

NBER WORKING PAPER SERIES

THE UNEMPLOYMENT-INFLATION TRADE-OFF REVISITED:
THE PHILLIPS CURVE IN COVID TIMES

Richard K. Crump
Stefano Eusepi
Marc Giannoni
Ayşegül Şahin

Working Paper 29785
<http://www.nber.org/papers/w29785>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2022

This paper was initially presented at the January 3, 2021 ASSA session on “Measuring the Unemployment Gap.” We thank Emmanuel Saez and Pascal Michaillat for organizing the session and the discussant, Regis Barnichon. Jin Yan and Charles Smith provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Bank of New York, Federal Reserve Bank of Dallas, the Federal Reserve System, or the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Richard K. Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Unemployment-Inflation Trade-off Revisited: The Phillips Curve in COVID Times
Richard K. Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin
NBER Working Paper No. 29785
February 2022
JEL No. D84,E24,E31,E32,J11,J3

ABSTRACT

We estimate the natural rate of unemployment, often referred to as u^* , in the United States using data on labor market flows, short-term and long-term inflation expectations and a forward-looking New-Keynesian Phillips curve for the 1960-2021 period. The natural rate of unemployment was at around 4.5% before the onset of the pandemic and increased to 5.9% by the end of 2021. This pronounced rise was primarily informed by strong wage growth rather than changes in inflation expectations. Despite the rise in the natural rate of unemployment, the secular trend of unemployment continued to fall and stands at around 4.2% reflecting ongoing secular developments which have been pushing down the unemployment rate over the last 30 years. Our model forecasts strong wage growth to moderate only sluggishly continuing to put upward pressure on inflation in the medium-run. We project underlying inflation to remain 0.5 percentage points above its long-run trend by the end of 2023 even if long-run inflation expectations remain well anchored.

Given the importance of wage growth for the inflation outlook, we examine detailed micro data on job-filling rates, posted wages for vacant positions, and workers' reservation wages. In particular, we construct a composition-bias free measure of wage growth at the employer-job level using Burning Glass Technologies data and document strong wage growth for both teleworkable and non-teleworkable jobs. Moreover, we find that workers' reservation wages increased substantially after the pandemic. Our empirical analysis suggests that the strong wage growth is likely not a one-time adjustment of additional compensation for jobs that pose health risks to workers but rather reflects a tight labor market accompanied with a changing work-leisure trade-off.

Richard K. Crump
Capital Markets Function
Federal Reserve Bank of New York
33 Liberty Street
New York, NY 10045
richard.crump@ny.frb.org

Stefano Eusepi
Department of Economics
University of Texas at Austin
eusepi73@yahoo.com

Marc Giannoni
Federal Reserve Bank of Dallas
Research Department
2200 N Pearl St.
Dallas, TX 75201
United States
Marc.Giannoni@dal.frb.org

Ayşegül Şahin
Department of Economics
University of Texas at Austin
2225 Speedway
Austin, TX 78712
and NBER
aysegul.sahin@austin.utexas.edu

1 Introduction

The longest expansion in the postwar period came to an abrupt end in February 2020 with the emergence of the novel coronavirus and the broad implementation of lockdowns across the US. The unemployment rate climbed from its fifty-year low level of 3.5% to 14.7% in a matter of weeks. Despite being the deepest recession, the COVID-19 recession was the briefest downturn in postwar history as the acute disruptions in the macroeconomy reversed rapidly. Economic activity picked up briskly in 2021: real GDP grew at an annualized rate of 6.9% in the fourth quarter, job openings reached 11 million in December, and the unemployment rate declined to 4.0% in January 2022. This rapid pickup in economic activity was also accompanied by a steep rise in inflation. In January 2022, CPI inflation rose 7.5% on a year-over-year basis—a rate not seen in almost 40 years.

This brief but deep recession and its brisk recovery has brought the US economy to unfamiliar territory. Not since the 1960s has the US economy experienced such high levels of price inflation accompanied with such low levels of unemployment. The 1970s, generally referred to as the stagflation period, ended with both the unemployment rate and price inflation reaching double digits. The 1980s, the Volcker disinflation period, was a decade where both inflation and the unemployment rate declined precipitously. These tumultuous decades were followed by three decades of quiescent inflation, despite substantial fluctuations in the unemployment rate. This thirty-year long experience of subdued inflation, seemingly immune to labor market fluctuations, was interpreted as the weakening of the unemployment-inflation trade-off—often referred to as the death of the Phillips curve. The developments of the last two years—the dramatic decline in the unemployment rate and the sharp increases in inflation—have brought discussions of a changing unemployment-inflation trade-off back to the fore.

A useful construct to gauge the unemployment-inflation trade-off is the so-called *natural rate of unemployment*, which is defined as the unemployment rate such that, controlling for supply shocks, inflation remains stable. The natural rate of unemployment, which we also refer to as u_t^* , is affected both by business-cycle fluctuations and secular factors. Furthermore, the unemployment-inflation trade-off is linked by the classical determinants of inflation such as inflation expectations. To accommodate all of these facets, a comprehensive framework is required. We developed such a framework in [Crump et al. \(2019\)](#) where the natural rate is informed by wage and price inflation, inflation expectations, and changing secular factors. This Micro-Macro Phillips Curve framework not only creates a clear link between the labor market and inflation, it also directly incorporates the movements in

survey-based inflation expectations.

We estimate the natural rate of unemployment, u_t^* , over the period 1960 to 2021. We find that our Phillips curve captures the joint behavior of unemployment, wage and price inflation, and inflation expectations very well with a time-invariant slope—estimated to be quite flat. Our approach is flexible enough to overcome common empirical obstacles such as the unavailability of a universally accepted measure of wages free of composition bias. Instead we utilize five different measures of wage inflation representing noisy observations of latent underlying wage growth. The estimation relies on two key ingredients: First, we propose a measure of the secular trend in the unemployment rate, which we refer to as \bar{u}_t , obtained from separation (unemployment inflow) rates and job-finding (unemployment outflow) rates. We exploit the rich cross-sectional variation in the flow rates of different demographic groups to obtain more precise estimates of the underlying trends in a state-space framework. This directly aids in the measurement of the unobserved natural rate of unemployment.¹ Second, we use survey-based forecasts to measure the term structure of inflation expectations, that is, the forward-looking component of the Phillips curve. We conclude that it is also vital to account for the behavior of expectations to reconcile the observed behavior of inflation and slack over time as in [Del Negro, Giannoni, and Schorfheide \(2015\)](#), [Crump et al. \(2019\)](#), and [Carvalho, Eusepi, Moench, and Preston \(2021\)](#).

We find that the natural rate of unemployment has risen appreciably since the start of the pandemic, from 4.5% at the end of 2019 to 5.9% at the end of 2021. We conduct counterfactual analyses and find that strong wage growth accounts for the vast majority of the rise in u_t^* from 2019 to 2021. Other factors such as rising short-term inflation expectations or the behavior of inflation itself only plays a minor role. The behavior of u_t^* stands in stark contrast to the secular trend in unemployment which continued its downward drift unabated through the pandemic to 4.2%—its lowest level in 60 years. The last time we observed both the natural rate well above the secular trend along with high inflation was in the 1970s. Another notable similarity between the 1970s and 2021 is the existence of the large negative unemployment gap, the difference between the actual unemployment rate and the natural rate of unemployment. In both late 1970s and in the last quarter of 2021, this gap stood at around -1.5 percentage points. These two factors point to inflationary pressures for the US economy.

We forecast the evolution of price and wage inflation to the end of 2023. Our forecast is charac-

¹The main drivers of this downward trend can be traced to *grand gender convergence* and *dual aging* of workers and firms. See [Section 2.1](#) for a brief discussion and [Crump et al. \(2019\)](#) for an in-depth analysis using detailed micro data.

terized by elevated price and wage inflation which only gradually reverts to pre-pandemic levels. Our model projects underlying inflation to remain 0.5 percentage point above its long-run trend by the end of 2023 even if long-run inflation expectations remain well anchored. Importantly, the forecast path is directly linked to the expected path of the unemployment gap which is, in turn, is measured using information from all of our different sources including inflation expectations, the secular trend in the unemployment rate, five different wage series, and price inflation.

The reemergence of high inflation readings has drawn comparisons to the US economy in the 1970s. However, we see two key differences between the 1970s and the current state of the economy. First is the notable difference in the level *and* the bearing of the secular trend of unemployment. The secular trend of unemployment is now low following a period of persistent declines whereas in the 1970s it was high following persistent escalations. This indicates that the secular pressures on the unemployment rate that contributed to the high unemployment environment in the 1970s are no longer present. Second, the behavior of long-term inflation expectations over the two periods is markedly different. [Carvalho et al. \(2021\)](#) show that inflation expectations were unmoored over the 1970s, but have remained firmly anchored over the past twenty years. The anchoring of expectations, as measured by their sensitivity to short-term inflation forecast surprises, depends on the size and persistence of current and past forecast errors. Large forecast errors indicate instability in the long-run mean of inflation: in response, long-term inflation expectations become closely linked to recent forecast errors, as agents track a new inflation regime. Conversely, anchored expectations are fairly unresponsive and stable around the central bank's inflation objective. The expansionary policies during the mid-1960s and the 1970s then produced an upward trend in inflation expectations which eventually led to stagflation and, through the 1980s, to the costly Volcker disinflation. The current set of expansionary policies are, for the moment, associated with fairly stable long-run inflation expectations, reflecting the higher degree of credibility afforded to the Federal Reserve.

It might appear that with this 1.5 percentage point negative unemployment gap, we would expect inflation to rise further in the absence of a sharp deterioration in labor market outcomes—a remake of the 1970s and 1980s. Said differently, wouldn't we expect a high unemployment cost to reducing inflation with a flat Phillips curve and a negative unemployment gap? The anchoring of expectations is tightly connected with the costs of reducing inflation back to target. Given the forward-looking nature of our estimated Phillips curve, a flat slope does not necessarily imply high costs of disinflation. Instead, there is a much more important role for long-term inflation expectations along with the

expected future path of the unemployment gap. Monetary policy is the key factor driving these expectations. An (expected) credible monetary policy response would both keep inflation anchored (i.e. disconnected from the current surge in inflation) and induce a faster closing of the unemployment gap. This, in turn, would lead to a reduction in inflation for a given level of the current unemployment gap. In this buoyant scenario the cost of disinflation in terms of unemployment would be small. In sharp contrast, a loss of central bank credibility would require a large increase in the current unemployment gap to offset the sluggish adjustment of expectations.

Our analysis of the Phillips curve highlights the key role of wage growth for the evolution of u_t^* and the inflation outlook. We further investigate the current and prospective path of wages using a variety of cross-sectional data sources to provide external corroboration for our analyses and forecast. First, we show a strong historical linkage between underlying wage growth and the job-filling rate. At the end of 2021, the job-filling rate was near its all-time low and broadly in-line with its downward path through the previous expansion, suggesting that underlying wage growth is likely to continue to be firm. Second, we study wage growth based on vacancies posted by the same firm for the same job over time, to minimize the role of composition bias and influence of worker characteristics. This measure also provides a forward-looking measure of wage growth since it captures new hires' wages as documented by [Hazell and Taska \(2020\)](#). We find that posted wage growth shows similar patterns across occupations with different exposures to the pandemic, suggesting broad-based changes, rather than additional compensation for jobs that pose health risks to workers. Finally, we show that survey responses of reservation wages have been rising broadly since 2018, and through the pandemic, providing further evidence that some key aspects of wage dynamics represent the continuation of pre-pandemic behavior as well as a change in willingness to work. Taken in sum, this cross-sectional evidence is consistent with the time-series evidence we present from our model.

The structure of the paper is as follows. Section [2.1](#) estimates the secular trend of unemployment using detailed information for unemployment inflows and outflows by demographic groups and provides a brief discussion of its drivers. Section [2.2](#) introduces a forward-looking wage Phillips curve, discusses its theoretical underpinning, and details the estimation methodology. Section [3](#) presents the time series for the natural rate of unemployment, u_t^* for 1960-2021 including a detailed discussion of the pandemic economy and presents forecasts for wage and price inflation. Section [4](#) provides micro-data based evidence on the state of the labor market and wage growth using a variety of

additional data sources. Section 5 concludes.

2 A Micro-Macro Phillips Curve Framework

We estimate u^* in two steps, succinctly summarized in the following equation:

$$u_t = \bar{u}_t + \underbrace{(u_t - u_t^*)}_{x_t} + \underbrace{(u_t^* - \bar{u}_t)}_{z_t}. \quad (1)$$

building on the methodology we introduced in [Crump et al. \(2019\)](#). In the first step we extract \bar{u}_t , the secular trend in the unemployment rate, from in the inflow and outflow rates using a linear unobserved factor model. In the the second step, we combine this trend estimate, together with measures of price inflation, wage inflation, and inflation expectations to infer the natural rate of unemployment, u_t^* along with the unemployment gap, x_t , from a New-Keynesian Phillips curve. The natural rate of unemployment is defined as the unemployment rate such that, controlling for supply shocks, inflation remains stable. It is therefore conceptually distinct from \bar{u}_t and its deviations from the secular trend measured by the variable z_t .

2.1 The Secular Trend in the Unemployment Rate, \bar{u}_t

Our main premise is that the flow origins of unemployment rate movements help us better connect to the underlying drivers of unemployment fluctuations and trends. Therefore we start with the evolution of the unemployment stock from month t to month $t + 1$

$$dU/dt = s_t(L_t - U_t) - f_t U_t \quad (2)$$

where L_t denotes the labor force, s_t is the separation rate (inflow rate) to unemployment and f_t is the job-finding rate (outflow rate) from unemployment. While s_t is generally referred to as the separation rate and f_t as the job-finding rate, we will use the inflow-outflow terminology as in [Elsby, Michaels, and Solon \(2009\)](#) and [Elsby, Hobijn, and Şahin \(2010\)](#). This terminology creates a clear differentiation between s_t and f_t and employment-to-unemployment and unemployment-to-employment flow rates based on gross flows data computed using longitudinally matched monthly CPS microdata.²

²It is important to note that we focus on a two-state representation of unemployment where we do not explicitly differentiate between the source of unemployment inflows and destination of unemployment outflows following [Shimer](#)

The unemployment rate, u_t is defined as the fraction of the labor force L_t that is unemployed, $u_t = U_t/L_t$. We follow [Shimer \(2005, 2012\)](#) and calculate the outflow probability F_t using the observation that

$$U_{t+1} - U_t = U_{t+1}^S - F_t U_t$$

where U_{t+1}^S is the number of unemployed who report having been unemployed for less than one month. Solving for F_t ,

$$F_t = 1 - \frac{U_{t+1} - U_{t+1}^S}{U_t}$$

which can be mapped into a Poisson outflow hazard rate $f_t = -\log(1 - F_t)$. The idea behind this calculation is intuitive: individuals who reported being unemployed for less than one month were not in the unemployed pool in the previous month and therefore subtracting them out from this month's unemployment pool leaves us with the unemployed who failed to exit unemployment between month t and month $t + 1$. Solving the differential equation (2) forward as in [Shimer \(2012\)](#), we can solve for the unemployment inflow rate s_t

$$U_{t+1} = \frac{(1 - e^{-[s_t+f_t]})s_t}{s_t + f_t} L_t + e^{-[s_t+f_t]} U_t.$$

Given the fast transitional dynamics of the unemployment rate in the U.S., as noted by [Shimer \(2005\)](#), [Elsby, Michaels, and Solon \(2009\)](#) and others, the unemployment rate is closely approximated by its flow steady-state value given by

$$\frac{s_t}{s_t + f_t}. \tag{3}$$

The Current Population Survey (CPS) provide us with monthly measures of stock of unemployment, short-term unemployment and labor force. We calculate monthly unemployment inflow and outflow hazard rates using the methodology described above. We estimate the slow-moving trend in the inflow and outflow rates using six different demographic groups for each rate: the interaction between gender and age grouped by 16-24, 25-54, and 55 and above.³

Figure 1 shows the aggregate inflow rate, outflow rate and unemployment rate along with their corresponding estimated secular trends, \bar{s}_t , \bar{f}_t and \bar{u}_t for the whole sample 1960-2021. The secular trend of the inflow rate shows a decline of about 50% since the 1980s. In contrast, the secular trend

(2005, 2012), [Hall \(2005\)](#), [Elsby, Michaels, and Solon \(2009\)](#), [Elsby, Hobijn, and Şahin \(2010\)](#), [Davis, Faberman, Haltiwanger, Jarmin, and Miranda \(2010\)](#) and [Şahin, Song, Topa, and Violante \(2014\)](#).

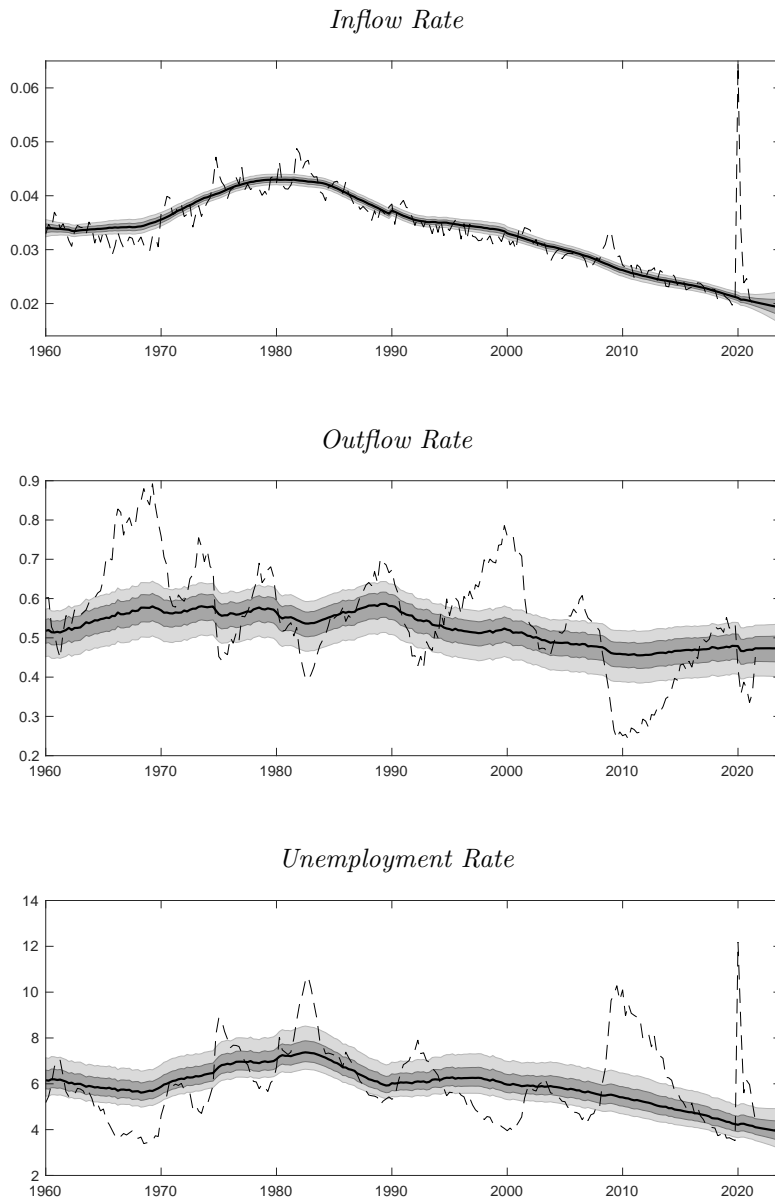
³We correct for the effects of CPS redesign on duration of unemployment using the correction factors in [Elsby, Hobijn, and Şahin \(2010\)](#).

in the outflow rate is generally stable, but has fallen since the 1990s consistent with the evidence presented in [Davis, Faberman, Haltiwanger, Jarmin, and Miranda \(2010\)](#). Finally, the secular trend in the unemployment rate, \bar{u}_t , can be constructed using \bar{s}_t and \bar{f}_t and the steady-state approximation to the unemployment rate, via

$$\bar{u}_t = \frac{\bar{s}_t}{\bar{s}_t + \bar{f}_t}, \tag{4}$$

and is shown in the bottom panel of [Figure 1](#). The trend unemployment rate was about 6% in 1960 and increases to over 7% in 1983. Since then it has displayed a clear downward trend, reaching about 4.2% by the end of 2021. Interestingly, this downward trend continued even after the dramatic job losses of the Great Recession and the pandemic recession, underscoring the importance of secular trends in the labor market. Since the outflow rate shows little trending behavior we observe from [equation \(4\)](#) that the overall downward trend is driven by the numerator, \bar{s}_t . The secular trend in the unemployment rate is estimated with a reasonably high degree of precision; for example, the 68% confidence interval at the end of the sample is comfortably less than one percentage point.

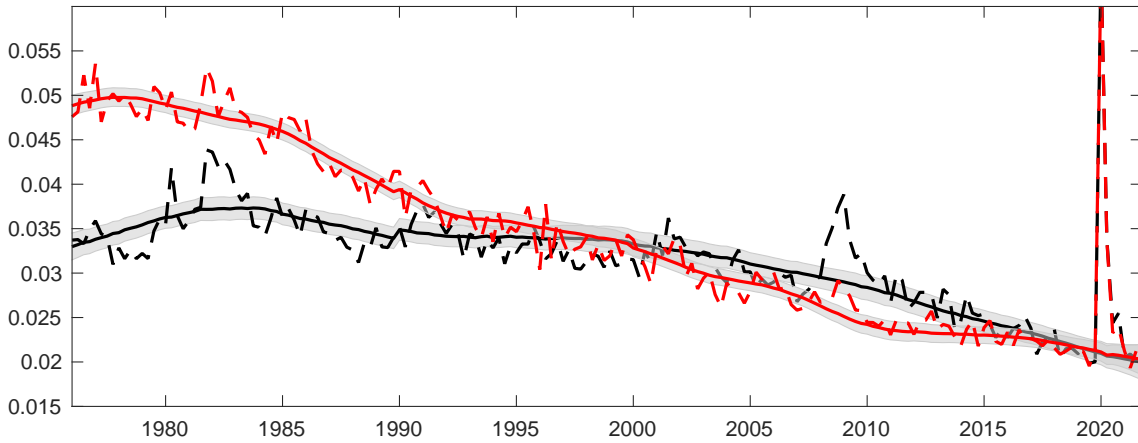
Figure 1. Inflow Rate, Outflow Rate and Unemployment Rate along with Estimated Secular Trend Actual rates denoted by dashed lines, median estimates of secular trend (\bar{s}_t , \bar{f}_t , and \bar{u}_t) denoted by solid black lines. Shading denotes 68% and 95% coverage intervals.



Our earlier work, [Crump et al. \(2019\)](#), has identified important changes in the labor market in the last 40 years as the drivers of the declining incidence of unemployment: grand gender convergence and dual aging. We provide a brief summary of these sources since the drivers of the trend are informative about the evolution of the unemployment rate in the medium-run.

Grand Gender Convergence The U.S experienced *Grand Gender Convergence* in the 20th century with female labor participation increasing from around 47% in 1976 to approximately 60% in

Figure 2. Unemployment Inflows by Gender. This figure displays the realized inflow rate along with the estimated secular trend for men (black line) and women (red line). Actual rates are denoted by dashed lines whereas solid lines indicate the median estimates of the secular trend. Grey shaded areas denote the 95% coverage interval.



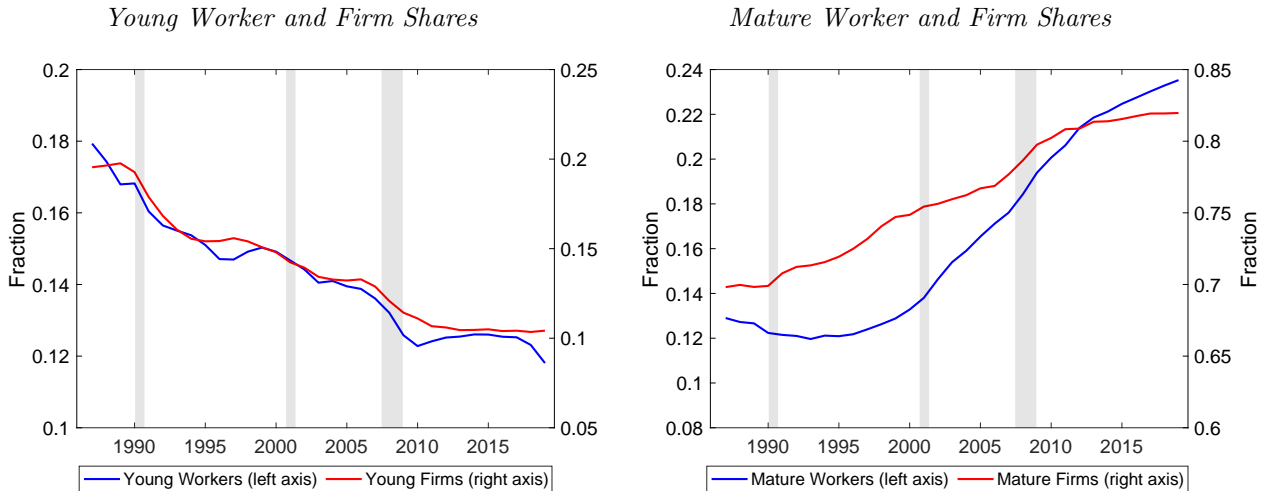
2000 (Goldin (2006)). The main driver of the rise in the female labor force participation rate was the increase in participation of married women with children. Women started to work longer into their pregnancy and started working after childbirth sooner than their counterparts in the 1960s likely due to changes in social norms, more widespread availability of maternity leave, and advances in maternal health and childcare. As labor market interruptions declined, women’s labor force attachment gradually increased. Having uninterrupted employment spells allowed women to build more stable employment relationships. This reduced frictional unemployment through a decline in the incidence of job loss and incidence of unemployment during re-entry into the labor force. Figure 2 shows the unemployment inflow rate by gender and its estimated trend. By the late 1990s the unemployment inflow rate for women converged to men, driving down the secular trend of unemployment in the 1980s and 1990s. The importance of gender convergence was relatively minor after 2000. This is when another prominent trend—dual aging—took over.

Dual Aging The U.S. economy has been experiencing a striking shift towards older workers and older firms since the mid-1990s which we refer to as *dual aging*. While the change in worker demographics is directly attributable to the baby boom, the drastic increase in births following World War II, the emphasis on aging of firms is relatively new as data have only recently become available. The intuition is very similar for firms: declining firm births almost fully account for the shift of employment towards older firms. Moreover, Karahan et al. (2022) establish a clear link from worker

to firm demographics. They show that the origin of the decline in firm entry is the decline in labor supply growth arising from the aging of the baby boom cohort and the flattening out of the female labor force participation rate.

The aging pattern is stark. Around 18% of the labor force was comprised of workers between 16 to 24 years old (*young workers*) in 1987. By 2019, this fraction declined to comfortably below 10%. *Young firms*' (firms younger than 5 years old) employment share also followed a similar pattern with their employment share declining from around 20% to 10%. On the flip side, in 1987, firms 11 or more years old—which we refer to as *mature firms*—used to employ around two thirds of the workers in the economy. By 2019, that fraction increased to over 80% as seen in Figure 3.

Figure 3. Dual Aging of Workers and Firms. This figure shows the employment share of 16-24 year old workers and the employment share of 5 year old and younger firms (left plot) along with the employment share of 55+ year old workers and the employment share of firms 11 or more years old (right plot). The sample period is 1987–2019.



Younger workers are four times more likely to flow into unemployment than prime-age workers. Similarly, firms aged between one and five years old are twice as likely to destroy jobs than their older counterparts. These patterns suggest that a direct consequence of dual aging is a decline in unemployment and job destruction. While the shift in worker and firm age composition falls short of accounting for the decline in the inflow rate, aging also affects the economy by affecting age-specific outcomes. Put differently, in economies with older workers and firms, unemployment and job destruction is lower even for all workers. Using state-level variation and an instrumental variables approach, in Crump et al. (2019) we show that a 1 percentage point increase in mature firm share lowers the job destruction rate by 0.60 percentage points for younger firms.

While grand gender convergence was important in accounting for the secular decline in the un-

employment rate until 2000, dual aging stands out as an important driver of the declining secular trend rate of unemployment in the last two decades. Given that the workforce and firms are still aging, the decline in the secular trend of unemployment is likely to continue in the medium-run.

2.2 Wage Phillips Curve

To measure the natural rate of unemployment we combine a simple statistical model for the evolution of $z_t = u_t^* - \bar{u}_t$ and the unemployment gap $x_t = u_t - u_t^*$ in (1) with a New Keynesian Phillips curve connecting wage (π_t^w) and price (π_t) inflation to the unemployment gap.⁴ Following Galí (2011), both prices and wages are set in an environment where firms and workers have some market power. While prices are set in the absence of nominal rigidities, nominal wages are sticky. Wage and price inflation evolve according to

$$\pi_t^w = g_w + \pi_t^* + \gamma(\pi_{t-1} - \pi_{t-1}^*) - \kappa x_t - \kappa\beta\mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} x_{s+1} + \varphi\Delta a_t \quad (5)$$

$$\pi_t = \pi_t^w - g_w - \Delta a_t \quad (6)$$

$$\Delta a_t = \rho_a \Delta a_{t-1} + \sigma_a \epsilon_t^a. \quad (7)$$

Wage and price inflation are determined by five core factors. Real wage growth is driven by productivity and price markup shocks captured by the process⁵ Δa_t . Nominal wages are partially indexed to past inflation measured by the parameter $0 \leq \gamma < 1$. Because of nominal rigidities, wage setting is forward-looking and depends on the expected path of the unemployment gap (discounted at the rate $0 < \beta < 1$). The slope of the Phillips curve, measured by $\kappa > 0$, regulates inflation's responsiveness to current and expected unemployment gaps. A second crucial forward-looking component is agents' estimate of the long-run inflation mean, π_t^* , which serves as a proxy for the degree of expectations' anchoring. This process, modeled as a random walk, induces shifts in the relationship between inflation and the unemployment gap. Reflecting the forward-looking nature of the New Keynesian Phillips curve, the future expected path of the unemployment gap is a key driver of wage inflation. In fact, inflation expectations contain information about expected future unemployment

⁴In this simple model the natural rate of unemployment is driven by market distortions captured by shifting market power of workers and other factors.

⁵The autocorrelation coefficient ρ_a is restricted to be between zero and one. The parameter $\varphi > 0$ depends on both ρ_a and β as it measures the expected discounted path of a_t .

gaps: using equations (5) through (7) we get

$$\mathbb{E}_t \pi_{t+1} = g_w + \pi_t^* + \gamma(\pi_t - \pi_t^*) - \kappa \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} x_{s+1} + (\varphi - 1) \rho_a \Delta a_t. \quad (8)$$

2.3 Estimating the Wage Phillips curve

Assume under ideal conditions we can perfectly observe price and wage inflation; long-run inflation expectations (π_t^*); short-term inflation expectations ($\mathbb{E}_t \pi_{t+1}$); and the model’s key parameters. Then it is straightforward to see that the unemployment gap can be obtained just by “inverting” the wage Phillips curve using equations (5), (6), and (8).

We strive to get as close as possible to this ideal scenario by collecting a wealth of information on each the model’s components. In addition to the unemployment rate (and its estimated trend \bar{u}_t) and a measure of CPI inflation, we use five different measures of wage inflation, together with short- and long-term inflation expectations from professional forecasters⁶. However, we face two challenges. First, information about wage growth and inflation expectations contains significant measurement errors. In the case of wages, this is evident from the fact that we use multiple measures for the same variable. This limitation implies that we can only infer the unemployment gap with some degree of uncertainty. Additionally, we do not have strong prior information about a key parameters such as the slope of the Phillips curve, κ , which needs to be estimated. In fact, a large literature focused on estimating the Phillips Curve—see, for example, [Mavroeidis et al. \(2014\)](#) for a comprehensive discussion.

Estimation of the model comprising equations (1) and (5)–(7) requires additional identifying assumptions: in particular, a law of motion for the joint behavior of x_t and z_t . Similar to [Laubach \(2001\)](#), [Galí \(2011\)](#), and [Laubach and Williams \(2003\)](#) we model these unobserved components as exogenous process:

$$z_t = \rho_z z_{t-1} + \sigma^z \epsilon_t^z \quad (9)$$

$$x_t = a_{x,1} x_{t-1} + a_{x,2} x_{t-2} + \sigma_x \epsilon_t^x. \quad (10)$$

This specification allows for persistent deviations of u_t^* from the secular trend, but imposes that, over the longer run, these deviations shrink toward zero. In addition to producing an estimate for the the natural rate of unemployment, these additional modeling assumptions also allow us to conduct

⁶Details about the dataset can be found in the Appendix.

the simple inflation forecasting exercise discussed in Section 3.2. Using the model it is straightforward to produce time- t forecasts at horizon $n > 1$ (and the associated forecast distribution) for inflation and the natural rate of unemployment:

$$\pi_{t+n|t} = \pi_{t|t}^* + \phi_{\{\pi,n\}}(\pi_t - \pi_{t|t}^*) + \phi_{\{x1,n\}}x_{t|t} + \phi_{\{x2,n\}}x_{t-1|t} + \phi_{\{a,n\}}\Delta a_{t|t} \quad (11)$$

$$u_{t+n|t}^* = \rho_z^n z_{t|t} + \bar{u}_{t+n|t} \quad (12)$$

where the coefficients $\phi_{\{i,n\}}$ capture the model solution consistent with the data generating process for the output gap in equation (10). Variables $y_{\cdot|t}$ denote estimates of the unobserved states using information up to the current period. A useful object that we can construct based on equation (11) is our measure of “underlying inflation,” defined as $\pi_{t+n|t} - \phi_{\{a,n\}}\Delta a_{t|t}$, linking inflation only to the unemployment gap. Another measure that we consider in the next section is the inflation gap, or the difference between underlying inflation and the estimated inflation trend, $\pi_{t|t}^*$.

The model is estimated with Bayesian methods over the sample 1960:Q1–2019:Q4 using quarterly data. Details about the estimation approach can be found in [Crump et al. \(2019\)](#). It is important to emphasize that our estimate of the unemployment gap $x_{t|t}$ reflects *all available information* through model linkages. This includes the trend in unemployment (that is, its estimate) even though it does not directly appear in the Phillips curve. It is also useful to discuss three key parameters that greatly affect the behavior of inflation and the estimate natural rate of unemployment. The slope of the Phillips curve κ is precisely estimated and in the range 0.02-0.04 which implies a fairly flat curve, as often found in the literature (for recent papers see, for example, [Negro et al. 2020](#) and [Hazell et al. \(2021\)](#)). We find little evidence for inflation inertia, with an estimate of $\gamma \in (0, 0.1)$, so that the behavior of inflation is highly forward-looking. Finally, our estimate of $\rho_z \in (0.96, 0.99)$ indicates persistent deviations of the the natural rate of unemployment from its historical trend, suggesting that changes in medium-term productivity or the structure of the labor market play an important role beyond the slow-moving demographic factors captured by our estimate of \bar{u}_t .

3 The Phillips Curve and the Natural Rate: 1960-2021

Our estimated framework allows us to examine the evolution of the natural rate of unemployment, u_t^* since 1960. Figure 4 shows that the natural rate hovers slightly below 6% and starts rising in the early 1970s. u_t^* continues to rise, reaching comfortably above 7% by the late 1970s, before falling to

about 7% in 1983. The natural rate then declines sharply throughout the 1980s. The time period spanning the 1990s to the Great Recession is characterized by a slight upward drift in the natural rate of unemployment, with a range of 4.8% to 5.8%. In the Great Recession the natural rate increased to above 6.2% but after reaching its peak, began a steady descent to a little above 4% in 2017. As the expansion matured, u_t^* reversed course and began to rise and was at around 4.5% at the end of 2019. During the COVID-19 pandemic, u_t^* has risen appreciably to 5.9% at the end of 2021.

We first discuss the roles of the key ingredients of our Phillips curve framework, the secular trend of unemployment and inflation expectations, in Section 3.1 and end with a detailed discussion of the effects of the pandemic on the natural rate and the outlook for inflation in Section 3.2.

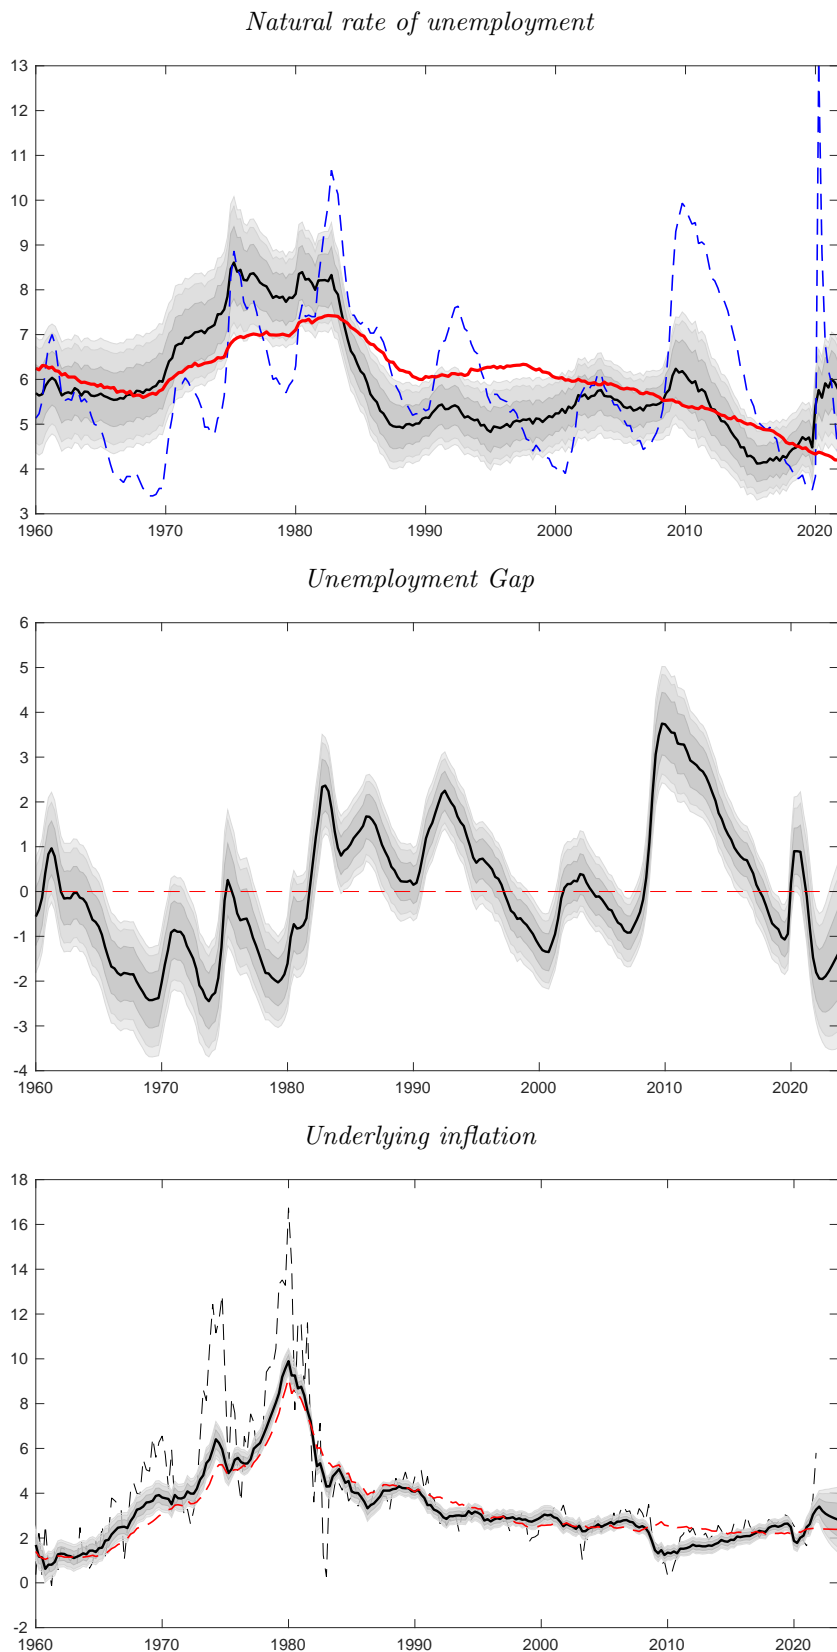
3.1 The Role of Inflation Expectations and Secular Unemployment Trends

The main novelty of our framework is its explicit use of data on inflation expectations and incorporation of micro-data based secular trend of unemployment into a forward looking New Keynesian Phillips curve framework. We find that allowing for these two ingredients capture the joint evolution of unemployment and inflation very well in the last 60 years. We discuss how our framework provides a reconciliation of seemingly different unemployment-inflation trade-offs in different episodes. Figure 4 plots the time series for the whole sample period and Table 1 summarizes our findings focusing on five time periods separately for ease of comparison: (1) 1970-79 (*Stagflation*), (2) 1980-89 (*Volcker disinflation*); (3) 1990-2006 (*Great Moderation*); (4) 2007-2019 (*Great Recession*); (5) 2021Q4 (the *Covid Pandemic*).⁷

Role of Inflation Expectations A key advantage of our framework is our explicit use of data on short-term and long-term inflation expectations. If one ignores the de-anchoring in the 1970s and subsequent anchoring of inflation expectations in the 1990s one might conclude that the unemployment-inflation trade-off was time varying. Put differently, the Phillips Curve would appear steeper in the 1970s and the 1980s before becoming flatter thereafter. This is particularly relevant after the pandemic recession since we observe sharp movements in both unemployment and inflation. However, our Micro-Macro framework is consistent with periods of large slack in the labor market and relatively stable inflation. This is perfectly illustrated in Table 1 with the comparison of the Volcker disinflation period with the Great Recession period (2007-2019). The unemployment gap, x_t aver-

⁷Since the COVID-19 pandemic triggered a deep but brief recession in the US labor market, averaging over 2020 and 2021 masks stark movements in the measures we report. Instead, we report our most recent estimates in Table 1 for the pandemic recession.

Figure 4. Phillips Curve Estimates This figure shows the main model outputs from our estimated Phillips curve. The top panel shows the estimate of the natural rate of unemployment, u_t^* (black line) along with the observed unemployment rate (blue dotted line) and the median secular trend, \bar{u}_t (red line). The middle panel shows the corresponding unemployment gap (black line). The bottom panel shows the model estimate of underlying inflation (black line) along with the Federal Reserve Bank of Cleveland's median CPI inflation (grey dashed line) and the long-run component of inflation, π_t^* (red dashed line). Grey shaded areas denote 68%, 90% and 95% coverage intervals. The sample period is 1960-2021.



aged at 0.75 during 1980-89 while it was 1.3 over 2007-2019. While core CPI inflation declined by 6.3 percentage points during the Volcker disinflation period, it remained largely unchanged and increased mildly by 0.2 percentage points in the 2007-2019 period. This observation is often attributed to the flattening of the Phillips curve since shifting inflation expectations are not incorporated in the estimation. However, as Table 1 shows most important for the stability of inflation is the fact that inflation expectations declined only modestly during and after the Great Recession, while the decline was stark in the 1980s. As indicated in our Phillips curve, equation (8), inflation expectations reflect the expected path of future unemployment gaps, and so the near-stability of inflation expectations in the aftermath of the Great Recession suggests that the unemployment gap was expected to close. This is consistent with the attenuated response of inflation to the large and persistent unemployment gap.

Our analysis of the Great Recession, through the lens of our estimation results, does not, however, imply that inflation is necessarily insensitive to the unemployment gap. In fact, we see that a somewhat smaller rise in the unemployment gap in the early 1980s caused a much more significant drop in inflation. The key determinant is the behavior of long-run inflation expectations, which dropped much more sharply in the 1980s (by 4.5 percentage points) than was the case following the Great Recession (only a decline of 0.20 percentage points). The comparison of the early 1980s with the Great Recession period stresses the importance of accounting for inflation expectations in explaining the behavior of inflation and the unemployment gap, and hence to estimate u_t^* .

Table 1. The Phillips Curve Over Time This table summarizes the behavior of key objects related to the unemployment-inflation trade-off. The rows labelled \bar{u}_t , u_t^* , x_t report the average value for the median estimated secular trend of unemployment, the natural rate of unemployment, and the unemployment gap over the specified time periods. The row labelled $\Delta\pi_t^{\text{xfe}}$ reports the change in year-over-year core CPI growth from the beginning to the end of the specified time periods. The row labelled $\Delta\pi_{t|t}^*$ reports the change in the median estimated inflation trend. For the column labelled “2021Q4” the last two rows represent the change as compared to 2020Q4.

	<i>Stagflation Period</i> 1970–1979	<i>Volcker Disinflation</i> 1980–1989	<i>Great Moderation</i> 1990–2006	<i>Great Recession</i> 2007–2019	<i>COVID Pandemic</i> 2021Q4
\bar{u}_t	6.63	6.86	6.07	5.11	4.20
u_t^*	7.57	6.52	5.27	5.01	5.86
x_t	-1.35	0.75	0.22	1.29	-1.56
$\Delta\pi_t^{\text{xfe}}$	4.85	-6.32	-1.71	0.23	3.39
$\Delta\pi_{t t}^*$	5.63	-4.48	-1.75	-0.20	0.08

Role of Secular Trend of Unemployment There are long-lasting social and demographic changes in the economy which affect the unemployment rate, such as women’s labor force attachment and aging of workers and firms as we discussed earlier.⁸ While these trends are arguably not affected by the business cycle, they affect the long-term behavior of the unemployment rate and the natural rate of unemployment. We have captured these secular trends through the variable \bar{u}_t introduced in Section 2.1. This “anchor” to the natural rate of unemployment then also becomes a relevant variable for assessing the state of the labor market and, in particular, achievable targets for sustainable unemployment levels. For example, the 1970s, which were characterized by high unemployment and inflation, also coincided with a rising \bar{u}_t when the secular trend of unemployment averaged at 6.6% as noted in Table 1.⁹ Figure 4 shows that our estimate of the natural rate of unemployment was consistently above \bar{u}_t over this period as a consequence of highly accommodative monetary policy. This pattern reverses starting in the early 1980s as the natural rate of unemployment remained below its secular trend for almost three decades. This is a reflection of anchoring inflation expectations, higher productivity, and the disinflationary effect of rising import penetration.¹⁰ Against this backdrop, the secular trend in the unemployment rate had a local peak in the late 1990s and has steadily declined since. Despite the dramatic rise in the unemployment rate during the Great Recession, the natural rate of unemployment gradually declined roughly in line with its secular trend.

Figure 4 also makes clear that secular trend of unemployment is, as of the end of 2021, at its lowest level in the sample at 4.2%. This suggests that despite the fact that u_t^* is above \bar{u}_t and inflation is high, the labor market is fundamentally different now than in the 1970s. In the 1970s, \bar{u}_t had risen from about 6% to almost 7.5% by the end of the decade. The high unemployment rate in the 1970s, which coexisted with accommodative monetary policy, partially reflects the secular rise in the unemployment rate.

⁸In Crump et al. (2019), we provide a causal link between these factors and the decline in the incidence of unemployment using micro-data.

⁹It also coincided with a heated academic debate about the level of the natural rate of unemployment. Going back to earlier papers such as Hall (1970b,a), Gordon (1972, 1982), Perry (1978), or Tobin (1974), there appears to have been a consensus that the natural rate of unemployment increased over this period. Interestingly these insightful analyses did not get much traction in policy circles and the Humphrey-Hawkins Full Employment and Balanced Growth Act of 1978 set an unemployment target of 4 percent for 1983. For a general discussion of unemployment rate benchmarks that are frequently used by policymakers, see Crump et al. (2020).

¹⁰See for example, Heise et al. (2022) for the disinflationary effect of rising import competition in late 1990s and early 2000s.

3.2 The Unemployment-Inflation Trade-Off During the COVID-19 Pandemic

Our estimation has shown that the natural rate of unemployment, u_t^* , has risen from around 4.5% in 2019Q4 to 5.9% in 2021Q4. While there was a brief period of a positive unemployment gap at the onset of the pandemic in early 2020, the unemployment gap exhibited a sharp reversal.¹¹ As of 2021Q4, the unemployment gap—the difference between the actual and natural rates of unemployment—stood at -1.6 percentage points. This is the highest negative unemployment gap in our sample since late 1970s.

Our framework allows us to isolate the effect of each margin in accounting for the notable increase in the natural rate of unemployment. Implementing counterfactuals in our framework, we trace the rise in u_t^* to the behavior of wages. We also report results from a forecasting exercise summarizing the expected path of key variables for the 2022-2023 period.

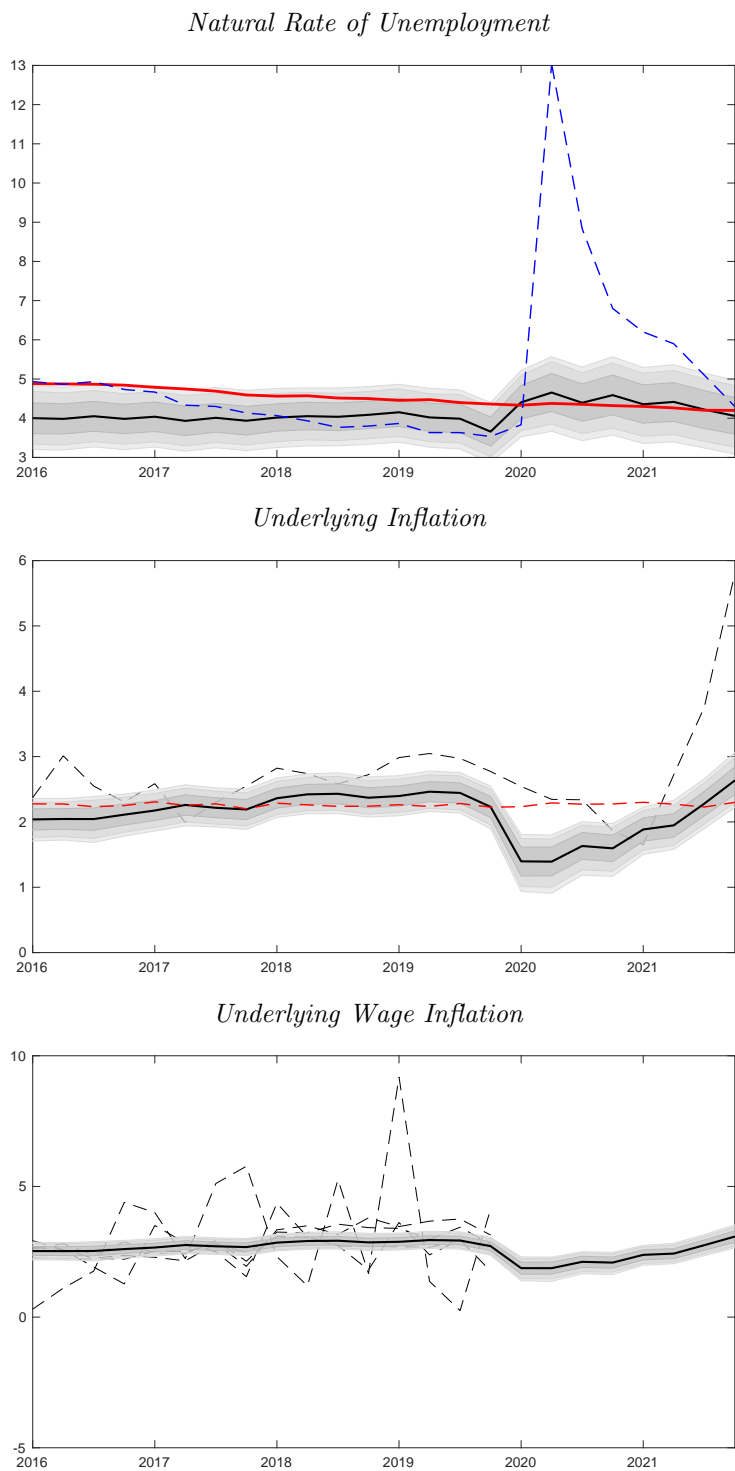
Tracing the rise in u^* We illustrate that the behavior of observed wages in the aftermath of COVID is the dominant signal for a the increase in u_t^* and, correspondingly, the increase in underlying inflation. To do so, we implement a simple counterfactual exercise where we omit any wage information past the fourth quarter of 2019 to assess the role of wage inflation.

Figure 5 shows the corresponding model outputs using this smaller information set. The top panel of the figure shows that the rise in u_t^* is now much more muted, hewing closely to the secular trend in the unemployment rate, with a near-zero unemployment gap at the end of 2021. Consequently, underlying inflation does not deviate significantly from longer-term inflation expectations despite the large rise in realized inflation and short-term inflation expectations. In the bottom panel, we show the counterfactual path of underlying wage growth. In this exercise, underlying wage growth recovers more gradually with a peak of 3.1% at the end of 2021 which is closely aligned with its pre-pandemic behavior.

Taken in sum, we can draw two conclusions from these results. First, the rise in u^* is driven primarily by the behavior of observed wages rather than changes in inflation or inflation expectations. Second, in the absence of wage information after 2019, both underlying inflation and underlying wage inflation recover to about the same level as observed before the pandemic.

¹¹This pattern is in stark contract to the Great Recession period where the large and positive unemployment gap persisted for almost a decade.

Figure 5. Phillips Curve Estimates Without Wage Information This figure shows the main model outputs from our estimated Phillips curve omitting wage information after 2019. The top panel shows the counterfactual estimate of the natural rate of unemployment, u_t^* (black line) along with the observed unemployment rate (blue dotted line) and the median secular trend, \bar{u}_t (red line). The middle panel shows the counterfactual estimate of underlying inflation (black line) along with the Federal Reserve Bank of Cleveland's median CPI inflation (grey dashed line) and the long-run component of inflation, π_t^* (red dashed line). The bottom panel shows the counterfactual estimate of underlying wage growth (black line) along with the realized wage series (grey dashed lines). Grey shaded areas denote 68%, 90% and 95% coverage intervals.



Forecasting u^* , inflation, and wages We now explore the implications for the near-term evolution of inflation through the lens of our Phillips curve. The top panel of Figure 6 shows the evolution of underlying inflation from 2014 on to the end of our forecasting horizon (2023Q4). The black line is the median prediction from the model whereas the grey shaded area depicts the 68%, 90%, and 95% coverage intervals. Underlying inflation rose steadily through the second-half of the previous expansion and clearly moved above model-implied long-run inflation expectations (red dotted line) by 2019. This expansion was interrupted by the COVID pandemic, but over the last year we again observe underlying inflation increasing and comfortably above long-term inflation rates. Underlying inflation is centered at a range of 3.0% to 3.6% – the highest reading we have observed since the early 1990s. Beyond 2021Q4, the figure shows that the unconditional model forecast is characterized by a sluggish decline in underlying inflation which remains about a half of a percentage point above long-run inflation even at the end of 2023.

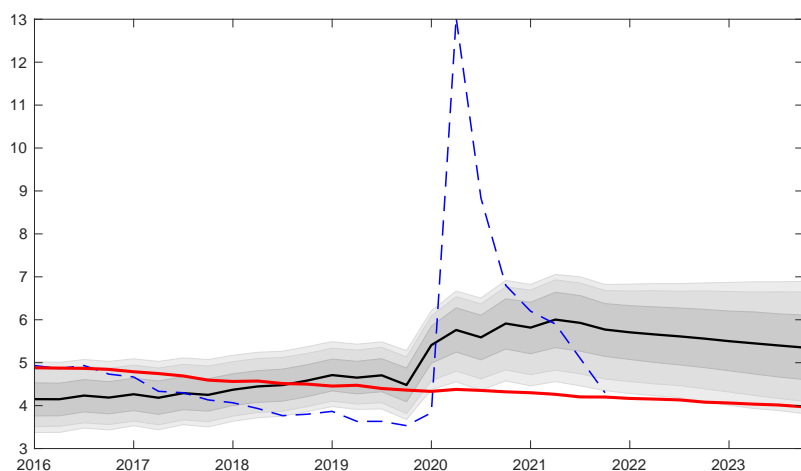
In the middle panel of Figure 6, we show the four-quarter moving average of the inflation gap – defined as the difference between underlying inflation (black line in top chart) and its long-term expectation (red line in top chart). For comparison, we add alternative forecasts from both private forecasters and the SEP from the FOMC.¹² Although private forecasters also anticipate short-term inflation to remain well above their long-term forecast, their path for the gap sits near the bottom of our 68% forecast distribution. Thus, inflation stabilizes more rapidly in their forecast than that of our model. In contrast, the FOMC’s December 2021 SEP is more consistent with the model’s median path (black line). That said, the median path is above the median participant’s projection (red diamonds) and toward the upper end of the FOMC’s central tendency (red lines).

Recall that the forecast path for inflation is directly linked to the expected path of the unemployment gap. This gap, in turn, is measured using information from different sources including inflation expectations, the secular trend in the unemployment rate, and wages. The bottom panel of Figure 6 shows our measure of underlying wage growth along with the forecasted path beyond 2021. Underlying wage growth has more than fully recovered the pre-COVID level and is about 3.5% at the end of 2021. For reference, this is about the same level as observed on the eve of the 2007-2009 recession but meaningfully higher than at the end of 2019. After the initial drop in 2020, the series rises steeply and is forecasted to peak at 3.9% before slowly declining.

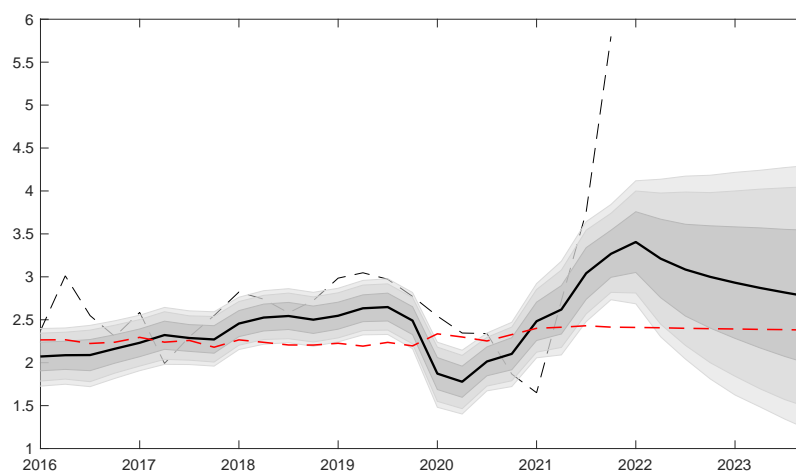
¹²The SEP provides projections for the Q4 to Q4 PCE price index growth, rather than the CPI; however, since we plot the gap between the near-term projection and the longer-run value we can make direct comparisons to our model. Moreover, Q4 to Q4 growth is well approximated by a four-quarter moving average of quarterly annualized growth rates (Crump et al. 2014).

Figure 6. Phillips Curve Based Forecasts This figure shows the main model outputs from our estimated Phillips curve along with their unconditional forecast after 2021Q4. The top left panel shows the median estimate of the natural rate of unemployment, u_t^* (black line) along with the observed unemployment rate (blue dotted line) and the median secular trend, \bar{u}_t (red line). The top right panel shows the median estimate of underlying inflation (black line) along with the Federal Reserve Bank of Cleveland's median CPI inflation (grey dashed line) and the long-run component of inflation, π_t^* (red dashed line). The bottom left panel shows the four-quarter moving average of underlying inflation with survey responses from Consensus Economics (blue square), Blue Chip Economic Indicators (cyan circle), Survey of Professional Forecasters (green square), along with the median and central tendency from the FOMC SEP (red diamond and line). The bottom right panel shows the median estimate of underlying wage inflation (black line) along with realized wages (grey dashed line). Grey shaded areas denote 68%, 90% and 95% coverage intervals.

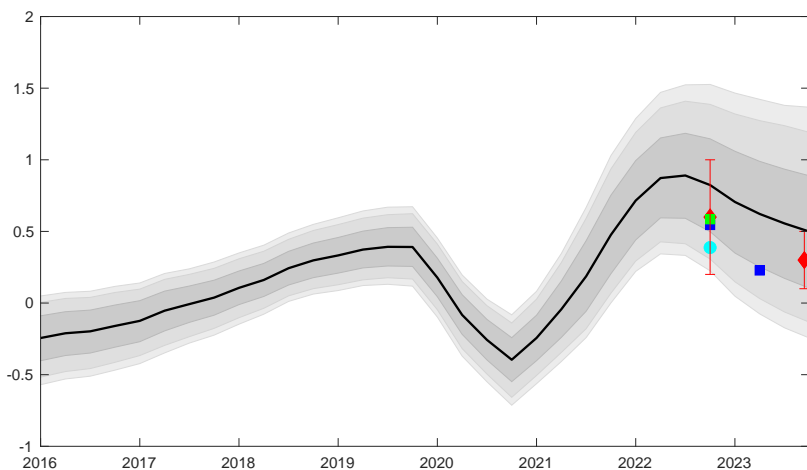
Natural Rate of Unemployment



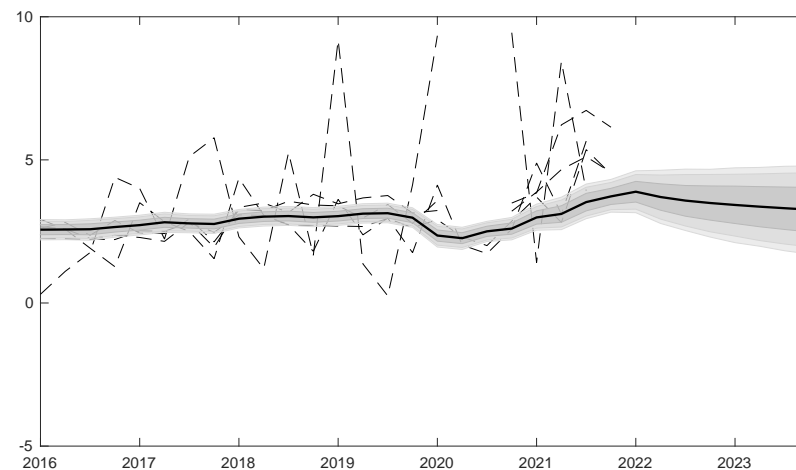
Underlying Inflation



Underlying Inflation Relative to Long-Run Forecast (4-quarter moving avg.)



Underlying Wage Inflation



Risks to the Inflation Forecast: Year over year core CPI inflation rose 5.5% in December 2021—the highest reading in 30 years. The model’s forecasts along with the forecasts of professional forecasters all anticipate a gradual reduction in the rate of inflation over the forecast horizon. A key feature of the unconditional forecast of the model is that long-term inflation expectations remain anchored by assumption.

This stability of long-term expectations in our projection is consistent with the current behavior of surveys of professional forecasters whereas the evidence from household surveys shows a more consistent rise over the last year or so. Further in the future inflation expectations from the NY Fed’s Survey of Consumer Expectations and the University of Michigan’s Survey of Consumer Sentiment have risen between half and a full percent over the last year. If long-term inflation expectations began to rise, this would likely lead to higher underlying inflation than what is currently forecasted. This scenario would be particularly worrisome given that the current estimated degree of slack, as measured by the unemployment gap, is comparable with the 1970s.

The model does not allow for a mechanism which relates recent inflation forecast errors to revisions in long-term forecasts. [Carvalho et al. \(2021\)](#) introduce a structural model of expectations anchoring, where the sensitivity of long-run expectations to short-term forecast errors is endogenous and depends on the history of forecast errors. This mechanism can accurately predict the behavior of long-term inflation expectations of both professional forecasters and households, in the US and other developed countries over the past forty years. [Carvalho et al. \(2022\)](#) present recent evidence that while long-term expectation in the U.S. and a few other OECD countries remain anchored, the risk of un-anchoring has increased notably in recent years.

The anchoring of expectations is also tightly connected with the costs of reducing inflation back to target. Given the highly forward-looking nature of our estimated Phillips curve, a flat slope does not necessarily imply high costs of disinflation. In fact, much more important is the role of long-term inflation expectations and the expected future path of the output gap. Monetary policy is the key factor driving these expectations. An (expected) credible monetary policy response would both keep inflation anchored (i.e. disconnected from the current surge in inflation) and induce a faster closing of the unemployment gap. This, in turn, would lead to a more rapid reduction in inflation *for a given level of the current unemployment gap*. In this buoyant scenario the cost of disinflation in terms of unemployment would be small. Conversely, a loss of central bank credibility would require a large increase in the *current* unemployment gap to offset the sluggish adjustment of expectations. The

simple framework proposed here has no explicit role for policy, and our unemployment gap forecast reflects the historical dynamics of the estimated unemployment gap.

As evident in Figure 6, our measure of underlying wage growth lies at the bottom of the range of realized wage growth based on our 5 observable series. If these different measures of wage growth remain high, all else equal, we would expect underlying wage growth to be revised upward leading to an even higher value for u_t^* . Consequently, the forecast for inflation would also be revised upward remaining well above its pre-COVID baseline for a substantial amount of time.

4 The Outlook for Wage Growth

Our model analysis in the previous section identify strong wage growth as the main source of the information resulting in the rise in u^* and our forecast calls for only a sluggish decline in wage growth over the next two years. In this section, we investigate the current and prospective path of wages using a variety of cross-sectional data sources to provide external corroboration for our analysis and forecast. First, we examine the joint evolution of the job-filling rate and underlying wage growth in the aggregate and across industries to assess whether strong growth is concentrated in a small segment of the labor market or not. Second, we use wages advertised in newly posted vacancies by the same employers for the same jobs to isolate a composition-free measure of *forward-looking* labor costs. Third, we analyze the evolution of reservation wages to detect whether the pandemic caused a shift in willingness to work. The evidence we present does not appear to be consistent with a one-time health risk adjustment to the behavior of wages concentrated in the health-risk prone segments of the economy. Instead, we find broad-based wage growth consistent with a tight labor market which arguably is likely to be persistent.

4.1 Underlying Wage Growth and Job-Filling Rate

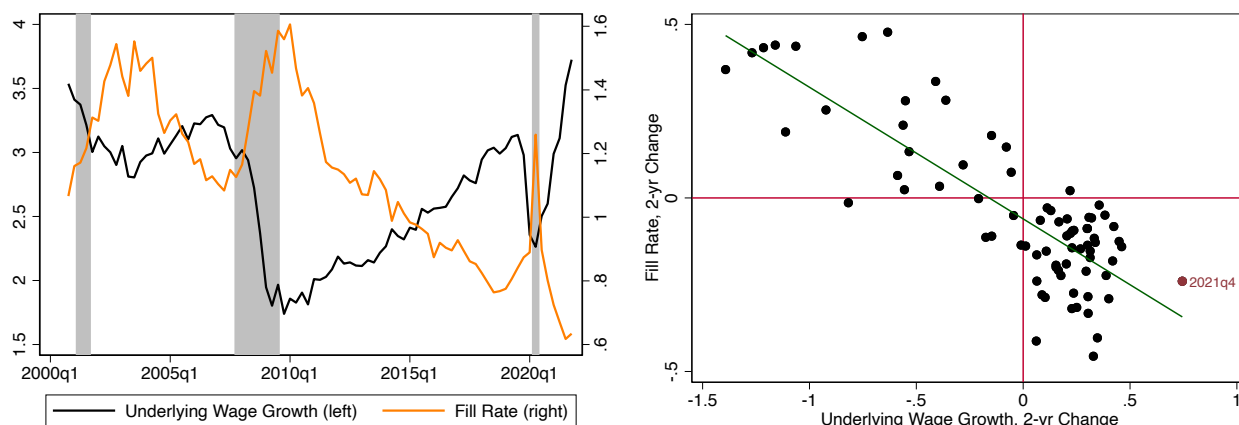
The COVID pandemic was highly disruptive to the labor market due to the health risk associated with labor market participation. The dramatic swings in labor market indicators since the onset of the pandemic has made assessing labor market conditions more challenging than usual. One indicator that is especially useful in light of the possible changing willingness to work is the so-called job-filling rate.¹³ Formally, the job-filling rate is defined as the the number of hires per vacancy and provides a reliable measure of how “easy” it is to fill open positions for firms. With the advent of the BLS’

¹³For example [Abraham et al. \(2020\)](#) advocate using hires-based measures of labor market tightness.

Job Openings and Labor Turnover Survey (JOLTS), we can calculate the aggregate job-filling rate since 2000.

The left chart in Figure 7 shows our underlying wage growth measure along with the economywide job-filling rate.¹⁴ Theory would predict a tighter labor market—characterized by a lower job filling rate—would coincide with higher wage growth. Or, as lucidly described in (Pissarides, 2000, p.7), “...firms with vacancies find workers more easily when there are more workers relative to available jobs.” Therefore labor shortages that arise from a decline in workers’ willingness to work, say due to looming health risks or career dissatisfaction, would lower the job filling rate and would engender wage growth. Figure 7 shows that this negative correlation is strongly borne out in the data. When the labor market is tight (when it is harder for firms to fill open positions) wage growth in the economy tends to be higher. In the right chart of Figure 7 we show the two-year changes in underlying wage growth and the job-filling rate. Remarkably, almost all of the data points reside in the top-left and lower-right quadrants, i.e., wage growth accelerates when the job-filling rate declines and wage growth decelerates when filling jobs becomes easier. The job-filling rate currently stands at a little above 0.6 which is even lower than its value in 2019 (the end of the previous expansion). In concert, underlying wage growth is 3.7% which is the highest value since 1991.

Figure 7. Underlying Wage Growth and Job-Filling Rate This figure displays the relation between the underlying wage growth introduced in Section 3 and the job-filling rate from JOLTS. The left plot of the figure shows the time series for each variable whereas the right plot of the figure shows a scatterplot of two-year changes in each variable along with the OLS fitted line.

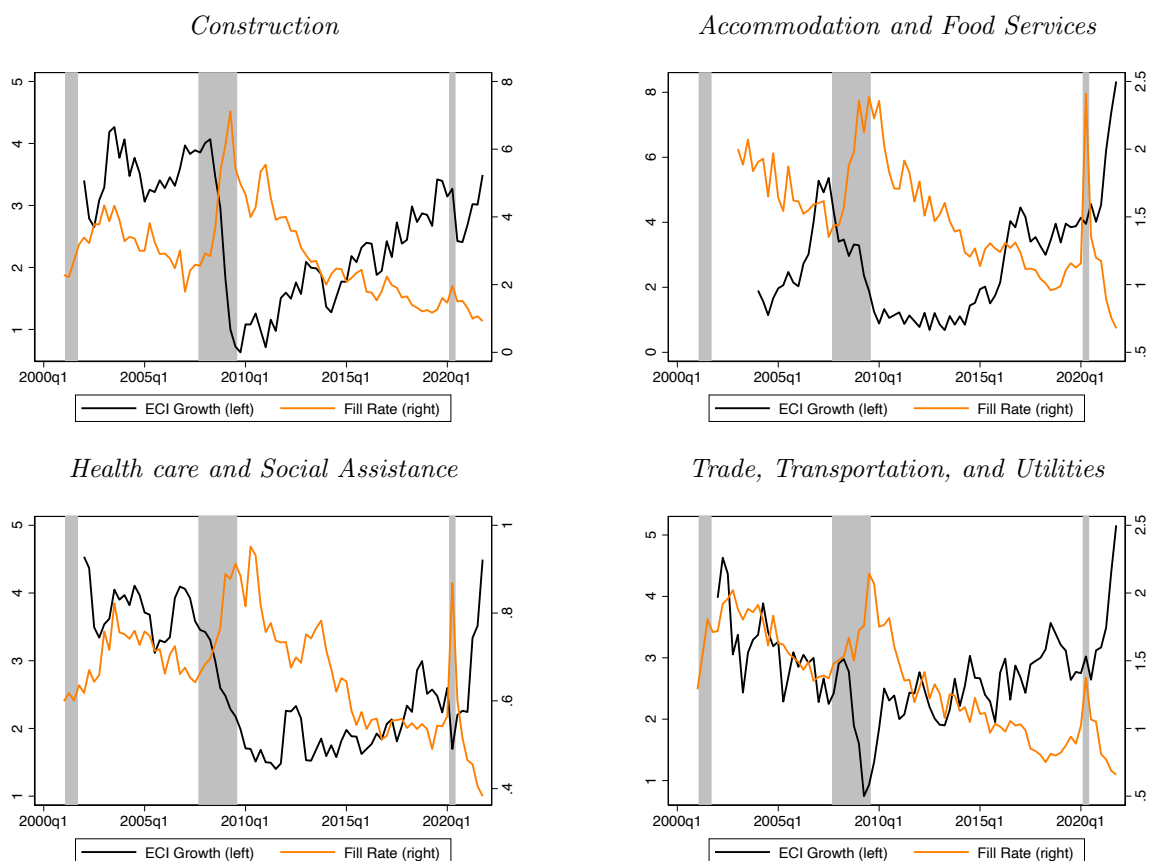


Although Figure 7 shows that this negative relationship has consistently held over the last twenty years, it is possible that the recent stark movements are driven by pandemic-specific factors. In Figure

¹⁴While the underlying wage growth measure is derived using all of our observables, as a robustness check, we also present a similar chart relying instead on ECI total compensation growth and find similar results. See Figure 13 in the Appendix.

8 we show the corresponding graphs for four different industries, namely, construction, accommodation and food services, health care and social assistance, and trade, transportation, and utilities. To obtain industry-specific wage growth we replace our underlying wage growth measure with data from the employment cost index (ECI). While we see sharper swings in both the fill rate and also wage growth in industries more exposed to health risk from COVID-19, the job-filling rate is at its series minimum for all industries shown.¹⁵ At first glance, this figure is suggestive that the strong movements in the economy-wide job-filling rate and wage growth is driven by potentially temporary factors related specifically to the pandemic. However, a clear drawback to decomposing the data by industry is that it is difficult to make temporal comparisons as there are very different occupations within each industry along with the well-known concerns about changing worker composition over time.

Figure 8. Wage Growth and Job-Filling Rate by Industry This figure displays the relationship between the ECI total compensation growth and the job-filling rate from JOLTS for selected industries.



¹⁵In fact, across the vast majority of industries in JOLTS, the job-filling rate is at, or near historic lows, at the end of 2021.

4.2 Forward-Looking Labor Costs: Growth in Job-Level Posted Wages

To assess the degree of the pandemic’s role in the strong wage growth observed we utilize data from Burning Glass Technologies on posted job vacancies. These data have a number of noteworthy advantages. First, since they provide information about wages posted at the job level, rather than the worker, there are less concerns about unobserved heterogeneity generating patterns in the data. Second, they provide detailed information about each job vacancy—including information on firm, location, occupation, and posted wage. Third, the data set is very large reporting more than 4 million job openings as of December 2021. Finally, and most importantly, analyzing posted wage behavior at the job level, which are not subject to worker composition bias, allows us to identify emergent trends more accurately. Rising posted wages suggest that employers increase the wages they post to attract more workers. This in turn increases wages of new hires. [Hazell and Taska \(2020\)](#) who advocate using posted wages as a composition-bias free measure of wages show that posted wages move almost one-to-one with new hires wages in the CPS and the QWI, lending credibility to the measure we use to assess wage growth.

It is well known that changing worker composition makes it harder to assess the cyclical of wage growth ([Perry \(1972\)](#), [Bils \(1985\)](#), [Solon et al. \(1994\)](#) and, more recently, [Daly and Hobijn \(2022\)](#)). A better measure of inflationary wage growth would be to observe the change in remuneration for the same job in the same firm and location. This would alleviate concerns about unobserved worker characteristics that are not even controlled for in the employment cost index.¹⁶ If firms post a higher wage for the same job when the labor market is tight, we might expect to see inflationary pressures due to rising labor costs. To operationalize this concept, we consider vacancies posted by the same firm for the same job over time following a similar methodology to [Hazell and Taska \(2020\)](#) who use posted wages to analyze downward wage rigidity. To ensure that we are identifying the same position we match on firm, job title, location, and 5-digit occupation code (further details on the data are provided in the Appendix). While looking at within employer, within job posted wage growth is appealing it may not capture posted wage growth in positions that are not posted systematically. As a robustness check, we will also consider posted wage growth for average wages posted for the same 5-digit occupation in the same metropolitan statistical area (MSA) and draw similar conclusions (See [Figure 15](#)).¹⁷

¹⁶The ECI, unlike some other measures of wages, has a fixed composition of occupations which partially, but not fully, addresses these issues.

¹⁷We also show in the Appendix that posted wages align well with BLS wage data by occupation – especially for

We study two-year posted wage growth from the fourth quarter of 2019 to the fourth quarter of 2021 to minimize distortions related to the pandemic. As a comparison we also study the two year wage growth ending in the fourth quarter of 2019 – representing posted wage growth in the tightest labor market, as measured by the raw unemployment rate, since the 1960s. On average, posted wages for jobs with salaries below \$75,000 grew at a rate of about 12% from 2019 to 2021 as compared to about 8% from 2017 to 2019. Because of the large number of observations and underlying noise in the data, we utilize a binned scatterplot which presents a nonparametric estimate of the conditional median of posted wage growth as a function of the level of wages. The shaded areas indicates associated 95% confidence bands, all based on the methodology introduced in Cattaneo et al. (2021b). While average posted wage growth is somewhat higher over the second of these two periods, Figure 9 shows that over the 2019-2021 period there is a much stronger rise in posted wage growth for job openings which are below the 2019 median salary of \$35,000.¹⁸ Interestingly, from 2017-2019 posted wage growth was more even across this salary spectrum.¹⁹

The strong posted wage growth at lower salary positions over the last two years is consistent with the ample anecdotal evidence that these positions have become more difficult to fill.²⁰ Moreover, there is also evidence that lower paying jobs tend to be more consumer facing and less amenable to remote work than higher paying jobs (Dingel and Neiman 2020). Potentially the steep rise in posted wage growth in lower-paying positions could reflect required compensation for the additional health risk or a rapid tightening of the labor market for lower-paying jobs (or both). If the former explanation is dominant, we would expect to see strong wage growth in “non-teleworkable” jobs as compared to “teleworkable” jobs.²¹ In the bottom panel of Figure 9, we show the corresponding binned scatter plots (with associated confidence bands) for the conditional median function for either category. These two estimated relations are very similar across the salary spectrum, with confidence bands overlapping for almost the whole support, suggesting that the tightness in the labor market is uniform across these two categories of occupations. It might, at first glance, seem surprising that

salaries below \$75,000. See Figure 12.

¹⁸See <https://www.ssa.gov/oact/cola/central.html>. Moreover, Howard et al. (2022) show that composition-adjusted real wage growth has remained positive over the period spanning 2020 and 2021.

¹⁹These results are qualitatively consistent with results available from the Atlanta Fed Wage Tracker for wage growth by wage quartile. See Figure 17 in the Appendix.

²⁰Anecdotal evidence on employers’ difficulties finding and retaining workers can be found, for example, in recent Federal Reserve Beige Books. See for example, <https://www.federalreserve.gov/monetarypolicy/beigebook202110.htm> or <https://www.federalreserve.gov/monetarypolicy/beigebook202112.htm>. Survey evidence is presented in Figure 14 in the Appendix.

²¹Dingel and Neiman (2020) classify the feasibility of working at home for all occupations using the Standard Occupational Classification (SOC). We use their definition for the five-digit SOC occupation code of the job posting. The classification data is available at <https://brentneiman.com/research/DingelNeimanCSVs.zip>.

health risk does not appear to be driving the posted wage growth, however, the work/life disruptions caused by the COVID pandemic may have substantial indirect effects on individual’s ability and desire to supply labor.

4.3 Rise in Reservation Wages

A useful metric to summarize the trade-offs that affect labor supply decisions is the reservation wage of workers. Our measure of the reservation wage is obtained from the following question from the NY Fed’s Survey of Consumer Expectations²²:

Suppose someone offered you a job today in a line of work that you would consider. What is the lowest wage or salary you would accept (BEFORE taxes and other deductions) for this job?

Rises in reservation wage could reflect perceived health risk but also changes in preferences and willingness towards work. In Figure 9, the top left shows the reservation wage for two age categories (above and below 45). We observe a steep rise in reservation wages starting at the end of 2017 (vertical line) when the unemployment rate first fell below 4%. The pace of the increase of the reservation wage picked up further after the pandemic began. Notably, this rise was very similar for respondents in both age categories. This behavior is consistent with our earlier evidence that posted wage growth has evolved similarly for both teleworkable and non-teleworkable jobs.

In the top right plot of Figure 9, we examine reservation wages by educational attainment. Reservation wages started to rise for both workers with college education and for those without starting in 2017 but the rise was much steeper for workers without a college degree after the pandemic echoing the steep posted wage increases we have seen in the Burning Glass data. Finally, in the bottom two panels we show reservation wage split by gender and household income which show similar patterns.²³

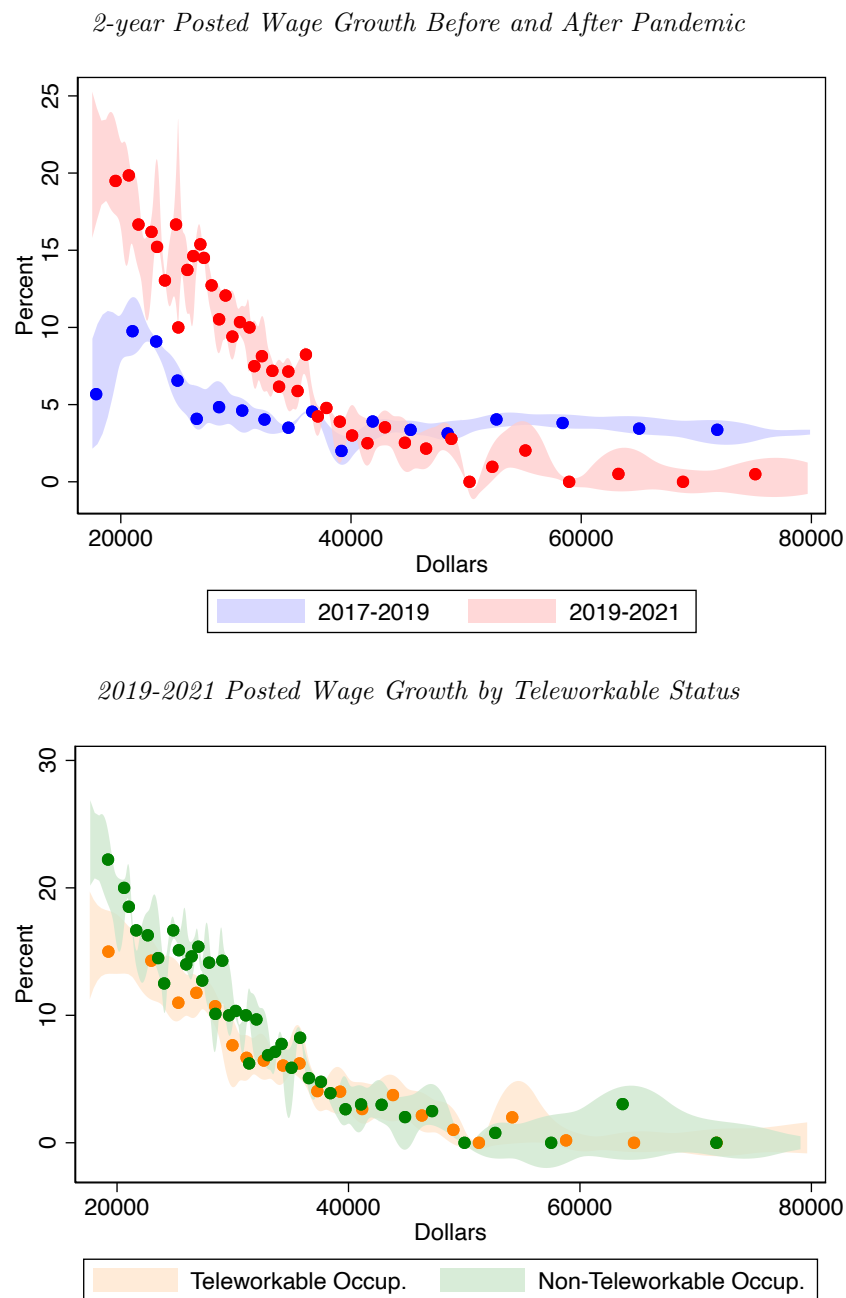
5 Conclusion

We revisited the unemployment-inflation trade-off by estimating the natural rate of unemployment using a Micro-Macro Phillips Curve framework. The natural rate of unemployment stood at 5.9%

²²In particular, we utilize responses on participants’ reservation wage which is a part of the survey’s Labor Market Module. This module is conducted three times per year in March, July, and November.

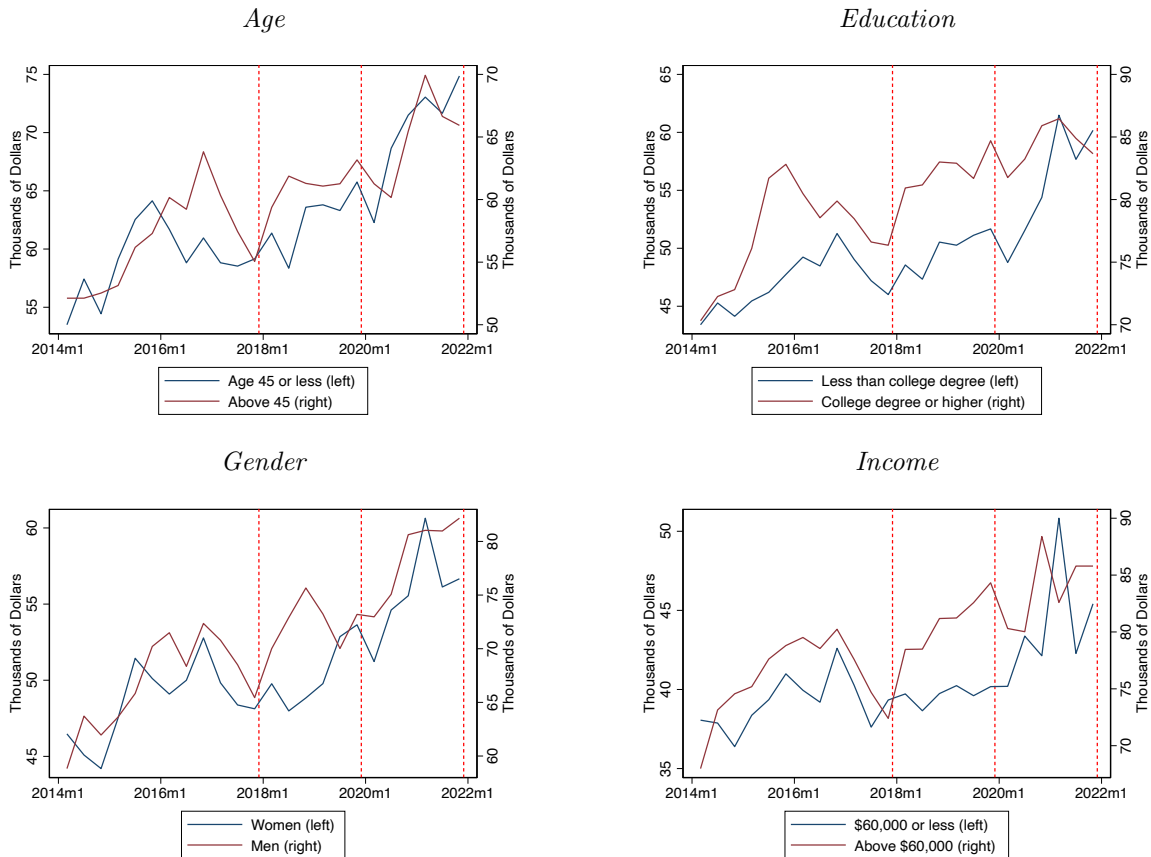
²³These results are consistent with Faberman et al. (2022) who show that willingness to work has declined after the pandemic both along the extensive and intensive margin of labor supply.

Figure 9. Posted Wage Growth Comparisons This figure presents nonparametric estimates of the conditional median function of two-year posted wage growth given initial wage level, based on data from Burning Glass Technologies. Posted wage growth is constructed by matching posted wages for the same job listings at two-year intervals. See Appendix for further details. The nonparametric curve estimates rely on Cattaneo et al. (2021b) and Cattaneo et al. (2021a). Shaded regions denote 95% confidence bands.



at the end of 2021—1.4 percentage points higher than at the end of 2019. This pronounced rise was mostly driven by strong wage growth rather than changes in inflation expectations. Despite the rise in the natural rate of unemployment, the secular trend of unemployment was around 4.2% reflecting the effect of ongoing secular trends that have been pushing down the unemployment rate in the last

Figure 10. Reservation Wages This figure presents the average reported reservation wage from survey respondents of the Labor Market Module from the Federal Reserve Bank of New York’s Survey of Consumer Expectations. Each panel presents average reported reservation wage for different subsets of respondents. The sample period is March 2014–November 2021. Vertical red dotted lines represent December 2017, 2019, and 2021, respectively.



30 years.

Our analysis has two important implications for the reconciliation of the unemployment-inflation trade-off in the medium-run. First, our model projects underlying inflation to remain 0.5 percentage point above its long-run trend even by the end of 2023 even if wage growth slows down to levels similar to its pre-pandemic levels and long-run inflation expectations remain well anchored. Second, the flat Phillips curve we estimate does not necessarily imply high unemployment costs of reducing inflation but rather underscores the importance of anchored inflation expectations and credible monetary policy.

Given the importance of wage growth for the inflation outlook, we examined detailed data on posted wages and workers’ reservation wages. Our findings suggest that strong wage growth is likely not a one-time adjustment for increased health risk but rather reflects a work-leisure reoptimization. Naturally, a positive labor supply shock, such as the entry of side-lined workers to the labor force at a

faster pace—would help alleviate wage pressures. Since the participation cycle lags the unemployment cycle, the cyclical adjustment of the labor force participation rate is not yet complete. However, as shown by [Hobijn and Şahin \(2021\)](#), these procyclical movements are not driven by entry of workers into the labor force but rather a consequence of employment stability which takes time to build. Given our analysis of different factors, we expect strong wage growth to moderate only sluggishly, continuing to put upward pressure on inflation.

Another noteworthy concept that our work connects to is the well-known misery index. The misery index, which was developed in the 1970s by Arthur Okun as a proxy to capture the high welfare cost of the high inflation-high unemployment economic environment of the 1970s, is calculated by adding the level of the unemployment rate and the annual inflation rate. While the current unemployment gap is similar in level to its 1970s estimate, in terms of the misery index the economies look very different—owing to the secular trends driven by slow moving demographic and social changes. Understanding the welfare implications of high inflation with a declining secular trend of unemployment in heterogeneous agent models (e.g., HANK models) is an important open research area.

Finally, we isolate the job-filling rate, posted wage growth at the job-level and changes in reservation wages as the three important measures of labor market to follow, beside the unemployment rate, to assess inflationary pressures. These measures are particularly important in a tight labor market since they provide forward-looking information about wage and inflation pressures.

References

- Abraham, K. G., Haltiwanger, J. C., Rendell, L. E., 2020. How tight is the US labor market? *Brookings Papers on Economic Activity* 51, 97–138.
- Bils, M. J., 1985. Real wages over the business cycle: Evidence from panel data. *Journal of Political Economy* 93, 666–689.
- Carvalho, C., Eusepi, S., Moench, E., Preston, B., 2021. Anchored inflation expectations. *American Economic Journal: Macroeconomics* Forthcoming.
- Carvalho, C., Eusepi, S., Moench, E., Preston, B., 2022. Measuring the risk of inflation expectations unanchoring during COVID: A structural approach, working paper.
- Cattaneo, M. D., Crump, R. K., Farrell, M. H., Feng, Y., 2021a. Binscatter regressions, working paper.
- Cattaneo, M. D., Crump, R. K., Farrell, M. H., Feng, Y., 2021b. On binscatter. Staff Report 881, Federal Reserve Bank of New York.
- Crump, R. K., Eusepi, S., Giannoni, M., Şahin, A., 2019. A unified approach to measuring u^* . *Brookings Papers on Economic Activity* 50, 143–214.
- Crump, R. K., Eusepi, S., Lucca, D., Moench, E., 2014. Which growth rate? It’s a weighty subject. *Liberty Street Economics Blog*.
- Crump, R. K., Nekarda, C. J., Petrosky-Nadeau, N., 2020. Unemployment rate benchmarks. *Finance and Economics Discussion Series 2020-072*, Federal Reserve Board of Governors.
- Şahin, A., Song, J., Topa, G., Violante, G. L., 2014. Mismatch unemployment. *American Economic Review* 104, 3529–3564.
- Daly, M. C., Hobijn, B., 2022. The importance of the part-time and participation margins for real wage adjustment. *Journal of Money, Credit and Banking* Forthcoming.
- Davis, S. J., Faberman, R. J., Haltiwanger, J., Jarmin, R., Miranda, J., 2010. Business volatility, job destruction, and unemployment. *American Economic Journal: Macroeconomics* 2, 259–287.
- Del Negro, M., Giannoni, M. P., Schorfheide, F., 2015. Inflation in the great recession and new Keynesian models. *American Economic Journal: Macroeconomics* 7, 168–196.
- Dingel, J. I., Neiman, B., 2020. How many jobs can be done at home? *Journal of Public Economics* 189, 1–8.
- Elsby, M. W. L., Hobijn, B., Şahin, A., 2010. The labor market in the great recession. *Brookings Papers on Economic Activity* 41, 1–69.
- Elsby, M. W. L., Michaels, R., Solon, G., 2009. The ins and outs of cyclical unemployment. *American Economic Journal: Macroeconomics* 1, 84–110.
- Faberman, R. J., Mueller, A. I., Şahin, A., 2022. Has the willingness to work fallen during the covid pandemic?, working paper.
- Galí, J., 2011. The return of the wage Phillips curve. *Journal of the European Economic Association* 9, 436–461.
- Goldin, C., 2006. The quiet revolution that transformed women’s employment, education, and family. *American Economic Review* 96, 1–21.
- Gordon, R. J., 1972. Wage-price controls and the shifting Phillips curve. *Brookings Papers on Economic Activity* 3, 385–421.
- Gordon, R. J., 1982. Inflation, flexible exchange rates, and the natural rate of unemployment. In: Baily, M. N. (ed.), *Workers, Jobs, and Inflation*, Brookings Institution.

- Hall, R. E., 1970a. The process of inflation in the labor market. *Brookings Papers on Economic Activity* 5, 343–393.
- Hall, R. E., 1970b. Why is the unemployment rate so high at full employment? *Brookings Papers on Economic Activity* 1, 369–410.
- Hall, R. E., 2005. Employment fluctuations with equilibrium wage stickiness. *American Economic Review* 95, 50–65.
- Hazell, J., no, J. H., Nakamura, E., Steinsson, J., 2021. The slope of the phillips curve: Evidence from u.s. states. Working Paper 28005, NBER.
- Hazell, J., Taska, B., 2020. Downward rigidity in the wage for new hires, working paper.
- Heise, S., Karahan, F., Şahin, A., 2022. The missing inflation puzzle: The role of the wage-price pass-through. *Journal of Money, Credit and Banking* 54, 7–51.
- Hobijn, B., Şahin, A., 2021. Maximum employment and the participation cycle. Working Paper 29222, NBER.
- Howard, S., Rich, R., Tracy, J., 2022. Real wages grew during two years of covid-19 after controlling for workforce composition. *Dallas Fed Economics*, February.
- Karahan, F., Pugsley, B., Şahin, A., 2022. Demographic origins of the startup deficit. *American Economic Review* Forthcoming.
- Laubach, T., 2001. Measuring the NAIRU: Evidence from seven economies. *Review of Economics and Statistics* 83, 218–231.
- Laubach, T., Williams, J. C., 2003. Measuring the natural rate of interest. *Review of Economics and Statistics* 85, 1063–1070.
- Mavroeidis, S., Plagborg-Møller, M., Stock, J. H., 2014. Empirical evidence on inflation expectations in the new Keynesian Phillips curve. *Journal of Economic Literature* 52, 124–188.
- Negro, M. D., Lenza, M., Primiceri, G. E., Tambalotti, A., 2020. Whats up with the phillips curve? *Brookings Papers on Economic Activity* 51, 301–373.
- Perry, G. L., 1972. Unemployment flows in the U.S. labor market. *Brookings Papers on Economic Activity* 3, 779–787.
- Perry, G. L., 1978. Slowing the wage–price spiral: The macroeconomic view. *Brookings Papers on Economic Activity* 9, 259–299.
- Pissarides, C. A., 2000. *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, second ed.
- Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review* 95, 25–49.
- Shimer, R., 2012. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics* 15, 127–148.
- Solon, G., Barsky, R., Parker, J. A., 1994. Measuring the cyclicity of real wages: How important is composition bias. *Quarterly Journal of Economics* 109, 1–25.
- Tobin, J., 1974. Monetary policy in 1974 and beyond. *Brookings Papers on Economic Activity* 5, 219–232.

A Data Description

In this section we summarize our data sources for the paper. Our observed measure of u_t is the civilian unemployment rate from the Bureau of Labor Statistics (BLS). Inflation is measured as the median CPI inflation in quarterly annualized percent changes available from the Federal Reserve Bank of Cleveland. For wage growth and inflation expectations we combine data from a variety of sources. For a full description of these sources, see [Crump et al. \(2019\)](#). Finally, all growth rates are expressed at a quarterly, annualized rate.

In Section 4, we introduce empirical evidence from a variety of sources. First, we utilize data from Burning Glass Technologies (BGT) which compiles job openings along with detailed and standardized characteristics of each opening such as a firm identifier or SOC occupation code. Figure 11 compares labor market tightness – the ratio of the level of vacancies to the level of unemployment from the BLS – using data from BGT and also from the Job Openings and Labor Turnover Survey (JOLTS). The figure shows that the time-series dynamics of these two measures are very similar which provides a natural robustness check to the data set.

Importantly, the BGT data include information on posted salaries for about 20% of all reported vacancies from 2017-2021. To compute posted salary growth, we compute the posting’s salary as the average of the minimum and maximum posted salary. We drop observations in the District of Columbia and in U.S. territories. We also drop observations if they do not report a salary or if they are missing any of the characteristics we use to match job postings across time. To construct the underlying data for Figure 9, we use the following data fields: employer, job title, SOC5 code, state and county FIPS code, and pay frequency. For each quarter of the BurningGlass data, we compute the average posted salary by the interaction of those six categories. We then restrict the data to observations in the fourth quarter of 2017, 2019, and 2021 and only consider matched jobs across time: from 2017 to 2019 and from 2019 to 2021. To determine posted salary growth by teleworkable status, we incorporate data from [Dingel and Neiman \(2020\)](#), which classifies whether a job is teleworkable or not based on its SOC5 code. To construct the underlying data for Figure 15 we follow a similar approach as in Figure 9, but instead match on SOC5 code, MSA, and pay frequency. We still drop observations that do not report a salary, an employer, a job title, an SOC5 code, an MSA, are identified as an internship, or report a pay frequency that is not hourly or annual (hourly and annual pay frequencies make up more than 90% of the data). We then compute the average posted salary by SOC5 code, MSA, pay frequency, and quarter; further, we restrict to observations

in the fourth quarter of each year and compute posted salary growth. Again, to determine wage growth by teleworkable status, we use the classification data from [Dingel and Neiman \(2020\)](#).

Finally, for Figure 16 we construct labor market outcomes by teleworkable status. Using the basic monthly data from the Current Population Survey (CPS), we use an individual's primary job's occupation code to determine whether their primary job is teleworkable. We use a crosswalk between PTIO1OCD and SOC3 code, which allows us to categorize the SOC3 code for individuals in the CPS. To compute whether an occupation is teleworkable using the SOC3 code, we compute the share of SOC5 codes within each SOC3 which are classified as teleworkable by [Dingel and Neiman \(2020\)](#). Any SOC3 category which has at least half of its SOC5 code categories as teleworkable are defined as teleworkable occupations. Finally, we aggregate the individual data to obtain the number of unemployed and employed by teleworkable status.

B Additional Figures

Figure 11. Labor Market Tightness: JOLTS vs. Burning Glass Technologies This figure shows labor market tightness – the ratio of job openings to the number of unemployed – based on job openings data from either JOLTS or Burning Glass Technologies.

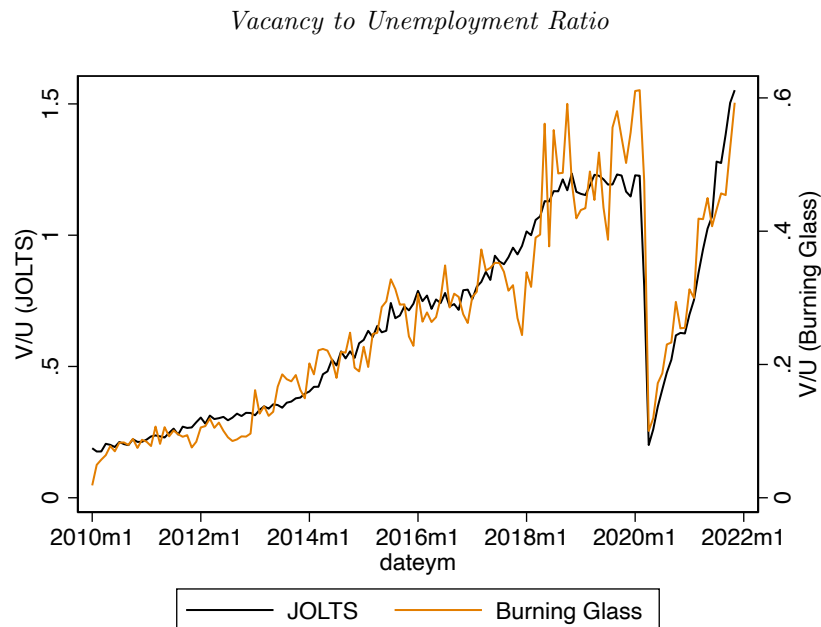


Figure 12. Posted Wages versus Realized Wages This figure shows a scatterplot of the median 2019 posted wage from Burning Glass Technologies compared to the median 2019 realized wage from the BLS National Occupational Employment and Wage Estimates at the 5-digit SOC level.

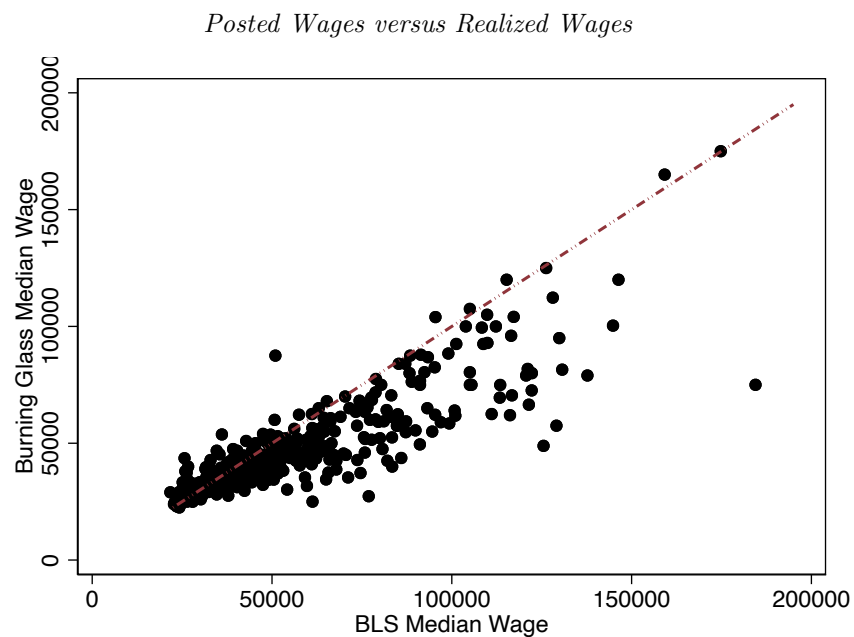


Figure 13. Total Compensation Growth and Job-Filling Rate This figure displays the relation between year-over-year total compensation growth from the ECI and the job-filling rate from JOLTS. Grey shaded regions denote NBER recessions.

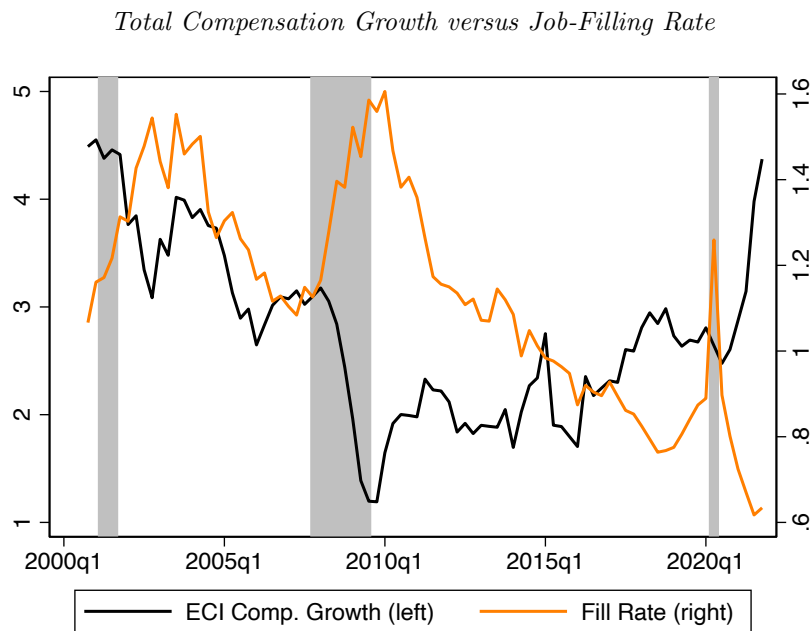


Figure 14. NFIB Survey Response on Filling Openings This figure shows the percentage of firms with at least one unfilled job opening from the NFIB Small Business Jobs Report.

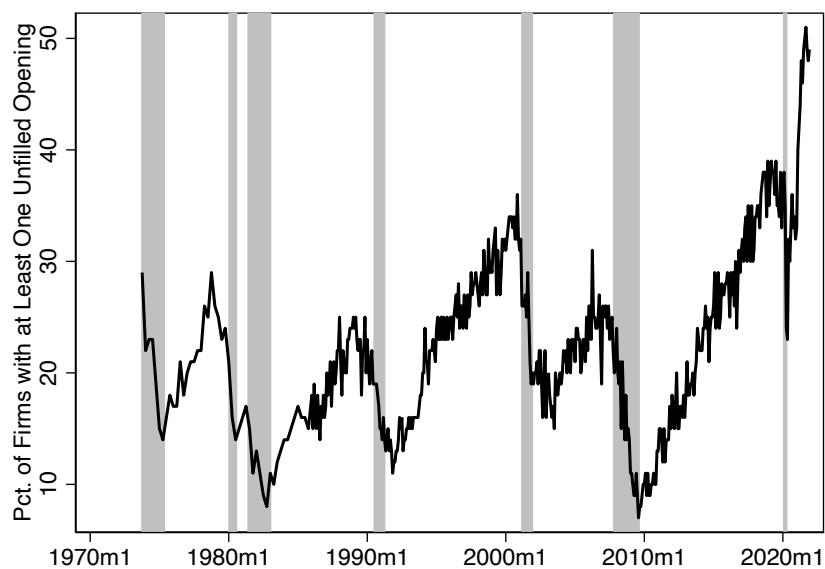
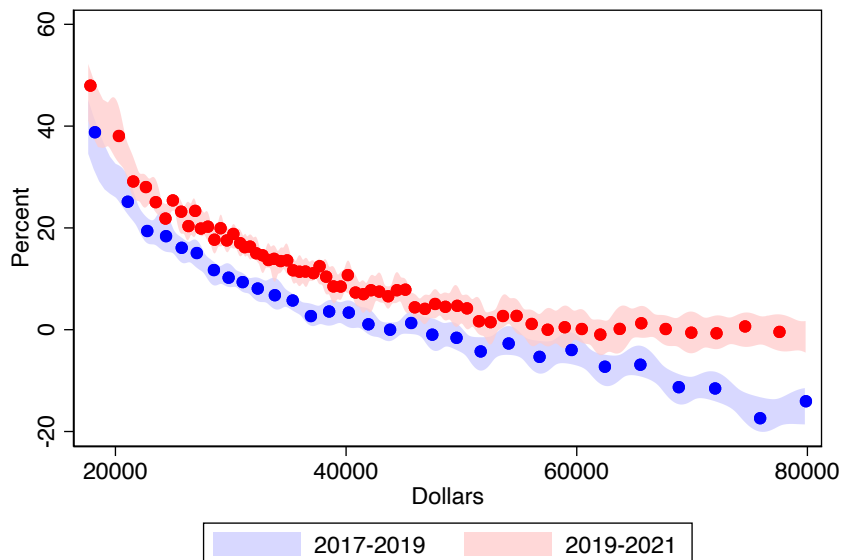


Figure 15. Posted Wage Growth Comparisons This figure presents nonparametric estimates of the conditional median function of two-year posted wage growth given initial wage level, based on data from Burning Glass Technologies. Posted wage growth is constructed with average posted wages in for the same MSA-SOC5 pair at two-year intervals. See Appendix for further details. The nonparametric curve estimates rely on [Cattaneo et al. \(2021b\)](#) and [Cattaneo et al. \(2021a\)](#). Shaded regions denote 95% confidence bands.

2-year Posted Wage Growth Before and After Pandemic



2019-2021 Posted Wage Growth by Teleworkable Status

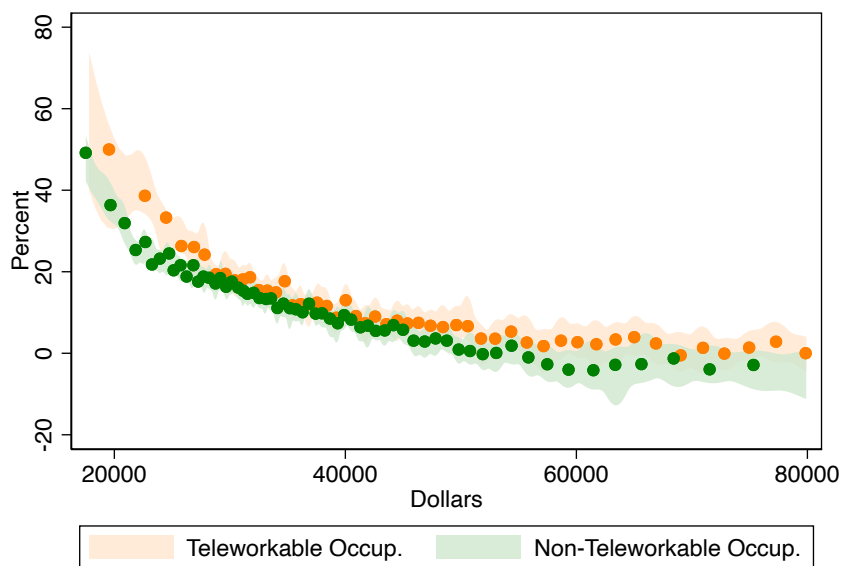


Figure 16. Labor Market Outcomes by Teleworkable Status This figure displays the number of unemployed workers and the ratio of unemployed to employed workers by teleworkable status using the occupational classification introduced in [Dingel and Neiman \(2020\)](#). Vertical red dotted lines represent December 2017, 2019, and 2021, respectively.

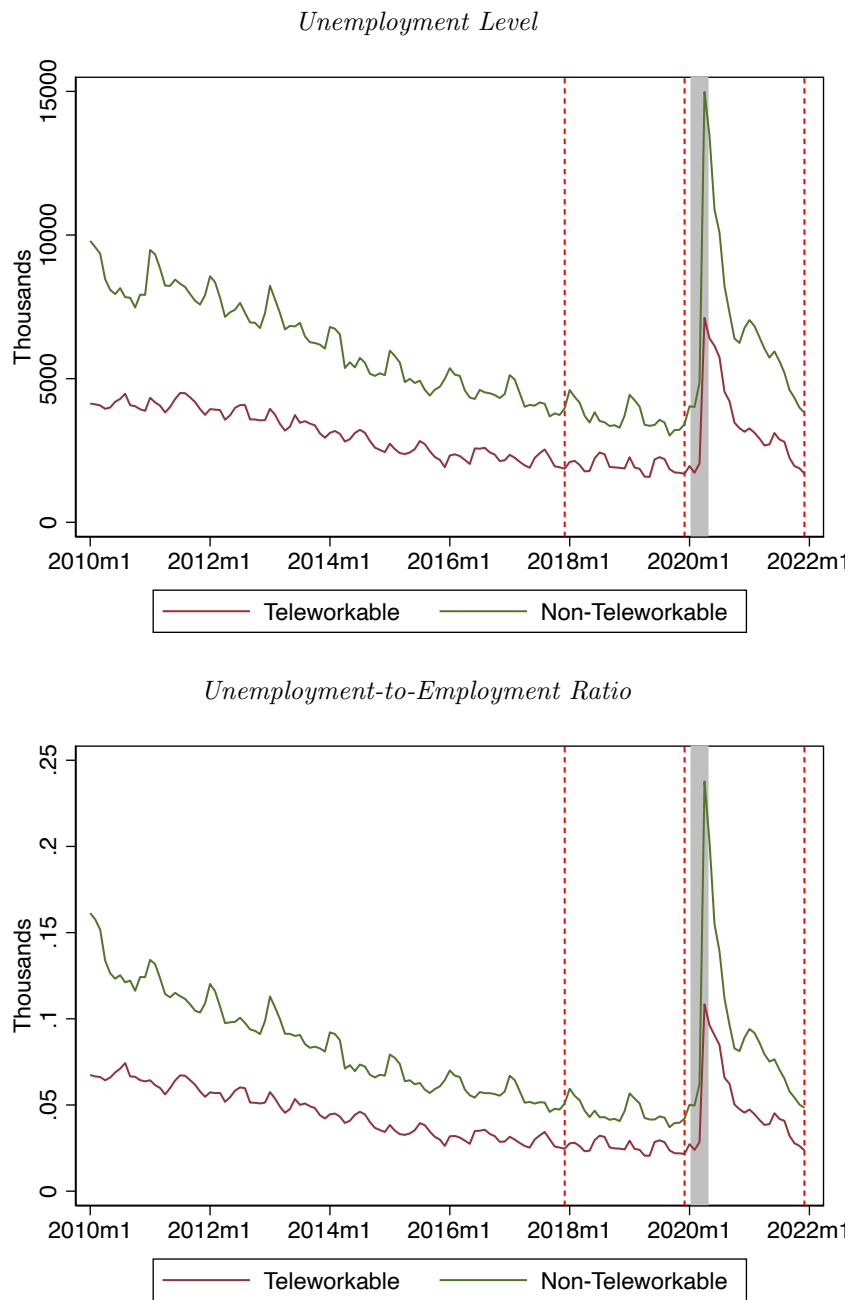


Figure 17. Atlanta Fed Wage Tracker by Wage Quartile This figure shows wage growth by wage quartile obtained from the Atlanta Fed Wage Tracker (available at <https://www.atlantafed.org/chcs/wage-growth-tracker>). The top plot shows the full sample whereas the bottom plot shows the more recent sample. Grey shaded regions denote NBER recessions and vertical red dotted lines represent December 2017, 2019, and 2021, respectively.

