

NO. 1086
MARCH 2024

The Unemployment-Inflation Trade-off Revisited: The Phillips Curve in COVID Times

Richard K. Crump | Stefano Eusepi | Marc Giannoni |
Ayşegül Şahin

**The Unemployment-Inflation Trade-off Revisited:
The Phillips Curve in COVID Times**

Richard K. Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin
Federal Reserve Bank of New York Staff Reports, no. 1086

March 2024

<https://doi.org/10.59576/sr.1086>

Abstract

Using a New Keynesian Phillips curve, we document the rapid and persistent increase in the natural rate of unemployment, u_t^* , in the aftermath of the pandemic and characterize its implications for inflation dynamics. While the bulk of the inflation surge is attributed to temporary supply factors, we also find an important role for current and expected negative unemployment gaps. Through the lens of the model, the 2022-23 disinflation was driven by the expectation that the unemployment gap will close through a progressive decline in u_t^* and a rise in the unemployment rate. This implies that convergence to long-run price stability depends critically on expectations about labor market tightness. Using a variety of cross-sectional data sources, we provide corroborating evidence of unusually tight labor market conditions, consistent with our estimated rise in u_t^* .

JEL classification: D84, E24, E31, E32, J11

Key words: Phillips curve, unemployment, inflation, natural rate of unemployment, expectations

Crump: Federal Reserve Bank of New York (email: richard.crump@ny.frb.org). Eusepi, Şahin: University of Texas at Austin (email: stefano.eusepi@austin.utexas.edu, aysegul.sahin@austin.utexas.edu). Giannoni: Barclays (email: mpgiannoni@gmail.com). Marc Giannoni contributed to this paper while he was an employee of the Federal Reserve Bank of Dallas. This paper was initially presented at the January 3, 2021 ASSA session on “Measuring the Unemployment Gap.” The authors thank Emmanuel Saez and Pascal Michaillat for organizing the session and Regis Barnichon for his discussion. They also thank their editor Ricardo Reis, an anonymous referee, their discussants Jordi Galí and Jonathan Hazell, and the participants of the 2022 JME-SNB-SCG conference. In addition, they thank Edward Nelson for helpful comments about the interpretation of historical inflation dynamics. Jin Yan, Charles Smith, and Ignacio Lopez Gaffney provided excellent research assistance.

This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the author(s) and do not necessarily reflect the position of the Federal Reserve Banks of New York or Dallas or the Federal Reserve System, The University of Texas at Austin, or Barclays. Any errors or omissions are the responsibility of the author(s).

To view the authors’ disclosure statements, visit
https://www.newyorkfed.org/research/staff_reports/sr1086.html.

1 Introduction

In the aftermath of the COVID-19 pandemic the US economy experienced a swift recovery accompanied by a sharp rise in inflation, unseen since the late 1960s. This episode has re-ignited the debate about the trade-off between inflation and unemployment and its implications for inflation stability in the long-run. One camp argued that the surge in inflation was driven primarily by transitory factors, such as global supply chain disruptions and demand shifts, with little negative growth consequences of disinflation for the US economy.¹ A more pessimistic view embraced by others envisioned a costlier disinflation process leading to a recession (e.g., [Blanchard et al. 2022](#)).

In this paper we evaluate the unemployment-inflation trade-off through the lens of a simple New Keynesian Phillips curve focussing on the post-pandemic economy. In the model, inflation is determined by transient supply factors and by labor market conditions, as reflected by both the current and expected future unemployment gaps. The latter is the difference between the unemployment rate and the *natural rate of unemployment*, u_t^* , which is defined as the unemployment rate such that, controlling for supply shocks, inflation remains stable.

Crucially, we allow the natural rate of unemployment rate to vary over time, as argued by Friedman in his AEA Presidential address in 1968: *To avoid misunderstanding, let me emphasize that by using the term “natural” rate of unemployment, I do not mean to suggest that it is immutable and unchangeable...Improvements in employment exchanges, in availability of information about job vacancies and labor supply, and so on, would tend to lower the natural rate of unemployment.*

We use a wealth of labor market and inflation data to infer the evolution of both u_t^* and of economic agents’ expectations about the future path of the unemployment gap based on the framework in [Crump et al. \(2019\)](#). On the one hand, job market flows help identify the demographic factors driving the secular trend in the natural rate of unemployment. On the other hand, multiple nominal wage measures, coupled with inflation and survey-based inflation expectations provide information about higher frequency shifts in u_t^* and expectations about future labor market tightness. We find that the Phillips curve, covering the 1960–2023 period, captures the joint behavior of unemployment, wage and price inflation, and inflation expectations well with a time-invariant slope—estimated to be quite flat.

Our estimates point to a more nuanced view of post-pandemic inflation dynamics. While the bulk

¹For example, <https://www.nytimes.com/2021/09/10/opinion/transitory-inflation-covid-consumer-prices.html> and <https://rooseveltinstitute.org/2023/11/15/a-victory-lap-for-the-transitory-inflation-team/>.

of inflation is attributed to the direct effects of temporary factors, the more persistent component of inflation, driven by expected future unemployment gaps, has increased sharply from its long-run mean of 2% to 4% in 2022, a shift unseen since the early 1970s. This increase in underlying inflation in the aftermath of the pandemic is explained through a rapid and persistent increase in the natural rate of unemployment. After hovering around its secular trend of 4% before the pandemic, u_t^* reaches 7% in 2021 and stands at 6.5% at the end of 2023. This sizable increase surpasses the recent peak of 6.5% during the Great Recession, which was accompanied by persistent labor market disruptions.² Furthermore, it is comparable with the experience of the 1970s, once we take into account the higher secular trend in the unemployment rate during that period.

Model-implied underlying inflation has been falling since 2022 but remains above its pre-pandemic level. This decline in inflation was brought about by the expectation that the unemployment gap would close through a progressive decline in u_t^* and a rise in the unemployment rate. Despite the flat slope of our estimated Phillips curve, disinflation was not accompanied by a notable increase in the unemployment rate which remained below 4%. Instead, the *expectation* of the closing of the unemployment gap reduced underlying inflation from 4% at the peak to 3.2% at the end of 2023. To validate this key model prediction, we show that survey-based forecasts in 2022 and 2023 anticipated a rise in the unemployment rate. Although these data are not used in the estimation, the forecast paths were aligned with the corresponding model-based forecast.

In sum, our model implies that disinflation in 2022 and 2023 is accounted for by two components: the first is the abating of the transitory supply shocks and their direct impact on prices and the second is the expectation of the closing of the unemployment gap through a rise in unemployment and a decline in u_t^* . We find that while overall inflation has come down notably from its peak, the decline in underlying inflation was more modest suggesting that only half of the rise in the persistent component of inflation has been reversed. The final convergence to long-run price stability then depends critically on expectations of diminished labor market tightness.

The model's implications for the behavior of inflation is strongly linked to the stark and persistent rise in u_t^* . In the last section of our paper, we further investigate other labor market indicators and provide independent evidence for this persistent rise in u_t^* in the aftermath of the pandemic. We first show that reservation wages (the lowest wage or salary a worker would accept) have risen considerably and there has been a decline in willingness to work using the New York Fed's Survey

²Crump et al. (2019) show evidence of a decline in labor market matching efficiency during that period.

of Consumer Expectations (SCE). These labor supply constraints imply difficulty in hiring and an increase in labor costs which we examine directly. We find that the job-filling rate was at its all-time low in the first part of 2022—coinciding with the peak in inflation—and remained substantially below its pre-pandemic levels throughout 2023. Moreover, we show a strong historical linkage between our measure of underlying inflation and the job-filling rate suggesting that declines in inflation are tied closely to easing of hiring difficulties. Motivated by this finding, we study wage growth based on vacancies posted by the same firm for the same job over time. This measure also provides a forward-looking measure of wage growth since it captures new hires’ wages as documented by [Hazell and Taska \(2020\)](#).³ We find that posted wage growth has picked up substantially in the 2020–2022 period relative to 2017–2019, especially at low-paid jobs. Even though we document some moderation in wage growth in the 2021–2023 period, which coincided with the monetary policy tightening cycle, posted wage growth remains high compared to the pre-pandemic labor market. Taken in sum, this cross-sectional evidence is consistent with the time-series evidence we present from our model: the labor market has remained tight since 2021 and it is only gradually reverting to pre-pandemic conditions.

Related Literature. Our paper is related to a recent and growing literature focused on decomposing the drivers of inflation dynamics in the post-COVID period and the macroeconomic effects of supply chain disruptions and the shift of consumption from services to goods.⁴ This literature has emphasized that shocks to import prices and supply chains contributed to higher inflation, especially in the goods sector in 2021 and 2022. We also find an important role of supply shocks for the rise in inflation and the subsequent decline in the 2022–2023 period. In addition, we provide an estimate for underlying inflation and show that while it did not spike as starkly as total inflation, its increase has been more persistent and it has remained elevated throughout 2023.

We also relate to recent work re-examining the Phillips curve by offering modifications to the slack measure. For example, some papers have argued for replacing the unemployment rate in the Phillips curve with the vacancy-to-unemployment ratio such as [Ball et al. \(2022\)](#), [Bernanke and Blanchard \(2023\)](#) and [Benigno and Eggertsson \(2023\)](#). This choice is motivated by the observation that job-openings and inflation surged together after the pandemic. However, the performance of the vacancy-to-unemployment ratio in capturing inflation is not better than the unemployment rate

³Moreover, it minimizes the role of composition bias and influence of worker characteristics.

⁴See for example, [di Giovanni et al. \(2022\)](#), [Baqae and Farhi \(2022\)](#), [Cavallo and Kryvtsov \(2023\)](#), [Amiti et al. \(2023\)](#), [Shapiro \(2022\)](#), [Rubbo \(2024\)](#), [Comin et al. \(2023\)](#), [Ferrante et al. \(2023\)](#)

in the pre-pandemic data as shown in [Furman and Powell \(2021\)](#) and, in particular, in the 1970s as argued by [Şahin \(2022\)](#). Moreover, as [Barlevy et al. \(2023\)](#) argue, the ratio between vacancies and unemployment (instead of the unemployment rate) should not matter if the Beveridge curve is stable. When the Beveridge curve shifts, relying on the vacancy-to-unemployment ratio without examining the reasons behind the shift could be misleading.

A second type of modification to the Phillips curve is to allow for a time-varying slope which goes back to [Ball et al. \(1988\)](#) and more recently [Cerrato and Gitti \(2022\)](#), [Boehm and Pandala-Nayar \(2022\)](#) and [Harding et al. \(2023\)](#). The main idea is that when the economy gets closer to its capacity constraints, wage and price pressures might build up rapidly leading to a nonlinear Phillips curve. While it is tempting to consider a shift in the slope of the Phillips curve when the observed relationship between the unemployment rate and inflation changes, the reason for the assumed change in the slope is often traced back to an omitted variable bias. For example, as argued by [Crump et al. \(2019\)](#) and [Hazell et al. \(2022\)](#), the flattening of the unemployment-inflation relationship in the 1990s and 2000s reflected the omission of the role of inflation expectations which became anchored after the Volcker disinflation.

We, instead, take a different approach and continue relying on the unemployment rate as our measure of labor market tightness and complement it with multiple measures of labor compensation. This approach is informed by multiple reasons. The first and most important reason is that the unemployment rate has been consistently measured starting with the introduction of the Household Survey starting in 1948. Second, the drivers of the trends in the unemployment rate are much better understood than the trends in the job openings rate which is based on the JOLTS survey starting in 2000 and the Help Wanted Index before 2000—two different datasets in scope and source. We estimate and control for the secular trend in the unemployment rate carefully in our analysis, which is an arguably more reliable process because of the availability of rich cross-sectional and time-series data on labor-market flows. Third, our framework allows for the natural rate of unemployment to vary over time. Furthermore, our detailed analysis of other labor market indicators lend independent support for the rise in the natural rate instead of relying on different indicators for different periods and assuming a different slope for each episode (e.g., [Ball and Mazumder \(2019\)](#)).

The structure of the paper is as follows. Section 2 introduces the modeling framework and presents the time series for the natural rate of unemployment, u_t^* and underlying inflation. Section 3 discusses the unemployment-inflation trade-off for the pandemic period. Section 4 presents external

evidence on the state of labor market conditions in support of the model results. Section 5 concludes. An online Supplemental Appendix (hereafter, “SA”) provides additional results.

2 A Micro-Macro Phillips Curve Framework

Following Crump et al. (2019), we estimate u_t^* in two steps, succinctly summarized in the following equation:

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

In the first step we extract \bar{u}_t , the secular trend in the unemployment rate, from the inflow and outflow rates using a linear unobserved factor model. In 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 Estimating the Unemployment Trend from Labor Market Flows

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.⁵

⁵It 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 (2005, 2012), Hall (2005), Elsby, Michaels, and Solon (2009), Elsby, Hobijn, and Şahin (2010), Davis, Faberman,

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

$$u_t \approx \frac{s_t}{s_t + f_t}. \quad (3)$$

The Current Population Survey (CPS) provides monthly measures of unemployment, short-term unemployment, and the 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.⁶ In particular, we follow [Crump et al. \(2019\)](#) and use an unobserved components model with a slow-moving trend and serially correlated cyclical dynamics to estimate the secular trends in the inflow and outflow rates and, therefore, the unemployment rate.

Haltiwanger, Jarmin, and Miranda (2010) and [Şahin, Song, Topa, and Violante \(2014\)](#).

⁶We correct for the effects of CPS redesign on duration of unemployment using the correction factors in [Elsby, Hobijn, and Şahin \(2010\)](#).

2.2 The Secular Trend in the Unemployment Rate

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 sample, 1960Q1–2023Q3. The secular trend of the inflow rate shows a decline of about 50% since the 1980s. In contrast, the secular trend 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

$$\bar{u}_t = \frac{\bar{s}_t}{\bar{s}_t + \bar{f}_t}, \quad (4)$$

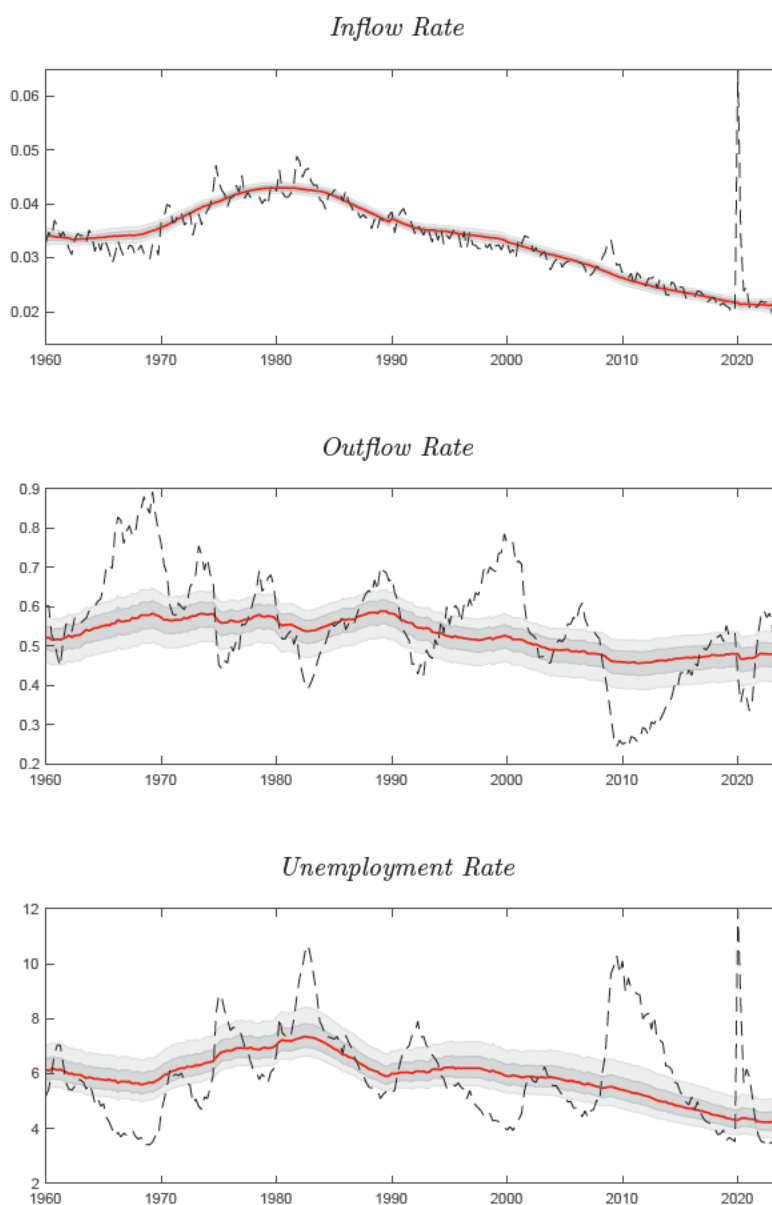
and is shown in the bottom panel of Figure 1. The trend unemployment rate was about 6% in 1960 and increases to over 7% in 1983. Since then it has displayed a clear downward trend, reaching about 4.25% toward the end of 2023. Interestingly, this downward trend continued even after the dramatic job losses 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.

[Crump et al. \(2019\)](#) 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. 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 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.

While grand gender convergence was important in accounting for the secular decline in the unemployment rate until 2000, the shift towards older workers and older firms since the mid-1990s,

which Crump et al. (2019) refer to as dual aging, stands out as an important driver of the declining secular trend rate of unemployment in the last two decades. Younger workers are more likely to flow into unemployment than prime-age workers. Similarly, young firms, aged between one and five years old, are also more 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. Given that worker and firm demographics tend to be slow moving, we would not expect the secular trend of unemployment to start rising in the medium term.

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 red lines. Shading denotes 68% and 95% posterior coverage intervals.



2.3 The 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.⁷ We base our empirical specification of the New Keynesian Phillips curve on Galí (2011). In this setting, 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

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

$$\pi_t^w = g_w + \pi_t - \beta^{-1} \frac{1 - \beta \rho_\eta}{1 - \rho_\eta} \eta_t \quad (6)$$

$$\eta_t = \rho_\eta \eta_{t-1} + \sigma_\eta \epsilon_t^\eta, \quad (7)$$

where $\beta \in (0, 1)$ is the discount factor and $\rho_\eta \in (0, 1)$.

Wage and price inflation are determined by five key factors. Labor productivity and price markup shocks are captured by the process η_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 discounted expected path of the unemployment gap. 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 mean of inflation, 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 determinant of wage inflation. In fact, inflation expectations contain information about expected future unemployment gaps: using equations (5) and (7) we obtain

$$\mathbb{E}_t \pi_{t+1} = \pi_t^* + \gamma(\pi_t - \pi_t^*) - \kappa \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} x_{s+1} + \rho_\eta \eta_t. \quad (8)$$

⁷In this simple model the natural rate of unemployment is driven by market distortions captured by shifting market power of workers and other factors.

2.4 Estimating the Wage Phillips curve

Suppose we are under the ideal conditions where 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 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 of the model’s components. In addition to the unemployment rate, 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 forecasters.⁸ That said, we face two challenges. First, information about wage growth and inflation expectations contains significant measurement errors.⁹ In the case of wages, this is evident from the fact that we use multiple measures for underlying nominal wage growth π_t^w . This limitation implies that we can only infer the unemployment gap with some degree of uncertainty. Second, we do not have strong prior information about key parameters such as the slope of the Phillips curve, κ , which needs to be estimated. In fact, a large literature is 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 \(2001\)](#), [Laubach and Williams \(2003\)](#), and [Galí \(2011\)](#) we model these unobserved components as exogenous processes:

$$z_t = \rho_z z_{t-1} + \sigma_z \epsilon_t^z \tag{9}$$

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

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 allow for the construction of forecasts which we utilize in our out-of-sample forecasting exercise in Section 3. In particular, time- t forecasts at horizon $n > 1$ (and the associated forecast distribution) for inflation and the natural

⁸Details about the dataset can be found in Appendix A.

⁹We assume our measure of CPI inflation and inflation expectations are noisy measures of the model’s counterparts. We therefore impose *i.i.d.* Gaussian measurement errors. For the different wage measures we assume measurement errors that are independent first-order autoregressive processes (AR (1)).

rate of unemployment are:

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

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

$$u_{t+n|t} = u_{t+n|t}^* + x_{t+n|t} \quad (13)$$

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 key model output is a measure of “underlying” or “fundamental” inflation. Underlying inflation is defined as

$$\tau_t \equiv \pi_t^* + \gamma(\pi_{t-1} - \pi_{t-1}^*) - \kappa \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} x_s \quad (14)$$

Intuitively, underlying inflation represents the component which solely depends on the long-run trend and the sequence of current and future unemployment gaps (Del Negro et al. 2015). Said differently, this measure excludes the direct influence of productivity or price markup shocks on the inflation process.¹⁰ From equations (6) and (14) we can re-express inflation and our underlying nominal wage inflation measures as

$$\pi_t = \tau_t + \eta_t \quad (15)$$

$$\pi_t^w \approx g_w + \tau_t, \quad (16)$$

where the approximation in the second line stems from our assumption that the discount rate β is close to one.¹¹ Information about τ_t can be extracted from the multiple observed measures of nominal wage growth, price inflation, and inflation expectations at different horizons. Inflation is then described by the sum of underlying inflation and the direct effects from supply shocks.

The model is estimated with Bayesian methods over the sample 1960Q1–2019Q4 using quarterly data. Details about the estimation approach can be found in Crump et al. (2019). We restrict the estimation period to the pre-COVID sample to ensure that our parameter estimates are not driven

¹⁰Note that in a structural model, general equilibrium effects can drive co-movement between these two variables, as supply shocks can affect the path of the unemployment gap. Here our goal is not to decompose the evolution of inflation and the unemployment gap in terms of underlying structural shocks. Instead, we focus on measuring the role of the persistent component τ_t .

¹¹The parameter is set to $\beta = 0.99$ so that shocks measured by η_t have minimal effects on underlying nominal wage growth—see equation (6).

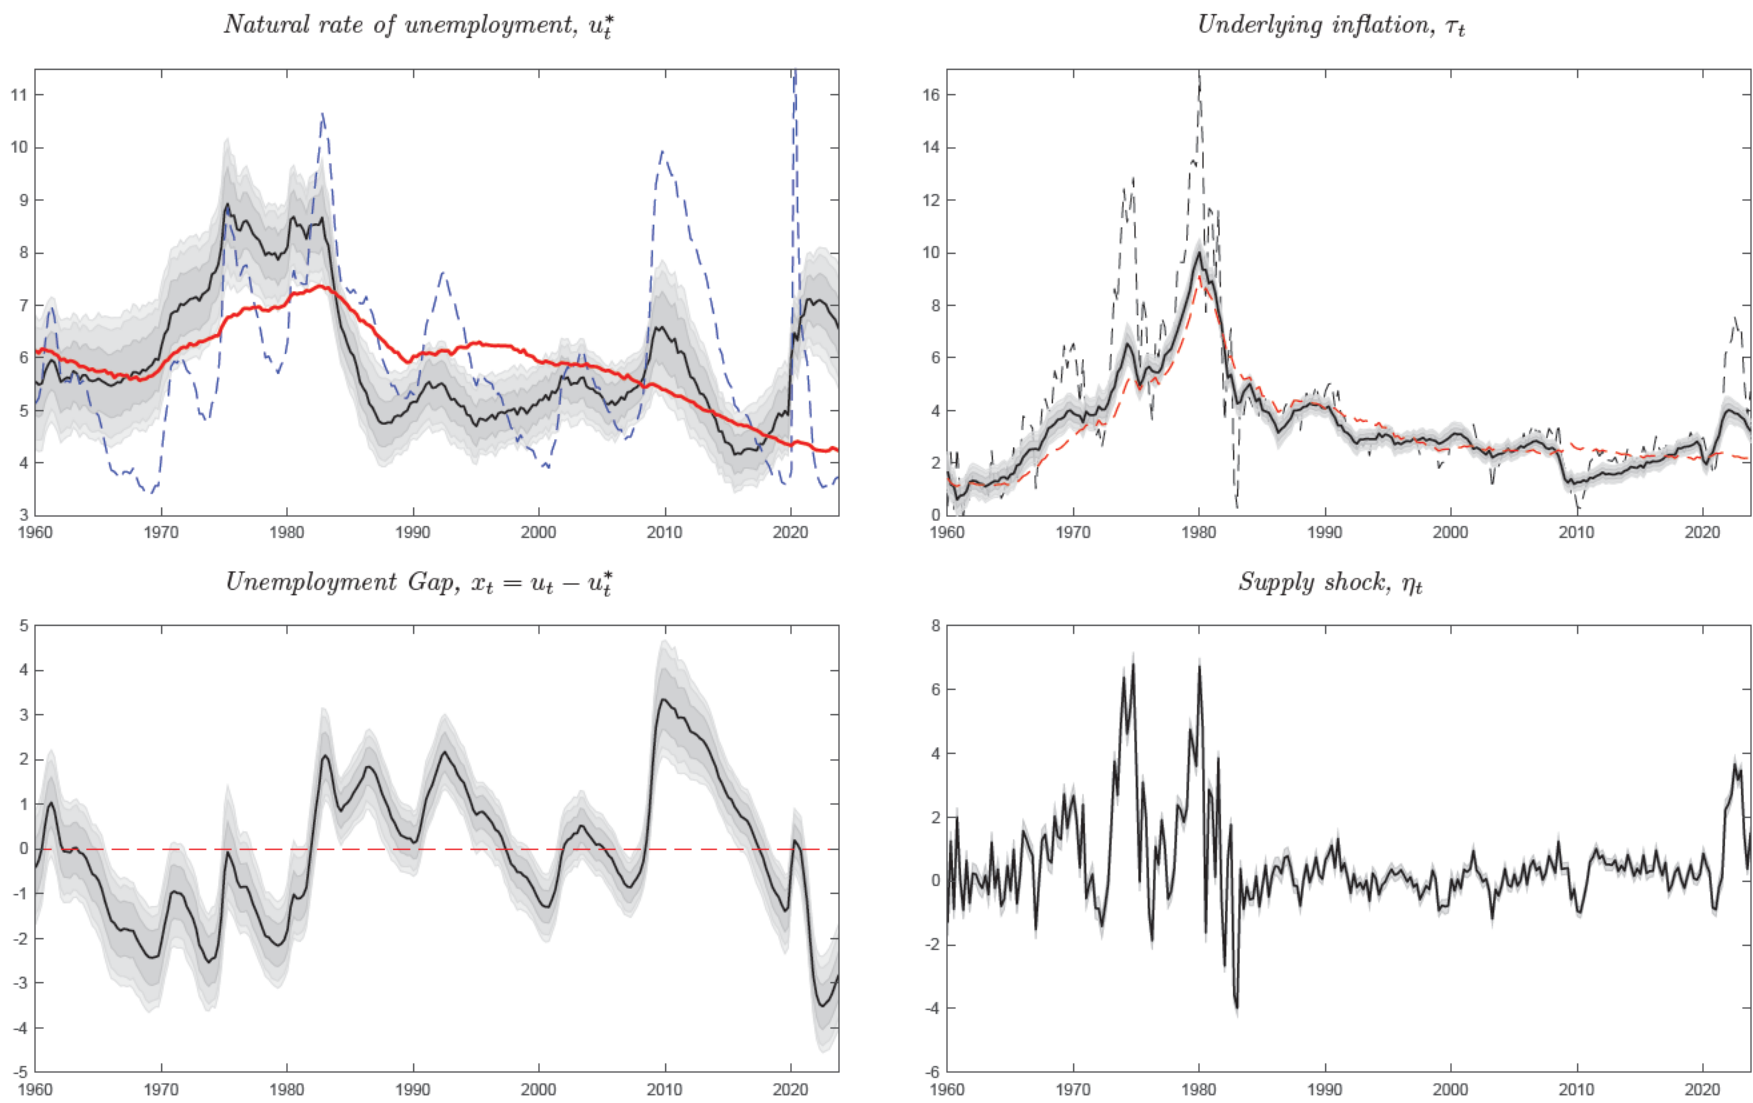
by the large changes in some series that occurred in 2020 and 2021. We then use the obtained posterior distribution of the parameters to estimate the model unobservables up to 2023Q4.¹² It is important to emphasize that our estimate of the unemployment gap, $x_{t|t}$, reflects all available information through model linkages. This includes the secular 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 of the natural rate of unemployment. The slope of the Phillips curve, κ , is estimated with a median of 0.03 and a range of 0.02–0.06 which implies a fairly flat curve, as is often found in the literature (for recent papers see, for example, [Del Negro et al. 2020](#) and [Hazell et al. 2022](#)). 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. The persistence of supply shocks, ρ_η is estimated in the range 0.3–0.4. 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 labor market conditions play an important role beyond the slow-moving demographic factors captured by our estimate of \bar{u}_t .

2.5 The Natural Rate of Unemployment and Underlying Inflation

The estimated model allows us to examine the evolution of two key unobserved variables, the natural rate of unemployment, u_t^* and the underlying rate of inflation, τ_t , since 1960. The top left chart in [Figure 2](#) shows that the natural rate exhibits persistent fluctuations around its long-run trend \bar{u}_t . After hovering around 6% through the 1960s, u_t^* starts rising in the early 1970s, reaching almost 9% by the late 1970s. The natural rate then declines sharply throughout the 1980s and over the time period spanning the 1990s to the Great Recession, it stabilizes in a range of 4.8% to 5.8%. In the Great Recession the natural rate increased to about 6.5% but, after reaching its peak, began a steady descent to a little above 4% in 2017, catching up with its long-term trend. As the expansion matured, u_t^* reversed course and rose to just below 5.0% at the end of 2019. During the COVID-19 pandemic, u_t^* increased appreciably to 7% at the end of 2022 and has declined to 6.6% toward the end of 2023. Such a sharp and persistent increase of u_t^* has only previously occurred in the 1970s.

¹²We utilize the one-period forecast of \bar{u}_t to obtain a value for 2023Q4. Wage series which were not yet available for 2023Q4 are treated as missing observations.

Figure 2. Phillips Curve Estimates This figure shows the main model outputs from our estimated Phillips curve. The top-left panel shows the estimate of u_t^* (black line) along with the observed unemployment rate (blue dashed line) and the median secular trend, \bar{u}_t (red line). The bottom-left panel shows the unemployment gap (black line). The top-right panel shows the model-implied estimate of underlying inflation (black line), τ_t , 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-right panel shows the estimated supply shock, η_t (black line). Grey shaded areas denote 68%, 90% and 95% posterior coverage intervals.



The top right chart of Figure 2 presents underlying inflation, τ_t , alongside realized inflation and the agents' estimate of trend inflation. Underlying inflation rose dramatically throughout the 1970s and then fell precipitously in the Volcker disinflation period. After a period of relative stability it declines during the financial crisis before gradually returning toward its long-run trend. During the pandemic, underlying inflation increased markedly (at a pace matched only in the 1970s) and has partially reversed course during 2023. Underlying inflation stood at 3.2% at the end of 2023—about a percentage point above its long-run trend. The bottom right chart of Figure 2 shows the time series of the supply shock, η_t . The direct effect of pandemic-related supply shocks accounts for the bulk of the difference between realized inflation and underlying inflation. Again, the only precedent for such behavior is the 1970s.

Through the lens of the model the relation between the current unemployment rate and wage/price inflation is time-varying and depends on the behavior of inflation expectations and the natural rate of unemployment. Thus, the joint behavior of these key objects is vital to analyze inflation dynamics historically. Consider the following two examples: (i) If one ignores the de-anchoring in the 1970s and subsequent anchoring of inflation expectations in the 1990s the Phillips Curve would appear steeper in the 1970s and the 1980s before becoming flatter thereafter. (ii) Inflation expectations, and therefore the expected path of future unemployment gaps, played an important role in explaining the “missing” disinflation during the Great Recession. The subdued response of inflation to the persistently high unemployment rate reflected two factors: first, the natural rate of unemployment is estimated to have risen in the aftermath of the financial crisis, attenuating the rise in the unemployment gap; second, there was an expectation of the unemployment gap narrowing in the future (Del Negro et al. 2015, Crump et al. 2019).

In the next section we discuss how inflation expectations and shifts in the natural rate of unemployment have shaped the unemployment-inflation trade-off during and after the pandemic.

3 The Unemployment-Inflation Trade-Off During the Pandemic

The large increase in inflation in the aftermath of the pandemic has triggered a debate about its costs for the US economy. One side of the debate viewed the rise in prices as temporary, driven primarily by supply disruptions, and with little to no consequences for long-run inflation stability. The other side feared a scenario of more persistent inflation and a costly disinflation. Figure 3 shows the model-implied path for τ_t since 2016. After narrowly fluctuating around its long-run trend,

π_t^* , underlying inflation sharply increases by two percentage points in 2022 and gradually declines afterwards. While the direct effects of temporary supply shocks η_t account for the bulk of the rise in realized inflation (see Figure 2), the increase in τ_t is persistent and, even at the end of 2023, remains about one percentage point above its long-run trend. These conditions suggest that achieving the so-called final mile (Daly 2023) of disinflation then depends critically on expectations of diminished labor market tightness. To see this note that by equation (14) we have that,

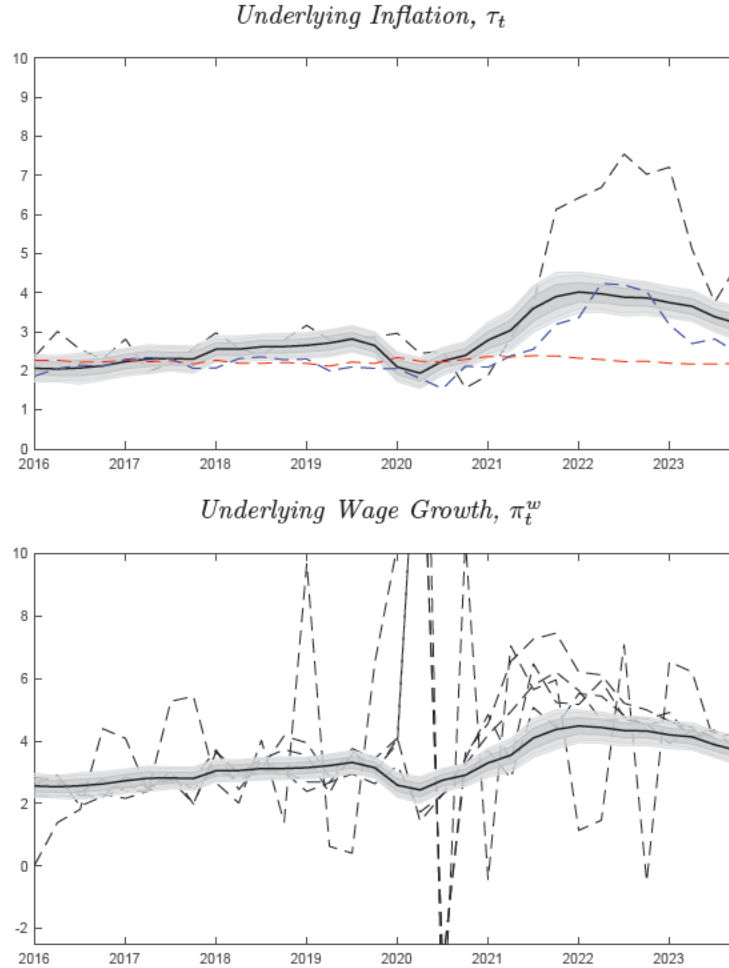
$$\tau_t \approx \pi_t^* - \kappa \mathbb{E}_t \sum_{s=t}^{\infty} \beta^{s-t} x_s. \quad (17)$$

Given that long-run inflation expectations have remained stable ($\pi_t^* \approx \pi^*$ in this period), the persistent rise in underlying inflation must reflect both current and expected discounted future unemployment gaps.¹³ As shown in the previous section these unobserved variables are inferred from price inflation and inflation expectations (top panel of Figure 3) and nominal wage growth measures (bottom panel of Figure 3).

The unemployment gap x_t widened sharply in the pandemic and was estimated at -3.5 percentage points in the first half of 2022 (bottom left panel of Figure 2). This was the largest negative unemployment gap in our sample. As of 2023Q4, the unemployment gap stood at -2.8 percentage points. Given the evolution of the unemployment rate, the measured gap reflects a sizable increase in the natural rate of unemployment, u_t^* , from 4.9% in 2019Q4 to 6.6% in 2023Q4.

¹³Here the approximation uses the result that inflation inertia is estimated to be small and thus has little effect on underlying inflation.

Figure 3. Underlying Price and Wage Inflation The top panel shows underlying inflation τ_t (black line) along with short-term inflation expectations (blue dashed line), 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 underlying wage growth $\pi_t^w \approx g_w + \tau_t$ (black line) along with the individual wage series used in the estimation (grey dashed lines). Grey shaded areas denote 68%, 90% and 95% posterior coverage intervals.

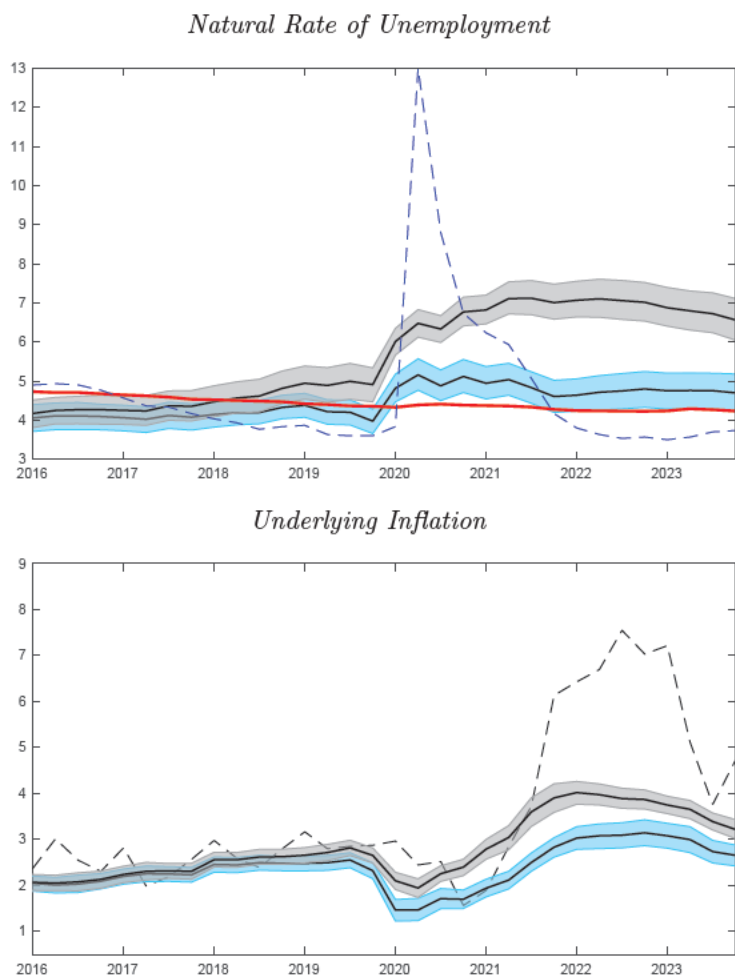


What signal drives the estimate of x_t and, consequently, u_t^* ? Figure 4 shows that the behavior of observed wages in the aftermath of COVID is the dominant signal for the increase in u_t^* . The blue shaded area shows the coverage interval for the model predictions without wage growth information, while the grey shaded area shows our baseline estimate.¹⁴ The top panel of the figure shows that this alternative estimate of u_t^* is now much more muted, hewing closely to the secular trend in the unemployment rate, \bar{u}_t , with a near-zero unemployment gap in mid-2022. Consequently, underlying inflation, τ_t , does not deviate significantly from π_t^* despite the large rise in realized inflation and short-term inflation expectations. Instead, the increase in inflation observed during the period is entirely attributed to the direct effects of temporary supply shocks. This corresponds to the most optimistic

¹⁴In this exercise we omit the five measures of wage growth from the observation equation starting in 2019Q4.

scenario where the surge in inflation was a purely temporary phenomenon. The key informational role of nominal wage measures for the rise in u_t^* suggests that changes in the labor market have played an important role in shaping the behavior of inflation during and after the COVID-19 pandemic.

Figure 4. 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 alternative estimate of the natural rate of unemployment, u_t^* (black lines) along with the observed unemployment rate (blue dashed line) and the median secular trend, \bar{u}_t (red line). The grey shading corresponds to the coverage interval of the baseline model, while the light blue shading corresponds to the coverage interval of the model without wage information. The bottom panel shows the alternative estimate of underlying inflation (black line) along with the Federal Reserve Bank of Cleveland’s median CPI inflation (grey dashed line). Shaded areas denote 68% posterior coverage intervals.



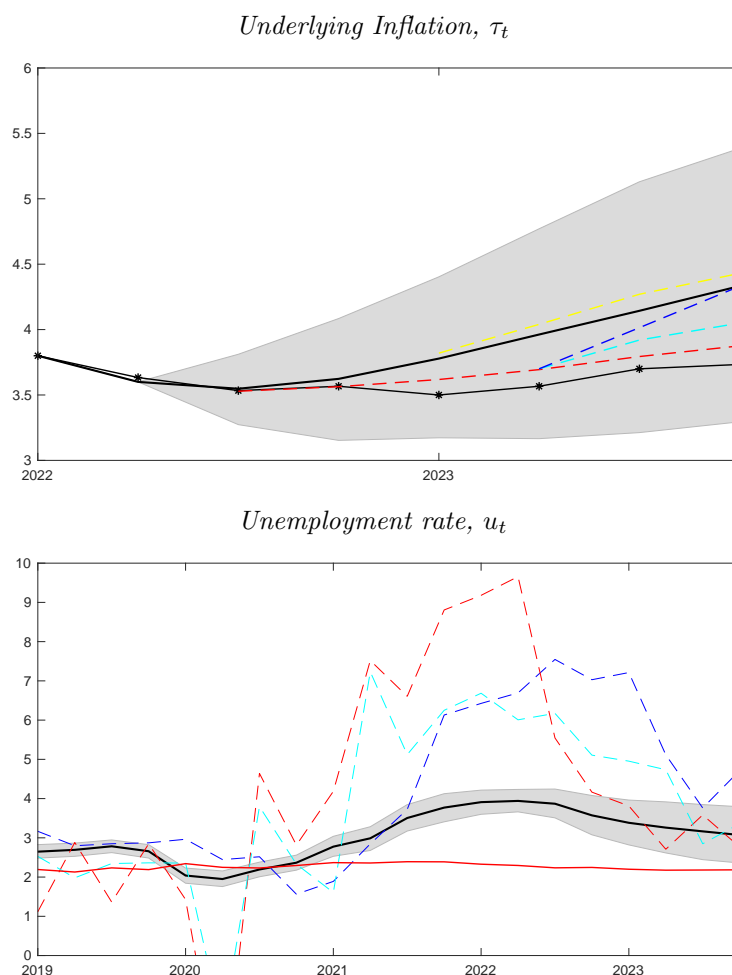
While the large negative unemployment gap x_t accounts for higher underlying inflation, how do we explain its gradual decline over 2023? Here it is useful to recall that the current unemployment gap in this model plays only a minor role in determining τ_t , especially given that our estimated Phillips curve is quite flat. The key driver of underlying inflation is the expectation about *future* unemployment gaps. Through the lens of our model, the gradual reduction of τ_t is achieved through

a modest expected increase in the unemployment rate (a so-called soft landing) and an expected decline in the natural rate of unemployment.

To validate this model prediction Figure 5 shows the real-time, out-of sample forecasts for underlying inflation and the unemployment rate from the model as of 2022Q2. The top panel of Figure 5 shows a gradual increase in the unemployment rate (median) forecast to above 4% at the end of 2023 (black solid line). While this is decisively higher than the realized unemployment rate over the period (the starred-black line is flat over the forecasting horizon), what matters for the dynamics of inflation is the *expectation* about future unemployment. We can compare this model forecast of the unemployment rate with those of professional forecasters. By mid-2022 both the Survey of Professional Forecasters (SPF) and the Blue Chip Economic Indicators (BCEI) survey display upward-sloping paths for the unemployment rate in line with the model forecast. This expectation of a gradual increase in the unemployment rate then drives the gradual decline in underlying inflation (see equation (14)).

The bottom panel of Figure 5 shows that common measures of realized inflation converge to the model forecast for underlying inflation by the end of 2023. Importantly, this out-of-sample exercise shows that underlying inflation, τ_t , captures the medium-term dynamics of inflation whereas the direct effect of the supply shock, η_t , and the long-run trend, π_t^* play the key role for short-term and long-term dynamics, respectively.

Figure 5. Out-of-Sample Forecast Comparison The figure shows the real-time, model-implied forecast distribution for the unemployment rate (top panel) and underlying inflation (bottom panel), based only on information available up to 2022Q2. The black solid line denotes the median prediction, while the grey area shows the 68% posterior coverage interval. In the top panel, the line with asterisks denotes the realized unemployment rate through 2023Q4, while the dashed lines display professional forecasts of the unemployment rate taken at different times. The red and blue lines show the Blue Chip Economic Indicators forecast path as of 2022Q2 and 2023Q1, respectively. The yellow and cyan lines denote the forecast path from the Survey of Professional Forecasters as of 2022Q4 and 2023Q1, respectively. In the bottom panel, the dashed lines represent realized CPI inflation (red), the Federal Reserve Bank of Cleveland’s median CPI inflation (blue) and core CPI inflation (cyan).



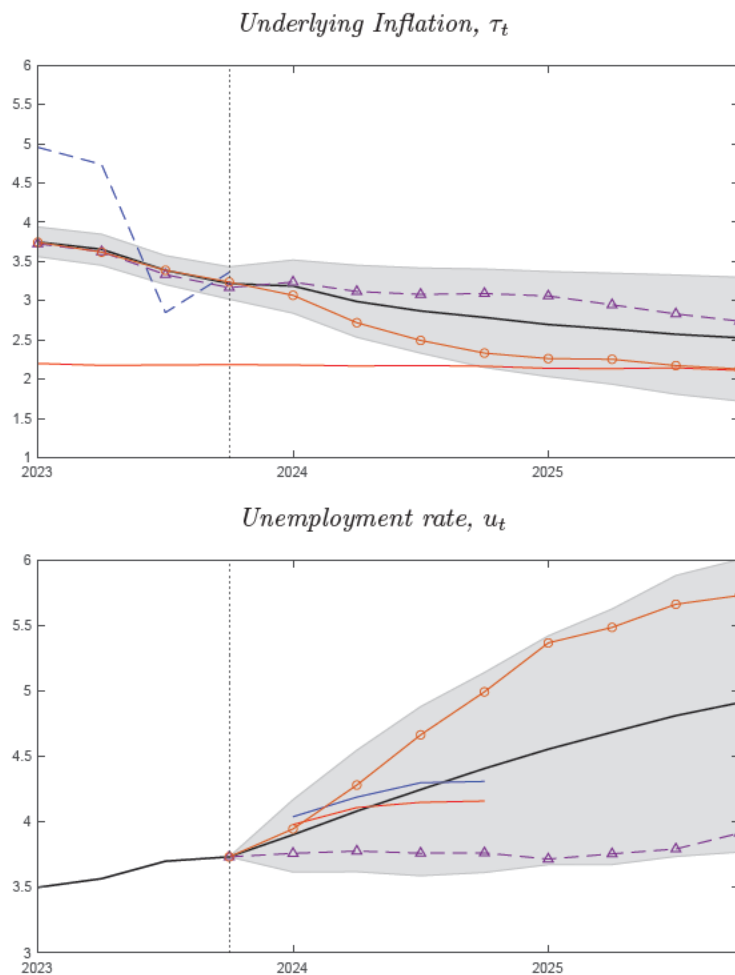
Why is it different from the 1970s? The reemergence of high inflation readings has drawn comparisons to the US economy in the 1970s and early 1980s where sharp monetary tightening led to economic contraction and a period of disinflation. However, we see two key differences between this period and the state of the economy after 2020 that significantly worsened the unemployment-inflation trade-off in the former period. First, there is a 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 and early 1980s 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 and early 1980s are no longer present. Second, the behavior of long-term inflation expectations over the two periods is markedly different. [Carvalho et al. \(2023\)](#) show that inflation expectations were unmoored during the 1970s, but have remained firmly anchored over the past twenty years.

What is the current forecast? Figure 6 shows the forecast distribution for underlying inflation and the unemployment rate based on information up to 2023Q4. Consistent with model predictions from the previous year (Figure 5), the top panel shows τ_t gradually converging toward its long-run mean through 2025, driven by a moderate rise in the unemployment rate. As can be gleaned from the bottom panel, 2023Q4 survey responses from professional forecasters (red and blue lines) showed a continued expectation of a rise in the unemployment rate over the forecast horizon. Looking at the overall forecast distribution, there is considerable uncertainty about the final mile in the disinflation process. The orange and purple lines correspond to possible scenarios conditioning on the top 75th and bottom 25th percentile of the unemployment distribution.¹⁵ The purple line denotes a scenario where the unemployment rate barely increases over the forecasting horizon. Consequently, disinflation is slower. The orange line indicates a scenario where inflation reaches its long-run trend before 2025. The unemployment rate path that delivers this outcome crosses 5% by mid-2024. Through the lens of the model, we observe that labor market conditions define the path to long-run price stability.

¹⁵We sort the projected unemployment rate paths in 2025Q1 and average 50 paths above (below) the path corresponding to the 25th (75th) percentiles.

Figure 6. Current Forecast The figure shows the model-implied forecast distribution for underlying inflation (top panel) and the unemployment rate (bottom panel), based on information available up to 2023Q4 (delineated by the vertical line). The black solid line denotes the median prediction, while the grey area shows the 68% posterior coverage interval. The orange (purple) lines denote inflation forecasts conditional on a steeper (shallower) increase in forecasted unemployment relative to the median forecast. In the top panel, the red line denotes estimated π_t^* , while the dashed blue line is core CPI inflation. In the bottom panel the red and blue lines show the Survey of Professional Forecasters and Blue Chip Economic Indicators forecast paths as of 2023Q4.



4 The Post-Pandemic Labor Market and the Rise in u^*

The natural rate of unemployment plays the key role in our model and we view the unemployment gap as an appealing measure of labor market tightness. Our estimate of the unemployment gap implies that the post-pandemic period has seen the tightest labor market in our sample. At the end of 2023 the negative unemployment gap remains unprecedentedly large. In contrast, the unemployment rate would suggest a labor market about as tight as at the end of 2019. From the end of 2019 to the end of 2023, the unemployment gap widened by 1.5 percentage points—almost entirely driven by the rise

in the natural rate of unemployment.

In this section, we investigate the sources of the rise in the natural rate of unemployment in the 2020–2023 period using a variety of cross-sectional data to provide external corroboration for our analysis. First, we analyze the evolution of reservation wages and desired work hours and find that the pandemic caused a shift in reservation wages and willingness to work. Second, we examine the joint evolution of the job-filling rate and underlying wage growth and find that the job-filling rate had declined to its lowest point in the first quarter of 2022 and has only increased modestly in 2023, consistent with the emergence of labor shortages and increased difficulties in filling jobs.¹⁶ Last, 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 which connects to the evolution of underlying wage growth.

4.1 Rise in Reservation Wages and Decline in Willingness to Work

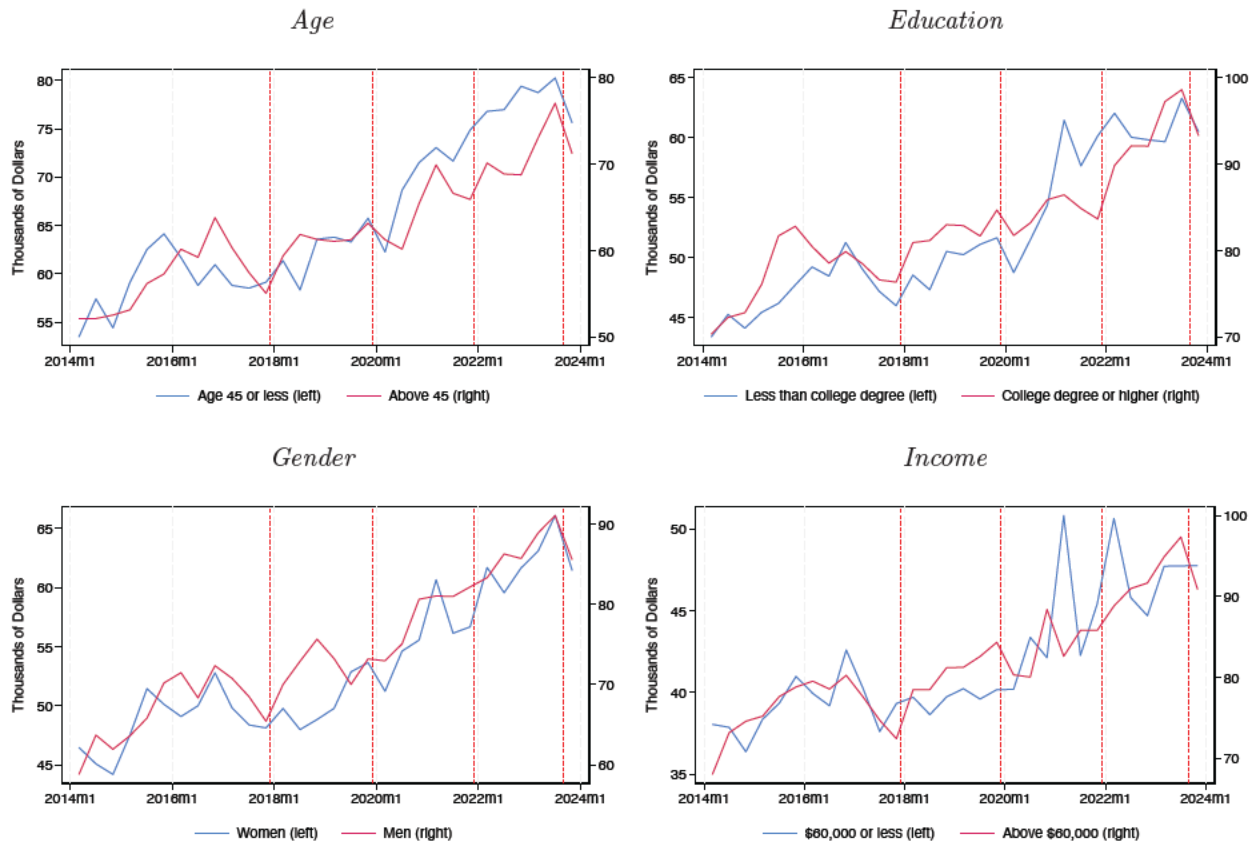
The labor shortages that emerged following the COVID pandemic have been widely discussed and several factors have been attributed to this phenomenon such as health risks, persistent effects of infections, changes in immigration flows, and a re-evaluation of work-life balance. A useful metric to summarize 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 Expectations:

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 wages could reflect perceived health risk but also changes in preferences and willingness to work. In Figure 7, the top left chart 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.

¹⁶Anecdotal evidence on employers’ difficulties finding and retaining workers can be found, for example, in <https://www.federalreserve.gov/monetarypolicy/beigebook202110.htm> or <https://www.federalreserve.gov/monetarypolicy/beigebook202112.htm>. See also <https://www.cbsnews.com/news/full-transcript-fed-chair-jerome-powell-60-minutes-interview-economy/>. Survey evidence is presented in Figure B.3 in the SA.

Figure 7. 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 2014M3–2023M11. Vertical red dotted lines represent December 2017, 2019, 2021, and September 2023, respectively.



In the top right chart of Figure 7, 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 findings of Autor et al. (2023) who document compression in the US wage distribution data. Finally, in the bottom two panels we show reservation wage split by gender and household income which show similar patterns.

These results are consistent with Faberman et al. (2022) who show that willingness to work has declined after the pandemic along both the extensive and intensive margins of labor supply. Figure 8 shows the series they construct using the SCE data.¹⁷ The series is based on the following question that is asked to all respondents who searched or said they wanted work in the last 4 weeks:


Assuming you could find suitable/additional work, how many hours PER WEEK would

¹⁷Data are available at <https://docs.google.com/spreadsheets/d/1r-0ujZDDIrt91FV8Fj11HCxT2zmlM7Lr/edit?usp=sharing&oid=107507111193001876864&rtmpof=true&sd=true>.

you prefer to work on this new job?

The figure shows that there was a drastic drop in the total number of hours that individuals would prefer to work at the onset of the pandemic which only slowly recovered. Despite this gradual recovery in willingness to work, desired hours remain below their pre-COVID levels.

Figure 8. Desired Work Hours This chart shows the potential work hours series from [Faberman et al. \(2022\)](#). The sample period is 1994M1–2023M6. Grey shaded regions denote NBER recessions.



4.2 Job-Filling Rate, Matching Efficiency and Wage Growth

One indicator that is especially useful in assessing labor market tightness in light of the changing willingness to work is the so-called job-filling rate.¹⁸ 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. This is likely also exacerbated by what is often referred to as the “Great Resignation.” The quits rate, which was around 2.4% in 2019, peaked at 3% in early 2022 as more workers re-evaluated their current jobs and decided to quit in search of other opportunities at a higher rate than before the pandemic.¹⁹

¹⁸For example [Abraham et al. \(2020\)](#) advocate using hires-based measures of labor market tightness.

¹⁹See [Bagga et al. \(2024\)](#) for a general equilibrium model which links the decline in the job-filling rate to the rise in quits.

Formally, the job-filling rate is defined as 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’ Job Openings and Labor Turnover Survey (JOLTS), we can calculate the aggregate job-filling rate since 2000. The left chart in Figure 9 shows our underlying wage growth measure from equation (16), π_t^w , along with the economy-wide job-filling rate.²⁰

Figure 9. Underlying Wage Growth and Job-Filling Rate This figure displays the relation between the underlying wage growth introduced in Section 2 and the job-filling rate from JOLTS. The left plot of the figure shows the time series for each variable whereas the right plot shows a scatterplot of two-year changes in each variable along with the OLS fitted line for the pre-COVID sample (solid line) and the full sample (dashed line). The sample period is 2000Q4–2023Q4. Grey shaded regions denote NBER recessions.

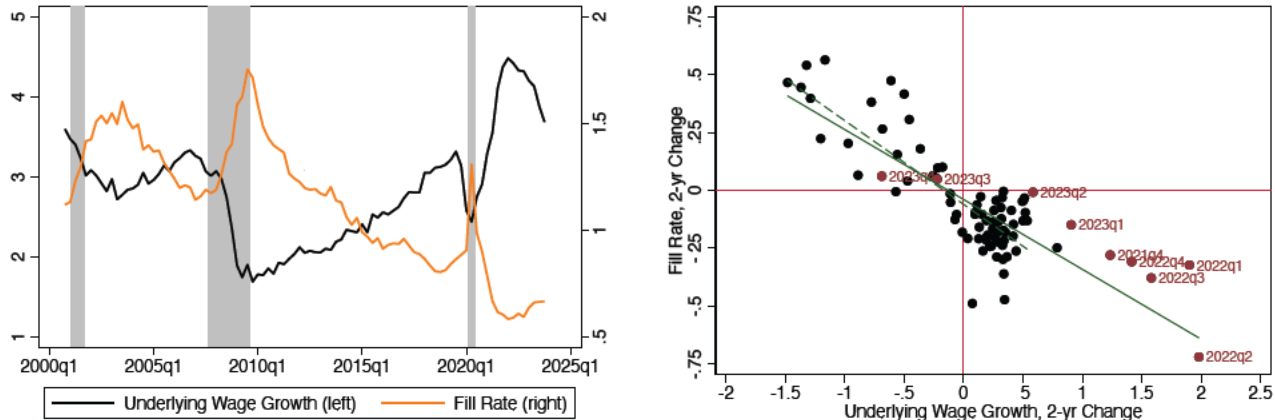


Figure 9 shows that when the labor market is tight (when it is harder for firms to fill open positions) wage growth in the economy tends to be higher. In the right chart of Figure 9 we show the two-year changes in underlying wage growth and the job-filling rate. Remarkably, almost all of the data points 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. It is important to recall that the job filling rate is not used in the model estimation. We also overlay the OLS fitted lines for the pre-pandemic period (solid line) and full sample (dashed line) which have very similar slopes. The job-filling rate stood at a little below 0.6 in early 2022 when the monetary policy tightening cycle began. In concert, underlying wage peaked at 4.5% which was the highest value since 1990. Since early 2022, as the job-filling rate started to increase, underlying wage growth also started to moderate but remains comfortably above its pre-pandemic level.

One might argue that a simple metric for tightness, the vacancy-to-unemployment ratio, as

²⁰While 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 B.1 in the SA.

typically defined in search-theoretic models of the labor market, would accurately capture the state of the labor market (e.g., [Ball et al. 2022](#), [Bernanke and Blanchard 2023](#), [Benigno and Eggertsson 2023](#)). This is because in the Diamond-Mortensen-Pissarides (DMP) framework, the vacancy-to-unemployment ratio and the job-filling rate co-move exactly. However, this is not the case in practice. As we show below, the job-filling rate has been even lower than what would be implied by the behavior of the vacancy-to-unemployment ratio.

In the DMP framework, the hiring process is summarized by a Cobb-Douglas matching function. Let inputs to the matching function at time t be the v_t vacancies posted by firms looking to hire and u_t unemployed workers looking for jobs. The Cobb-Douglas matching function can be written as

$$h_t = \Phi v_t^\alpha u_t^{1-\alpha} \quad (18)$$

where h_t is the total hires and $\alpha \in (0, 1)$ is the vacancy share. Φ is the aggregate matching efficiency parameter. It is easy to see that the job-filling rate then follows as

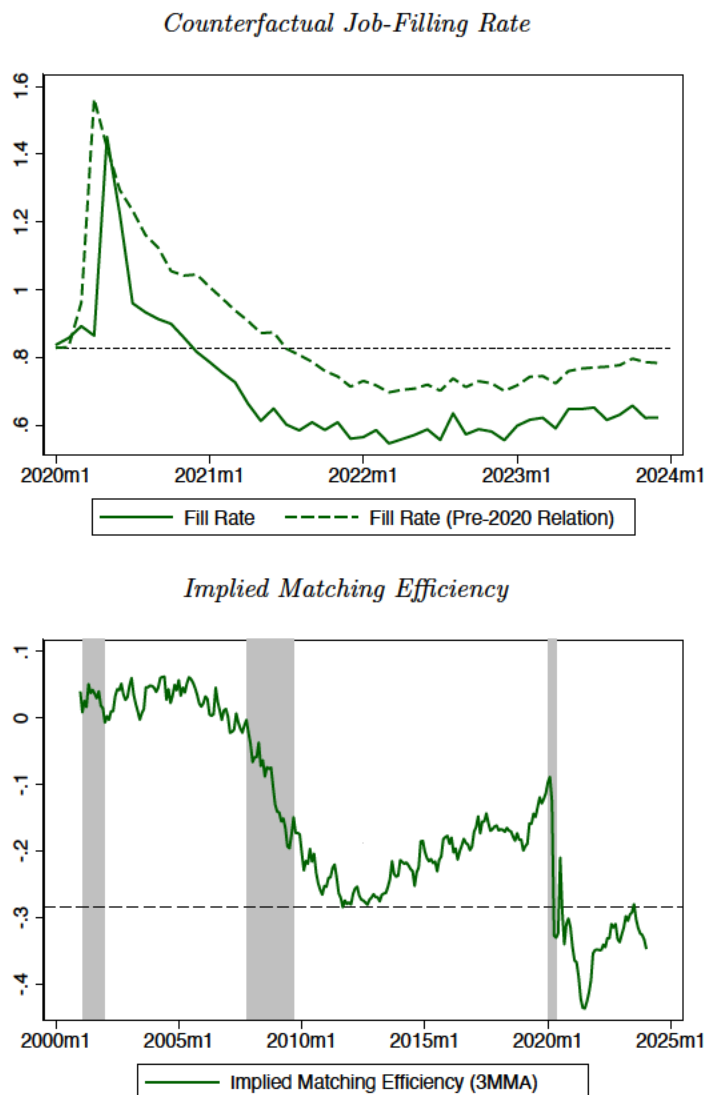
$$\frac{h_t}{v_t} = \Phi \cdot \left(\frac{v_t}{u_t} \right)^{\alpha-1} \quad (19)$$

As the equation shows, if the matching efficiency, Φ , does not vary over time, the job-filling rate would be only a function of the vacancy-to-unemployment ratio. However, mismatch between vacant jobs and unemployed workers, the search effort of workers and recruiting intensity of firms, reservation wages, and job-to-job transitions all affect the matching efficiency in the economy (see, e.g., [Crump et al. 2019](#), [Barlevy et al. 2023](#)). Changes along these dimensions result in a wedge between the job-filling rate and the vacancy-to-unemployment ratio and also cause a shift in the Beveridge curve.

We can use the matching function framework to estimate a log-linear relation between the job-filling rate and v_t/u_t using the pre-COVID sample, 2000–2019. We include year dummies in the regression as an informal way to allow for time variation in matching efficiency. We then use the estimated relation along with the realized v_t/u_t ratio to predict the job-filling rate for the 2020–2023 period. This predicted series is plotted along with the actual job-filling rate in the top chart of [Figure 10](#). The realized job-filling rate is even lower than implied by the behavior of the unemployment-to-vacancy ratio. In fact, the alternative job-filling rate series, implied by the evolution of v_t/u_t , is almost back to its pre-pandemic level (see the horizontal line in the top chart of [Figure 10](#)). An

important reason for this deviation is the steep drop in matching efficiency (which we estimate via $\log(h_t/v_t) - (\hat{\alpha} - 1) \cdot \log(v_t/u_t)$) and plot in the bottom chart of Figure 10. While matching efficiency typically improves as the labor market becomes tighter, in this cycle, the recovery in matching efficiency was substantially muted. The leading explanation for the decline in matching efficiency is the record high quits as argued by Barlevy et al. (2023) and Bagga et al. (2024). As workers started quitting their jobs at a higher rate than usual, they left behind vacancies for firms to fill at a time when willingness to work declined. Consequently, the rate of job filling has decreased, while wage growth has accelerated.

Figure 10. Fill Rate and Matching Efficiency This figure presents results based on estimation of equation (19). The estimate of α is obtained using a log-linear regression of equation (19) with year dummies for the sample ending in 2019. The predicted fill rate shown in the top chart is obtained using $\hat{\alpha}$ and the pre-2020 average log matching efficiency. The log matching efficiency shown in the bottom chart is the three-month moving average of $\log(h_t/v_t) - (\hat{\alpha} - 1) \cdot \log(v_t/u_t)$. The sample period is 2001M1–2023M12. Grey shaded regions denote NBER recessions.



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

The behavior of real wages starting in 2020 has been a topic of intense debate with some economists arguing for a decline in real wages²¹ while some arguing the opposite position.²² A simple approach to calculate real wages is to consider the difference between nominal wage growth and the rate of inflation. This definition is theoretically sound if we consider a single firm which only uses labor to produce. If we know how fast unit labor costs and prices are growing for the firm, the difference between the two tells us how the firm’s marginal cost and its workers’ real labor income is changing over time. While the theoretical definition of real wage is intuitive, in practice, it suffers from composition bias.²³ 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.²⁴ 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 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 fewer 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. One downside of this data source is that the time series only begins after the Great Recession.

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.

²¹See, for example, <https://twitter.com/jasonfurman/status/1468602144844513280>.

²²See, for example, <https://twitter.com/paulkrugman/status/1468552448424005634>.

²³It is well known that changing worker composition makes it harder to assess the cyclical effect of wage growth as well as its inflationary effect ([Perry \(1972\)](#), [Bils \(1985\)](#), [Solon et al. \(1994\)](#), [Daly and Hobijn \(2022\)](#)).


²⁴The ECI, unlike some other measures of wages, has a fixed composition of occupations which partially, but not fully, addresses these issues.

To ensure that we are identifying the same position we match on firm, job title, location, 5-digit occupation code, tax term, and salary type (further details on the data are provided in the Appendix).


We study two-year posted wage growth after the onset of the pandemic. As a comparison we study the two year wage growth ending in the third quarter of 2019 – representing posted wage growth in the tightest labor market, as measured by the raw unemployment rate, since the 1960s. We present results from the third quarter of 2020 to the third quarter of 2022 as well as from the third quarter of 2021 to the third quarter of 2023 to assess whether the moderation of underlying wage growth is also visible in posted wages. 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 indicate the associated 95% confidence bands, all based on the methodology introduced in [Cattaneo et al. \(2024\)](#).

Figure 11. 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. (2024) and Cattaneo et al. (2023). Shaded regions denote 95% confidence bands.

2-Year Posted Wage Growth Before and After Pandemic (Job Match)



2-Year Posted Wage Growth Before and After Pandemic (SOC5-MSA Match)



On average, posted wages for jobs with salaries below \$75,000 grew at a rate of about 18% from 2020 to 2022 and 15% from 2021 to 2023 as compared to about 7% from 2017 to 2019. Notably, Figure 11 shows that over the 2020–2023 period there is a much stronger rise in posted wage growth for job openings which are below the 2019 median salary of \$35,000. Instead, from 2017–2019 posted wage growth was more even across this salary spectrum.²⁵ This observation is consistent with wage

²⁵These results are qualitatively consistent with results available from the Atlanta Fed Wage Tracker for wage growth

compression documented in the CPS by [Autor et al. \(2023\)](#) and also with the ample anecdotal evidence that these positions have become more difficult to fill. Finally, we find that there has been some moderation in wage growth in the 2021–2023 period which coincided with the monetary policy tightening cycle. However, at lower salary positions, posted wage growth remains comfortably above the pre-pandemic level which is consistent with the behavior of underlying wage growth from our model. Moreover, the shape of posted wage growth across salaries from 2021 to 2023 is different than 2020 to 2022, with moderation for lower salaries and an upward shift in wage growth at higher salaries, partially reversing the wage compression observed earlier. As a robustness check, we 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 the bottom chart in [Figure 11](#)).²⁶

5 Conclusion

Using the Micro-Macro Phillips Curve framework of [Crump et al. \(2019\)](#), we document the rapid and persistent increase in the natural rate of unemployment in the aftermath of the pandemic and its implications for inflation dynamics. While the bulk of the inflation surge during that period is attributed to the direct effects of temporary supply factors like supply-chain disruptions, the more persistent component of inflation, driven by wide current and expected future negative unemployment gaps, has increased sharply in the post-pandemic years, a shift unseen since the early 1970s.

We show the disinflation process, started in 2022, is driven by the expectation that the unemployment gap will close through a progressive decline in u_t^* and a rise in the unemployment rate. For this reason the observed disinflation was not accompanied by a notable increase in the unemployment rate, despite our Phillips curve being quite flat. To validate this key model prediction, we show that survey-based forecasts in 2022 and 2023 anticipated a rise in the unemployment rate. Although these data are not used in the estimation, their forecast paths were aligned with the corresponding model-based forecast.

Finally, using a variety of cross-sectional data sources we provide external evidence of unusually tight labor market conditions, consistent with our estimated rise in u_t^* . We use survey data showing an increase in reservation wages, coupled with a lower willingness to work. In addition, we examine

by wage quartile. See [Figure B.4](#) in the Appendix.

²⁶We also show that posted wages align well with BLS wage data by occupation—especially for salaries below \$75,000. See [Figure B.2](#) in the SA.

alternative measures of labor market tightness such as the job-filling rate, and show that it tightly comoves with the model-based measure of expected unemployment gaps. Last, we use wages advertised in newly posted vacancies by the same employers for the same jobs to isolate a composition-free measure of labor costs. The evidence we present is consistent with persistently tight labor market conditions in the aftermath of the pandemic.

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.
- Amiti, M., Heise, S., Karahan, F., Şahin, A., 2023. Inflation strikes back: The role of import competition and the labor market. In: *NBER Macroeconomics Annual*, University of Chicago Press, vol. 38.
- Autor, D., Dube, A., McGrew, A., 2023. The unexpected compression: Competition at work in the low wage labor market. Working Paper 31010, National Bureau of Economic Research.
- Bagga, S., Mann, L., Şahin, A., Violante, G. L., 2024. Job amenity shocks and labor reallocation, mimeo.
- Ball, L., Leigh, D., Mishra, P., 2022. Understanding U.S. inflation during the COVID era. *Brookings Papers on Economic Activity* 53, 1–54.
- Ball, L., Mankiw, N. G., Romer, D., 1988. The new Keynesian economics and the output-inflation trade-off. *Brookings Papers on Economic Activity* 19, 1–82.
- Ball, L., Mazumder, S., 2019. A Phillips curve with anchored expectations and short-term unemployment. *Journal of Money, Credit and Banking* 51, 111–137.
- Baqae, D., Farhi, E., 2022. Supply and demand in disaggregated Keynesian economies with an application to the COVID-19 crisis. *American Economic Review* 112, 1397–1436.
- Barlevy, G., Faberman, R. J., Hobijn, B., Şahin, A., 2023. The shifting reasons for Beveridge-curve shifts. Working Paper 31783, National Bureau of Economic Research.
- Benigno, P., Eggertsson, G. B., 2023. It's baaack: The surge in inflation in the 2020s and the return of the non-linear Phillips curve. Working Paper 31197, National Bureau of Economic Research.
- Bernanke, B., Blanchard, O., 2023. What caused the U.S. pandemic-era inflation? Working Paper 86, Brookings Institution Hutchins Center.
- Bils, M. J., 1985. Real wages over the business cycle: Evidence from panel data. *Journal of Political Economy* 93, 666–689.
- Blanchard, O., Domash, A., Summers, L., 2022. Bad news for the Fed from the Beveridge space. Policy Briefs PB22-7, Peterson Institute for International Economics.
- Boehm, C. E., Pandalai-Nayar, N., 2022. Convex supply curves. *American Economic Review* 112, 3941–69.
- Carvalho, C., Eusepi, S., Moench, E., Preston, B., 2023. Anchored inflation expectations. *American Economic Journal: Macroeconomics* 15, 1–47.
- Cattaneo, M. D., Crump, R. K., Farrell, M. H., Feng, Y., 2023. Binscatter regressions, working paper.
- Cattaneo, M. D., Crump, R. K., Farrell, M. H., Feng, Y., 2024. Nonlinear binscatter methods, working paper.
- Cavallo, A., Kryvtsov, O., 2023. What can stockouts tell us about inflation? evidence from online micro data. *Journal of International Economics* 146, Article 103769.
- Cerrato, A., Gitti, G., 2022. Inflation since COVID: Demand or supply. Graduate Student Bravo Working Paper 2022-005, Brown University.
- Comin, D. A., Johnson, R. C., Jones, C. J., 2023. Supply chain constraints and inflation. Working Paper 31179, National Bureau of Economic Research.
- Crump, R. K., Eusepi, S., Giannoni, M., Şahin, A., 2019. A unified approach to measuring u^* . *Brookings Papers on Economic Activity* 50, 143–214.
- Daly, M. C., 2023. Monetary policy: Progress is not victory. <https://www.frbsf.org/news-and-media/speeches/mary-c-daly/2023/10/monetary-policy-progress-is-not-victory/>, Remarks by Mary C. Daly at the Economic Club of New York, New York, NY.

- 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* 54, 89–111.
- 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.
- Del Negro, M., Lenza, M., Primiceri, G. E., Tambalotti, A., 2020. What’s up with the Phillips curve? *Brookings Papers on Economic Activity* 51, 301–373.
- di Giovanni, J., Kalemli-Özcan, Ş., Silva, A., Yildirim, M. A., 2022. Global supply chain pressures, international trade, and inflation. Working Paper 30240, National Bureau of Economic Research.
- 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? *Labour Economics* 79, Article 102275.
- Ferrante, F., Graves, S., Iacoviello, M., 2023. The inflationary effects of sectoral reallocation. *Journal of Monetary Economics* 140, Supplement, S64–S81.
- Furman, J., Powell, W., 2021. What is the best measure of labor market tightness? <https://www.piie.com/blogs/realtime-economic-issues-watch/what-best-measure-labor-market-tightness>, Petersen Institute for International Economics Blog.
- 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.
- Hall, R. E., 2005. Employment fluctuations with equilibrium wage stickiness. *American Economic Review* 95, 50–65.
- Harding, M., Lindé, J., Trabandt, M., 2023. Understanding post-COVID inflation dynamics. *Journal of Monetary Economics* 140, Supplement, S101–S118.
- Hazell, J., Herreño, J., Nakamura, E., Steinsson, J., 2022. The slope of the Phillips curve: Evidence from U.S. states. *Quarterly Journal of Economics* 137, 1299–1344.
- Hazell, J., Taska, B., 2020. Downward rigidity in the wage for new hires, working paper.
- 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.
- Mavroidis, 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.
- Perry, G. L., 1972. Unemployment flows in the U.S. labor market. *Brookings Papers on Economic Activity* 3, 779–787.
- Pissarides, C. A., 2000. *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, second ed.
- Rubbo, E., 2024. What drives inflation? lessons from disaggregated price data, working paper.

- Şahin, A., 2022. Comment on “Understanding U.S. inflation during the COVID era”. *Brookings Papers on Economic Activity* 53, 65–76.
- Şahin, A., Song, J., Topa, G., Violante, G. L., 2014. Mismatch unemployment. *American Economic Review* 104, 3529–3564.
- Shapiro, A. H., 2022. Decomposing supply and demand driven inflation, forthcoming.
- 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 cyclical behavior of real wages: How important is composition bias. *Quarterly Journal of Economics* 109, 1–25.

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.

Inflation expectations. We obtain a range of inflation expectations from different surveys of professional forecasters. For short-term inflation expectations we combine six-month ahead expectations, averaged across forecasters, from the Livingston survey (available at a semi-annual frequency throughout our sample) and the Survey of Professional Forecasters (SPF, available quarterly since 1981Q3). For long-term inflation expectations we combine five-to-ten year ahead forecasts from Blue Chip Economic Indicators, Blue Chip Financial Forecasts and the SPF. For the years 1975–1977 we also use five-to-ten year ahead inflation expectations from the University of Michigan Consumer Sentiment survey.

Nominal wage growth. For wage growth and inflation expectations we use five measures of labor compensation that vary in their coverage and data sources. Two of these series are growth in wages and salaries for private industry workers and the growth in total compensation for all civilian workers from the Employment Cost Index (ECI) release. One advantage of the ECI is that it provides a wage measure which is free from the influence of employment shifts among occupations and industries. From the Establishment Survey, as part of the Employment Situation release, we use growth in average hourly earnings of all private sector employees and growth in average hourly earnings of production and nonsupervisory employees. Finally, we use growth in compensation per hour of the nonfarm business sector. All growth rates are expressed at a quarterly, annualized rate.

Additional labor market indicators. In Section 4, we introduce empirical evidence from a variety of sources. First, we utilize responses on participants' reservation wage from the Labor Market Module of the Survey of Consumer Expectations (SCE). This module is conducted three times per year in March, July, and November.²⁷ To construct the job-filling rate and the vacancy-to-unemployment ratio we use total nonfarm job openings and hires from the JOLTS survey along with the level of unemployment from the BLS. Finally, we utilize data from Burning Glass Technologies

²⁷Data are available at <https://www.newyorkfed.org/microeconomics/sce/labor#/>.

(BGT) which compiles job openings along with detailed and standardized characteristics of each opening such as a firm identifier or SOC occupation code. 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 11, we use the following data fields: employer, job title, SOC5 code, state and county FIPS code, pay frequency, tax term (e.g., employee, contractor), and job hours (e.g., full-time). For each quarter of the Burning Glass data, we compute the average posted salary by the interaction of those six categories. We then restrict the data to observations in the third quarter of 2017, 2019, 2020, 2021, 2022, and 2023. We only consider matched jobs across time: from 2017 to 2019, from 2020 to 2022, and from 2021 to 2023. To construct the underlying data for the bottom chart of Figure 11 we follow a similar approach as in Figure 11, 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.

Supplemental Appendix:

**“The Unemployment-Inflation Trade-off Revisited:
The Phillips Curve in COVID Times”**

Richard K. Crump
(New York Fed)

Stefano Eusepi
(UT-Austin)

Marc Giannoni
(Barclays)

Ayşegül Şahin
(UT-Austin & NBER)

1 Additional Figures

Figure B.1. 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. The sample period is 2000Q4–2023Q4. Grey shaded regions denote NBER recessions.




Figure B.2. 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 B.3. 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. The sample period is 1973M10–2024M1. Grey shaded regions denote NBER recessions.




Figure B.4. 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. The sample period is 1997M12–2023M12. Grey shaded regions denote NBER recessions and vertical dotted lines represent December 2017, 2019, and 2021, respectively.

