, if the vacancy filling cost is the relevant forcing variable, then matching efficiency should have an effect on inflation above and beyond the effect of the V/U ratio on inflation. This is what column (5) shows.

We can run a similar analysis using the MSA level data. While we are not able to construct a filling cost variable by geography, we can back out the MSA-level residual, $\hat{\mu}_t$, which captures shifts in the Beveridge curve. Specifically, to construct $\hat{\mu}_t$ 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 columns 2 and 3 include additional proxies for shifts in the Beveridge curve: either the inclusion of the log V/U

ratio or the inclusion of $\hat{\mu}$. The results show that both $\hat{\theta}$ and $\hat{\mu}$ are statistically significant, over and above the inclusion of $\hat{u}$.

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.

6 Non-linear effects of slack on inflation

In light of the post-COVID outburst in inflation, a number of recent work has argued that slack has non-linear effects of inflation; that the Phillips curve can steepen substantially in tight labor markets (Benigno and Eggertsson, 2023; Gitti, 2024).

Using our Beveridge curve decomposition (16), we can explore the sources of that non non-linearity. Specifically, for a Phillips curve with the log V/U ratio as forcing variable —our preferred specification—-, we have

$$\begin{aligned} \pi_t &= \mathrm{E}_t\pi_\infty + \beta_\theta(\theta_t)\hat{\theta}_t + \omega_t \\ &= \mathrm{E}_t\pi_\infty + \beta_u(\theta_t)\hat{u}_t + \beta_\mu(\theta_t)\hat{\mu}_t + \omega_t \end{aligned} \tag{17}$$

Clearly, if β_θ —the slope of the Phillips curve— depends on $\hat{\theta}_t$, then the Phillips curve is non-linear —a genuine non-linearity. However, equation (17) suggests another possibility: that the Beveridge curve shifts (μ_t) systematically in tight labor markets; when the V/U ratio is high.[13] 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 a non-linear Phillips curve. To test between these different possibilities, we run the regression:

$$\pi_t = \beta_\pi \mathrm{E}_t\pi_\infty + \beta_u\hat{u}_t + \gamma_u\hat{u}_t\mathbb{1}_{\theta_t>\bar{\theta}} + \beta_\mu\hat{\mu}_t + \gamma_\mu\hat{\mu}_t\mathbb{1}_{\theta_t>\bar{\theta}} + \omega_t \ , \tag{18}$$

using as threshold variable $\bar{\theta}$ the median of θ_t.[14]

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

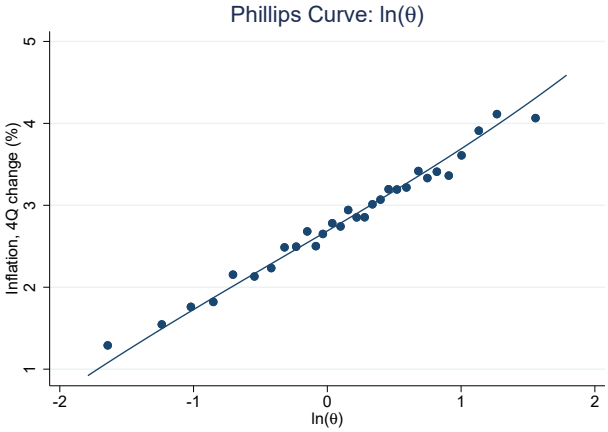
We can use the MSA-level variation to further explore the presence of non-linearity. However, the evidence for non-linearity is much weaker as the MSA level. In Table 7, we report tests for nonlinearities in $\hat{\theta}$ (column 2) as well as $\hat{\mu}$ (column 4).[15] The results show that the

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

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

[15] The threshold for the nonlinearity is 50th percentile by MSA, but results are generally similar when altering the thresholds as shown in Figure 6.

Figure 6: 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.

impact of $\hat{\theta}_t$ on inflation is quite linear, where the interaction term is positive but insignificant. Similarly, the effect of $\hat{\mu}_t$ shows minimal nonlinearity. This can be seen more clearly in Figure 6 which plots a binned scatter plot of the inflation rate (4-quarter change) and $\hat{\theta}$. There is little evidence of a non-linear 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 US 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.

As last word, we note an important caveat to this last finding: while we showed that Beveridge curve shifts correlate strongly with (and predict) inflation movements, and in particular that a decline in matching efficiency correlates with higher inflation, we did not establish a causal link. For that purpose, one would need to find instrumental variables that move matching efficiency and are independent of the other determinants of inflation. To the extent that the large decline in matching efficiency owes to the post-Covid Great Resignation and reconsideration of career choices and work-life balance, the post-2022 decline in matching efficiency could be interpreted as a convincing case in point of the effect of lower matching efficiency on inflation. Identifying the causal effect of matching efficiency on inflation is an important topic for future studies.

References

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

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

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

Barlevy, G., Faberman, R. J., Hobijn, B., and Şahin, A. (2023). The shifting reasons for beveridge-curve shifts. Technical report, National Bureau of Economic Research.

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

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

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. (2020). Measuring employer-to-employer reallocation. Technical report, National Bureau of Economic Research.

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 us 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.

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*, pages 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*.

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. (2008). Estimating the new keynesian phillips curve: a vertical production chain approach. *Journal of Money, Credit and Banking*, 40(4):627–666.

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

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)
u_t	-0.19^{***} (0.08	-0.84^{***} (0.12)	—	—	—	—	—	—
$\hat{\theta}_t$	—	—	0.28^{***} (0.08)	1.08^{***} (0.10)	—	—	—	—
$\hat{\chi}_t$	—	—	—	—	0.30^{***} (0.09)	1.07^{***} (0.11)	—	—
$\hat{\theta}^*_t$	—	—	—	—	—	—	0.35^{***} (0.09)	1.17^{***} (0.10)
Inflation	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
Adjusted R^2	0.223	0.614	0.267	0.755	0.265	0.731	0.291	0.790

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

Table 2: Philips Curve Estimates, 1960-2023

	(1)	(2)	(3)	(4)	(5)	(6)
u_t	-0.27^{***} (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.

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

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

	(1)	(2)	(3)	(4)	(5)
$\hat{u}_t$	-0.91^{***} (0.13)	1.45^{*} (0.75)	-0.42 (0.30)	-0.55^{**} (0.28)	—
$\hat{\theta}_t$	—	1.36^{***} (0.39)	—	—	0.52^{***} (0.13)
$\hat{\mu}_t$	—	—	-1.36^{***} (0.40)	—	—
$\hat{s}_t$	—	—	—	1.03^{*} (0.54)	—
$\hat{m}_{0t}$	—	—	—	-1.43^{***} (0.40)	-0.67^{***} (0.32)
$E_t\pi_\infty$	1.01^{***} (0.07)	0.88^{***} (0.07)	0.88^{***} (0.07)	0.89^{***} (0.07)	0.91^{***} (0.07)
Adjusted R^2	0.817	0.826	0.826	0.826	0.826

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)	
$\hat{\mu}$			-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)
$\hat{\mu}_t$	—	—	0.02
			(0.10)
$\hat{\mu}_t \mathbb{1}_{\theta>\bar{\theta}}$	—	—	−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)
$\hat{\mu}$			-0.122***	-0.0991*
			(0.0384)	(0.0541)
$\hat{\mu}$ 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

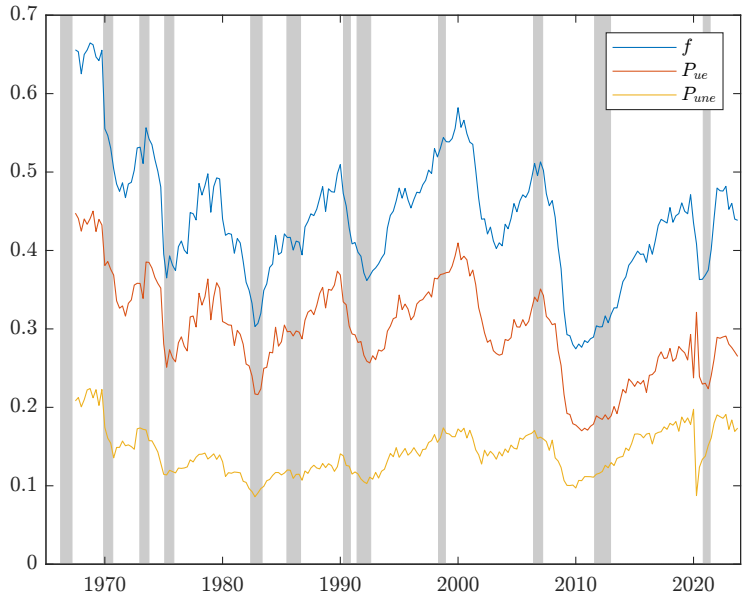
All variables z-scored (demeaned and normalized to unit standard-deviation) for comparability across columns. Controls include lagged inflation, the lagged ratio of the goods and services price level, and threshold dummies. Thresholds are the 50th percentile by MSA. Standard errors clustered by MSA

8 Appendix

UE and UIE rates

Figure 7 plots the UE and UIE rate. We can see that the two flow rates are highly correlated.[16] In fact, a matching function (going back to Pissarides, 1985) does a great job at capturing the behavior of f_t.

Figure 7: Unemployment outflow probabilities



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 $s_t = -\ln(1 - S_t)$.

[16] The raw correlation of 0.75 and a correlation of 0.90 after detrending with a quadratic polynomial.

Raw series for inflation and candidate forcing variables

Figure 8: Raw series 1995-2023

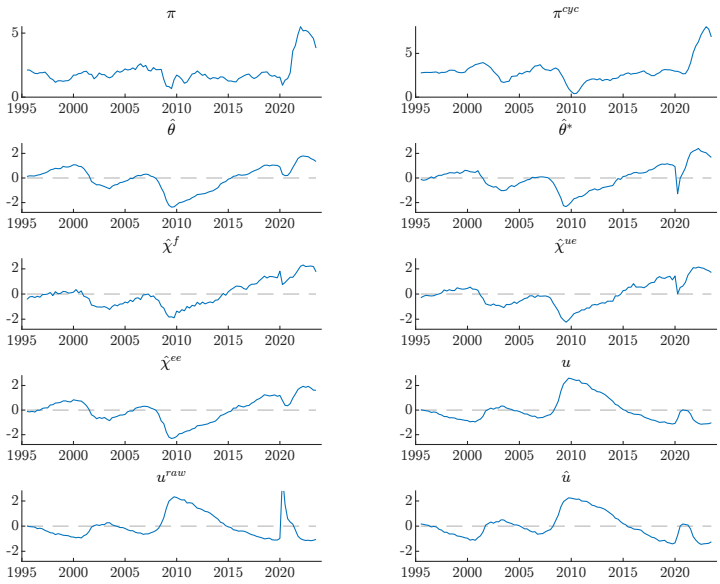


Figure 9: Raw series 1960-2023

