RICHARD K. CRUMP
Federal Reserve Bank of New York

MARC GIANNONI
Federal Reserve Bank of Dallas

STEFANO EUSEPI
University of Texas at Austin

AYŞEGÜL ŞAHIN
University of Texas at Austin

A Unified Approach to Measuring u*

ABSTRACT This paper bridges the gap between two popular approaches to estimating the natural rate of unemployment, u_t^*. The first approach uses detailed labor market indicators, such as labor market flows, cross-sectional data on unemployment and vacancies, and various measures of demographic changes. The second approach, which comprises reduced-form models and dynamic stochastic general equilibrium models, relies on aggregate price and wage Phillips curve relationships. We combine the key features of these two approaches to estimate the natural rate of unemployment in the United States, using both data on labor market flows and a forward-looking Phillips curve linking inflation to current and expected deviations of unemployment from its unobserved natural rate. We estimate that the natural rate of unemployment was about 4.0 percent toward the end of 2018 and that the unemployment gap was roughly closed. Identification of a secular downward trend in the unemployment rate, driven solely by the inflow rate, facilitates the estimation of u_t^*. We identify the increase in labor force attachment of females, the decline in job destruction and reallocation intensity, and the dual aging of workers and firms as the main drivers of the secular downward trend in the inflow rate.

Conflict of Interest Disclosure: Richard Crump is vice president and function head of Capital Markets at the Federal Reserve Bank of New York; Stefano Eusepi is associate professor of economics at the University of Texas at Austin; Marc Giannoni is senior vice president and director of research for the Federal Reserve Bank of Dallas; and Ayşegül Şahin is the Richard J. Gonzalez Regents Chair in Economics at the University of Texas at Austin and an adviser to the Federal Reserve Bank of Dallas. Beyond these affiliations, the authors did not receive financial support from any person for this paper or from any firm or person with a financial or political interest in this paper. They are currently not officers, directors, or board members of any organization with an interest in this paper. No outside party had the right to review this paper before circulation except the Federal Reserve Bank of New York and the Federal Reserve Bank of Dallas. The views expressed in this paper are those of the authors, and do not necessarily reflect those of the Federal Reserve Bank of New York, the University of Texas at Austin, or the Federal Reserve Bank of Dallas.

The unemployment rate in the United States peaked at 10.2 percent in October 2009. Since then, it has declined gradually, reaching below 4 percent for the first time in almost 20 years. A debate has arisen about how sustainable these low levels are and how monetary policy should respond. Starting with Milton Friedman (1968) and Edmund Phelps (1967, 1968), both academics and policymakers have endeavored to measure a sustainable level of unemployment and what implications deviations from this level have for price and wage inflation. This natural rate of unemployment, u_t^*, is broadly defined as the unemployment rate such that, controlling for supply shocks, inflation remains stable. The measure u_t^* is thought to vary over time with changes in the economy, such as demographic shifts, changes in the structure of the labor market, and technological advances.[1]

There are two popular approaches to estimating u_t^* in the literature. The first approach uses detailed labor market data, such as changes in demographics (Perry 1970; Summers 1986; Shimer 1998; Brauer 2007; Barnichon and Mesters 2018), labor market flows and job vacancies (Blanchard and Diamond 1989; Daly and others 2012), firms' recruiting intensity (Davis, Faberman, and Haltiwanger 2013) and skills mismatch (Şahin and others 2014). One potential limitation of this approach is the absence of information from inflation to infer u_t^*; moreover, these measures are not additive, because they cannot be considered as independent from each other and thus they are not conclusive as to the *level* of the natural rate of unemployment. Finally, there is the need for detailed data sets (for example, to build mismatch indexes) that are available only for the more recent period.

The second approach—which comprises reduced-form models (Staiger, Stock, and Watson 1997; Laubach 2001; Orphanides and Williams 2002) and dynamic stochastic general equilibrium (DSGE) models (Galí 2011; Gertler, Sala, and Trigari 2008; Galí, Smets, and Wouters 2012)—relies mainly on price and wage Phillips curve relationships, together with model-specific assumptions on aggregate demand. This approach, in

1. For example, in the face of a spike in unemployment during the Great Recession, the modest decline in inflation was, in part, attributed to increases in mismatch unemployment, decline in firms' recruiting intensity, the extension of unemployment benefits, and uncertainty in economic conditions. This time variation in u_t^* is reflected in the time series of forecasted longer-run unemployment observed in survey data and projections by the Federal Open Market Committee; see figure C.1 in the online supplemental appendix. The online appendixes for this and all other papers in this volume may be found at the *Brookings Papers* web page, www.brookings.edu/bpea, under "Past BPEA Editions."

contrast, makes little use of detailed labor market information, and has been subjected to two sets of criticism. First, the natural rate estimates obtained from these models tend to be surrounded by a considerable degree of uncertainty, hampering their use for policy decisions. Second, the relationship between "economic slack" and inflation has been called into question since the financial crisis of 2007–8, because the strong rise in unemployment did not lead to a sizable and persistent decline in inflation.

We combine the key features of these two approaches and estimate u_t^* using a forward-looking Phillips curve linking inflation to current and expected deviations of unemployment from its *unobserved* natural rate. This estimation relies on two key pieces of information. First, we propose a measure of the secular trend in the unemployment rate 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 an estimate of the trends. Our analysis of unemployment flows identifies the downward trend in the inflow rate as the main driver of the secular unemployment trend. The identification of such trends aids the measurement of the unobserved natural rate of unemployment. Second, we use survey-based professional forecasts to measure the term structure of inflation expectations, that is, the forward-looking component of the Phillips curve. We find that it is vital to account for the behavior of expectations to reconcile the observed behavior of inflation and slack over time consistent with the research of Marco del Negro, Marc Giannoni, and Frank Schorfheide (2015) and of Carlos Carvalho and others (2017).

We estimate the natural rate of unemployment for the United States over the period 1960–2018. As of the third quarter of 2018, we estimate that u_t^* was about 4 percent; in particular, using only information from price inflation, we estimate that u_t^* stood at 4.0 percent, with a 68 percent confidence interval of 3.5 to 4.5 percent. When we add information from wage inflation, the estimate shifts down slightly, to 3.8 percent, with an associated confidence interval of 3.5 to 4.2 percent. We find that the unemployment gap was roughly closed by the end of 2018, as short-term inflation expectations approximately converged toward their long-run mean. More generally, we find that the natural rate of unemployment, estimated using both price and wage inflation, was steady, at just below 6 percent, in the 1960s; rose sharply in the 1970s, to over 8 percent; and then fell steadily, to below 5 percent, in 2000. During the 2000s up until the Great Recession, the natural rate of unemployment was range-bound. In the Great Recession, we document a rise in u_t^* of about 1 percentage point relative to its prerecession levels. We demonstrate that this estimate aligns well with estimated contributions to the unemployment rate attributed to mismatch unemployment and changes in recruiting intensity.

We trace the long-term decline in u_t^* over the last 40 years to a secular downward trend in the rate at which workers become unemployed—the inflow rate. The decline in the inflow rate reflects three important changes in the labor market: (1) the rise in participation and labor force attachment of females, which coincided with fewer labor force interruptions related to maternity and childbirth and culminating in the closing of the gender unemployment gap; (2) the shift of the labor force from younger workers, who frequently become unemployed, to older workers, who are less likely to become unemployed; (3) the aging of firms, as older firms tend to have reduced rates of job destruction (layoffs and firings). The second and third changes are connected, and we refer to them as the *dual aging* of the U.S. economy, which has resulted in less job destruction and unemployment incidence in the labor market, not only through a composition effect but also by reducing unemployment incidence (job destruction) for workers (firms) in all age groups. Dual aging stands out as an important driver of the lower trend rate of unemployment, especially in the last two decades. Together, these secular changes have reduced the overall flow rate into unemployment and, consequently, the unemployment rate itself.

The structure of the paper is as follows. Section I presents an overview of the paper and discusses its contributions relative to the extensive literature on the natural rate of unemployment. Section II estimates the secular trend in unemployment, using detailed information for unemployment inflows and outflows by demographic group. Section III introduces a simple forward-looking Phillips curve, discusses its theoretical underpinning, and details the estimation methodology. Section IV presents the time series for the natural rate of unemployment, u_t^*, for the sample 1960–2018. Section V provides a quantitative evaluation of three factors driving the trend decline in the unemployment inflow rate: the increase in female labor force attachment; the decline in job destruction and reallocation; and the dual aging of workers and firms in the economy. Section VI concludes.

I. Overview and Relation to the Literature

The object we seek to estimate is "the natural rate of unemployment," u_t^*, which is defined as the unemployment rate such that, controlling for supply shocks, inflation remains stable.[2] Although the relation between inflation and unemployment is a perennial topic in macroeconomics (Humphrey 1991), the concept of the natural rate is often attributed to Friedman (1968) and Phelps (1967, 1968), and the notation u_t^* can be traced back to Phelps. As originally suggested by Friedman, u_t^* is generally assumed to vary over time, possibly as a function of demographic shifts, changes in the structure of the labor market, or technological advances. Friedman, in his 1968 presidential address to the American Economic Association, wrote:

> 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. On the contrary, many of the market characteristics that determine its level are man-made and policy-made. . . . 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. (Friedman 1968, 9)

Friedman clearly pointed out changes in labor supply behavior and in the efficiency of the matching process in the labor market arising from better matching technology as shifters of the natural rate. However, despite this key insight, an ongoing assumption of the time was that the natural rate was about 4 percent, which caused policymakers to underestimate how tight the labor market was. Various influential papers in the inaugural volumes of *Brookings Papers on Economic Activity* in the early 1970s studied the rise in the *natural* rate of unemployment, such as those by Robert Hall (1970a, 1970b), Robert Gordon (1970a, 1970b), George Perry (1970, 1972), and Charles Schultze (1971). These papers emphasized the role of the changing demographic structure of the economy and the importance of labor market flows in assessing the natural rate in real time. Here, we expand on these enduring insights and estimate the secular trend in unemployment and integrate it into the New Keynesian Phillips curve.

2. The extensive literature on the natural rate of unemployment used long run, frictional, average, equilibrium, normal, steady state, lowest sustainable, Hodrick–Prescott trend, nonaccelerating inflation rate of unemployment (NAIRU) and the unemployment at full employment to refer to related, perhaps the same, object that we are trying to estimate. An insightful article by Richard Rogerson (1997), titled "Theory Ahead of Language in the Economics of Unemployment," discusses the confusion and uncertainty around the language used.

Our point of departure is a simple decomposition of the unemployment rate,

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

where $\bar{u}_t$ is the *secular trend in unemployment* and u_t^* is the *natural rate*. The secular trend in unemployment, $\bar{u}_t$, captures the elements of the unemployment rate that are driven by slow-moving factors such as demographics and social change. The unemployment gap, x_t, measures the deviation of the observed unemployment rate from the natural rate and is the primary input to monetary policy considerations (for example, the goal of maximum employment). The natural rate of unemployment is defined as the sum of the secular trend component and a cyclical component z_t. Conceptually, we would expect the natural rate of unemployment to converge to $\bar{u}_t$ over time in the absence of shocks.

Although it is tempting to use traditional filtering techniques to eliminate the higher-frequency fluctuations in the unemployment rate, we instead rely on rich cross-sectional variation in unemployment flow rates by demographic groups to assess the extent of the secular trend in unemployment, $\bar{u}_t$. We do so for three main reasons: (1) the inherent asymmetry in the unemployment rate (Montgomery and others 1998; Hamilton 2005) makes it challenging to directly estimate its secular, slow moving trend; (2) the inflow/outflow dynamics of the unemployment rate—which is the source of the underlying asymmetry—by itself provides a better characterization of the evolution of the unemployment rate (Blanchard and Diamond 1990; Barnichon and Nekarda 2012; Şahin and Patterson 2012); and (3) extensive cross-sectional information on these flow rates enables us to better distinguish and analyze the underlying common and group-specific trends.

In estimating the secular trend in unemployment, we allow trends in unemployment inflows and outflows to vary by age and gender. This follows a long-standing body of literature dating back to George Perry's influential *Brookings Paper* in 1970, which recognized age and gender as the main demographic characteristics that need to be taken into account in assessing the natural rate of unemployment. In particular, Perry suggested an adjustment to account for the rising share of teenagers and females in the labor force that is often referred to as the *Perry-adjusted* unemployment rate. This adjustment—which assigns a lower weight to the unemployment rate of demographic groups with lower hours and wages—has been used in the literature in estimations of the Phillips curve such as those made by

Gordon (1982) and Lawrence Summers (1986); and it provides a basis for different measures of labor market underutilization, such as U-1 and U-6 (for definitions of these measures, see BLS 2018). Robert Shimer (1998) built on the research by Perry (1970), Gordon (1982), and Summers (1986), and provided a critical evaluation of the underlying assumption of applying demographic adjustments to the unemployment rate: demographic shifts in the labor market only affect the aggregate unemployment rate through the changing labor force shares without affecting group-specific unemployment rates. Shimer (1998) argued that this assumption is adequate with respect to changes in the age structure but is violated when there are changes in educational attainment. More recently, Regis Barnichon and Geert Mesters (2018) revisited the demographic adjustment of the unemployment rate and proposed a new demographic adjustment based on gross flows data. We build on the research of Barnichon and Mesters (2018) and examine the relationship between demographics and unemployment flows instead of focusing directly on the unemployment rate.

To connect inflation to the state of the labor market, we employ a forward-looking Phillips curve linking inflation to expected inflation and the unemployment gap. Following Friedman (1968) and Phelps (1967, 1968), and building on the rational expectations school of thought in the 1970s (Sargent 1971; Lucas 1972), it has become common to link the gap between the unemployment rate and a natural rate of unemployment to the inflation rate, through an expectations-augmented Phillips curve. According to this relationship, whenever the unemployment rate is equal to its natural rate, inflation and inflation expectations should settle to their long-run value in the absence of supply shocks. For this reason, the natural rate of unemployment is sometimes called the nonaccelerating inflation rate of unemployment (NAIRU).[3] Moreover, for given unemployment, inflation, and an assumption about inflation expectations, this relation allows for the estimation of u_t^*. We utilize survey-based expectations of inflation at different horizons to provide noisy signals of true inflation expectations and impose that the secular trend act as an anchor for the natural rate, although accommodating the possibility of persistent deviations.

A Phillips curve by itself is, however, not a panacea to estimate the natural rate of unemployment. Indeed, as many authors have emphasized,

3. Modigliani and Papademos (1975) defined the noninflationary rate of unemployment, which they referred to as NIRU, as a rate such that, as long as the unemployment rate is above it, inflation can be expected to decline, and they estimated it to be somewhat over 5.5 percent in 1975.

the estimates of the response of inflation to the unemployment gap in conventional backward-looking Phillips curves—that relate current inflation to a measure of economic slack and lags of inflation to proxy for inflation expectations—appear to have diminished substantially since the late 1980s (Hall 2011; Ball and Mazumder 2011). This raises two issues. First, an instability in key parameters of the Phillips curve renders the estimating of u_t^* more difficult. Second, relatively flat Phillips curves may result in uncertain estimates of u_t^*.

Several researchers (Ball and Mazumder 2011; Hall 2011; Blanchard 2016) have also questioned the Phillips curve relationships on the grounds that the dramatic increase in the unemployment rate and the collapse in economic activity recorded during the Great Recession should have implied a very large drop in the inflation rate, or even deflation, in contrast to the relatively modest decline in inflation registered in the aftermath of the Great Recession. However, recent research has shown that while the criticism of backward-looking Phillips curves is well justified, it does not apply to forward-looking Phillips curves linking inflation to the unemployment gap and expectations of *future* inflation.[4] For example, del Negro, Giannoni, and Schorfheide (2015) show that a relatively standard monetary DSGE model with forward-looking expectations and financial frictions can account remarkably well for the joint evolution of inflation and economic activity during and after the Great Recession. This is because forward-looking agents in the model understand that monetary policy will be more accommodative in the future the more activity contracts, thereby helping to anchor inflation expectations. Carvalho and others (2017) provide further evidence that it is possible to reconcile the observed behavior of inflation with the level of slack during the crisis and its aftermath. In particular, they show that inflation expectations have remained "anchored" over that period, which has contributed to promoting price stability.

We build on this insight and use a forward-looking Phillips curve linking inflation to the unemployment gap and expectations of future inflation, which is based on the model of Jordi Galí (2011). Importantly, the forward-looking nature of this Phillips curve implies that inflation depends not only on the contemporaneous unemployment gap but also on the entire path of

4. Early derivations of the New Keynesian Phillips curve that features expectations of future inflation were calculated by Calvo (1983), Rotemberg (1987), Roberts (1995), Fuhrer and Moore (1995), Yun (1996), Goodfriend and King (1997), Rotemberg and Woodford (1997), Galí and Gertler (1999), and Sbordone (2002).

expected future unemployment gaps. This tighter link, along with data on inflation expectations at various horizons, coupled with the secular trend in the unemployment rate, helps us identify u_t^*.

Finally, our examination of the determinants of the secular trend in the unemployment rate links to the recent but growing literature on the decline in labor market dynamism. Worker and job reallocation have declined substantially in the recent decades, as initially documented by Steven Davis and others (2006) and recently analyzed by Davis and John Haltiwanger (2014) and by Raven Molloy and others (2016). Our analysis links the decline in labor market fluidity to movements in the inflow rate and assesses its quantitative impact on the natural rate of unemployment.

II. The Secular Trend in Unemployment

We estimate u_t^* in two steps. In the first step, described in this section, we extract the slow-moving trends in the inflow and outflow rates using a linear state-space model and obtain the unemployment rate trend $\bar{u}_t$. 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 from a New Keynesian Phillips curve.

It could be argued that, from a statistical standpoint, it is more efficient to jointly estimate the unemployment trend and the natural rate. Our choice reflects two main considerations. First, the current approach is simple to implement and is transparent. Single-step estimation would add significant complexity because it would require conducting inference with a nonlinear state-space model of a reasonably large dimension (see equation 3 below). Second, as we argue in section V, the evolution of the unemployment secular trend is driven by forces such as changes in labor supply behavior reflecting social change or slow-moving demographic changes. This is broadly consistent with the assumption that $\bar{u}_t$ evolves exogenously to the state of the business cycle or changes in monetary and fiscal regimes during our sample period.

Subsection II.A introduces and summarizes the flow origins of the unemployment rate. Subsection II.B characterizes overall flows into and out of unemployment, whereas subsection II.C focuses on these flows for specific demographic subgroups. Finally, in subsection II.D, we introduce a state-space model to estimate the slow-moving components of the inflows and outflows to unemployment that maps directly to the slow-moving component of the unemployment rate.

II.A. The Flow Dynamics of the Unemployment Rate

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 $t + 1$:

$$dU/dt = s_t(L_t - U_t) - f_t U_t, \tag{2}$$

where L_t denotes the labor force, s_t is the inflow rate (separation rate) to unemployment, and f_t is the outflow rate (job-finding rate) from unemployment. Although s_t is generally referred to as the separation rate and f_t as the job-finding rate, we use the inflow/outflow terminology used by Michael Elsby, Ryan Michaels, and Gary Solon (2009) and by Elsby, Bart Hobijn, and Ayşegül 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 Current Population Survey (CPS) micro data.

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

$$U_{t+1} - U_t = U^S_{t+1} - F_t U_t,$$

where U^S_{t+1} is the number of unemployed who report having been unemployed for less than five weeks. Solving for F_t,

$$F_t = 1 - \frac{U_{t+1} - U^S_{t+1}}{U_t},$$

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 five weeks were not in the unemployed pool in the previous month, and therefore subtracting them out from this month's unemployment pool leaves us with those unemployed persons who failed to exit unemployment between month t and month $t + 1$. Solving the differential equation 2 forward, as done by 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 United States—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

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

It 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 the destination of unemployment outflows, following Shimer (2005, 2012); Hall (2005); Elsby, Michaels, and Solon (2009); Elsby, Hobijn, and Şahin (2010); Davis and others (2010); and Şahin and others (2014). The inflow and outflow rates we use are estimated from CPS time series, rather than the longitudinally matched monthly CPS micro data. This abstraction simplifies the framework and better connects to the literature on unemployment dynamics. Although we maintain the two-state abstraction throughout section IV, we explicitly consider the role of the participation margin for females when we examine the drivers of the trends in unemployment flows in section V.

II.B. Unemployment Inflows and Outflows

The CPS provides us with monthly measures of the stock of unemployment, short-term unemployment, and the labor force. We calculate monthly unemployment inflow and outflow hazard rates using the methodology described above and plot quarterly averages of monthly s_t and f_t for the 1976:Q1–2018:Q4 period, as shown in figure 1.[5] Visual examination of inflow and outflow rates confirms the findings of the earlier literature regarding the cyclical properties of these flows. The inflow rate

5. Online supplemental appendix B provides details on the data sources used in this paper.

Figure 1. Inflow and Outflow Rates, 1976–2018[a]

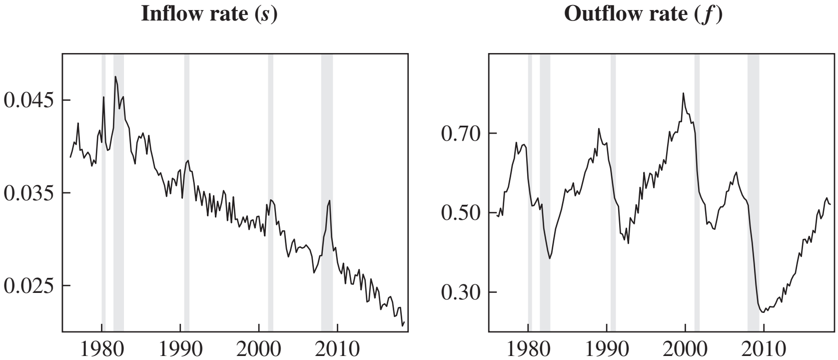

Source: Current Population Survey.

a. This figure shows the unemployment inflow rate (left panel) and outflow rate (right panel) for the sample 1976:Q1–2018:Q4. Inflow and outflow rates are weighted averages of flow rates using counts of employed and not in the labor force from the six demographic categories described in subsection II.D.

is characterized by sharp, short-lived spikes during recessions.[6] The outflow rate from unemployment is strongly procyclical, with persistent downswings during recessions.

Figure 1 also reveals that these two flows that shape the evolution of the unemployment rate over time exhibit differential long-run trends. The inflow rate has a striking downward trend, declining gradually to 0.02, half its level preceding the twin recessions of the early 1980s. In contrast, there is less evident trending behavior in the outflow rate.

Although it is tempting to use traditional filtering techniques to filter out the trends in the inflow and outflow rates, it is well known that the presence of a severe downturn—such as the Great Recession—at the end of the sample is likely to affect the estimate of the underlying trend. Instead, we rely on rich cross-sectional variation in the flow rates to assess the extent of the trends. In addition, cross-sectional information allows us to analyze the underlying drivers of the trends in the flow rates.

6. As emphasized by Shimer (2005, 2012) and Hall (2005), the response of the inflow rate was relatively muted during the mild recessions in 1990–91 and 2001. However, the inflow rate, without exception, exhibited sharp increases during severe recessions, including the most recent 2007–9 recession, as emphasized by Elsby, Hobijn, and Şahin (2010) and by many others.

Figure 2. Labor Force Shares, 1976–2018[a]

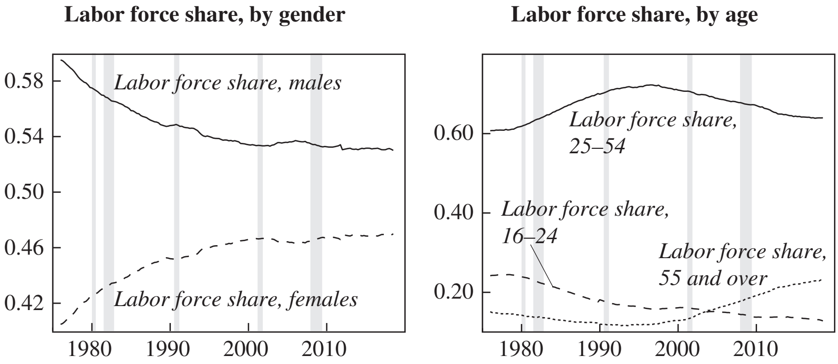

Source: Current Population Survey.
a. This figure shows the labor force shares by gender (left panel) and age (right panel) for the sample 1976:Q1–2018:Q4.

II.C. The Demographics of Unemployment Inflows and Outflows

We start with a visual examination of the flows by gender and age before we move on to our state-space setting to estimate the secular trends. Figure 2 summarizes the changes in the gender and age composition of the labor force from 1976 to 2018. With the rise in the female labor force participation rate, the labor force share of females increased to about 46 percent by 2000 and has stabilized. The age composition shifted from younger workers to prime age workers as the baby boom cohort entered the labor force and gradually aged, with the share of prime age workers peaking in the late 1990s. Since 2000, workers age 55 and older have constituted an increasing share of the labor force—and this has come about not just because of the aging of the population but also because of the differential trend in the participation rates of younger and older workers.

The left panels of figure 3 reveal the drastic convergence of males and females, both in terms of their unemployment inflow and outflow rates. A clear implication of this pattern is the disappearance of the gender unemployment gap, as discussed by Shimer (1998) and by Stefania Albanesi and Şahin (2018). Although the inflow rate has a downward trend for both male and female workers, the downtrend is more pronounced for female workers in the earlier part of the sample. The right panels of figure 3 show the importance of age composition. Workers younger than 25 years experienced an unemployment inflow hazard that was about five

Figure 3. Inflow and Outflow Rates by Gender and Age, 1976–2018[a]

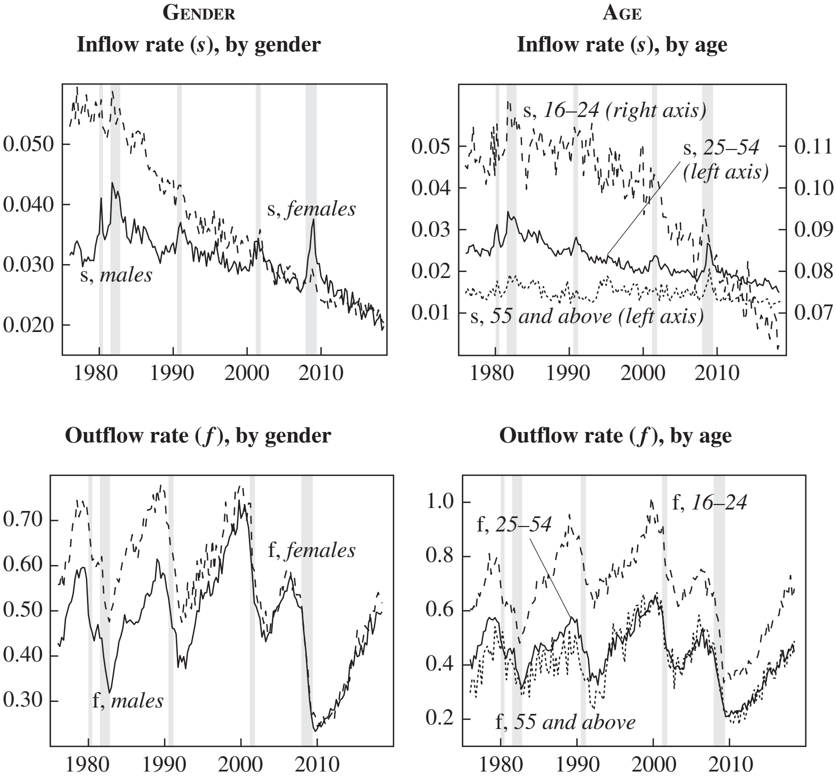

Source: Current Population Survey.

a. This figure shows inflow rates (top row) and outflow rates (bottom row) by gender (left panels) and by age (right panels).

times that of those of prime age workers in the early 1980s. The recent decades show a partial convergence in their inflow rate as well as a decline in their shares.

II.D. Secular Trends in Unemployment Flows

In this subsection, we estimate the slow-moving trend in the inflow and outflow rates. We build on the research by Barnichon and Mesters (2018) and cast these trends in job market flows as latent processes in a linear state-space setting (Tasci 2014; Hornstein and Kudlyak 2019). Each flow is described by this set of equations:

$$s_t^i = \theta_i^s \varphi_t^s + \tau_t^{s,i} + \epsilon_t^{s,i} \text{ and} \tag{4}$$

$$f_t^i = \theta_i^f \varphi_t^f + \tau_t^{f,i} + \epsilon_t^{f,i}, \tag{5}$$

for $i = 1, \ldots 6$, our six demographic subgroups, and we normalize one element of θ_i^s and θ_i^f to 1. The individual flow rates are mapped into the aggregate rates s_t and f_t, using a choice of weights giving

$$s_t = \sum_{i=1}^{6} \omega_{it}^s s_t^i \quad \text{and} \quad f_t = \sum_{i=1}^{6} \omega_{it}^f f_t^i. \tag{6}$$

We use as weights for the inflow rate each group's share in the employment and out-of-the-labor-force pool because flows into unemployment originate from these two stocks.[7] For the outflow rate, we use each group's share in the unemployment pool. The trends in the inflow and outflow rate evolve according to

$$\tau_t^{s,i} = g_t^{s,i} + \tau_{t-1}^{s,i},\ g_t^{s,i} = g_{t-1}^{s,i} + \eta_t^{s,i} \text{ and} \tag{7}$$

$$\tau_t^{f,i} = \tau_{t-1}^{f,i} + \eta_t^{f,i}. \tag{8}$$

We assume that the slow-moving trend for the inflow rate, $\tau_t^{s,i}$, is characterized by an integration of order-two, $I(2)$, process, to accommodate the apparent secular trend in these flows. The trend for the outflow rate, $\tau_t^{f,i}$, is instead a random walk. The common component, $\varphi_t \equiv (\varphi_t^s, \varphi_t^f)'$, follows a second-order vector autoregressive process:

$$\varphi_t^f = \phi_{11}\varphi_{t-1}^f + \phi_{12}\varphi_{t-1}^s + \phi_{13}\varphi_{t-2}^f + \phi_{14}\varphi_{t-2}^s + \zeta_t^f \text{ and} \tag{9}$$

$$\varphi_t^s = \phi_{21}\varphi_{t-1}^f + \phi_{22}\varphi_{t-1}^s + \phi_{23}\varphi_{t-2}^f + \phi_{24}\varphi_{t-2}^s + \zeta_t^s. \tag{10}$$

The common cyclical component, φ_t, accommodates joint business cycle behavior between the inflow rate and the outflow rate for each group. This allows us to capture the specific lead-lag relationship around business cycle turning points (Fujita and Ramey 2009; Elsby, Hobijn, and Şahin 2013).

7. Recall from the text above that s is the solution to a nonlinear equation and that unlike f, it is not linear in group-specific weights. In unreported results, we verify that using employment shares and a more complex weighting method that corrects for time aggregation does not alter our main findings. However, small differences can arise between the overall series computed using aggregate data and the series constructed as a weighted average.

Finally, the flow-specific components follow a first-order autoregressive process

$$\epsilon_t^{s,i} = \psi_i^s \epsilon_{t-1}^{s,i} + \varepsilon_t^{s,i} \text{ and} \tag{11}$$

$$\epsilon_t^{f,i} = \psi_i^f \epsilon_{t-1}^{f,i} + \varepsilon_t^{f,i}, \tag{12}$$

representing idiosyncratic, possibly persistent, movements in the individual flow rates. The innovations $(\eta_t^{s,1}, \ldots, \eta_t^{s,6}, \eta_t^{f,1}, \ldots, \eta_t^{f,6}, \zeta_t^s, \zeta_t^f, \varepsilon_t^{s,1}, \ldots, \varepsilon_t^{s,6}, \varepsilon_t^{f,1}, \ldots, \varepsilon_t^{f,6})'$ are mutually independent, Gaussian random variables. The initial conditions $(\varphi_0, \varphi_{-1}, \epsilon_0^s, \epsilon_0^f)'$ are normally distributed with mean zero and unconditional variance implied by equations 9 through 12. Equations 4 through 6 represent the observation equations in the state-space model, and equations 7 through 12 are the transition equations. The model is estimated using Bayesian methods, utilizing the Gibbs sampler approach proposed by del Negro and others (2017).[8] (Also see Carter and Kohn 1994; Kim and Nelson 1999.) We estimate the model using quarterly data on labor market flows for the sample 1960:Q1–2018:Q3. We focus on six demographic subgroups: the interaction of three age groups—16–24 years, 25–54 years, and 55 years and over—with gender. Because individual flow rates are available only starting in 1976, we use aggregate flows for the remaining sample period, together with the weights, ω_{it}^s and ω_{it}^f, in order to estimate the unobserved trends for the entire sample.

The priors for the coefficients and variances of the vector autoregression with two lags, VAR(2), common components have the standard form

$$p\left(\text{vec}(\Phi)|\Sigma_\zeta\right) = \mathcal{N}\left(0, \Sigma_\zeta \otimes \Omega\right) \quad \text{and} \quad p\left(\Sigma_\zeta\right) = \mathcal{IW}\left(\kappa_\zeta, \Psi_\zeta\right), \tag{13}$$

where Φ is the autoregressive matrix corresponding to equations 9 and 10. The priors for the variance-covariance term of the innovations, Σ_ζ, is a fairly diffuse inverse Wishart with just enough degrees of freedom ($\kappa_\zeta = 4$) to have a well-defined prior mean for the innovations to the VAR. For simplicity of exposition, in this subsection only, we work with the flow rates multiplied by 100. The choices of priors in equation 13 imply standard deviations of 0.45 and 1.4 for the innovations to φ_t^s and φ_t^f, and also imply the mutual independence of these innovations.

8. We thank the authors for helpful discussions and for sharing their code.

The prior for Φ is a standard Minnesota prior with the hyperparameter for the overall tightness $\lambda = 2$ (which regulates the matrix Ω); this parameterization reflects a relatively loose prior. We implement similar priors for the two first-order autoregressive processes described by equations 11 and 12. The prior on the innovations corresponds to an inverse-gamma with shape and scale parameters implying a diffuse prior consistent with a well-defined mean.[9] The prior mean delivers a standard deviation of 0.45 for innovations to ε_t^s and 1 for innovations to ε_t^f. The prior on the autocorrelation coefficient is normally distributed with zero mean and variance determined using the same parametrization, with $\lambda = 2$ as in the VAR Minnesota prior. Also, in this case the prior is fairly diffuse. Finally, the priors on the innovations in the trend variables, $\tau_t^{s,i}$ and $\tau_t^{f,i}$, have inverse-gamma distributions. The parameters are chosen to guarantee a well-defined mean and deliver a standard deviation at the mean prior of 0.1 for the innovation to the trend in the outflow rate and of 0.01 to the innovation of the inflow rate.[10] The priors on the loadings θ^s and θ^f in equations 4 and 5 are defined as independent normal densities with mean 1 and standard deviation 0.5. Finally, we should note that in order to assess the role of the choice of priors, we have reestimated the model with uninformative priors and found broadly similar results.

Figure 4 shows the six inflow rates and their estimated secular trend with associated confidence intervals, for the part of the sample where such flows are observable.[11] These trends differ substantially by gender and age group. Females age 16–24 and 25–54 show a pronounced downward trend starting in the 1980s, about halving their level from early in the sample. The trend for males age 16–24 displays a clear hump-shaped pattern, peaking in the first half of the 1990s and then falling by about 30 percent by the end of the sample. The remaining three groups, prime age males and females and males older than 55, demonstrate a milder secular trend. However, prime age males experience a notable decline in the inflow rate over the last decade or so. In contrast to the inflow rates, figure 5 shows that outflow rates have fairly stable trends, with the exception of females age 25–54. For this latter group, the outflow rate has fallen since the early 1990s. All other groups show little evidence of a secular trend over our sample.

9. The distribution has a shape parameter of 1.5 and scale parameters of 0.1 and 1 for ϵ_t^s and ϵ_t^f, respectively.

10. The gamma parameters correspond to a shape factor of 1.5 and a scale factor of 0.01 for the outflow rate and of $(0.01)^2$ for the inflow rate.

11. Individual flow rates going back to 1960 are estimated with a considerable degree of uncertainty. For this reason, we report below only the aggregate trend for the whole sample.

Figure 4. Inflow Rates by Gender and Age Subgroups, 1976–2018[a]

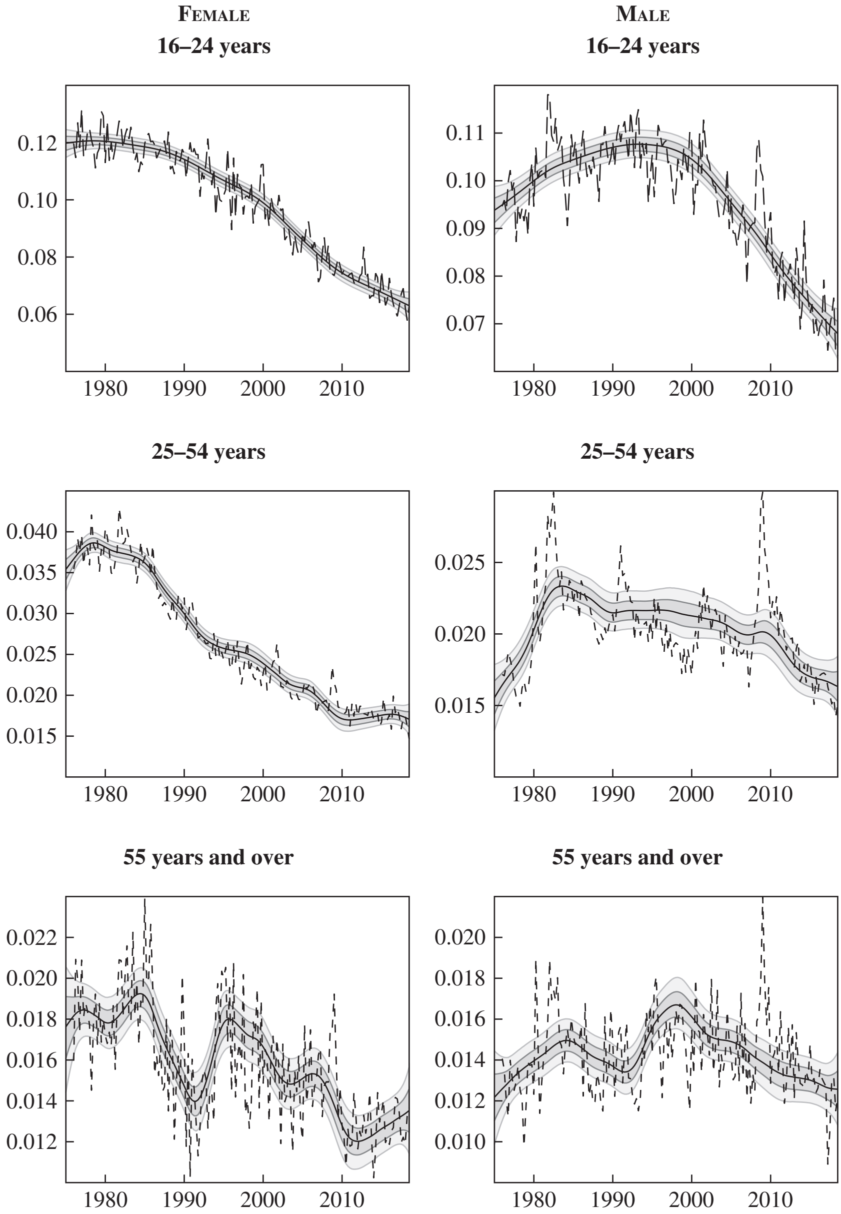

Sources: Current Population Survey; authors' calculations.

a. This figure shows observed inflow rates (dashed line) along with median estimates of the secular trend ($\tau_t^{s,i}$, solid line) for the inflow rate for each age and gender subgroup. Shading denotes 68 percent and 95 percent confidence intervals.

Figure 5. Outflow Rates by Gender and Age Subgroups, 1976–2018[a]

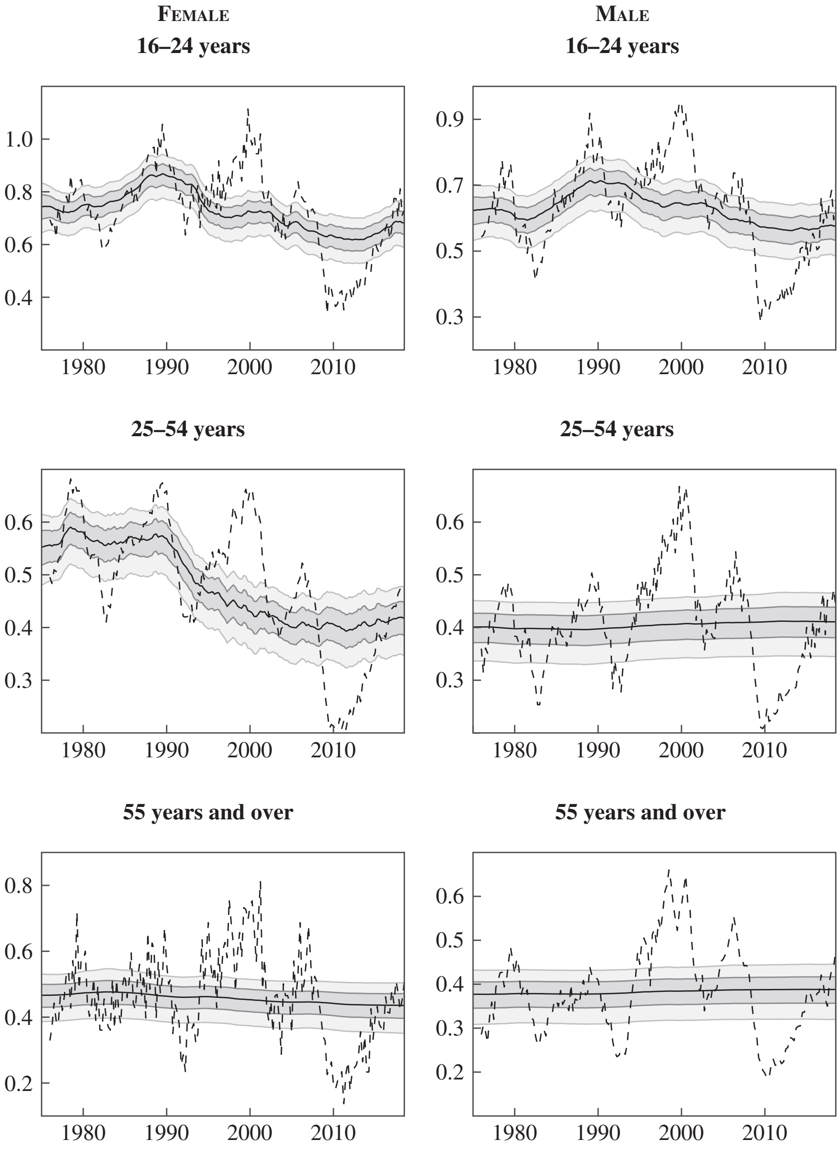

Sources: Current Population Survey; authors' calculations.

a. This figure shows observed outflow rates (dashed line) along with median estimates of the secular trend ($\tau_t^{f,i}$, solid line) for the outflow rate for each age and gender subgroup. Shading denotes 68 percent and 95 percent confidence intervals.

II.E. The Secular Trend in the Unemployment Rate

We map the individual secular trends for each subgroup using appropriate weights to obtain $\bar{s}_t$ and $\bar{f}_t$ as:

$$\bar{s}_t = \sum\nolimits_{i=1}^{6} \omega_{i,t}^{s} \hat{\tau}_t^{s,i}, \qquad \bar{f}_t = \sum\nolimits_{i=1}^{6} \omega_{i,t}^{f} \hat{\tau}_t^{f,i}. \tag{14}$$

Figure 6 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, 1960–2018. The secular trend in the inflow rate shows a decline of about 50 percent 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 by Davis and others (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

$$\bar{u}_t = \frac{\bar{s}_t}{\bar{s}_t + \bar{f}_t}, \tag{15}$$

and is shown in the bottom panel of figure 6. The trend unemployment rate was about 6 percent in 1960 and increases to over 7 percent in 1983. Since then, it has displayed a clear downward trend, reaching about 4.5 percent by the end of the sample. Because the outflow rate shows little trending behavior, we observe from equation 15 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 percent confidence interval at the end of the sample is comfortably less than 1 percentage point.

To illuminate interesting features of the trends in labor market flows over the last 60 years, we perform a number of counterfactual exercises using the model introduced above. Although this analysis is mostly descriptive, in section V we complement this analysis using more detailed micro data to analyze the economic drivers of these changes.

THE ROLE OF THE DECLINE IN THE OUTFLOW RATE As we observed in the previous subsection, $\bar{f}_t$ shows only a very modest decline in our sample. However, this decline has a nonnegligible role in the behavior of the trend in the unemployment rate. In the top panel of figure 7, we display two counterfactuals for $\bar{u}_t$:

(1) The measure $\bar{u}_t$ when $\bar{f}_t$ is set constant to its sample mean (the dotted line). We observe that, starting in 1990, we would have observed about half

Figure 6. The Inflow Rate, Outflow Rate, and Unemployment Rate, Along with the Estimated Secular Trend, 1960–2018[a]

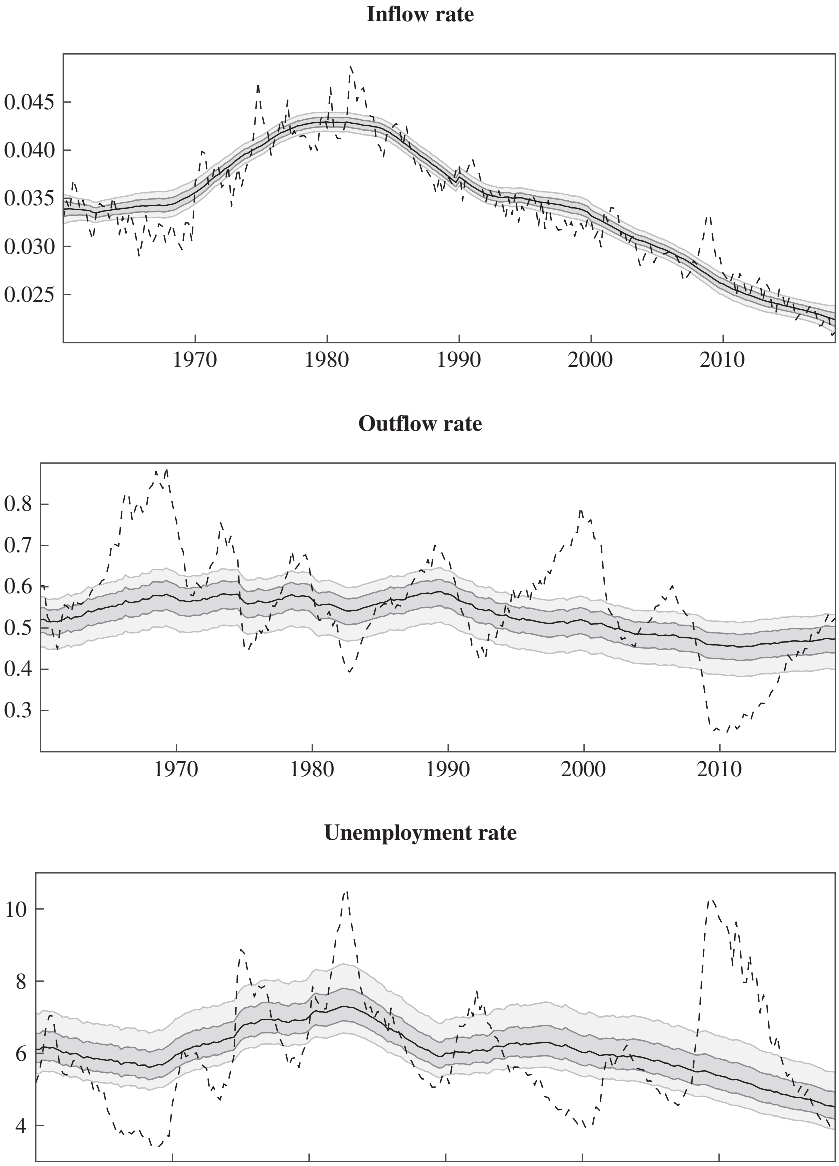

Sources: Current Population Survey; authors' calculations.

a. Actual rates denoted by dashed lines, median estimates of secular trend ($\bar{s}_t$, $\bar{f}_t$, and $\bar{u}_t$) denoted by the solid lines. Shading denotes 68 percent and 95 percent confidence intervals.

Figure 7. Counterfactual Exercises, Various Periods[a]

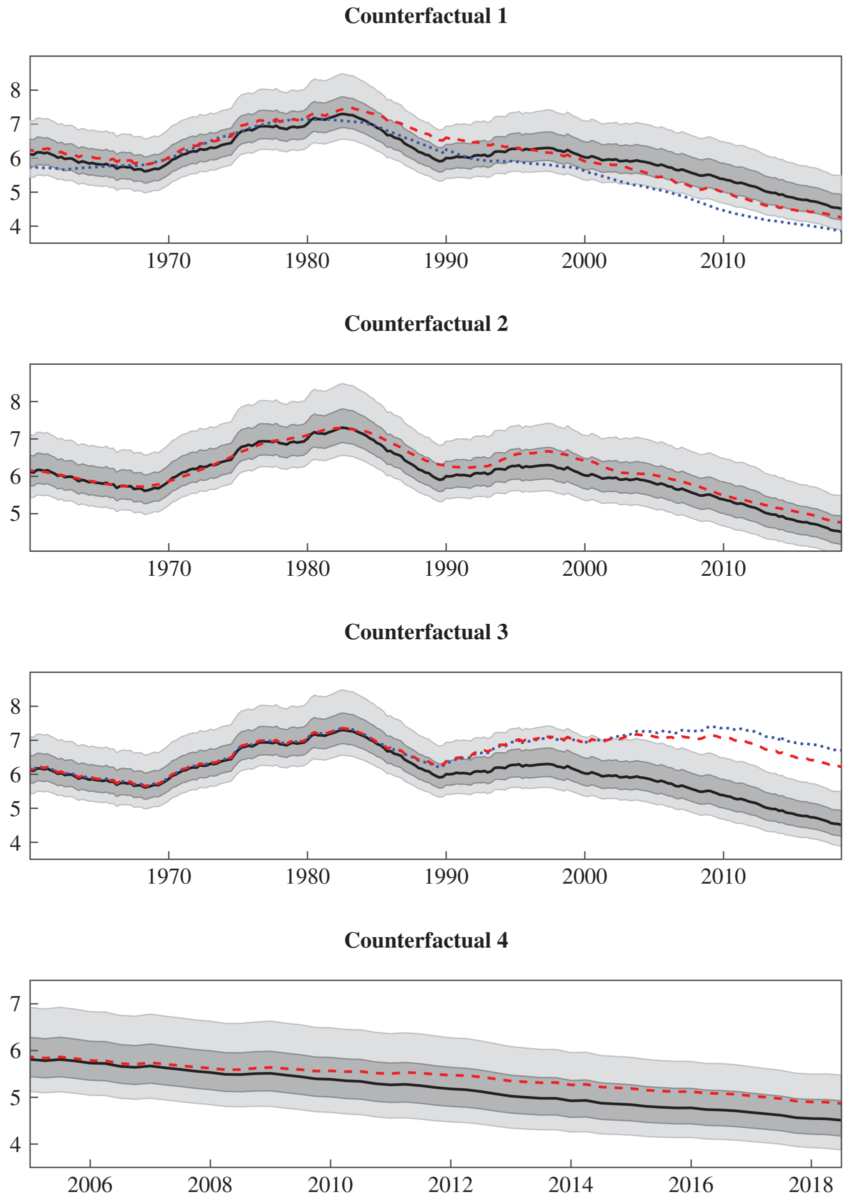

Sources: Current Population Survey; authors' calculations.

a. This figure shows the baseline estimate of $\bar{u}_t$ along with different counterfactual series (the dashed and dotted lines) based on the described scenarios. Shading denotes 68 percent and 95 percent confidence intervals.

a percentage point lower value of $\bar{u}_t$. This exercise implies that the entirety of the downward trend in $\bar{u}_t$ is driven by the inflow rate decline because it more than offsets the decline in the overall outflow rate.

(2) The measure $\bar{u}_t$ when $\bar{f}_t^i$ for $i = 1 \ldots, 6$ (the dashed line) are reestimated under the assumption that they are time invariant. We observe that this alternative tracks our baseline $\bar{u}_t$ closely, suggesting that $\bar{f}_t$ varies primarily through changes in the composition of the unemployed pool.

THE ROLE OF AGE AND GENDER COMPOSITION Section II summarized the sweeping demographic changes over the last 60 years that caused substantial shifts in age and gender composition. In this exercise, we capture the direct effect of the changing composition:

(1) The measure $\bar{u}_t$ when the weights $\omega_{i,t}^s$ and $\omega_{i,t}^f$ are fixed at their 1976 level (the dashed line). Between 1976 and the late 1990s, the change in the shares account for some of the decline. After 2000, the counterfactual series are very close to our baseline $\bar{u}_t$ estimate, implying that the majority of the secular decline after 2000 reflects the trends in the group-specific flows rather than changes in the composition.

THE ROLE OF YOUNG WORKERS AND FEMALES The bulk of the decline in the secular trend in the unemployment rate is accounted for by the group-specific trends of young workers (age 16–24) and prime age females until 2000. After the late 2000s, the decline is more broad based and goes beyond young workers and females. We conduct two counterfactual exercises:

(1) The measure $\bar{u}_t$ when the inflow rate of young and prime age females is held constant at its 1985 level (the dashed line). The secular trend in the unemployment rate would have experienced only a modest decline by 2000, highlighting the importance of females' role in the secular trend. After the late 2000s, the counterfactual series exhibits a substantial decline, pointing to a more broad-based decline.

(2) The measure $\bar{u}_t$ when, in addition to point 1, young male workers also experienced constant inflow rates at their 1985 level (the dotted line). The secular trend in the unemployment rate would have been about constant up until 2010, before declining to slightly above 6.5 percent.

THE IMPACT OF THE GREAT RECESSION Prime age males and females age 55 and older are the only two demographic groups that experienced a differential trend in their inflow rate after 2007. We interpret this change as the effect of the Great Recession and carry out a counterfactual exercise to capture its effect:

(1) The measure $\bar{u}_t$ when we fix the inflow rate for prime age males and females over 55 at their 2006 level, right before the onset of the Great Recession (the dashed line; the graph begins in 2005). The declines in the

inflow rate for these subgroups have only a modest effect on the secular trend, accounting for less than 50 basis points at maximal impact. This suggests that the declining trend in the inflow rate is primarily driven by forces in place well before the beginning of the last recession.

This set of counterfactuals illustrates clearly that most of the steady decline in the secular trend in inflow rates is associated with strong declines in the group-specific, trend inflow rate of young workers and prime age females observed from the early 1980s up to the late 2000s. The secular decline in the last two decades or so is more broadly based and cuts across demographic groups.

III. A Simple, Forward-Looking Phillips Curve

Thus far, we have focused on the secular trend in the unemployment rate implied by trends in labor market flows. This is, however, conceptually distinct from the "natural rate of unemployment," u_t^*, which is defined as the unemployment rate such that, controlling for supply shocks, inflation remains stable. Although the New Keynesian model—which has become a popular framework for monetary policy analysis and the core structure in many monetary DSGE models—features a forward-looking Phillips curve, it is silent on u_t^*. Instead of relating inflation to the unemployment gap, it typically relates inflation to the output gap or real marginal costs (Woodford 2003; Galí 2015). However, Galí (2011) reintroduces the unemployment rate into a New Keynesian model by rewriting the wage inflation equation, albeit with the assumption of a constant natural rate of unemployment. As described next, this motivates our formulation of the New Keynesian Phillips curve that connects inflation, π_t, to the unemployment gap, $u_t - u_t^*$, and retains forward-looking inflation expectations.

III.A. A Stylized New Keynesian Phillips Curve with Unemployment

We motivate our empirical specification of the New Keynesian Phillips curve with a stylized model based on the work of Galí (2011). In this framework, unemployment arises as a result of the market power of workers, which is reflected in positive wage markups. In particular, we assume a large representative household with a continuum of members specializing in different types of labor services and whose members experience different levels of disutility from working. Prices are fully flexible, but nominal wages are sticky; in each period, workers of a given type get to reset their wages with probability $1 - \theta_w$, similarly to that done by Christopher Erceg, Dale Henderson, and Andrew Levin (2000) and by Guillermo Calvo (1983).

Monopolistically competitive firms have access to a linear production function and produce using labor as the only input. Optimizing firms equate their marginal revenue and marginal costs:

$$a_t = w_t - p_t, \tag{16}$$

where the exogenous process a_t is the combination of (log) productivity and markup shocks to firms, w_t denotes the log nominal wage, and p_t is the log of the good's price. Given the wage, firms choose the quantity of workers employed, and the household assigns the workers with the lowest disutility of working. Because labor supply is elastic along the extensive margin, higher wage markups result in higher participation and therefore higher unemployment in the economy.

When they get to reset their wages, workers choose new wages that are a markup $\mu_{w,t}$ over a weighted average of current and expected future price-adjusted marginal rates of substitution. This results in a log-linearized wage Phillips curve of the form

$$\pi_t^w = -\kappa_w \left(\mu_{w,t} - \mu_{w,t}^*\right) + \beta E_t \pi_{t+1}^w, \tag{17}$$

where $\kappa_w = (1 - \theta_w)(1 - \beta\theta_w)/(\theta_w((1 + \varphi\varepsilon_w)) > 0$, φ is the steady state labor supply elasticity; $-\varepsilon_w$ is the steady state elasticity of demand for labor of different types; $\pi_t^w = w_t - w_{t-1}$ denotes nominal wage inflation; $\mu_{w,t}$ is the cross-sectional average wage markup over the economy's average marginal rate of substitution; and $\mu_{w,t}^*$ is an exogenous process capturing the time variation in the markup in the labor market, which in turn depends on firms' elasticity of demand for labor of different types, as well as that of the labor supply elasticity. Iterating equation 17 forward, we obtain

$$\pi_t^w = -\kappa_w \left(\mu_{w,t} - \mu_{w,t}^*\right) - \kappa_w E_t \sum_{T=t}^{\infty} \beta^{T-t+1} \left(\mu_{w,T+1} - \mu_{w,T+1}^*\right). \tag{18}$$

When average wage markups are below their desired level, workers who reset their wage will adjust it upward, resulting in positive wage inflation. Equation 18 reveals the central feature of the New Keynesian Phillips curve. The current gap is only one driver of inflation, and it might be a small contributor in the empirically relevant case if the slope of the curve, κ_w, is relatively flat. However, the discounted expected future path of the gap is a determinant of inflation as well. For a given level of the current gap, shifts in expectations have important implications for wage inflation—an insight lacking in the traditional backward-looking Phillips curve. The implication is then that it is important to take into account expectations when analyzing the relation between wage inflation and the markup gap.

Workers participate in the labor market only if their real wage is above their disutility from working, and Galí (2011) shows that this implies that the unemployment gap is proportional to the markup gap, so wage inflation can be expressed as

$$\pi_t^w = -\kappa x_t + \beta E_t \pi_{t+1}^w, \tag{19}$$

where $\kappa = \kappa_w \varphi > 0$, and x_t denotes the unemployment gap; here, the natural rate of unemployment (in deviation from its trend), defined as $z_t = u_t^* - \bar{u}_t = \varphi^{-1}\mu_{w,t}^*$, captures, in turn, time variation in firms' elasticity of demand for different types of labor, as well as in labor supply elasticity. Finally, using the identity relationship between price and wage inflation, we have

$$\pi_t^w = \pi_t + (w_t - p_t) - (w_{t-1} - p_{t-1}) = \pi_t + \Delta a_t, \tag{20}$$

where for the last equality, we use firms' profit-maximizing condition (equation 16). Using this expression to replace wage inflation in the wage Phillips curve, we obtain the price inflation New Keynesian Phillips curve, expressed in terms of the unemployment gap:

$$\pi_t = -\kappa x_t + \beta E_t \pi_{t+1} + \beta E_t (\Delta a_{t+1} - \Delta a_t), \tag{21}$$

where the last component is an exogenous term measuring expected wage growth, which depends on productivity and price markup shocks.

III.B. An Empirical Model

The model just described can be generalized in a variety of ways. In particular, we allow wages that are not optimally reset to be indexed to a combination of lagged inflation and the inflation target to better capture features of the data. Assuming rational expectations, the Phillips curve we consider in our empirical model thus takes the form

$$\pi_t - \pi_t^* = \gamma\left(\pi_{t-1} - \pi_{t-1}^*\right) - \gamma\sigma_{\pi^*}\varepsilon_t^{\pi^*} - \kappa E_t \sum_{T=t}^{\infty} \beta^{T-t} x_T - \beta \frac{1-\rho_a}{1-\beta\rho_a} \Delta a_t. \tag{22}$$

Here, π_t is determined by five core components: (1) π_t^*, which represents the drift in long-term inflation expectations, and therefore the degree of anchoring, and is assumed to evolve as

$$\pi_t^* = \pi_{t-1}^* + \sigma_{\pi^*}\varepsilon_t^{\pi^*}; \tag{23}$$

(2) π_{t-1}, which captures inertia in the inflation process; (3) $x_t = u_t - u_t^*$, which denotes the current unemployment gap, which evolves as

$$x_t = a_{x,1} x_{t-1} + a_{x,2} x_{t-2} + \sigma_x \varepsilon_t^x, \tag{24}$$

and here we adopt the common assumption of an exogenous data-generating process for the unemployment gap (Laubach 2001; Galí 2011); (4) the discounted expectation of future unemployment gaps, discounted at the rate β; and (5) the shock Δa_t, which we assume evolves as

$$\Delta a_t = \rho_a \Delta a_{t-1} + \sigma_a \varepsilon_t^a. \tag{25}$$

In terms of the structural model described above, this shock corresponds to real wage growth. Given the evolution of the unemployment gap, we can then rewrite equation 22 as

$$\pi_t - \pi_t^* = \gamma(\pi_{t-1} - \pi_{t-1}^*) - \gamma\sigma_{\pi^*}\varepsilon_t^{\pi^*} - \kappa w_{\pi,1} x_t - \kappa w_{\pi,2} x_{t-1} + \varsigma_t, \tag{26}$$

where $\varsigma_t = -\beta \dfrac{1-\rho_a}{1-\beta\rho_a} \Delta a_t$; $w_{\pi,1} = (1 - \beta(a_{x,1} + \beta a_{x,2}))^{-1}$; and $w_{\pi,2} = \beta a_{x,2} \cdot w_{\pi,1}$. As a result, observed inflation is measured as

$$\prod\nolimits_t = (\pi_t - \pi_t^*) + \pi_t^*, \tag{27}$$

and inflation expectations at different horizons j can be written as

$$\mathbb{E}_t \prod\nolimits_{t+j} = \pi_t^* + \ell_\pi' F^j X_t, \tag{28}$$

where $X_t = (\pi_t - \pi_t^*, x_t, x_{t-1}, \varsigma_t)'$, $\ell_\pi = (1, 0, 0, 0)'$ and $F = F(a_{x,1}, a_{x,2}, \gamma, \rho_a)$ is determined by equations 24 through 26.

The unemployment rate, u_t, may be expressed in terms of

$$u_t = x_t + z_t + \bar{u}_t, \tag{29}$$

with $z_t = u_t^* - \bar{u}_t$, the deviation of the natural rate of unemployment from its secular trend, which follows:

$$z_t = \rho_z z_{t-1} + \sigma_{z,\varsigma}\sigma_\varsigma \varepsilon_t^z. \tag{30}$$

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.

The model can be cast in state-space form. Equations 23 through 26, together with equation 30, represent the transition equations in the state-space model, and equations 27 through 29 are the observation equations. We estimate the model over the sample 1960:Q1–2018:Q3 using quarterly data. Our observed measure of u_t is the civilian unemployment rate from the Bureau of Labor Statistics (BLS). Inflation is measured as core CPI inflation in quarterly annualized percentage changes, which are also available from the BLS. We obtain a range of inflation expectations from various surveys by professional forecasters. For short-term inflation expectations, we combine six-month-ahead expectations, averaged across forecasters, from the Livingston Survey (available at semiannual frequency through our sample) and the Survey of Professional Forecasters (SPF, available since 1981:Q3). For long-term inflation expectations, we combine 5- to 10-year-ahead forecasts from Blue Chip Economic Indicators, Blue Chip Financial Forecasts, and the SPF. For the years 1975–77, we also use 5- to 10-year-ahead inflation expectations from the University of Michigan's Consumer Sentiment survey (see online supplemental appendix B for additional details about each series). Using equation 2, model-consistent, six-month-ahead inflation expectations are given by

$$\pi_t^* + (1/2)\,\ell_\pi' \sum_{j=1}^{2} F^j X_t, \tag{31}$$

while 5- to 10-year-ahead expectations can be expressed as

$$\pi_t^* + (1/20)\,\ell_\pi' (1 - F)^{-1}(1 - F^{20})F^{20} X_t. \tag{32}$$

We include independent measurement errors for both short-term and long-term forecasts with standard deviation parameters σ_{12Q} and σ_{510Y}. All parameters are estimated using Bayesian methods, with the exception of the discount rate, β. This parameter is set to $\beta = 0.99$, a value commonly used in the literature (Woodford 2003). We split the parameters in two vectors; $\Theta^1 = (\gamma, \kappa, a_{x,1}, a_{x,2}, \rho_a, \rho_z, \sigma_{z,\varsigma})'$ and $\Theta^2 = (\sigma^2_x, \sigma^2_\varsigma, \sigma^{*2}_\pi, \sigma^2_{12Q}, \sigma^2_{510Y})'$. Conditional on observing $\bar{u}_t$, this linear model can be estimated using the Gibbs sampler. In the first step, the Metropolis Hastings algorithm is used to draw parameters from Θ^1 for which we do not know the posterior distribution (this is due to the fact that the matrices F^j are nonlinear functions of the underlying model parameters).[12] In the second step, the Kalman smoother is used to draw the unobserved states, including initial conditions. In the third step, conditional on the drawn unobserved states, parameters from Θ^2 are drawn using known posterior distributions.[13] Because we observe a full distribution of paths of $\bar{u}_t$, we have to add an additional step in the estimation to account for this uncertainty. We first draw a path for $\bar{u}_{1:T}$, obtained from the estimation in section II, and conditional on this draw, we then proceed with the Gibbs sampler as described. We repeat this estimation procedure for a number of $\bar{u}_t$ paths selected at random.[14] This approach is motivated by the assumption, discussed in section II, that the unemployment trend $\bar{u}_t$ is exogenous to the variables in the Phillips curve model.

III.C. Adding Information from Wages

Although subsection III.A characterizes a simple wage Phillips curve, the discussion so far has not characterized the information available from observed wages. The importance of wages for assessing the unemployment gap has been emphasized by, for example, Robert Solow (1964), Olivier Blanchard and Peter Diamond (1989), and Galí (2011). Here we consider a second specification that includes *both* price and wage inflation. The goal here is to evaluate the impact of this additional information on our estimates of the natural rate of unemployment. Given that wages are measured with a considerable degree of noise, we extract information from five alternative data sources.[15] From the Employment Cost Index

12. For a comprehensive review of these Bayesian methods, see Herbst and Schorfheide (2015).

13. The estimation method follows quite closely del Negro and others (2017), who provide full details.

14. In more detail, we draw 250 paths from the distribution of $\bar{u}_t$. For each of these draws, we run a chain of 10,000 draws of the model parameters and states with the Gibbs sampler. Importantly, each parameter's chain is initialized by looking for a set of parameters close to the mode. We then randomly select 250,000 draws that we use to compute the joint distribution of parameters and states.

15. We thank our discussant, Steven Davis, for this suggestion.

release, we use growth in wages and salaries for private industry workers, along with growth in total compensation for all civilian workers (both starting in 2001:Q1). 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 (starting in 2006:Q1 and 1964:Q1, respectively). From the Productivity & Costs release, we use growth in compensation per hour of the nonfarm business sector (starting in 1947:Q1). All data are available from the BLS, and growth rates are expressed at a quarterly, annualized rate.

The relation between wage and price inflation implied by the model in equation 20 implies

$$\begin{aligned}(33) \qquad \pi_t^w = g_w + \pi_t^* + \gamma(\pi_{t-1} - \pi_{t-1}^*) - \gamma\sigma_{\pi^*}\varepsilon_t^{\pi^*} \\ - \kappa w_{\pi,1} x_t - \kappa w_{\pi,2} x_{t-1} - \frac{\beta^{-1} - 1}{1 - \rho_a}\varsigma_t,\end{aligned}$$

where π_t^w denotes nominal wage growth and g_w is the (constant) mean growth rate of real wages. This can be used to obtain this additional measurement equation to the model

$$\prod_t^{w,(j)} = \Theta^{(j)}\pi_t^w + \xi_t^{(j)}$$

where $j = 1, \ldots, 5$ denotes the individual nominal wage growth measures introduced above and where $\xi_t^{(j)}$ are first-order autoregressive measurement errors with autocorrelation coefficient $\rho_{\xi,(j)}$ and innovation standard deviation $\sigma_{\pi_{w,(j)}}$. We normalize the first loading coefficient $\Theta^{(1)} = 1$ and then estimate the remaining loadings $(\Theta^{(2)}, \ldots, \Theta^{(5)})$ along with $(g_w, \rho_{\xi,(1)}, \ldots, \rho_{\xi,(5)}, \sigma_{\pi_{w,(1)}}, \ldots, \sigma_{\pi_{w,(5)}})$. We view this as a particularly compelling exercise, given the relative stability of the wage Phillips curve over the sample, as shown by Galí (2011). In fact, the equation given above has a form similar to the one estimated by Galí (2011)—with a few key differences. Our specification includes an inflation trend, π_t^*, and, importantly, the original specification assumes a constant level of the natural rate of unemployment, which we eschew. We estimate this equation jointly with the rest of the model described in subsection III.B.

Table 1 shows the assumptions on the priors along with the properties of the posterior distribution for both model specifications. Notice first that the

Table 1. Parameter Estimates[a]

	Prior			Posterior (inflation only)			Posterior (inflation and wage inflation)		
Variable	*Distribution*	*Mean*	*Standard deviation*	*Mean*	*5%*	*95%*	*Mean*	*5%*	*95%*
$a_{x,1}$	Gamma	1.25	0.200	1.59	1.50	1.68	1.53	1.45	1.61
$a_{x,2}$	Normal	0.000	1.00	–0.618	–0.703	–0.529	–0.568	–0.639	–0.491
θ_w	Beta	0.750	0.100	0.879	0.847	0.906	0.852	0.820	0.878
κ	—	0.113	0.114	0.019	0.011	0.031	0.028	0.018	0.041
γ	Beta	0.500	0.265	0.049	0.004	0.196	0.025	0.003	0.060
ρ_{u*}	Beta	0.950	0.035	0.982	0.961	0.994	0.980	0.962	0.992
ρ_a	Beta	0.500	0.200	0.366	0.204	0.449	0.350	0.279	0.411
σ^2_x	InvGamma	0.12	—	0.043	0.032	0.054	0.043	0.032	0.055
σ^2_ς	InvGamma	1.00	—	1.44	1.23	1.68	1.34	1.13	1.56
$\sigma^2_{\pi*}$	InvGamma	0.12	—	0.067	0.053	0.083	0.069	0.054	0.087
$\sigma_{z,\varsigma}$	InvGamma	0.150	0.050	0.103	0.077	0.132	0.120	0.090	0.152
$\sigma^2_{o,\pi^{510Y}}$	InvGamma	0.005	—	0.003	0.001	0.007	0.004	0.001	0.010
$\sigma^2_{o,\pi^{12Q}}$	InvGamma	0.500	—	0.153	0.122	0.188	0.165	0.128	0.208
g_w	Normal	0.4	0.05	—	—	—	0.439	0.363	0.521

Source: Authors' calculations.

a. The variable κ is calculated assuming $\kappa = (1 - \theta_w)(1 - \beta\theta_w)/\theta_w$.

priors for the innovations' variance report only the mean, as the standard deviation is not defined; for these priors, we use an Inverse-Gamma distribution with shape parameter of 1.5, enough to have a well-defined mean. Second, the posterior distribution of the parameters for the two model specifications is broadly similar, with a few small differences that are discussed below. The process for the unemployment gap shows a high degree of persistence, consistent with medium-frequency business cycle movements in the unemployment rate. Regarding the Phillips curve, the slope is precisely estimated and in the range of about 0.02–0.03 across specifications, and it implies a fairly flat curve, as is often found in the literature (del Negro, Giannoni, and Schorfheide 2015). The addition of wages as an observable delivers a slightly steeper slope, but it does not fundamentally alter the estimated link between the current gap and inflation. The Phillips curve does not display significant inflation inertia, given the estimate of γ in the range about 0–0.2. The process for z_t is estimated to be highly persistent consistent with prolonged deviations of u_t^* from the secular trend in unemployment. The signal-to-noise ratio, $\sigma_{z,\varsigma}$ is tightly estimated in the range of about 0.12–0.15. Also, in this case, the introduction of wage inflation implies a somewhat more volatile natural rate of unemployment, as we discuss in the next section. Finally, the measurement errors for the survey-based measures of expectations are estimated to have small variances, allowing a tight mapping from observed expectations to the unobserved unemployment gap.

Priors for the five first-order autoregressive measurement errors (not shown in table 1) in the wage equation are as follows. The prior on the innovations corresponds to an inverse-gamma with shape and scale parameters implying a diffuse prior consistent with a well-defined mean of 1. The prior on the autocorrelation coefficient is normally distributed, with zero mean and variance determined using the same parameterization, with $\lambda = 0.1$ as in the VAR Minnesota prior. The priors are fairly diffuse. Finally, the priors on the loadings $\Theta^{(2)} \ldots \Theta^{(5)}$ (also not shown in the table) are defined as independent normal densities with mean 1 and standard deviation 0.5.

Before moving on to the main empirical results, it is useful to consider a simplified version of the model in order to make concrete the appropriate interpretation of κ, the slope of the Phillips curve, in this forward-looking model. As mentioned above, the estimated slope is fairly small; however, this does not necessarily imply a weak link between the unemployment gap and inflation. For example, consider a simpler model, where the unemployment gap, x_t, is a first-order autoregressive process (that is, $a_{x,2} = 0$) and set $\gamma = 0$. Then, solving forward for expectations delivers this relation between inflation and the contemporaneous unemployment gap,

$$\pi_t - \pi_t^* = -\frac{\kappa}{1 - a_{x,1}\beta} x_t + \varsigma_t .$$

Given that $\beta = 0.99$ and that we would expect a reasonably high degree of persistence in the unemployment gap, so that $a_{x,1}$ is near 1, then the coefficient relating actual inflation (in deviation from the trend) and the unemployment gap would be much larger than κ. The same intuition applies to the more general model introduced in this section.

IV. Measuring u_t^*

We estimate u_t^* since 1960 and show its evolution in the two panels of figure 8. The top panel refers to the model using price inflation only, and the bottom panel shows the results where both prices and wages are used. Both model specifications yield comparable predictions in general, with some differences, which are highlighted below. Overall, the natural rate of unemployment is estimated quite precisely with a 95 percent confidence interval of about 2 percentage points in the model using price inflation only, and with even narrower bands when wage information is added. When discussing ranges of the possible values of u_t^* at any particular point in time, we refer to the 68 percent interval. In the first decade of the sample, the natural rate hovers slightly below 6 percent and starts rising in the early 1970s, reaching comfortably above 7 percent by the late 1970s before falling to about 7 percent in 1983. The increase in the natural rate was the subject of a heated debate during the 1970s. Going back to earlier papers—such as those by Hall (1970a, 1970b), Gordon (1972, 1982), Perry (1978), and James Tobin (1974)—there appears to have been a consensus that the natural rate of unemployment increased to somewhere between a low of 5.0 percent and a high of 7.0 percent. 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. Subsequent research devoted substantial effort to understanding this period. For example, Summers (1986, 340) states that "where Kennedy-Johnson economists set 4 percent as an interim full-employment target, contemporary policymakers would regard even the temporary achievement of 6 percent unemployment as a great success." The natural rate then declines throughout the 1980s, consistently below

Figure 8. The Natural Rate of Unemployment, u_t^*, 1960–2018[a]

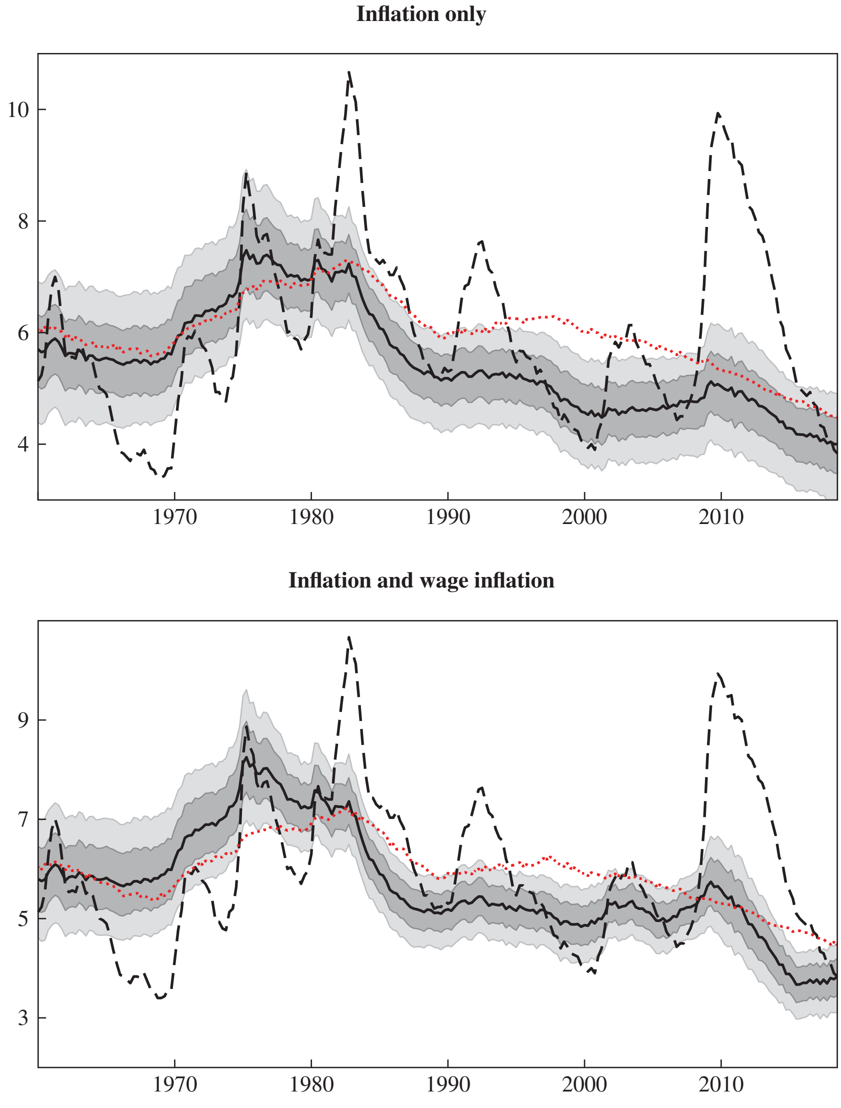

Source: Authors' calculations.

a. This figure shows the estimate of u_t^* for the inflation-only specification (top panel) and the inflation and wage inflation specification (bottom panel). The dotted line denotes the median $\bar{u}_t$. Shading denotes 68 percent and 95 percent confidence intervals.

the median secular trend in the unemployment rate (the dotted line). More recent analysis by Laurence Ball and N. Gregory Mankiw (2002) estimated the natural rate to be about 5.4 percent in 1960, and rising to 6.8 percent in 1979 and decreasing to 4.9 percent in 2000. Douglas Staiger, James Stock, and Mark Watson (1997) also have similar estimates.

One of the key differences between the top and bottom panels concerns the behavior of the natural unemployment rates during the 1970s. While the model with prices only estimates the natural rate to increase along with the secular trend in unemployment, the richer specification including wages estimates a further increase in the natural rate. One possible explanation for this discrepancy is the wage-price controls implemented in the early 1970s and their relative effects on wage growth and price inflation. As shown in the top panel of figure 9, while inflation dips in the early 1970s, wage inflation remains robust, signaling a strongly negative unemployment gap. The Nixon administration imposed wage and price controls in August 1971 that lasted until April 1974. The control program went through four phases. The first two phases were more strict and accomplished only a slight reduction in wage growth but a marked decline in the rise in prices between 1971:Q3 and 1972:Q2.[16] Phase II (which lasted until January 11, 1973) was followed by phases III and IV, but controls were generally relaxed in the last phases. Inflation started picking up in late 1972 while wage growth moderated. By the time wage-price controls were dismantled in April 1974, U.S. inflation had reached double digits. In fact, both panels of figure 10 show that, regardless of model specification, a substantial negative unemployment gap remained until the early 1980s.

The period spanning the 1990s to the Great Recession was characterized by a fairly stable natural rate of unemployment, which remained range-bound between 4.5 and 5.5 percent. During this period, the median u_t^* remained consistently below its secular trend. To speculate, some of this decline might have been due to the rapid growth of technological progress during the period.[17] As shown in figure 9, the unemployment gap had been consistently positive throughout the 1980s, around the deep monetary contractions of the Volcker disinflation period. It turned negative briefly in

16. See Gordon (1972, 1973).

17. As then–chair of the Federal Reserve Alan Greenspan said in a speech in 1998 on the New Economy: “Coupled with the quickened pace of productivity growth, wage and benefit moderation has kept growth in unit labor costs subdued in the current expansion. This has both damped inflation and allowed profit margins to reach high levels” (Greenspan 1998).

Figure 9. Inflation, Inflation Expectations, and Wages, 1960–2018[a]

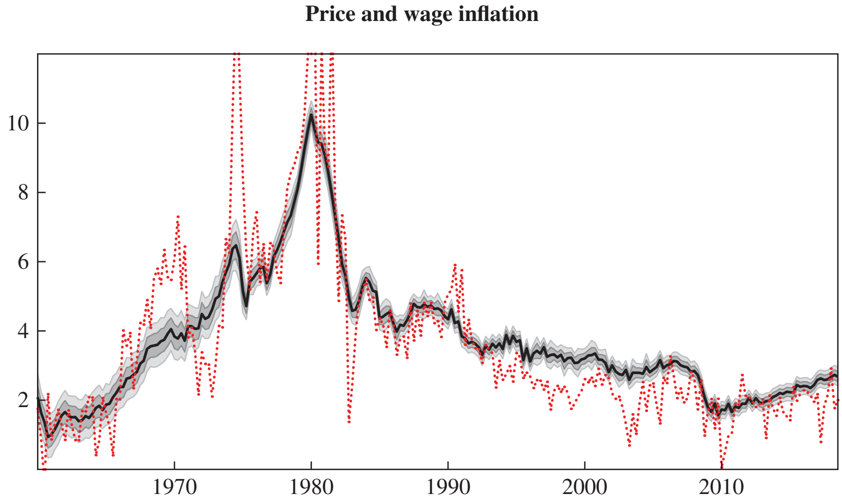

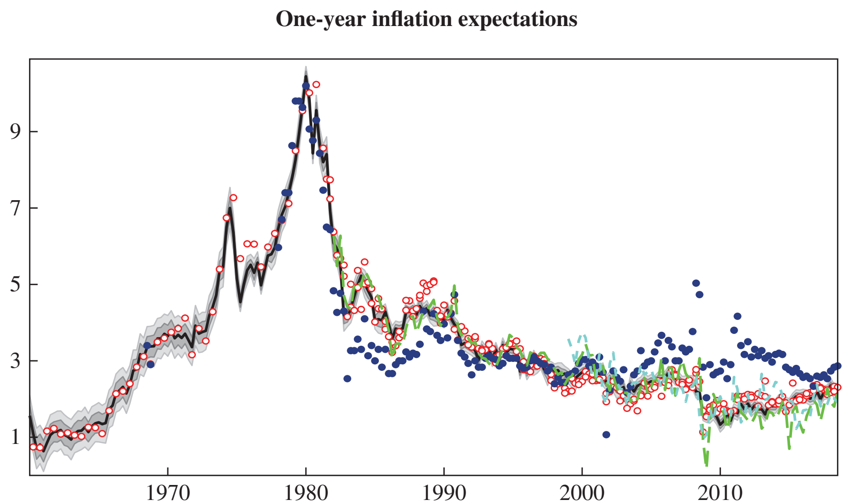

Long-term inflation expectations

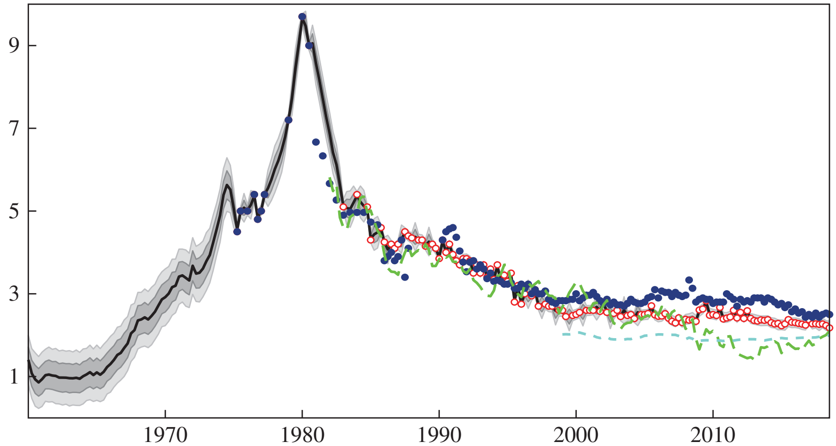

Sources: Bureau of Labor Statistics; Survey of Professional Forecasters; Livingston Survey; Blue Chip Economic Indicators; Blue Chip Financial Forecasts; University of Michigan; Abrahams and others (2016); Haubrich, Pennacchi, and Ritchken (2012).

a. The top panel shows realized quarterly annualized inflation (the dotted line) and the model predicted quarterly nominal wage inflation distribution (the solid black line and gray shading). The middle panel shows survey-based one-year inflation expectations of professional forecasters (the hollow dots) and households (the solid dots), and model-implied expectations (black line and gray shading). The short- and long-dashed lines show inflation expectations extracted from market prices from Abrahams and others (2016) (2-year) and Haubrich, Pennacchi, and Ritchken (2012) (one year), respectively. The bottom panel shows the same series from the middle panel but for the 5-year horizon beginning in 5 years. Shading denotes 68 percent and 95 percent confidence intervals.

Figure 10. The Unemployment Gap, 1960–2018[a]

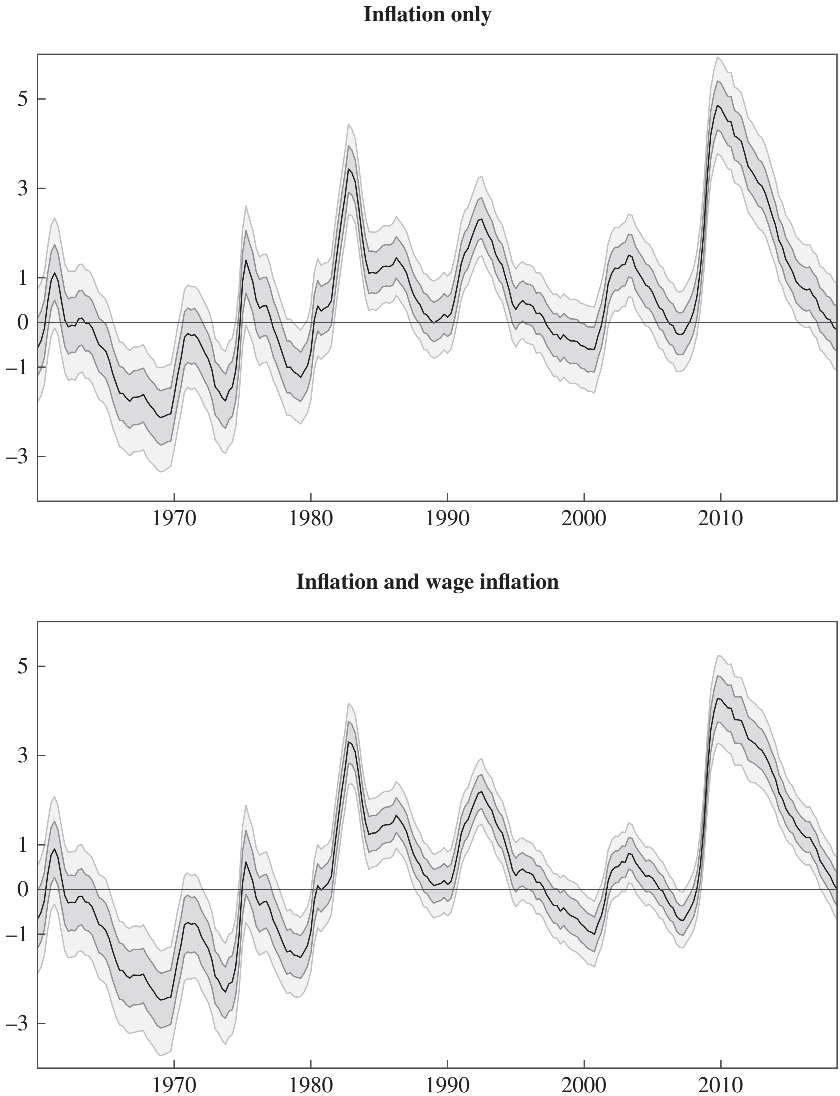

Source: Authors' calculations.

a. This figure shows the estimated unemployment gap, $u_t - u_t^*$, for the inflation-only specification (top panel) and the inflation and wage inflation specification (bottom panel). Shading denotes 68 percent and 95 percent confidence intervals.

the late 1990s, but this dip was preceded and followed by the 1990–91 and 2001 recessions.

Finally, during the prerecession years 2005–6, the natural rate of unemployment began increasing toward its long-run trend. This period presents the second important difference between the two model specifications. Including both prices and wages in the estimation leads to a higher estimate of u_t^*, which ends up overshooting its long-run trend, with its median estimate peaking in 2009–10 at just under 6.0 percent. Conversely, the model specification employing only price inflation predicts a milder increase (with a median that peaks at about 5.1 percent).[18] The different estimates reflect a pickup in wage growth in the period 2005–6, which we do not see in the measure of price inflation. Subsection IV.B. discusses the possible driving forces behind this increase. In the aftermath of the Great Recession, the natural rate of unemployment gradually declined, roughly in line with its secular trend. This finding implies that the fear of hysteresis after the Great Recession did not materialize, as we discuss in the next subsection. Both model specifications deliver estimates of the natural rate toward the end of 2018 in the range of 3.5 to 4.5 percent, which is consistent with the current unemployment gap of about zero.

Importantly, as shown in figures 9 and 10, the estimated Phillips curve is consistent with periods of large slack in the labor market and relatively stable inflation. This is perfectly illustrated by the Great Recession, which displays the largest unemployment gap in the sample, of about 4 percentage points, while price inflation declined only modestly. Indeed, while core CPI inflation was averaging 2.3 percent from 2005 to 2007, it declined to 1.4 percent on average from 2009 to 2011. Most important for the stability of inflation is the fact that inflation expectations declined only modestly both during and after the Great Recession, as shown in figure 9. As indicated in our Phillips curve, equation 22, 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 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

18. This is in line with the estimates by Daly and others (2012), who estimated the natural rate to be between 5.5 and 6.6 percent, with a preferred estimate at 6 percent at the end of 2011.

unemployment gap in the early 1980s caused a much more significant drop in inflation, with average core CPI inflation falling from 10.9 percent in the 1979–81 period to 4.5 percent in 1983–84. The key determinant is the behavior of inflation expectations, which dropped much more sharply in the early 1980s than was the case after the Great Recession. A 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 for estimating u_t^*.

The middle panel of figure 9 shows the model-implied predictive distributions (the gray shaded area) for the one-year-ahead inflation forecast, together with measured expectations from professional forecasters (the hollow dots). It is worth pointing out that as inflation and inflation expectations have been reverting to their long-term trend (see the figure's bottom panel), the unemployment gap has been steadily closing. The figure's middle panel also shows that alternative measures of one-year-ahead expectations display a roughly similar pattern as those of professional forecasters. For example, measures of inflation expectations extracted from asset prices (the dashed lines) are broadly in line. The median of households' expectations from the Michigan survey (the solid dots) behave differently since the early 2000s, predicting considerably higher inflation. Olivier Coibion and Yuriy Gorodnichenko (2015) show how the difference between the findings of household and professional forecasters over this period can be explained by oil prices. This difference, however, shrinks significantly when one looks at long-term inflation expectations, shown in the figure's bottom panel. These measures are of particular importance because they show the degree of anchoring of inflation expectations. As can be gleaned from the figure, all measures display a stable pattern after 1998, albeit at different levels, providing additional evidence to why the large unemployment gaps over these years were not associated with deflation.

IV.A. The Information Content of Inflation Expectations and the Secular Trend in the Unemployment Rate

In this subsection, we assess the role of observed inflation expectations and the secular trend in the unemployment rate in our estimated u_t^*. Figure 11 shows the results from two different estimation exercises. The top panel shows the estimate of u_t^* (for the price inflation only specification) when only information about $\overline{u}_t$ is provided, along with the realized unemployment rate and inflation rate. In this case, the estimated u_t^* is essentially identical to $\overline{u}_t$ (the dashed line), emphasizing the key role

Figure 11. Phillips Curve Models without Key Inputs, 1960–2018[a]

The model without inflation expectations

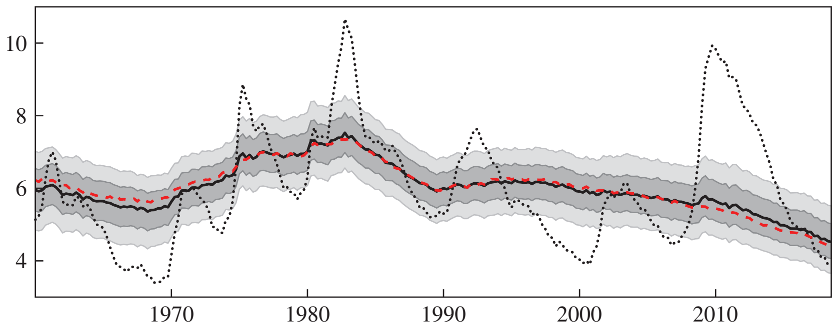

The model without trend or inflation expectations

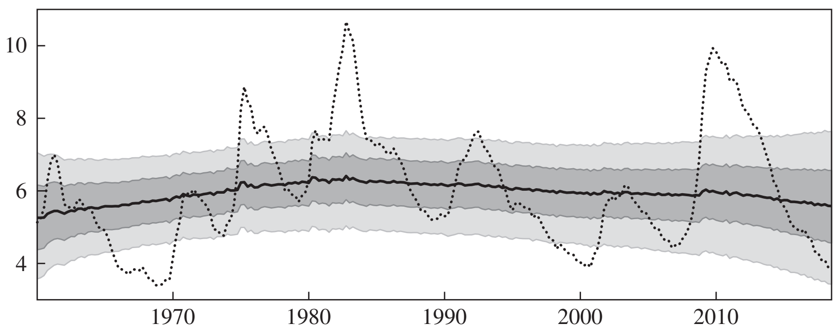

Source: Authors' calculations.

a. This figure shows results from an inflation-only Phillips curve model without key inputs. The top panel shows the estimated u_t^* (the black line) without inflation expectations as inputs. The dashed line represents the median estimate of $\bar{u}_t$. The bottom panel shows the estimated u_t^* (the black line) without information from inflation expectations or the secular trend in the unemployment rate. Shaded areas denote 68 percent and 95 percent confidence intervals.

that inflation expectations play in identifying movements in the unemployment gap across the state of the business cycle (the same result is obtained including both price and wage inflation in the estimation). In the bottom panel, we show the resulting u_t^* estimates when we further remove the secular trend from the observables. For this specification, we supply a stochastic process for the evolution of the trend, following the well-known model developed by Thomas Laubach (2001). The natural rate of unemployment follows the process

$$\bar{u}_t = \bar{u}_{t-1} + g_t \text{ and} \tag{34}$$

$$g_t = g_{t-1} + \sigma_g \varepsilon_t^g. \tag{35}$$

The model features no forward-looking behavior. Setting $\beta = 0$ delivers the Phillips curve

$$\pi_t - (1-\gamma)\pi_t^* - \gamma\pi_{t-1} = -\kappa x_t + \varsigma_t, \tag{36}$$

where ς_t is assumed to be independent and identically distributed under this specification. We also fix the standard deviation σ_g to deliver a smooth estimate of the trend, similar to our earlier analysis. In doing so, we considerably reduce the estimation uncertainty. In particular, we set $\sigma_g = 0.02 \times \sigma_u^*$. As is clear from the bottom panel of figure 11, the exercise shows that there is very little information about the natural rate of unemployment once we focus only on the joint behavior of inflation and unemployment. We conclude that inflation expectations and the secular trend in unemployment are therefore critical for assessing u_t^*.

IV.B. The Great Recession and Factors Affecting Matching Efficiency

It is of special interest to focus on the behavior of the natural rate of unemployment during the Great Recession and its aftermath. The Great Recession was not only the deepest postwar downturn in the labor market; it was also followed by an unprecedented period of high unemployment rates. The unemployment rate remained stubbornly high, printing at about 9 percent in January 2011, while many measures of economic activity had recovered by then. This disconnect triggered increased disagreement about the nature of the rise in the unemployment rate and whether the recession permanently affected the workings of the labor market. For example, figure C.3 in the online supplemental appendix summarizes the Federal Reserve Board's and Federal Reserve Banks' estimates of the NAIRU for three different periods: before the financial crisis in 2007, the current period (at the time), and 2015, as well as the increase between the first two periods as of January 25–26, 2011.[19] The figure shows that in 2011, there was increased disagreement not only about the current level of the natural rate but also about its level through 2015, suggesting that some participants viewed the natural rate of unemployment as higher even in the medium run due to hysteresis, as described by Blanchard and Summers (1986).

A careful examination of worker flows into and out of unemployment has shown that, though the inflow rate quickly returned to its prerecession level and gradually trended down, the persistently low outflow rate accounted for the high unemployment rate. Therefore, various explanations were suggested in the applied macroeconomics literature that operated through a long-lasting decline in the outflow rate, such as a rising mismatch, declining recruiting intensity, and a declining search effort by unemployed workers. This literature relied on rich micro data from various surveys, administrative data sources, and online data sources to quantify the effect of these factors. We next provide a simple framework derived from the search and matching literature to summarize these measures and then compare and contrast them with our measure of the natural rate of unemployment.

The point of departure is the matching function that characterizes the technology that firms and workers match with each other, building on the research of Diamond (1981), and of Dale Mortensen and Christopher Pissarides (1994), and on the work of Blanchard and Diamond (1989), who argue that changes in matching efficiency that shift the Beveridge curve may shed light on the Phillips curve. In its basic form, the inputs to the matching function at time t are the v_t vacancies posted by firms looking to hire and u_t unemployed workers looking for jobs. To accommodate the intensity of the recruiting and search effort, we denote the *recruiting intensity* of firms as q_t and the *search intensity* of workers as ι_t. A generalized Cobb–Douglas matching function that allows for shifts along the intensive margins of the firm and worker search effort can then be written as

$$h_t = \Phi_t (q_t v_t)^{\alpha} (\iota_t u_t)^{1-\alpha}, \tag{37}$$

where h_t is the total hires and $\alpha \in (0, 1)$ is the vacancy share. The term Φ_t is the aggregate matching efficiency parameter. As the specification shows,

19. See FOMC (2011a).

changes in q_t and ι_t would show up as a decline in the measured match efficiency. In addition, the mismatch between vacant jobs and unemployed workers (idle workers seeking employment in sectors, occupations, or industries different from those where the available jobs are) would manifest itself as a decline in Φ_t. Such a misalignment between the distribution of vacancies and unemployment along with a decline in the recruiting intensity and search effort would lower the aggregate outflow rate, which is defined as $f_t = h_t/u_t$.

MISMATCH Şahin and others (2014) formalize the notion of mismatch by defining the economy as a large number of distinct labor markets segmented by industry, occupation, and geography.[20] Each labor market i is frictional—that is, the hiring process within a labor market is governed by a matching function. To assess the existence of mismatch, they examine whether, given the distribution of vacancies observed in the economy, it would be feasible to reallocate unemployed workers across markets in a way that reduces the aggregate unemployment rate. This involves comparing the actual allocation of unemployed workers across labor markets with an optimal allocation that assumes costless worker mobility across these markets. Because the only frictions in such an environment are the ones embodied within each market-specific matching function, unemployment arising in this environment is purely frictional. The difference in unemployment between the observed allocation and the allocation implied by the optimal environment provides an estimate of the effect of mismatch.

Şahin and others (2014) calculate mismatch unemployment at the industry level using vacancy data from the Job Openings and Labor Turnover Survey (known as JOLTS), which provides survey-based measures of job openings and hires at a monthly frequency, starting from December 2000, and at the occupation level using vacancy data from the Help Wanted Online (known as HWOL) data set provided by the Conference Board. We plot an update of their occupation and industry mismatch unemployment measures in figure 12.

RECRUITING INTENSITY Recent research by Davis, Jason Faberman, and Haltiwanger (2013) has stressed the importance of channels other than a vacancy posting in the search and matching process. They argue that channels that affect how quickly firms fill those vacancies should be taken into account as determinants of the hiring process. A variety of factors—such as

20. In online supplemental appendix A, we present a simplified version of the derivation done by Şahin and others (2014).

Figure 12. A Comparison of u_t^* with Micro-Data-Based Estimates, 2006–18[a]

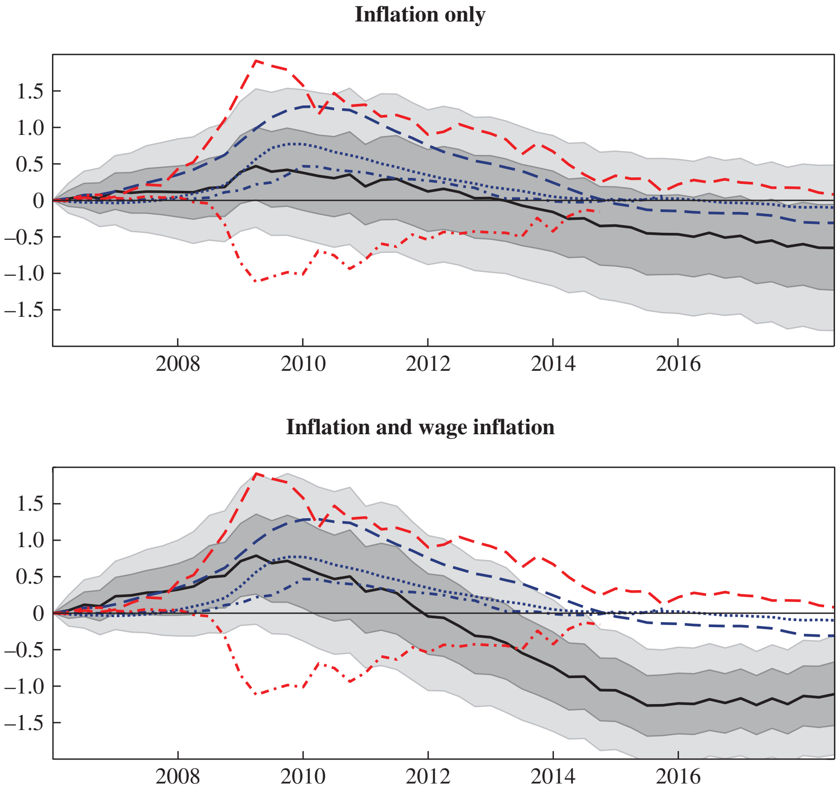

Sources: Current Population Survey; Job Openings and Labor Turnover Survey; American Time Use Survey; Conference Board.

a. Estimated u_t^*, occupation and industry mismatch unemployment rates (the lowest dashed line and dotted line), the recruiting-intensity-adjusted unemployment rate (the highest dashed line), the search-effort adjusted unemployment rate (the lowest dash-dotted line), and the unemployment insurance extension adjusted unemployment rate (the highest dash-dotted line). All series show changes relative to their level in 2006:Q1. Shading denotes 68 percent and 95 percent confidence intervals.

variations in hiring standards, wages offered that differ from those of competitors, variations in the amount of screening effort, and the propensity to use informal hiring methods—all contribute to what these authors refer to as *recruiting intensity.* They generate an aggregate time series of their measure of recruiting intensity using a generalized version of a standard matching function and their derivation of the monthly evolution of hiring and vacancies in the JOLTS data. Davis, Faberman, and Haltiwanger's (2013) recruiting intensity index provides us a normalized measure of q_t in the generalized Cobb–Douglas matching function.[21] In figure 12, we plot the counterfactual unemployment rate that is computed as the difference between the actual unemployment rate and a counterfactual unemployment rate that holds recruiting intensity constant at its mean value over the sample period replicated from the "Reserve Bank Report on Structural Unemployment" by Faberman and Şahin (2011).[22] This difference reflects the effect of changes in recruiting intensity on the unemployment rate.

THE WORKER SEARCH EFFORT AND EXTENSION OF UNEMPLOYMENT INSURANCE BENEFITS Another margin that is likely to be affected by aggregate conditions is unemployed workers' search effort. One often-discussed policy that is linked to the worker search effort is the extension of unemployment insurance (UI) benefits. In theory, receiving UI benefits for a longer period reduces the incentive of the unemployed to look for work. Similarly, it also increases unemployed workers' reservation wage, so that they may reject job offers that they would otherwise have accepted in the absence of these extended benefits. During the Great Recession, unemployment insurance benefits were extended to record lengths, with individuals in most states being eligible for up to 99 weeks of UI (and, at a minimum, 60 weeks). The Federal Reserve Board's Tealbook estimated that the extension of benefits raised the natural rate of unemployment, which we plot in figure 12.[23]

However, extension of UI is not the only channel that affected the worker search effort. Toshihiko Mukoyama, Christina Patterson, and Şahin (2018) showed that during the Great Recession, the unemployment pool shifted toward workers who are more attached to the labor force, who typically search harder for jobs. They showed that, as a consequence of this shift in the composition of the unemployed as well as the increased search

21. We thank Steven Davis for providing updated data of the recruiting intensity index and Jason Faberman for sharing his replication code.

22. See FOMC (2011b).

23. See FOMC (2012).

effort in response to declining household wealth, the aggregate search effort in the economy increased during the Great Recession. They find that the increase in search intensity during and after the Great Recession moderated the increase in the unemployment rate. Absent this increase, the unemployment rate would have peaked at about 11 percent and would have been consistently higher by about 0.5 to 1 percentage point during the recovery. Figure 12 shows the effect of the rise in workers' search effort on the unemployment rate.

IV.C. The Evolution of u^* and Factors Affecting the Matching Efficiency

In figure 12, we show the evolution of our estimated u_t^*, along with empirical estimates of the effects on unemployment of the persistent factors discussed in this section. All series are plotted in deviation from their level in 2006:Q1. There are two important features of these factors: (1) These persistent factors can only be measured as counterfactual gaps relative to an unemployment rate without additional assumptions; and (2) these measures are not additive, as they cannot be considered as independent from each other. That said, looking at the predictive distributions, we first observe that our measure of u_t^* aligns surprisingly well in terms of timing and magnitude with the evolution of these factors. The natural rate estimated using prices only is more aligned with industry mismatch unemployment and the effects of the extension of unemployment benefits (the dotted and dash-dotted lines in the figure). Conversely, u_t^* measured including both prices and wages in the estimation initially, is more aligned with the rise in occupational mismatch and decline in recruiting intensity. Though it displays a stronger increase relative to the inflation-only estimate, it still falls short of the observed spikes in these two series. Finally, job search intensity moved sharply, but in the opposite direction, moderating some of the effects of other factors. We also notice that these factors return to their 2006 levels in about 2014, while u_t^* continues to decline in line with the falling secular trend $\bar{u}_t$. Put differently, though the effect of the Great Recession on the unemployment rate persisted for almost a decade, these factors did normalize eventually. As such, we focus on driving forces that predate the Great Recession when studying the secular trend in the unemployment rate in the next section.[24]

24. This finding does not preclude persistent effects of the Great Recession on worker career paths (Davis and von Wachter 2011).

V. Changes in the U.S. Labor Market and Flow Dynamics

We have shown that the behavior of unemployment flows, especially the ongoing downward trend in the inflow rate, is the driver of the low levels of the natural rate of unemployment that the U.S. economy has been experiencing. In this section, we identify three important changes in the structure of the U.S. economy as the main drivers of the downward trend in the inflow rate: an increase in labor force attachment of females, a decline in job destruction and job reallocation, and the dual aging of workers and firms.

As a prelude to our analysis in investigating the economic changes that affected the evolution of the inflow rate, we report the changes in the inflow rate for the 1976–96 and 1996–2018 periods and decompose the total changes into changes accounted for by each gender and age group for the inflow rate

$$s_t \approx \sum_i \omega^s_{i,t} s^i_t, \tag{38}$$

where group i is defined as the interaction of gender and age. We consider three age groups for each gender, workers between 16 and 24 years, workers between 25 and 54, and workers older than 55 years. Table 2 decomposes the change in the inflow rate into changes accounted for by each demographic group using this simple approximation:

$$\Delta s(t, t') \approx \sum_i \left(\omega^s_{i,t'} s^i_{t'} - \omega^s_{i,t} s^i_t\right). \tag{39}$$

The table shows that females account for the majority of the decline in the inflow rate in the 1976–96 period, which coincides with the dramatic rise in the female labor force participation rate. In addition, during this period the baby boom cohort proceeded from younger ages to prime ages, reducing the aggregate inflow rate. Interestingly, the decline in the 1996–2018 period is very similar among males and females, suggesting a common factor after 1996.[25]

We also calculate the counterfactual contribution of each group, fixing their weights at their 1976 levels, a calculation that is often referred to as a shift-share analysis:

$$\Delta^C s(t, t') \approx \sum_i \left(\omega^s_{i,1976} s^i_{t'} - \omega^s_{i,1976} s^i_t\right). \tag{40}$$

25. The changes in the outflow rate are harder to interpret since the outflow rate is persistent and strongly procyclical. We report the corresponding changes in f in the online supplemental appendix.

Table 2. Shift Share with Fixed Demographic Composition[a]

Period	Aggregate change	Females			Males		
		16–24 years	*25–54 years*	*55+ years*	*16–24 years*	*25–54 years*	*55+ years*
A. Actual contribution							
1976–96	–0.80	–0.55	–0.24	–0.04	–0.17	0.13	0.01
1996–2018	–1.24	–0.34	–0.33	0.00	–0.34	–0.20	0.03
1976–2018	–2.04	–0.90	–0.57	–0.03	–0.51	–0.06	0.05
B. Counterfactual contribution with weights fixed in 1976							
1976–96	–0.36	–0.20	–0.32	–0.03	0.12	0.05	0.02
1996–2018	–1.29	–0.45	–0.23	–0.05	–0.42	–0.12	–0.02
1976–2018	–1.63	–0.65	–0.55	–0.08	–0.29	–0.06	0

Source: Current Population Survey.

a. Inflow rate changes are for 1976–96 and 1996–2018, and for the full sample, 1976–2018; and the actual and counterfactual contribution of each demographic group to the changes in the aggregate inflow rate. The counterfactual contributions are calculated using weights fixed at 1976 averages.

This counterfactual calculation shows that about half the decline in the 1976–96 period can be accounted for by the changing weights. In other words, if the shares of each age and gender group remained at their 1976 level, the inflow rate would have only declined by 0.36 percentage point, about 50 percent of the actual decline. The message is very different for the 1996–2018 period. Changes in the age and gender composition played no role in accounting for the decline in the inflow rate in this period.[26] Shift-share analyses—though informative—do not necessarily capture the full effects of demographic change. However, we find it useful because it helps us guide our examination of drivers of the decline in the inflow rate.

In light of our accounting exercise, and building on the earlier literature on female labor supply and firm dynamics, we now turn to the analysis of three channels that we show are important drivers of the downward trend in the inflow rate.[27]

V.A. The Increased Labor Force Attachment of Females

The United States experienced *grand gender convergence* in the 20th century, with female labor participation increasing from about 47 percent in 1976 to about 60 percent in 2000 (Goldin 2006).[28] The main driver of the rise in the female participation rate was the increase in participation of married females with children. Females started to work longer into their pregnancy, and they 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 child care. As labor market interruptions related to childbearing declined, females' labor force attachment gradually

26. This observation is consistent with the findings of Davis and Haltiwanger (2014), who show using various data sources that similar patterns apply to broader measures of worker allocation such as hires and separations.

27. It is possible that secular changes in factors affecting matching efficiency could have also played a role; however, these factors primarily affect the outflow rate which exhibits only a mild secular trend.

28. As discussed extensively by Goldin (2006, and the references therein), a large body of literature examines the drivers of the dramatic rise in female labor force participation. Various drivers of this change have been identified such as technological change, contraceptive innovation (for example, the birth control pill), a shift to a unilateral divorce legal framework, changes in social attitudes and norms toward married females working, advances in home production technology, a decline in maternal mortality, and the introduction of infant formula.

Figure 13. Labor Force Participation and Flows, 1976–2015[a]

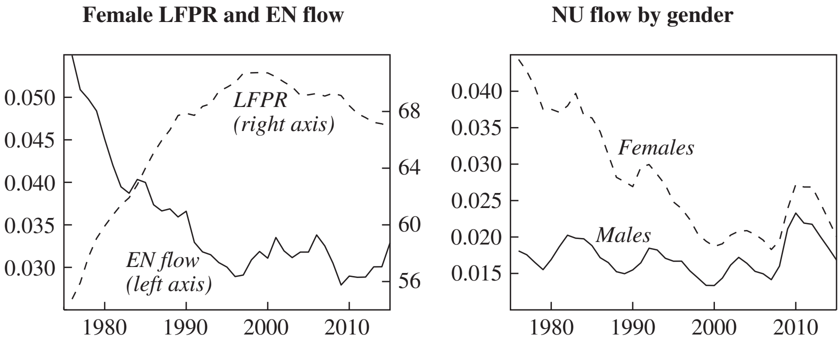

Source: Barnichon and Mesters (2018).

a. LFPR = labor force participation rate; EN = employment-to-nonparticipation; NU = nonparticipation-to-unemployment. Female labor force participation rate and EN flow rate (left) and NU flow rate normalized by the labor force by gender (right) for age 16–65 years.

increased.[29] The left panel of figure 13 shows that employed females left the labor force at a much lower rate in the 1990s than in the late 1970s. Though we do not have labor market flows before 1976, tabulations by Stephen Marston (1976) showed that the employment-to-nonparticipation flow rate for white females age 25–59 was 4.76 percent, while it was only 0.37 percent for white males in the same age category for the period 1967–73. Marston (1976) argues that the high rate at which employed females leave the labor force was the main factor in the higher unemployment rates they experienced. Marston referred to this as possibly a consequence of *participation instability*, almost antonymous to *increased labor force attachment*.

This decline in labor force exits has important implications for unemployment. Having uninterrupted employment spells allows workers to build more stable employment relationships, which is likely to reduce frictional unemployment through a decline in the incidence of job loss and the incidence of unemployment during reentry into the labor force. Examination of gross-flows data from the CPS based on longitudinally matched monthly CPS micro data confirms this intuition. As the right panel of figure 13 shows, as labor force departures became less common for females, entry from being out of the labor force into unemployment also became

29. For detailed statistics characterizing the changing patterns, see U.S. Census Bureau (2008).

increasingly rare. In the late 1970s, unemployment inflows from nonparticipation were about 4.5 percent of the labor force and declined by more than half, to about 2 percent by the late 1990s.[30]

Females also became less likely to leave unemployment for nonparticipation, increasing their duration of unemployment. Consequently, both the inflow and outflow gaps disappeared. On net, the decline in unemployment inflows dominated the rise in the duration of unemployment, causing full convergence of the unemployment rate of females to levels similar to males.[31] To summarize, even though declining exits from employment to nonparticipation will not have an immediate effect on the unemployment rate, they affected females' unemployment rate by lowering frictional unemployment, as shown by Katharine Abraham and Shimer (2002) using a flow decomposition and by Albanesi and Şahin (2018) using a three-state search and matching model.

GRAND GENDER CONVERGENCE IN THE CROSS-STATE DATA Evidence from U.S. states also confirms the relationship in the aggregate data: the rise in the female labor force participation rate was accompanied by an increase in labor force attachment, which in turn reduced frictional unemployment for females, generating a full convergence of unemployment rates by gender. We examine the evolution of the gender gaps in unemployment inflows and outflows at the state level. We first define the gender participation rate gap for each state at time t as

$$\frac{lfpr_{t,m} - lfpr_{t,f}}{lfpr_{t,m}}$$

and the unemployment inflow and outflow gaps as

$$\frac{s_{t,f} - s_{t,m}}{s_{t,m}} \quad \text{and} \quad \frac{f_{t,f} - f_{t,m}}{f_{t,m}},$$

with m denoting male outcomes and f denoting female workers' outcomes.

30. Abraham and Shimer (2002) and Albanesi and Şahin (2018) show nonparticipation to unemployment flows as a fraction of the stock of nonparticipation. We normalize these flows by the labor force because the unemployment rate is measured as a share of the labor force. An alternative is to compute the unemployment inflows by reason of unemployment, following Elsby, Michaels, and Solon (2009), who found that the inflow rate for labor force entrants declined starting in the early 1980s.

31. Albanesi and Şahin (2018) show that, though the gender unemployment gap has disappeared, the relative cyclicality of unemployment by gender has not changed.

Figure 14. The Convergence of the Inflow and Outflow Gaps

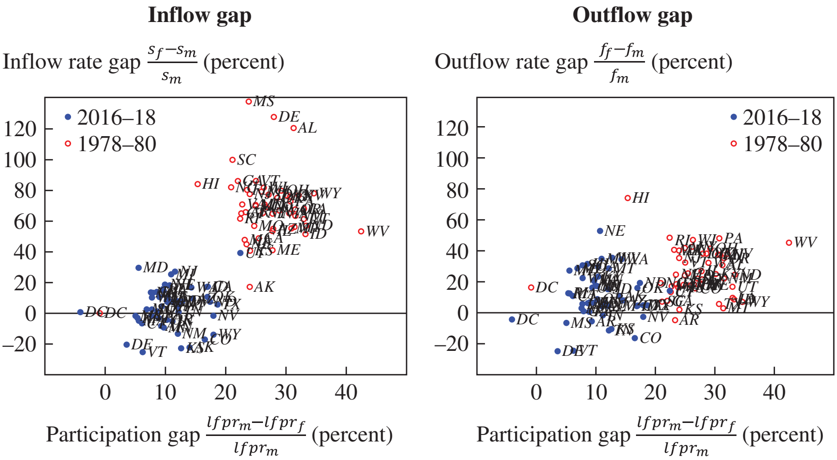

Source: Current Population Survey.
a. Participation and unemployment inflow/outflow gaps are for the 1978–80 and 2016–18 periods.

Figure 14 shows the state-level participation gaps and unemployment inflow and outflow gaps for the 1978–80 and 2016–18 periods. The convergence in labor market outcomes is clear. Moreover, unemployment inflows exhibit a starker convergence over time, consistent with the patterns in the aggregate data.

V.B. The Decline in Job Destruction Reallocation

Although females play an important role in the evolution of unemployment flows, almost all demographic groups' inflows declined over time. Especially after 1996, declines in group-specific inflow rates were the sole driver of the decline in the inflow rate, suggesting a common factor. Moreover, the rate employed workers transitioned into unemployment declined for both males and females, despite the dramatic job destruction at the onset of the Great Recession. This pattern suggests that changes in labor demand factors likely played a role.

The decline in unemployment inflows coincided with the decline in the volatility of firm-level growth rates and job destruction, as shown by Davis and others (2006); see figure 15. Search and matching models provide a natural link between the intensity of shocks that firms face and the incidence of unemployment. In this class of models with an endogenous job destruction margin, a decline in the intensity of firm-level

Figure 15. Job Destruction/Reallocation and Unemployment Inflow Rate, 1995–2018[a]

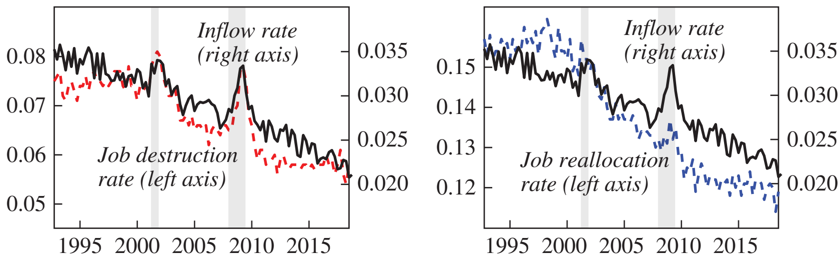

Sources: Current Population Survey; Business Employment Dynamics.
a. This figure shows job destruction and inflow rates (left) and job reallocation and inflow rates (right).

shocks would lower job destruction and the incidence of unemployment (Mortensen and Pissarides 1994). Davis and others (2010) formally examined this hypothesis, showing that industry-level movements in unemployment inflows are closely related to industry-level movements in several indicators for the intensity of idiosyncratic shocks for the 1990–2004 period. In this subsection, we extend Davis and others' (2010) analysis to the 1991–2017 period and evaluate the role of declining volatility on the trend decline in unemployment inflows and the employment-to-unemployment transition rate.

We use the Business Employment Dynamics (known as BED) data, which provide quarterly measures of job destruction at the industry level. We follow Davis and others (2010) and aggregate the data to these broad industry groups: construction, manufacturing, transportation and utilities, retail and wholesale trade, FIRE (finance, insurance, and real estate), and services; and we exploit within-industry time variation, the preferred specification of Davis and others (2010). The job destruction rate from quarter $t - 1$ to t is computed as the sum of job losses that are the result of contractions in employment at existing establishments and the loss of jobs at closing establishments, and it is expressed as a rate by dividing through by total employment.

We find economically and statistically significant effects of job destruction and job reallocation on the inflow rate and the employment-to-unemployment transition rate (table 3). The decline in job destruction and reallocation could be interpreted as declining firm level volatility and could arise from a changing nature of shocks or the declining responsiveness to shocks by firms as in Faberman (2017) and Ryan Decker and others (2017).

Table 3. Unemployment Inflow Rate and Employment-to-Unemployment Transition Rate Regressed on Job Destruction and Job Reallocation Rates, Quarterly Data[a]

	Inflow rate		*Employment-to-unemployment flow rate*	
Regressor	*(1)*	*(2)*	*(1)*	*(2)*
Job destruction rate	0.448***		0.382***	
	(0.0266)		(0.0202)	
Job reallocation rate		0.240***		0.197***
		(0.0198)		(0.0128)
Observations	618	618	618	618
R^2	0.935	0.927	0.950	0.933

Sources: Current Population Survey; Business Employment Dynamics.
a. Robust standard errors are in parentheses; ***$p < 0.01$, **$p < 0.05$, *$p < 0$. Quarterly data are from 1992:Q3–2018:Q1. The table includes time and sector fixed effects, with seven industry sectors.

V.C. Dual Aging

The U.S. economy has been experiencing a striking shift toward older workers and older firms since the mid-1990s, as we have discussed in section II. About 18 percent of the labor force was made up of workers between 16 to 24 years (*young workers*) in 1987. By 2017, this fraction declined to about 12 percent.[32] *Young firms'* (firms younger than 5 years old) employment share also followed a similar pattern, with their employment share declining from about 20 percent to 10 percent. On the flip side, in 1987, firms 11 or more years old—which we call *mature firms*, following Benjamin Pugsley and Şahin (2019)—used to employ about two-thirds of the workers in the economy. By 2017, this fraction increased to about 80 percent, as seen in figure 16.

Both worker age and firm age are widely recognized as important observables in accounting for differences in economic outcomes of workers and firms.[33] Table 4 shows the average unemployment inflow rates by worker age and job destruction rates by firm age. Younger workers are more than four times more likely to flow into unemployment than prime age workers. Similarly, firms between one and five years old are almost twice as likely to destroy jobs as their older counterparts. These patterns

32. In addition to the aging of the labor force, the ongoing decline in young workers' participation rate was a factor in this notable decline. As Krueger (2017) argues, the decline in participation of young workers was mostly offset by an increase in their college enrollment.

33. For worker age, see Perry (1970) and Shimer (1998); and for firm age, see Haltiwanger, Jarmin, and Miranda (2013) and Fort and others (2013).

Figure 16. The Dual Aging of Workers and Firms, 1987–2016[a]

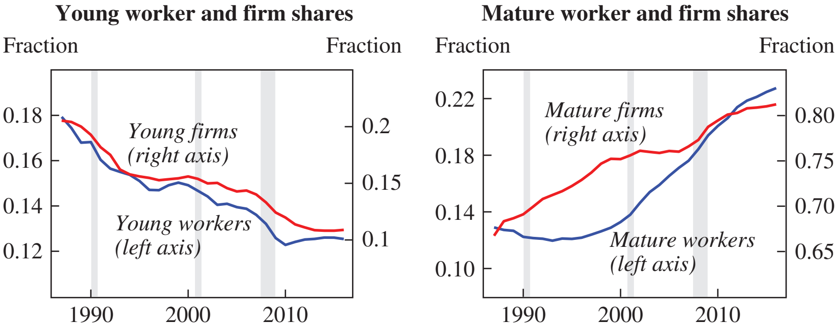

Sources: Current Population Survey; Business Dynamics Statistics.

a. This figure shows the employment share of 16- to 24-year-old workers and the employment share of 5-year-old and younger firms and employment shares of 55-and-above workers and employment share of firms 11 or more years old.

Table 4. Inflow and Job Destruction Rates, by Age[a]

Monthly inflow rates by worker age, 1978–2018		*Annual job destruction rates by firm age, 1987–2016*	
16–24	0.097	1–5	0.244
25–54	0.023	6–10	0.176
55+	0.015	11+	0.124

Sources: Current Population Survey (CPS); Business Dynamics Statistics (BDS).

a. This figure shows the average unemployment inflow rate by worker age in the CPS for 1976–2018 and the average job destruction rate by firm age in the BDS for 1987–2016. The inflow rates are calculated using the CPS, and the job destruction rates are calculated using the BDS.

suggest that a direct consequence of dual aging is a decline in unemployment inflows and job destruction.

We first conduct a simple worker-age-composition adjustment in the left panel of figure 17. We set the age composition of workers to their 1976 shares. We use three age groups for workers: 16–24, 25–54, and 55 or older. The shift toward an older population by itself accounts for about a quarter of the decline in the inflow rate, yet attributes a significant portion to the age-specific evolution of the inflow rate—a finding that resonates with our earlier analysis.

We repeat the same simple firm-age-composition adjustment in the right panel of figure 17, setting the age composition of firms to their 1987 shares using the Business Dynamics Statistics data set. We use three age groups of firms: 1–5 years, 6–10 years, and 11 or more years. Though it is hard

Figure 17. Shift Share Analyses with Fixed Age Composition, Various Periods[a]

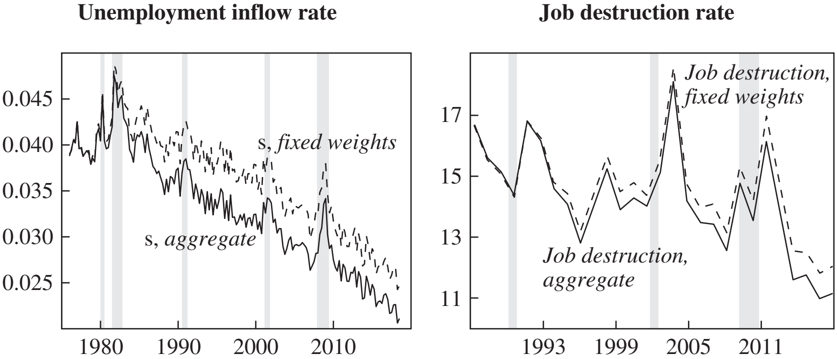

Sources: Current Population Survey; Business Dynamics Statistics; authors' calculations.
a. Unemployment inflow rate: aggregate and keeping the worker age composition unchanged at its 1976 shares (left) and job destruction rate: aggregate and keeping the firm age composition unchanged at its 1987 shares (right).

to assess the exact fraction that this calculation accounts for due to the pronounced countercyclicality of the job destruction rate, the change in the firm age compositions seems to be about as important as worker aging. However, the bulk of the decline still remains unaccounted for, similar to the inflow rate. This finding is consistent with those of Davis and Haltiwanger (2014) and of Haltiwanger, Ron Jarmin, and Javier Miranda (2013), who also show that though the shifts in the worker and firm age compositions go in the right direction, they still remain short of explaining the majority of the decline in the unemployment inflow and job destruction rates.

Although the shift in worker and firm age composition falls short of accounting for the decline in the inflow rate, recent research has emphasized that aging could also influence the economy by affecting age-specific outcomes, as noted by Shimer (2001), Fatih Karahan and Serena Rhee (2017), and Niklas Engbom (2019). Shimer (2001) refers to the direct effect as the effect of aging arising solely from changes in the age composition and to any additional effects as the indirect effect. These papers argue that the effect of aging goes beyond just shifting the composition of the economy in the context of unemployment, migration, and various measures of dynamism.

We build on this insight and show that the age composition of workers affects age-specific inflow rates and that the age composition of firms affects firm-age-specific job destruction rates, suggesting that indirect effects also play a substantial role.

DUAL AGING IN THE CROSS-STATE DATA We now turn to geographic variations to examine the direct and indirect effects of dual aging on unemployment inflows and job destruction rates using cross-state data.

Worker demographics and unemployment inflows. We should expect those states with larger changes in demographic makeup to experience the sharpest declines in inflow rates. Given the slow-moving nature of demographic changes, it is natural to compare long-horizon changes in these variables. To do so, we regress the change in the inflow rate for each state from its average value in 1978–82 to its average value in 1997–2001 on the change in the share of those age 15–24 relative to those age 15–64 from 1978 to 1998. We choose to take five-year averages of the inflow rates, and these years in particular, to ensure that our long differences are not unduly affected by the state of the business cycle (the first year in our sample is 1978, and the subsequent recession began in January 1980). We focus on the period up until the late 1990s because that is the period during which the share of young people in the population moved dramatically; since then, the changes have been relatively modest. In table 5, we show the results for this long-difference regression:

$$(41)\qquad s^i_{1997-2001} - s^i_{1978-1982} = \beta_0 + \beta_1\left[\left(\frac{\text{pop 15 to 24}}{\text{pop 15 to 64}}\right)_{i,1998} - \left(\frac{\text{pop 15 to 24}}{\text{pop 15 to 64}}\right)_{i,1978}\right] + \varepsilon_i .$$

The ordinary least squares (OLS) estimate suggests that a decline of 1 percentage point in the share of young people in a state corresponds to a fall of about 0.15 percentage point in the inflow rate. A common choice of instrument in regressions with age shares is to use lagged age shares adjusted for the deterministic aging that would be expected to occur.[34] This strategy relies on the idea that lagged age shares are not informative about current business conditions—such as a labor demand shock—that could potentially move both the age composition and unemployment inflows contemporaneously. In this case, we need to forecast the share of young people in 1998 as of 1978. We use 1978 births along with the rest of the age distribution at the time to do so. In particular, we forecast the long-horizon change by replacing

34. In a recent example, Davis and Haltiwanger (2014), estimate the effect of reallocation measures on employment and unemployment outcomes by age, gender, and education, using instruments based on age shares at the state level.

Table 5. Changes in Inflow Rates and Population Composition[a]

	Long-horizon change in inflow rate for:				
Specification	*Overall*	*All ages*	*16–24*	*25–54*	*55+*
Ordinary least squares	0.157	0.089	0.143	0.076	0.049
p value	(0.006)	(0.260)	(0.489)	(0.051)	(0.641)
IV	0.186	0.174	0.334	0.099	0.087
90% confidence interval	[0.05, 0.33]	[–0.00, 0.36]	[–0.17, 0.88]	[0.01, 0.20]	[–0.06, 0.24]
Observations	50	148	50	50	48

Sources: Current Population Survey; National Cancer Institute.

a. IV = instrumental variable. This table reports regression results for the specification of equation 41. The second row reports p values associated with the ordinary-least-squares estimate with robust standard errors; the fourth row reports weak-instrument robust confidence intervals constructed by inverting the Anderson–Rubin test (Mikusheva and Poi 2006). The "all ages" specification includes age effects. Results omit the District of Columbia.

$$\left(\frac{\text{pop 15 to 24}}{\text{pop 15 to 64}}\right)_{i,1998} - \left(\frac{\text{pop 15 to 24}}{\text{pop 15 to 64}}\right)_{i,1978}$$

with

$$\left(\frac{6\cdot(\text{1978 births}) + \text{pop 1 to 4}}{6\cdot(\text{1978 births}) + \text{pop 1 to 44}}\right)_{i,1978} - \left(\frac{\text{pop 15 to 24}}{\text{pop 15 to 64}}\right)_{i,1978}.$$

In words, to estimate the population of 15- to 24-year-olds in 1998 requires an estimate of the number of births in the period 1979–83 along with births and the population of 1- to 4-year-olds in 1978. We estimate the number of births in 1979–83 by assuming that births are constant at their 1978 level over that period. The instrumental variables (IV) estimate, also shown in table 5, is slightly larger than the OLS estimate, and weak-IV robust confidence intervals comfortably reject the null of a slope of zero.[35] These preceding results are similar in spirit to that of the shift-share analysis we have already conducted in table 2.

We next replace the long-horizon differences in inflow rates for all workers by the corresponding changes for workers below 25, age 25–54, and those 55 and over. This allows us to assess whether the maturing population—the decline in young people—is correlated with declines in

35. Although we employ weak-instrument robust confidence intervals, we note for reference that the first-stage F statistic is about 90.

the inflow rate for prime age and older workers. The results are suggestive that the change in the share of young people is associated with changes in separation rates across age groups (the second column of table 5). If we dig deeper and run these regressions by each age group separately, we observe that a higher share of young people is positively associated to the inflow rates of all groups with the effect declining by age (last three columns of the table). The most striking finding perhaps is the observation that prime age workers' unemployment inflow rate declines with the share of young workers in the economy. These results suggest that separation rates of prime age workers, in particular, were affected by the maturing population.

Firm demographics and job destruction. To investigate the role of firm aging, we follow a similar empirical strategy as in the case of worker demographics. We should expect those states with more substantial shifts toward older firms to experience the biggest declines in job destruction. We again consider long changes in job destruction due to the slow-moving firm demographics, as shown in figure 16. We compare the 3-year average of job destruction rates in 1987–89 with 2012–14 and examine how they are affected by the aging of firms, using the change in employment share of firms 11 years and older as a proxy for aging. This choice of regressor is motivated by the work of Haltiwanger, Jarmin, and Miranda (2013), who show that most of the young-firm dynamics continue throughout the first 10 years of firms' lives.

In table 6, we show the results for this long-difference regression, using OLS:

$$(42)\quad jd^{i}_{2012-2014} - jd^{i}_{1987-1989} = \beta_0 + \beta_1\left[\left(\frac{\text{emp}\,11+}{\text{emp}}\right)_{i,2014} - \left(\frac{\text{emp}\,11+}{\text{emp}}\right)_{i,1987}\right] + \varepsilon_i,$$

using job destruction data by state from Business Dynamics Statistics. The OLS estimate implies that an increase of 1 percentage point in the employment share of mature firms in a state corresponds to a fall of about 0.28 percentage point in the job destruction rate. This effect is both statistically and economically significant, and is more substantial quantitatively than the implication of the shift-share analysis. However, it is subject to the usual critique that firm demographics and job destruction could be affected by common shocks. To address this concern, we devise an IV strategy that parallels the IV approach that we have employed for worker demographics. To do so, we use the employment share of new firms (births) in

Table 6. Changes in Job Destruction Rates and Firm Age Composition[a]

	Long-horizon change in job destruction rate for:				
Specification	*Overall*	*All ages*	*1Y–5Y*	*6Y–10Y*	*11Y+*
Ordinary least squares	–0.284	–0.254	–0.261	–0.303	–0.197
p value	(0.008)	(0.002)	(0.146)	(0.025)	(0.035)
IV	–0.492	–0.571	–0.805	–0.595	–0.312
90% confidence interval	[–0.73, –0.33]	[–0.76, –0.41]	[–1.35, –0.45]	[–0.95, –0.34]	[–0.50, –0.16]
Observations	50	150	50	50	50

Source: Business Dynamics Statistics.

a. IV = instrumental variable. This table reports regressions results for the specification of equation 42. The second row reports *p* values associated with the ordinary-least-squares estimate with robust standard errors; the fourth row reports weak-instrument robust confidence intervals constructed by inverting the Anderson–Rubin test (Mikusheva and Poi 2006). The "all ages" specification includes age effects. Results omit the District of Columbia.

1979 as the instrument.[36] Petr Sedlácek and Vincent Sterk (2017) illustrate the strong persistence over time in the employment shares of startups. Because surviving firms that originated in 1979 did not become age 11 years until after 1987, these lagged employment shares of startups should forecast the subsequent long-horizon change of employment shares in old firms.[37] The bottom rows of table 6 show that the IV estimate is even stronger, suggesting a more substantial effect of firm aging on job destruction. In fact, the OLS estimate is outside the 90 percent weak-instrument robust confidence interval.

We next assess the indirect effect of firm aging on job destruction by investigating the relation between firm demographics and job destruction rates by firm age. In the second column of table 6, we pool the job destruction rates for all three age groups and include age effects. In concert with the overall results, the IV estimates suggest a larger magnitude of effect. In the table's last three columns, we consider individual specifications for each age group. We see clear negative effects of firm aging on job destruction for all firm age groups, suggesting that the shift in the age composition of firms is not the only effect of aging on job destruction. An older firm age distribution implies a lower overall job destruction rate by lowering job destruction for firms of all ages. We should also note that, across five specifications, the p values associated with the IV estimates are all smaller than 0.01. To ensure that our results are not driven by, for example, the types of industries that prevail in each state, we also report results for a panel version of equation 42, splitting the long-horizon change into observations of changes over two subperiods and taking the change (in the change) to account for unobserved heterogeneity (see table C.3 in the online supplemental appendix). Although the confidence intervals widen considerably, we continue to observe an estimated negative relationship and reject the null hypothesis at least at the 10 percent level, as judged by the weak-IV robust confidence intervals for all but the 6Y–10Y age category.

Our analysis showed that changes in worker and firm demographics, to which we refer as the dual aging of the U.S. economy, are important drivers of the decline in job destruction and unemployment inflows—two measures that we linked in the preceding subsection. Although the change in worker demographics is directly attributable to the baby boom, the drastic increase in births after World War II, the emphasis on aging of

36. Although the Business Dynamics Statistics data begin in 1977, we use 1979, due to measurement concerns discussed by Pugsley and Şahin 2019.

37. Although we employ weak-instrument robust confidence intervals, we note for reference that the first-stage F statistic is about 28.

firms is relatively new because data have only recently become available. However, Pugsley and Şahin (2019), using a firm dynamics framework, showed that the intuition is very similar for firms: declining firm births almost fully account for the shift of employment toward older firms. Moreover, Karahan, Pugsley, and Şahin (2018) show that the origin of the decline in firm entry is the decline in the labor supply growth arising from the aging of the baby boom cohort and the flattening out of the female labor force participation rate. These downward trends in unemployment inflows and job destruction pertain to a broader set of worker and job reallocation measures, as first documented by Davis and others (2006). Relatedly, recent research by Davis and Haltiwanger (2014) has shown that, in addition to shifts to older businesses and an aging workforce, policy developments that suppress reallocation—such as occupational labor supply restrictions, exceptions to the employment-at-will doctrine, the establishment of protected worker classes, and job lock associated with employer-provided health insurance—are among the policy factors that suppress labor market fluidity. Although analyses of these factors are beyond the scope of our paper, we believe that the interaction of policy decisions with labor market reallocation is an important issue for better understanding many important aggregates, such as unemployment, employment, productivity, and wages.

VI. Conclusion

We estimate the natural rate of unemployment in the 1960–2018 period by unifying two distinct estimation approaches that are popular in the literature. We exploit a rich set of labor market and inflation expectations data to provide tight estimates of the natural rate and study the underlying determinants of its movements. As of the third quarter of 2018, we estimate that u_t^* was about 4 percent; in particular, using only information from price inflation, we estimate that u_t^* stood at 4.0 percent with a 68 percent confidence interval of 3.5 to 4.5 percent. When we add information from wage inflation, the estimate shifts down slightly, to 3.8 percent, with an associated confidence interval of 3.5 to 4.2 percent. Our natural rate estimate is about 60 basis points lower than that of the CBO's estimate and 50 basis points lower than the median longer-run unemployment rate projection from the Federal Open Market Committee's "Summary of Economic Projections."[38] Importantly, our estimates imply that the unemployment gap was roughly closed toward the end of 2018.

38. See CBO (2019); FOMC (2019).

During the Great Recession, we find that the unemployment gap peaked at about 4 percentage points, which was far more severe than in any other downturn since the 1960s. Moreover, the closing of the unemployment gap has occurred only slowly, falling below 2 percentage points in 2014—about five years after the recession ended, and closing entirely only recently. We confront the micro data–based estimates of the rise in the natural rate during the Great Recession and find that our estimate of the rise is, remarkably, in agreement with the rise in mismatch unemployment and the decline in recruiting intensity. We view this similitude as an important success, given that these measures use almost completely separate sources of information.

Our analysis highlights a slow-moving secular trend that has been dragging down the unemployment rate since the early 1980s. This downward trend, until the late 1990s, was mostly driven by young workers and prime age females, while the secular trend in the last two decades is common across age and gender groups. We identify the rise in female labor force attachment, the decline in job destruction and reallocation intensity, and the dual aging of workers and firms in the economy as key drivers of this trend. Furthermore, we view these three developments as major changes that have had, and will continue to have, important and long-lasting effects on the economy.

The female labor force participation rate flattened in the late 1990s, and the unemployment rate for females fully converged with that of males. The participation gap has improved minimally since then, mostly on account of the deterioration of male participation outcomes and a gap of about 14 percentage points that still exists between prime age males and females. Our analysis of labor force attachment suggests that declining male attachment will be an upside risk for unemployment in the future, even though its effects are, thus far, more than being offset by the downward trend in job destruction. Another implication of our findings is that improvements in child care availability and maternity leave policy for females would also lower the natural rate of unemployment to the extent that they increase females' labor force attachment.

The aging of the population was predictable as early as the 1960s, and its consequences for innovation, productivity, government budgets, tax policy, Social Security, the labor market, and political economy have inspired an abundance of analyses and policy recommendations.[39] Although the discussion of the effects of aging goes back decades, there is still room

39. An insightful article by Cutler and others (1990) lays out various issues related to aging.

for further research on this topic, given that its effects on other parts of the economy—such as the decline in firm entry and the aging of firms—have taken shape. Another important implication of aging is the decline in workers' bargaining power, as recently analyzed by Andrew Glover and Jacob Short (2018). Using cross-sectoral variation, they find that older workers receive a smaller share of their marginal product than do younger workers, and they link the recent demographic trends in the United States to the declining bargaining power of workers. Because both worker and firm demographics are slow moving, and would likely take a long time to reverse, we expect these effects to persist.

Admittedly (and hopefully), our paper is not the last word on the natural rate of unemployment. However, we view our unified framework as a useful tool for future policy analyses, because it provides a bridge between the Phillips curve literature and the macro-labor literature, which focuses on measuring labor market efficiency by exploiting rich cross-sectional information. Moreover, the development of detailed micro data sources is a relatively recent development, and we expect that further progress harmonizing these two approaches will be made in future research.

ACKNOWLEDGMENTS The authors thank the *Brookings Papers* editors, Janice Eberly and James Stock, and commenters Steven Davis and Giorgio Primiceri. They also thank Olivier Blanchard, Jason Faberman, Bart Hobjin, Marco Del Negro, Domenico Giannone, Nora Traum, John Williams, and the Brookings conference participants for helpful comments. Réne Chalom, Brendan Moore, and Jin Yan provided excellent research assistance.

References

Abraham, Katharine G., and Robert Shimer. 2002. "Changes in Unemployment Duration and Labor Force Attachment." In *The Roaring Nineties*, edited by Alan Krueger and Robert Solow. New York: Russell Sage Foundation.

Abrahams, Michael, Tobias Adrian, Richard K. Crump, Emanuel Moench, and Rui Yu. 2016. "Decomposing Real and Nominal Yield Curves." *Journal of Monetary Economics* 84: 182–200.

Albanesi, Stefania, and Ayşegül Şahin. 2018. "The Gender/Unemployment Gap." *Review of Economic Dynamics* 30: 47–67.

Ball, Laurence, and N. Gregory Mankiw. 2002. "The NAIRU in Theory and Practice." *Journal of Economic Perspectives* 16, no. 4: 115–36.

Ball, Laurence, and Sandeep Mazumder. 2011. "The Evolution of Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity*, Spring, 337–405.

Barnichon, Regis, and Geert Mesters. 2018. "On the Demographic Adjustment of Unemployment." *Review of Economics and Statistics* 100, no. 2: 219–31.

Barnichon, Regis, and Christopher J. Nekarda. 2012. "The Ins and Outs of Forecasting Unemployment: Using Labor Force Flows to Forecast the Labor Market." *Brookings Papers on Economic Activity*, Fall, 83–131.

Blanchard, Olivier. 2016. *The U.S. Phillips Curve: Back to the '60s?* Policy Brief PB16-1. Washington: Peterson Institute for International Economics. https://piie.com/publications/policy-briefs/us-phillips-curve-back-60s.

Blanchard, Olivier J., and Peter Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity*, no. 1: 1–76.

———. 1990. "The Cyclical Behavior of the Gross Flows of U.S. Workers." *Brookings Papers on Economic Activity*, no. 2: 85–155.

Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual*, edited by Stanley Fischer. Cambridge, Mass.: MIT Press.

BLS (Bureau of Labor Statistics). 2018. "Labor Force Statistics from the Current Population Survey." https://www.bls.gov/cps/definitions.htm.

Brauer, David. 2007. "The Natural Rate of Unemployment." Working Paper 2007-06. Washington: Congressional Budget Office. https://www.cbo.gov/publication/18566.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12, no. 3: 383–98.

Carter, C. K., and R. Kohn. 1994. "On Gibbs Sampling for State Space Models." *Biometrika* 81, no. 3: 541–53.

Carvalho, Carlos, Sefano Eusepi, Emanuel Moench, and Bruce Preston. 2017. "Anchored Inflation Expectations." Working paper. https://dx.doi.org/10.2139/ssrn.3018198.

CBO (Congressional Budget Office). 2019. "Chapter 2: The Economic Outlook." In *The Budget and Economic Outlook: 2019 to 2029*. Washington: CBO. https://www.cbo.gov/system/files?file=2019-01/54918-Outlook-Chapter2.pdf.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7, no. 1: 197–232.

Cutler, David M., James M. Poterba, Louise M. Sheiner, and Lawrence H. Summers. 1990. "An Aging Society: Opportunity or Challenge?" *Brookings Papers on Economic Activity*, no. 1: 1–73.

Daly, Mary C., Bart Hobijn, Ayşegül Şahin, and Robert G. Valletta. 2012. "A Search and Matching Approach to Labor Markets: Did the Natural Rate of Unemployment Rise?" *Journal of Economic Perspectives* 26, no. 3: 3–26.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2013. "The Establishment-Level Behavior of Vacancies and Hiring." *Quarterly Journal of Economics* 26, no. 2: 581–622.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger, Ron Jarmin, and Javier Miranda. 2010. "Business Volatility, Job Destruction, and Unemployment." *American Economic Journal: Macroeconomics* 2, no. 2: 259–87.

Davis, Steven J., and John C. Haltiwanger. 2014. "Labor Market Fluidity and Economic Performance." NBER Working Paper 20479. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w20479.pdf.

Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda. 2006. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." In *NBER Macro Economics Annual*, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford. Cambridge, Mass.: MIT Press.

Davis, Steven J., and Till von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity*, Fall, 1–72.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2018. "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness." Working Paper 2018-007. Washington: Federal Reserve Board of Governors. https://www.federalreserve.gov/econres/feds/files/2018007pap.pdf.

Del Negro, Marco, Marc P. Giannoni, Domenico Giannone, and Andrea Tambalotti. 2017. "Safety, Liquidity, and the Natural Rate of Interest." *Brookings Papers on Economic Activity*, Spring, 235–94.

Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics* 7, no. 1: 168–196.

Diamond, Peter A. "Mobility Costs, Frictional Unemployment, and Efficiency." *Journal of Political Economy* 89, no. 4 (1981): 798–812.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2010. "The Labor Market in the Great Recession." *Brookings Papers on Economic Activity*, Spring, 1–69.

———. 2013. "Unemployment Dynamics in the OECD." *Review of Economics and Statistics* 95, no. 2: 530–548.

Elsby, Michael W. L., Ryan Michaels, and Gary Solon. 2009. "The Ins and Outs of Cyclical Unemployment." *American Economic Journal: Macroeconomics* 1, no. 1: 84–110.

Engbom, Niklas. 2019. “Firm and Worker Dynamics in an Aging Labor Market.” Working Paper 756. Minneapolis: Federal Reserve Bank of Minneapolis. https://www.minneapolisfed.org/research/wp/wp756.pdf.

Erceg, Christopher J., Dale W. Henderson, Andrew T. Levin, 2000. “Optimal Monetary Policy with Staggered Wage and Price Contracts.” *Journal of Monetary Economics* 46, no. 2: 281–313.

Faberman, R. Jason. 2017. “Job Flows, Jobless Recoveries, and the Great Moderation.” *Journal of Economic Dynamics & Control* 76: 152–70.

Faberman, R. Jason, and Ayşegül Şahin. 2011. “Reserve Bank Report on Structural Unemployment.” FOMC Memo, Meeting of January 25–26. https://fraser.stlouisfed.org/title/677/item/22545/content/pdf/fomc20110101memo01/toc/521590.

FOMC (Federal Open Market Committee). 2011a. “Material for FOMC Briefing on Structural Unemployment.” January 25–26. https://www.federalreserve.gov/monetarypolicy/files/FOMC20110126material.pdf

———. 2011b. “Reserve Bank Report on Structural Unemployment.” https://www.federalreserve.gov/monetarypolicy/files/FOMC20110101memo01.pdf.

———. 2012. “Domestic Economic Developments and Outlook.” https://www.federalreserve.gov/monetarypolicy/files/FOMC20121212tealbooka20121205.pdf.

———. 2019. “Economic Projections of Federal Reserve Board Members and Federal Reserve Bank Presidents under Their Individual Assessments of Projected Appropriate Monetary Policy, September 2018.” September 26. https://www.federalreserve.gov/monetarypolicy/files/fomcprojtabl20180926.pdf.

Fort, Teresa C., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2013. “How Firms Respond to Business Cycles: The Role of Firm Age and Firm Size.” *IMF Economic Review* 61, no. 3: 520–59.

Friedman, Milton. 1968. “The Role of Monetary Policy.” *American Economic Review* 58, no. 1: 1–17.

Fuhrer, Jeff, and George Moore. 1995. “Inflation Persistence.” *Quarterly Journal of Economics* 110, no. 1: 127–59.

Fujita, Shigeru, and Garey Ramey. 2009. “The Cyclicality of Separation and Job Finding Rates.” *International Economic Review* 50, no. 2: 415–30.

Galí, Jordi, 2011. “The Return of the Wage Phillips Curve.” *Journal of the European Economic Association* 9, no. 3: 436–61.

———. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton, N.J.: Princeton University Press.

Galí, Jordi, and Mark Gertler. 1999. “Inflation Dynamics: A Structural Econometric Analysis.” *Journal of Monetary Economics* 44, no. 2: 195–222.

Galí, Jordi, Frank Smets, and Rafael Wouters. 2012. “Unemployment in an Estimated New Keynesian Model.” In *NBER Macroeconomics Annual*, edited by Daron Acemoglu and Michael Woodford. Cambridge, Mass.: MIT Press.

Gertler, Mark, Luca Sala, and Antonella Trigari. 2008. “An Estimated Monetary DSGE Model with Unemployment and Staggered Nominal Wage Bargaining.” *Journal of Money, Credit and Banking* 40, no. 8: 1713–64.

Glover, Andrew, and Jacob Short. 2018. "Demographic Origins of the Decline in Labor's Share." Working paper. http://andyecon.weebly.com/uploads/5/1/3/8/5138005/ls_age_mar2018.pdf.

Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review* 96, no. 2: 1–21.

Goodfriend, Marvin, and Robert G. King. 1997. "The New Neoclassical Synthesis and the Role of Monetary Policy." In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Julio Rotemberg. Cambridge, Mass.: MIT Press.

Gordon, Robert J. 1970a. "Prices in 1970: The Horizontal Phillips Curve?" *Brookings Papers on Economic Activity*, no. 1: 8–47.

———. 1970b. "The Recent Acceleration of Inflation and Its Lessons for the Future." *Brookings Papers on Economic Activity*, no. 1: 8–47.

———. 1972. "Wage-Price Controls and the Shifting Phillips Curve." *Brookings Papers on Economic Activity*, no. 2: 385–421.

———. 1973. "The Response of Wages and Prices to the First Two Years of Controls." *Brookings Papers on Economic Activity*, no. 3: 765–79.

———. 1982. "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment." In *Workers, Jobs, and Inflation*, edited by Martin N. Baily. Washington: Brookings.

Greenspan, Alan. 1998. "Question: Is There a New Economy?" Speech given at Haas Annual Business Faculty Research Dialogue, Berkeley, September 4. https://www.federalreserve.gov/boarddocs/speeches/1998/19980904.htm.

Hall, Robert E. 1970a. "The Process of Inflation in the Labor Market." *Brookings Papers on Economic Activity*, no. 2: 343–93.

———. 1970b. "Why Is the Unemployment Rate So High at Full Employment?" *Brookings Papers on Economic Activity*, no. 3: 369–410.

———. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review* 95, no. 1: 50–65.

———. 2011. "The Long Slump." *American Economic Review* 101, no. 2: 431–69.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda. 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95, no. 2: 347–61.

Hamilton, James D. 2005. "What's Real About the Business Cycle?" *Federal Reserve Bank of St. Louis Review* 87, no. 4: 435–52.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken. 2012. "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps." *Review of Financial Studies* 25, no. 5: 1588–1629.

Herbst, Edward P., and Frank Schorfheide. 2015. *Bayesian Estimation of DSGE Models*. Princeton, N.J.: Princeton University Press.

Hornstein, Andreas, and Marianna Kudlyak. 2019. "Aggregate Labor Force Participation and Unemployment and Demographic Trends." Working Paper 2019-07. San Francisco: Federal Reserve Bank of San Francisco. https://www.frbsf.org/economic-research/files/wp2019-07.pdf.

Humphrey, Thomas M. 1991. "Nonneutrality of Money in Classical Monetary Thought." *Federal Reserve Bank of Richmond Economic Review* 18: 3–15.

Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin. 2018. "Demographic Origins of the Startup Deficit." Working paper. https://goo.gl/A3eX4r.

Karahan, Fatih, and Serena Rhee. 2017. "Population Aging, Migration Spillovers, and the Decline in Interstate Migration." Working Paper 699. New York: Federal Reserve Bank of New York. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr699.pdf.

Kim, Chang-Jin, and Charles R. Nelson. 1999. *State-Space Models with Regime Switching*. Cambridge, Mass.: MIT Press.

Krueger, Alan B. 2017. "Where Have All the Workers Gone? An Inquiry into the Decline of the U.S. Labor Force Participation Rate." *Brookings Papers on Economic Activity*, Fall, 1–87.

Laubach, Thomas. 2001. "Measuring the NAIRU: Evidence from Seven Economies." *Review of Economics and Statistics* 83, no. 2: 218–31.

Lucas, Robert E., Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4: 103–24.

Marston, Stephen T. 1976. "Employment Instability and High Unemployment Rates." *Brookings Papers on Economic Activity*, no. 1: 169–210.

Mikusheva, Anna, and Brian P. Poi. 2006. "Tests and Confidence Sets with Correct Size When Instruments Are Potentially Weak." *Stata Journal* 6, no. 3: 335–47.

Modigliani, Franco, and Lucas Papademos. 1975. "Targets for Monetary Policy in the Coming Year." *Brookings Papers on Economic Activity*, no. 1: 141–65.

Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding Declining Fluidity in the U.S. Labor Market." *Brookings Papers on Economic Activity*, Spring, 183–259.

Montgomery, Alan L., Victor Zarnowitz, Ruey S. Tsay, and George G. Tiao. 1998. "Forecasting the U.S. Unemployment Rate." *Journal of the American Statistical Association* 93, no. 422: 478–93.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61, no. 3: 397–415.

Mukoyama, Toshihiko, Christina Patterson, and Ayşegül Şahin. 2018. "Job Search Behavior over the Business Cycle." *American Economic Journal: Macroeconomics* 10, no. 1: 190–215.

Orphanides, Athanasios, and John C. Williams. 2002. "Robust Monetary Policy Rules with Unknown Natural Rates." *Brookings Papers on Economic Activity*, no. 2: 63–145.

Perry, George L. 1970. "Changing Labor Markets and Inflation." *Brookings Papers on Economic Activity*, no. 3: 411–48.

———. 1972. "Unemployment Flows in the U.S. Labor Market." *Brookings Papers on Economic Activity*, no. 2: 245–92.

———. 1978. "Slowing the Wage–Price Spiral: The Macroeconomic View." *Brookings Papers on Economic Activity*, no. 2: 259–99.

Phelps, Edmund S. 1967. "Phillips Curves, Expectations of Inflation, and Optimal Unemployment over Time." *Economica* 34, no. 135: 254–81.

———. 1968. "Money–Wage Dynamics and Labor-Market Equilibrium." *Journal of Political Economy* 76, no. 4: 678–711.

Polivka, Anne E., and Stephen M. Miller. 1998. "The CPS after the Redesign: Refocusing the Economic Lens." In *Labor Statistics Measurement Issues*, edited by John Haltiwanger, Marilyn E. Manser, and Robert Topel. Cambridge, Mass.: National Bureau of Economic Research.

Pugsley, Benjamin W., and Ayşegül Şahin. 2019. "Grown-Up Business Cycles." *Review of Financial Studies* 32, no. 3: 1102–47.

Roberts, John M. 1995. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit and Banking* 27, no. 4: 975–84.

Rogerson, Richard. 1997. "Theory Ahead of Language in the Economics of Unemployment." *Journal of Economic Perspectives* 11, no. 1: 73–92.

Rotemberg, Julio J. 1987. "The New Keynesian Microfoundations." In *NBER Macroeconomics Annual*, edited by Stanley Fischer. Cambridge, Mass.: MIT Press.

Rotemberg, Julio J., and Michael Woodford. 1997. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." In *NBER Macroeconomics Annual*, edited by Ben S. Bernake. Cambridge, Mass.: MIT Press.

Şahin, Ayşegül, and Christina Patterson. 2013. "The Bathtub Model of Unemployment: The Importance of Labor Market Flow Dynamics." Liberty Street Economics Blog, March. https://libertystreeteconomics.newyorkfed.org/2012/03/the-bathtub-model-of-unemployment-the-importance-of-labor-market-flow-dynamics.html.

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *American Economic Review* 104, no. 11: 3529–64.

Sargent, Thomas J. 1971. "A Note on the 'Accelerationist' Controversy." *Journal of Money, Credit and Banking* 3, no. 3: 721–25.

Sbordone, Argia M. 2002. "Prices and Unit Labor Costs: A New Test of Price Stickiness." *Journal of Monetary Economics* 49, no. 2: 265–92.

Schultze, Charles L. 1971. "Has the Phillips Curve Shifted? Some Additional Evidence." *Brookings Papers on Economic Activity*, no. 2: 452–67.

Sedláček, Petr, and Vincent Sterk. 2017. "The Growth Potential of Startups over the business Cycle." *American Economic Review* 107, no. 10: 3182–3210.

Shimer, Robert. 1998. "Why Is the U.S. Unemployment Rate So Much Lower?" In *NBER Macroeconomics Annual*, edited by Ben S. Bernake and Julio J. Rotemberg. Cambridge, Mass.: MIT Press.

———. 2001. "The Impact of Young Workers on the Aggregate Labor Market." *Quarterly Journal of Economics* 116, no. 3: 969–1007.

———. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95, no. 1: 25–49.

———. 2012. "Reassessing the Ins and Outs of Unemployment." *Review of Economic Dynamics* 15, no. 2: 127–48.

Solow, Robert M. 1964. *The Nature and Sources of Unemployment in the United States*. Stockholm: Almqvist & Wiksell.

Staiger, Douglas, James H. Stock, and Mark W. Watson, 1997. "How Precise Are Estimates of the Natural Rate of Unemployment?" In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer. Chicago: University of Chicago Press.

Summers, Lawrence H. 1986. "Why Is the Unemployment Rate So Very High Near Full Employment?" *Brookings Papers on Economic Activity*, no. 2: 339–96.

Tasci, Murat. 2014. "The Ins and Outs of Unemployment in the Long Run: Unemployment Flows and the Natural Rate." Working Paper 1233. Istanbul: Koç University-Tüsiad. https://eaf.ku.edu.tr/sites/eaf.ku.edu.tr/files/erf_wp_1233.pdf.

Tobin, James. 1974. "Monetary Policy in 1974 and Beyond." *Brookings Papers on Economic Activity*, no. 1: 219–32.

U.S. Census Bureau. 2008. "2008 Current Population Report on Maternity Leave and Employment Patterns of First-Time Mothers: 1961–2008." https://www.census.gov/prod/2011pubs/p70-128.pdf.

Woodford, Michael. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, N.J.: Princeton University Press.

Yun, Tack. 1996. "Nominal Price Rigidity, Money Supply Endogeneity, and Business Cycles." *Journal of Monetary Economics* 37 , no. 2: 345–70.

Comments and Discussion

COMMENT BY

STEVEN J. DAVIS Richard Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin make an excellent contribution to our understanding of unemployment and inflation dynamics, in two main parts: First, they estimate a trend decline in the U.S. unemployment rate of about 3 percentage points since the early 1980s. This development reflects a secular fall in the unemployment inflow rate, which the authors link to several deeper forces. This part of the paper brings together and extends a wide range of previous studies. Second, they develop an empirical model around a Phillips Curve with forward- and backward-looking elements, their measured unemployment trend, and survey data on inflation expectations. They use their empirical model to estimate the natural rate of unemployment, u^*, and to interpret the joint evolution of labor market slack and inflation. In their characterization, u^* moves over time due to the evolution of the unemployment trend and due to temporary forces.

Here, I first elaborate on the deeper forces behind the downward drift in unemployment and explain why this drift is important, quite apart from its implications for inflation. Then I express doubts about the practical usefulness of Crump and colleagues' Phillips Curve model, and the Phillips Curve concept more broadly, for the conduct of monetary policy. Relatedly, I conclude that continued efforts to precisely pin down the natural rate of unemployment and to estimate its impact on current and near-term inflationary pressures are unlikely to be fruitful. In closing, I suggest we can more readily advance our ability to assess current and near-term inflationary pressures by developing better measures of expected inflation, and a deeper understanding of how expected inflation behaves and feeds into current inflationary pressures.

A DOWNWARD DRIFT IN THE TREND UNEMPLOYMENT RATE Movements in the U.S. unemployment rate are well approximated by the steady state relation, $u \approx u^{SS} = s/(s + f)$, where s is the monthly unemployment inflow rate and f is the monthly outflow rate. In light of this observation, Crump and colleagues' measure the trend unemployment rate, $\bar{u}$, by extracting trends in s and f. Although they use more disaggregated data in measuring the trend, their figure 1 tells the story: The inflow rate, s, drifts down from the early 1980s, falling by roughly half over nearly four decades. The outflow rate, f, is highly procyclical and shows some indication of a downward drift after 2000. On this basis, the authors conclude that the trend unemployment rate, $\bar{u} = \bar{s}/(\bar{s} + \bar{f})$, reflects a large downward drift in the unemployment inflow rate. The proportionally modest decline in $\bar{f}$ works in the opposite direction.

WHAT IS BEHIND THE DOWNWARD DRIFT IN $\bar{s}$? Crump and colleagues identify four factors behind the downward drift in $\bar{s}$ in recent decades: an increased labor force attachment of females, the aging of the U.S. population ("worker aging"), a rightward shift in the employment-weighted age distribution of firms ("firm aging"), and secular declines in job destruction and reallocation rates. As the authors rightly note, these factors overlap, and we cannot simply add them up to get their combined contribution.

I agree that these four factors are important drivers of the downward drift in $\bar{s}$, but there is more to the story. Another important factor is the disappearance of short-duration employment relationships. Using data derived from administrative records, Henry Hyatt and James Spletzer (2017) show that more than half the drop in hiring and separation rates from 1996 to 2012 reflects a declining incidence of jobs that start and end in the same calendar quarter. They also find that the shifting composition of workers and employers accounts for only 22 percent of the declining incidence of short-duration employment relationships, mostly due to the aging of workers and firms.[1] This finding tells us that the disappearance of short-duration employment relationships is largely distinct from the first three factors that the authors stress. I suspect it is also largely distinct from the secular fall in job destruction and reallocation rates, given how job flows are measured.[2]

Another factor is falling labor participation rates among young adults. My figure 1 shows large declines since the late 1980s in the participation

1. Shifts by worker education and gender and by firm size and industry play a much smaller role or work in the opposite direction, according to Hyatt and Spletzer.

2. See Davis and Haltiwanger (1998) for an extended discussion of how worker flow measures (such as hiring and separation rates) relate to job flow measures.

Figure 1. U.S. Labor Force Participation Rates by Age Group, Annual, 1975–2018[a]

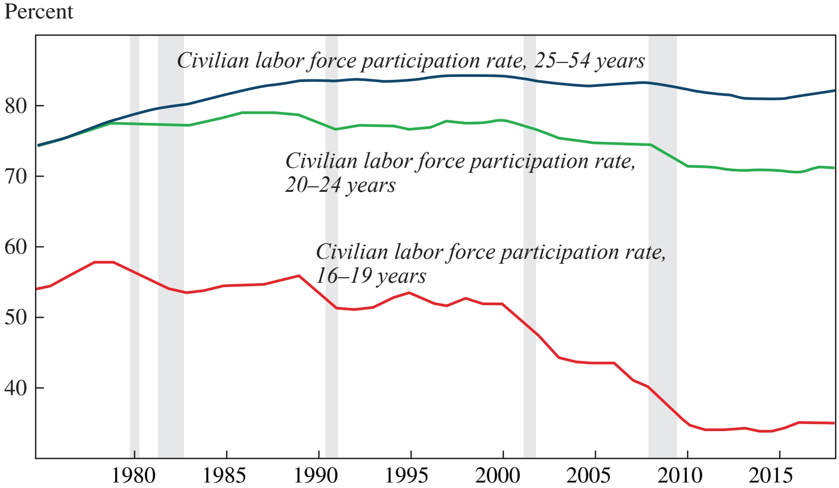

Sources: U.S. Bureau of Labor Statistics; Current Population Survey; FRED database of the Federal Reserve Bank of Saint Louis.

a. This figure shows data from the U.S. Bureau of Labor Statistics, specifically, the Civilian labor force participation rates from the Current Population Survey, retrieved from FRED on April 14, 2019: 16–19 years, LNU01300012; 20–24 years, LNU01300036; and 25 to 54 years, LNU01300060. Plotted data are annual averages of monthly values.

rates of persons who are 16–19 years old, more modest declines for those who are 20–24, and little change for those who are 25–54. My table 1 documents the well-known fact that young adults have relatively high unemployment rates. Unemployment inflow rates are also much higher among the young, as Crump and colleagues show in their figure 3. Taken together, these facts tell us that labor force participation rates by age shifted in a manner that contributes to the secular fall in $\bar{s}$ and $\bar{u}$. This factor is distinct from the role of population aging that the authors stress.

See my paper with John Haltiwanger (2015) for a discussion of developments that contributed to falling labor force participation rates among the young—and among the less-educated, who also have a relatively high propensity for unemployment. See Davis and others (2007), Decker and others (2014), and Davis and Haltiwanger (2015) on various factors behind the secular decline in job reallocation intensity.

THE SMALLER DOWNWARD DRIFT IN $\bar{f}$ The downward drift in the unemployment outflow rate is modest but has material implications for $\bar{u}$, as Crump and colleagues show. So, it is worth asking what deeper forces lie behind

Table 1. U.S. Civilian Unemployment Rates by Age Group and Time Period[a]

	Age group		
Time period	*25–54 years*	*20–24 years*	*16–19 years*
1975–2018	5.2	10.4	17.9
1975–99	5.2	10.5	17.8
2000–2018	5.1	10.2	18.1

Sources: U.S. Bureau of Labor Statistics; Current Population Survey; FRED database of the Federal Reserve Bank of Saint Louis.

a. This table shows data from the U.S. Bureau of Labor Statistics, specifically, the unemployment rates from the Current Population Survey, retrieved from FRED on April 14, 2019: 16–19 years, LNU04000012; 20–24 years, LNU04000036; and 25–54 years, LNS14000060. Table entries are averages of monthly values during the indicated period.

the downward drift in $\bar{f}$. In this respect, one noteworthy development is that geographic mobility has fallen in recent decades, even conditional on age (Molloy, Smith, and Wozniak 2014). As job losers and labor market entrants become less willing or able to migrate away from declining cities and regions, one likely effect is a fall in the unemployment outflow rate.

There are good reasons to think that falling geographic mobility is at least partly due to policy developments. As one example, the spread of occupational licensing (Kleiner and Krueger 2013) inhibits mobility across occupations and states. See the papers by Dick Carpenter and others (2012), the White House (2015), and Janna Johnson and Morris Kleiner (2017) for evidence. As a second example, Chang-Tai Hsieh and Enrico Moretti (2019) document a secular rise in the dispersion of nominal wages across U.S. cities from 1964 to 2009. They link this development to the adoption of land use restrictions in high-productivity coastal cities that reduced the elasticity of the housing supply and inhibited the in-migration of new workers from less productive cities and regions.

This raises interesting empirical questions that I have not seen addressed: To what extent is the downward drift in $\bar{f}$ concentrated in cities and regions with relatively low nominal wages? Is the downward drift present in cities with high nominal wages? Has the geographic dispersion of unemployment outflow rates risen in recent decades? If so, does the spatial pattern of rising dispersion in unemployment outflow rates align with the spatial pattern of rising nominal wage dispersion shown by Hsieh and Moretti (2019)?

FIVE TAKEAWAYS The authors make a compelling case that the trend component of the U.S. unemployment rate has fallen by roughly 3 percentage points since the early 1980s. They also identify several proximate drivers of this trend decline, and I have added to their list in the discussion above. These empirical results are quite helpful in assessing past U.S. economic performance and the future outlook for the U.S. economy, quite apart from any implications for inflation. In these respects, I see five important takeaways:

1. Much of the downward drift in U.S. unemployment rates over the past 35 years reflects good fortune rather than good macroeconomic policy. The "good fortune" includes the effects of population aging, increases in female labor force attachment, and declining business dynamism.

2. Some part of the downward drift in $\bar{u}$ probably reflects bad policy—that is, policies that drove younger, less educated, and other marginal workers out of the labor force.

3. Some past drivers of falling $\bar{s}$ have largely played out, and some may reverse. The increased labor force attachment of females, for example, seems to have largely played out. In all likelihood, we will be less fortunate with respect to the behavior of $\bar{s}$ and $\bar{u}$ in the coming decades.

4. The downward drift in $\bar{f}$, though modest and more recent, warrants concern. As the authors show (their figure 7, top panel), the fall in $\bar{f}$ raised $\bar{u}$ by 50 basis points. Lower values of $\bar{f}$ also slow recovery from the upward spikes in job destruction rates and s that typify the onset of recessions. That is, recessionary increases in unemployment take longer to unwind when $\bar{f}$ is lower, other things equal. Insofar as greater land use restrictions drove the downward drift in $\bar{f}$ by inhibiting migration to cities with better job opportunities, it will be politically challenging to reverse the decline in $\bar{f}$.

5. The trend component of the natural rate of unemployment has fallen substantially since the early 1980s, with 3 percentage points as a reasonable rough guess for the size of the fall.

THE CHALLENGE OF THE PHILLIPS CURVE Crump and colleagues integrate two very different approaches to estimating the natural rate of unemployment. This is a worthy ambition, and it is hard to take issue with the broad goal. Nevertheless, I see huge challenges in using the Phillips Curve to (1) sharpen our estimate of the natural rate of unemployment and (2) serve as a practical aid to monetary policymakers in assessing near-term inflationary pressures.

Justin Wolfers forcefully expresses one set of concerns about Phillips Curve modeling in his comments on a paper by Laurence Ball and Sandeep Mazumder (2011, 403–4) at a previous conference for the *Brookings Papers on Economic Activity*:

> That . . . the Phillips curve has not been proved false . . . may be because it is not falsifiable. . . . There are so many degrees of freedom to consider. . . . Inflation can be measured either as headline, core, or median, using either the PCE deflator,

> the CPI, or the GDP deflator. Inflation expectations can be modeled as rational, adaptive, or anchored. Data from different surveys can be utilized, such as the Livingston, the SPF, the Blue Chip, and the Michigan survey. Different measures of slack can be used, from the unemployment rate to the output gap to capacity utilization. The long-term unemployed can be included or not. Coefficients can be fixed or allowed to change over time. The lag structure can be adjusted, and nonlinearities can be assumed or ignored. Regime shifts can be invoked. Supply shocks can be included, including shocks to food, energy, and import prices, and price controls can be a factor in certain periods. Some economists in addition want to control for productivity or the labor share. In the end, there are more degrees of freedom than observations, which means that whatever path inflation might take, some researcher could plausibly claim to have found a Phillips curve that accounts for that path.

To this expansive list of Phillips curve variants, the paper at hand adds new degrees of freedom in the form of a richer, more flexible characterization of the natural rate of unemployment. Moreover, Crump and colleagues' account of puzzling U.S. inflation behavior during and after the global financial crisis leans very heavily on the paths of expected future inflation and expected future slack. We at least have multiple sources of data on expected future inflation, but expected future slack is essentially a free path variable, constrained only by the model and its functional form. Economists are very good at devising models with free path variables to fit nettlesome time series. From my vantage point, the authors' Phillips curve looks like the latest iteration in the long line of iterations that Wolfers summarizes. I do not think it will be the last iteration. More important, I do not see any reason to think it will prove a more useful practical guide to near-term inflationary pressures than many of its predecessors.

There is another view. At the same Brookings conference, James Stock responded to Wolfers as follows (Ball and Mazumder 2011, 404):

> The basic fact remain[s] that inflation in the United States and in other developed economies falls during periods of slack. This happened during the 1960s recession and again during the 1969 recession. The 1973 recession was different because of the oil price shock, but the pattern reappeared in the early-1980s and 1990 recessions, and again in 2000 for a while, except for a very interesting episode in 2004 and 2005. And much the same thing happened in 2007 and after, although the scale of it was in question. The issue then is not whether the pattern exists, but how to model it.

Stock's point about the "basic fact" is important and hard to deny. But it does not necessarily follow that the relationship between inflation and the unemployment gap (or other measures of slack) is sufficiently simple, stable, and predictable in its response to policy itself as to admit an empirical model that delivers confident predictions about near-term inflationary pressures.

In this respect, I am reminded of the view expressed by Olivier Blanchard (2016, 31): "Macroeconomists have learned, often painfully, that, while low unemployment creates inflation pressure, the form of the relation can change and has changed over time." Blanchard reviews some of these changes in U.S. unemployment-inflation dynamics since the 1960s.

In using the Phillips curve as a practical tool of monetary policy, perhaps the best we can do is to keep in mind Stock's "basic fact" and combine loose theorizing, simple statistical models, and informed judgment to obtain a very rough barometer of near-term inflationary pressures.

UNCERTAINTY ABOUT THE NATURAL RATE OF UNEMPLOYMENT The 68 percent confidence interval for u^* in the authors' preferred empirical model is about 1 percentage point. The 95 percent confidence interval is about 2 percentage points. Their figure 8 shows two estimated time series for u^*, one based on a price inflation Phillips curve, and one that relies on a Phillips curve specification that uses price and wage inflation data. As the authors discuss, their (median) estimated u^* differs a good deal between these two specifications during much of the 1970s and in the 2009–10 period. Eyeballing figure 8, it appears the peak difference is roughly 80 basis points. In addition to these sources of uncertainty about u^*, the estimated natural rate presumably depends on the choice of inflation expectations data to feed into the model, functional form choices, and more. Pulling these points together, I conclude that the authors' empirical undertaking does not yield much confidence about the value of the natural rate of unemployment at any point in time. Accounting for estimation uncertainty, specification uncertainty, and uncertainty about the appropriate data inputs, the range of reasonable values for u^* seems to be at least 250 basis points.

Here, I do not mean to suggest that the authors have done a poor job. Rather, I conclude that bringing the Phillips curve to the table helps little in sharpening our estimates for the natural rate of unemployment.

TAKING CRUMP AND COLLEAGUES' MODEL AT FACE VALUE: DOES SLACK MATTER? The authors estimate a flat Phillips Curve: the 90 percent confidence interval for the coefficient on the unemployment gap, $u_t - u_t^*$, is 0.011 to 0.031 in the model that uses price inflation data only and 0.018 to 0.041 in the model that uses both price and wage inflation data. Suppose these estimates are in the right ballpark. It follows that high uncertainty about the unemployment gap matters little for assessing current inflation pressures, so long as we have good data on expected inflation rates. To be concrete, suppose we misjudge the current value of u^* (and the unemployment gap) by 2 percentage points. Multiplying this misjudgment by a slope coefficient of 0.03 means that we misjudge current inflation pressures by only 6 basis points (annualized), conditional on expected future inflation. This is a tiny error. Indeed, it is probably smaller than the uncertainty about the current inflation rate.

The obvious corollary is that getting a sharp estimate for u^* matters very little for assessing current and near-term inflation pressures, provided that we have timely, high-quality measures of inflation expectations. In light of this corollary and 50 years of frustration in macroeconomists' efforts to develop a stable, reliable Phillips curve model, perhaps we should shift our focus to better measures of expected inflation, a deeper understanding of what causes expected inflation to move, and a better grasp on how expected future inflation feeds into current inflationary pressures.

REFERENCES FOR THE DAVIS COMMENT

Ball, Laurence, and Sandeep Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity*, Spring, 337–81; and comments and discussion, 382–405.

Blanchard, Olivier. 2016. "The Phillips Curve: Back to the 60s?" *American Economic Review* 106, no. 5: 31–34.

Carpenter, Dick M., Lisa Knepper, Angela C. Erickson, and John K. Ross. 2012. *License to Work: A National Study of Burdens from Occupational Licensing,* Arlington, Va.: Institute for Social Justice. https://ij.org/wp-content/uploads/2015/04/licensetowork1.pdf.

Davis, Steven J., and John C. Haltiwanger, 1998. "Measuring Gross Worker and Job Flows." In *Labor Statistics Measurement Issues*, edited by John Haltiwanger, Marilyn Manser, and Robert Topel. Chicago: University of Chicago Press.

———. 2015. "Labor Market Fluidity and Economic Performance." In *Economic Policy Symposium Proceedings: Re-Evaluating Labor Market Dynamics.* Jackson Hole, Wyo.: Federal Reserve Bank of Kansas City.

Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda. 2007. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." In *NBER Macroeconomics Annual 2006,* edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford. Cambridge, Mass.: MIT Press.

Decker, Ryan, John Haltiwanger, Ron S. Jarmin and Javier Miranda, 2014. "The Secular Decline in Business Dynamism in the U.S." Working paper. http://econweb.umd.edu/~haltiwan/DHJM_6_2_2014.pdf.

Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics* 11, no. 2: 1–39.

Hyatt, Henry R., and James R. Spletzer. 2017. "The Recent Decline of Single Quarter Jobs." *Labour Economics* 46: 166–76.

Johnson, Janna E., and Morris M. Kleiner. 2017. “Is Occupational Licensing a Barrier to Interstate Migration?” NBER Working Paper 24107. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w24107.pdf.

Kleiner, Morris M., and Alan B. Krueger. 2013. “Analyzing the Extent and Influence of Occupational Licensing on the Labor Market.” *Journal of Labor Economics* 31, no. 2: S173–S202.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak, 2014. “Declining Migration within the U.S.: The Role of the Labor Market.” NBER Working Paper 02138. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w20065.pdf.

White House. 2015. “Occupational Licensing: A Framework for Policy Makers.” Report prepared by Council of Economic Advisers, Department of the Treasury, and Department of Labor. https://obamawhitehouse.archives.gov/sites/default/files/docs/licensing_report_final_nonembargo.pdf.

COMMENT BY

GIORGIO E. PRIMICERI The goal of this paper by Richard Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin is to estimate the natural rate of unemployment, or u_t^*, in the postwar period. The authors combine two measurement approaches, one based on detailed data on flows into and out of unemployment for many demographic groups, and the other on the traditional Phillips curve relationship and data on aggregate unemployment, inflation, and inflation expectations. The paper provides three main takeaways. First, u_t^* is estimated to be about 4 percent in 2018. Second, u_t^* appears to have been trending down since the late 1980s. Third, this downward trend is due to the secular decline in the separation rate, which in turn was caused by the increased labor force attachment of females, the decline in job destruction and reallocation intensity, and the dual aging in the labor market of workers and firms. Overall, this paper is an impressive piece of work, with crucial policy implications. The most obvious of them is that the current low level of unemployment is roughly sustainable in terms of inflation, given that it is similar to the estimated natural rate, and the unemployment gap is thus close to zero.

My comments are organized around two main points. I first try to unpack the estimates of u_t^* presented by Crump and colleagues, to shed light on their essential drivers. I conclude that data on inflation expectations seem crucial for the measurement of u_t^* in the paper. In contrast, the detailed labor market flow data play a less central role for the measurement of u_t^*, although they are of course crucial for the interpretation of its secular trend. I then analyze the implications of this unpacking exercise for the New Keynesian Phillips curve, which seems in better shape than what some recent critics have suggested.

WHAT DRIVES CRUMP AND COLLEAGUES' ESTIMATE OF u_t^*? The authors' baseline estimate of u_t^* has three key features: (1) u_t^* has been trending down since the 1980s; (2) it is roughly equal to 4 percent in 2018; and (3) the uncertainty around its path is sizable but not overwhelming. My objective here is to understand what ingredients of the authors' complex empirical model are essential to these findings. To this end, I present a battery of u_t^* estimates, obtained using a sequence of progressively more complex models, the last of which corresponds to the authors' baseline model.

Model 1. The starting point of my unpacking exercise is the traditional backward-looking Phillips curve,

$$\pi_t - \pi^* = \gamma(\pi_{t-1} - \pi^*) - \kappa(u_t - u_t^*) + s_t. \tag{1}$$

According to this "Triangle model" (Gordon 1977, 2013), deviations of inflation, π_t, from its target value, π^*, depend on an inertial term, a "demand factor" represented by the gap between unemployment, u_t, and its natural rate, u_t^*, and a supply shock, s_t. In the estimation, u_t^* is treated as an unobservable variable and modeled as a random walk process, to capture the idea that its movements are very persistent. Like Crump and colleagues, I assume that the unemployment gap and the supply shock follow a first- and second-order autoregressive (AR(2), AR(1)) process respectively, although these two assumptions are not crucial for the results. The top left panel of my figure 1 presents the implied estimate of u_t^*, which is quite different from the authors' baseline estimate: u_t^* is relatively stable over time, it is roughly equal to 6 percent in 2018, and the uncertainty around its path is large.

Model 2. I augment model 1 with all the ingredients of Crump and colleagues' baseline setup, adding these components one at a time to understand their specific roles. The first step in this direction consists of turning the backward-looking Phillips curve of equation 1 into a more modern, New Keynesian, forward-looking Phillips curve, based on sticky wages and indexation to past and steady-state inflation:

$$\pi_t - \pi^* = \gamma(\pi_{t-1} - \pi^*) - \kappa E_t \sum\nolimits_{j=0}^{\infty} \beta^j \left(u_{t+j} - u_{t+j}^*\right) + s_t. \tag{2}$$

The main new feature of equation 2 is that inflation depends on the expected present discounted value of the future unemployment gaps, and

Figure 1. Estimates of u_t^* Based on Models 1 , 2, 3, and 4, 1960–2019

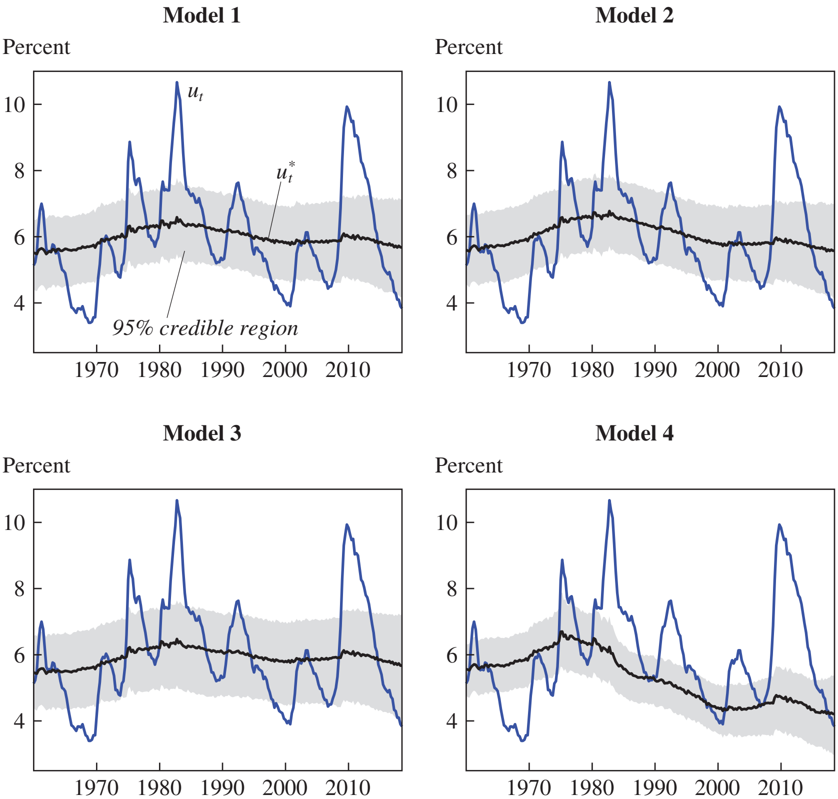

Sources: U.S. Bureau of Labor Statistics; author's calculations.

not just on the current level of this gap, as in equation 1. However, this change has a relatively minor impact on the measurement of u_t^*, as evident from comparing the top right and top left panels of my figure 1.

Model 3. The next step in the direction of Crump and colleagues' model is to introduce time variation in the inflation target. More precisely, π^* in equation 2 is replaced by π_t^*, which is modeled as a random walk to capture the low-frequency, hump-shaped behavior of inflation during the 1970s and 1980s. The role of this modification, however, is also relatively marginal, resulting again in nearly unchanged estimates of u_t^* relative to models 1 and 2 (the bottom left panel of my figure 1).

Model 4. Moving on, I now estimate the model by also using data on short- and long-term inflation expectations. As with Crump and colleagues' estimate, these data are equated to the short- and long-term forecasts of inflation implied by the model, up to a measurement error. These survey data help pin down the level of the time-varying inflation target, π_t^*, among other things. Notice that the implied estimate of u_t^*, presented in the bottom right panel of my figure 1, is now quite different from the previous ones: u_t^* clearly trends down starting in the 1980s, it is close to 4 percent in 2018, and its uncertainty is lower. Overall, the path of u_t^* is quite similar to the authors' baseline estimate, suggesting that the use of data on inflation expectations is a crucial component of their empirical model.

Model 5. The final ingredient of Crump and colleagues' model is the assumption that u_t^* cannot permanently deviate from the secular trend of the unemployment rate, $\bar{u}_t$. In particular, the authors rewrite u_t^* as

$$u_t^* = (u_t^* - \bar{u}_t) + \bar{u}_t,$$

where $\bar{u}_t$ is measured independently using disaggregated labor market flow data, and the term in parentheses—the distance between natural and secular unemployment—is assumed to follow an AR(1) process. I estimate this model using the authors' exact measure of $\bar{u}_t$, and I present the implied estimate of u_t^* in my figure 2. By construction, this estimate is identical to the authors' "inflation-only" estimate of u_t^*.

It is important to notice that the evolution of u_t^* shown in my figure 2 is similar overall to the path of u_t^* displayed in the bottom right panel of my figure 1. A possible interpretation of this finding is that disaggregated labor market flow data are not terribly useful to estimate u_t^*, because they do not drastically change our view about the time-series behavior of this variable. This interpretation, however, would probably be too literal and a bit naive. A more compelling view is that this consistency result—the fact that aggregate data on unemployment, inflation, and inflation expectations deliver estimates of natural unemployment in line with its secular trend—is remarkable, and provides an important external validation of the traditional Phillips curve framework.

IMPLICATIONS FOR THE NEW KEYNESIAN PHILLIPS CURVE The New Keynesian Phillips curve has recently been criticized because inflation fell little relative to the increase in unemployment during the Great Recession, the so-called missing disinflation phenomenon (Hall 2011). However—as stressed by Olivier Coibion and Yuriy Gorodnichenko (2015); Marco Del Negro, Marc Giannoni, and Frank Schorfheide (2016); and Carvalho and others (2017)—this somewhat puzzling behavior of inflation can be explained by the fact that inflation expectations were well anchored during the same

Figure 2. Estimate of u_t^* Based on Model 5, 1960–2019

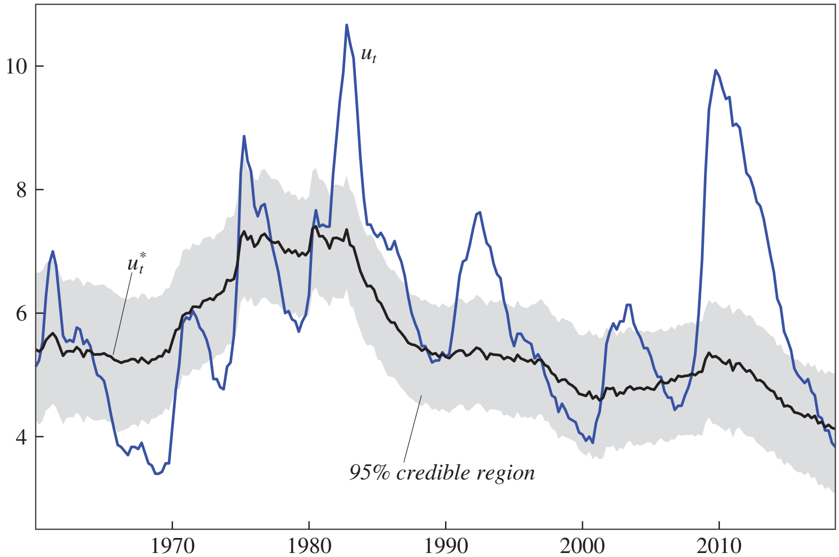

Sources: U.S. Bureau of Labor Statistics; author's calculations.

period. By this logic, the explicit use of inflation survey data should robustify inference in the context of the New Keynesian Phillips curve, and improve the estimation of u_t^*, which is exactly what the previous empirical results show.

To understand why these survey data play such a crucial role, let π_t^* replace π^* in the New Keynesian Phillips curve given in equation 2, and rewrite this equation as

$$
\begin{aligned}
(3) \qquad & \pi_t - \gamma\pi_{t-1} - (1-\gamma)\pi_t^* - \beta E_t\left[\pi_{t+1} - \gamma\pi_t - (1-\gamma)\pi_{t+1}^*\right] \\
& = -\kappa(u_t - u_t^*) + \tilde{s}_t,
\end{aligned}
$$

to make explicit the dependence of inflation on its expected future value.[1] The use of data on short- and long-term inflation expectations makes the variables $E_t\pi_{t+1}$ and π_t^* observable, up to some measurement error. As a consequence, it becomes easier for the econometrician to isolate the relationship between the left-hand side and the right-hand side of

1. The term $\tilde{s}_t$ is equal to $(1 - \beta\rho_s)s_t$, where ρ_s is the autocorrelation of the supply shock.

Figure 3. Post-1990 Estimate of u_t^* Based on Model 4 and Data on Inflation and Long-Term Inflation Expectations, 1990–2019

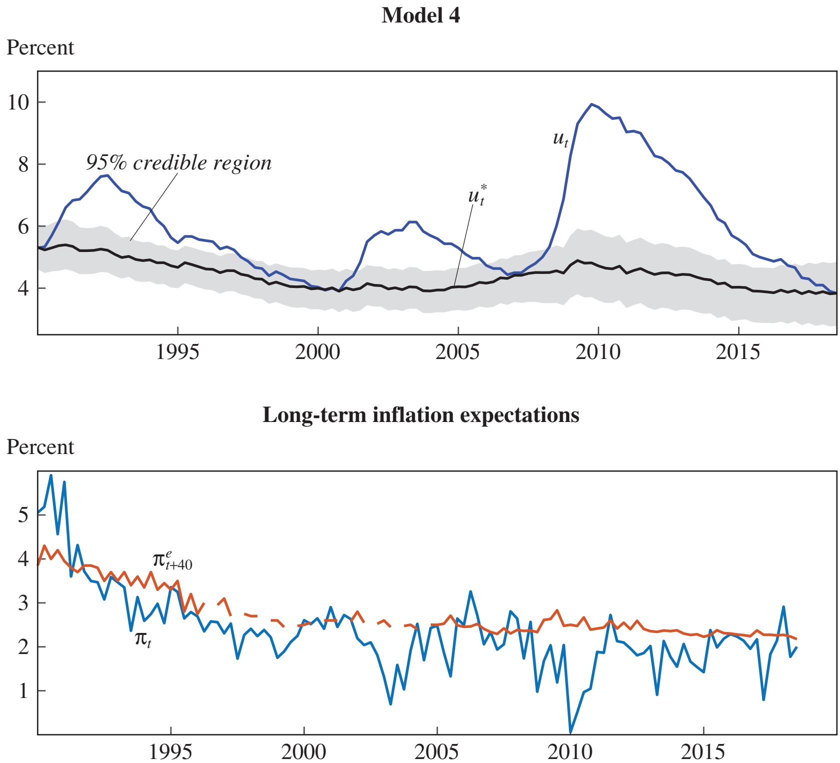

Sources: U.S. Bureau of Labor Statistics; Survey of Professional Forecasters; author's calculations.

equation 3, and to more precisely estimate the slope coefficients κ and the natural rate u_t^*. Furthermore, when using data on inflation expectations, inference about κ becomes also surprisingly stable over time. To illustrate this point, the top panel of my figure 3 presents the evolution of u_t^* according to model 4, when this model is estimated using only post-1990 data. Notice that the implied path of u_t^* is remarkably similar to the one based on the full-sample estimation of the same model, plotted in the lower right panel of my figure 1.

It is important to realize, however, that this powerful role of inflation expectation data in the estimation of the Phillips curve also comes with a disadvantage. The cost is the sensitivity of the estimate of u_t^* to the exact measurement of inflation expectations, which is notoriously difficult. For

Figure 4. Post-1990 Estimate of u_t^* Based on Model 4, Obtained Using Actual Inflation Data and Artificial Data on Inflation Expectations, 1990–2019

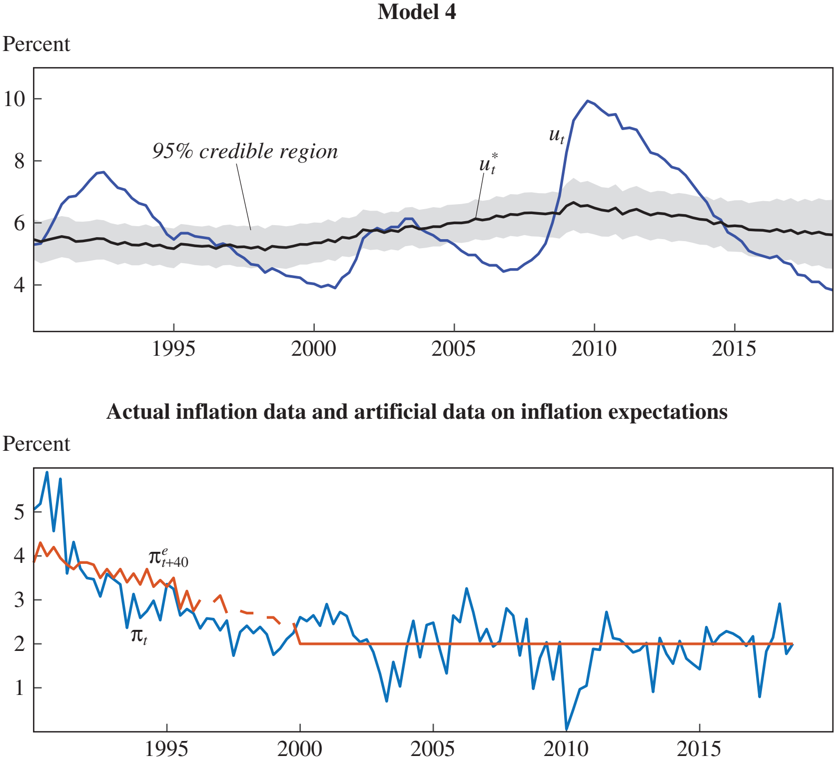

Sources: U.S. Bureau of Labor Statistics; Survey of Professional Forecasters; author's calculations.

example, the bottom panel of my figure 3 plots the time series of long-term inflation expectations used for the estimation of model 4 (and Crump and colleagues' baseline model). These expectations appear to be "upward biased," because agents systematically expect inflation in the long run to be higher than current inflation. Given that these survey data effectively pin down the level of π_t^*, this implies that actual inflation is almost always below target. In turn, this explains why the unemployment gap is almost always positive in the top panel of my figure 3.

To illustrate the sensitivity of u_t^* to the measurement of inflation expectations, the top panel of my figure 4 plots the implied evolution of u_t^* when I reestimate the model on post-1990 data, using a modified long-term inflation expectation series that is artificially set to 2 percent after 2000 (the bottom panel of my figure 4). The figure makes clear that the estimate of u_t^* would shift considerably in this counterfactual scenario, implying a substantially negative unemployment gap in 2018. This finding suggests that the Phillips curve is still relatively flat, despite being estimated quite robustly due to the explicit use of data on inflation expectations.

REFERENCES FOR THE PRIMICERI COMMENT

Carvalho, Carlos, Stefano Eusepi, Emmanuel Moench, and Bruce Preston. 2017. "Anchored Inflation Expectations." Working paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3018198.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7, no. 1: 197–232.

Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics* 7, no. 1: 168–96.

Gordon, Robert J. 1977. "Can the Inflation of the 1970s Be Explained?" *Brookings Papers on Economic Activity*, no. 1: 253–77.

———. 2013. "The Phillips Curve Is Alive and Well: Inflation and the NAIRU during the Slow Recovery." NBER Working Paper 19390. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w19390.pdf.

Hall, Robert E. 2011. "The Long Slump." *American Economic Review* 101, no. 2: 431–69.

GENERAL DISCUSSION David Romer started the discussion by complimenting the paper and the commenters. He expressed the view, however, that the paper might have overemphasized u^* (the natural rate of unemployment) relative to $\bar{u}$ (the secular trend of unemployment).

One reason to use either of these two concepts is to forecast inflation, for which he argued u^* would clearly be the preferred measure. However, it may not matter much which variable is used because inflation does not seem to be reacting to labor market tightness very strongly.

Another reason to use either variable would be as a baseline for where the economy should move over the medium term. In this context, because u^* eventually reverts to $\bar{u}$, then $\bar{u}$ may be the better measure to focus on. For example, during the Great Recession, policies such as unemployment insurance raised u^* somewhat, while $\bar{u}$, a secular construct, remained mostly unchanged. Because monetary policymakers, for example, target economic outcomes over a longer horizon, $\bar{u}$ would have been a better metric to use. Similarly, based on the paper's estimates, unemployment is currently at about its natural rate (u^*), but is about 0.5 percentage point below its secular trend ($\overline{u}$). If u^* is likely to rise to $\overline{u}$ over the next several years, then again $\overline{u}$ would be the more relevant variable for policymakers to consider.

In addition, Romer agreed with the presenters about the relevance of inflation expectations in the modern version of the Phillips curve, but noted that there are two components to inflation expectations in modern models. One piece is essentially the present value of future output gaps (or minus the departures of unemployment from the natural rate). This piece has not been a very good predictor of inflation and behaved similarly during the Great Recession and during the Volcker-era recession. The second piece, however, is the long-term expectations of inflation. Romer expressed his belief that this piece is crucial to "saving" the Phillips curve, but that it lacks any substantive microeconomic foundations. Macroeconomic models get around this problem by assuming that workers index their wages to some abstract concept of expected inflation, but this assumption has little basis in reality. As a result, he argued that economists are still a long way from explaining why and how inflation expectations actually matter for inflation, and therefore from understanding the behavior of inflation.

Robert Hall complimented the paper on its litany of data on entrance and exit rates from unemployment. He noted though that the paper's conceptual reliance on the Phillips curve goes against data showing it to be an unreliable predictor of inflation. He cited a paper by James Stock and Mark Watson, as well as Phillips curves used in work by Robert Gordon as better examples.[1] These formulations of the Phillips curve study changes in unemployment relative to changes in inflation rather than the relationship between the level of unemployment and inflation and do a better job of fitting the data. He also referenced a paper he wrote with Thomas Sargent on the failure of the modern Phillips curve due to its misunderstanding of Milton Friedman's original ideas about the construct.[2] He noted specifically that Sargent had been critical of it since the 1960s.

1. James H. Stock and Mark W. Watson, "Modeling Inflation after the Crisis," *Proceedings: Economic Policy Symposium—Jackson Hole, Federal Reserve Bank of Kansas City*, 173–220, https://www.kansascityfed.org/publicat/sympos/2010/Stock-Watson_final.pdf; Robert J. Gordon, "The Phillips Curve Is Alive and Well: Inflation and the NAIRU during the Slow Recovery," NBER Working Paper 19390 (Cambridge, Mass.: National Bureau of Economic Research, 2013), https://www.nber.org/papers/w19390.pdf.

2. Robert E. Hall and Thomas J. Sargent, "Short-Run and Long-Run Effects of Milton Friedman's Presidential Address," *Journal of Economic Perspectives* 32, no. 1 (2018): 121–34, https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.32.1.121.

Frederic Mishkin noted that good monetary policy that acts to stabilize inflation will induce a bias toward estimating a flatter slope of the Phillips curve than would otherwise be the case. When inflation goes up, a responsible central bank should raise interest rates, which will increase unemployment. This effect would partially offset the negative slope between inflation and unemployment contained in the structural construct of the Phillips curve. He argued that the flattening of the Phillips curve's slope since the 1980s can be explained by this phenomenon and may bias estimates of the curve. Although the structural Phillips curve may still be active, monetary policy could be masking its empirical identification. He argued this was an important issue to address, particularly because a policymaker such as National Economic Council director Larry Kudlow seemed to be implying that the Phillips curve was dead and therefore the economy could sustain low interest rates without suffering from inflation.[3] Ignoring these types of issues could result in bad policy results, he argued.

Robert Gordon complimented the paper and its commenters, and he stressed the importance for policy of the paper's estimate of the natural rate of unemployment. In particular, if the paper's estimates were correct, then the economy has been at full employment for the past year and has been fluctuating around it for even longer, which explains why there has been so little inflation. An alternative explanation for low inflation would be that the natural rate is higher, maybe 5 percent, but that the Phillips curve is very flat. The distinction between the two narratives is important for monetary policy, he argued, because pushing the unemployment rate down to 3 percent would lead to inflation in the first narrative but not the second. The paper comes out on the side of the first argument. To understand which argument is correct, Gordon referenced the period between 2009 and 2015, when unemployment increased but there was little change in inflation, a point that would argue in favor of the second narrative, that the Phillips curve is very flat and there is little actual evidence for what the natural rate of unemployment is.

Picking up on comments from Romer and from Giorgio Primiceri in his comment, Gordon commented on the increasing importance of inflation expectations in the modern Phillips curve. He noted that the relevance of expectations may have come about as a result of the Federal Reserve's explicit inflation target or as some other process that stabilized inflation.

3. Larry Kudlow and Stephen Moore, "Who's Afraid of Higher Wages?" *Wall Street Journal*, March 12, 2018, https://www.wsj.com/articles/whos-afraid-of-higher-wages-1520897733.

Including inflation expectations in the Phillips curve causes an increase in estimates of the natural rate of unemployment in the 1970s and 1980s, from about 6 percent to around 8 percent, something that did not occur in Gordon's own work, which did not model the Phillips curve using inflation expectations. He noted that this difference was due to the fact that his own models explicitly include shocks specific to the time, rather than inflation expectations more broadly.

Finally, Gordon commented on interesting trends shown in the paper's data on labor market entry and exit from unemployment. He noted the mounting pile of evidence on the declining rate of entry of workers into unemployment and the declining importance of layoffs. As evidence of this, he noted the similarity in the levels of new claims for unemployment insurance today compared with the 1960s. The levels are comparable today and are still quoted in thousands, even though the size of the labor force has doubled since the 1960s. This point further indicates how dramatically entry into unemployment has declined.

Wendy Edelberg compared estimates of the natural rate of unemployment in the paper with estimates from the Congressional Budget Office (CBO). She inquired as to how the authors thought about labor force participation when making their estimates, wondering if differences in the treatment of this variable could drive divergence in the paper's estimate of u^* and the CBO's estimates. Specifically, though the CBO considers the inflationary pressures arising from the Phillips curve when estimating the natural rate, they weight the construct referred to as $\overline{u}$ in the paper more heavily because it takes into account longer-term structural trends like participation. Instead, the CBO had started to focus on the concept of an employment gap rather than an unemployment gap. Estimates of $\overline{u}$ in the paper are indeed similar to CBO's estimate of the natural rate of unemployment, showing values of about 4.6 percent in recent data. She asked how the authors calibrated their measures of u^* to measures of slack in the labor market. She noted that based on rough calculations, the labor market slack implied by the paper shows a similar unemployment gap in 2016 to the CBO's estimate of labor market slack after incorporating weakness in labor force participation relative to its potential. She stressed the importance of thinking about participation when estimating u^*.

Justin Wolfers expressed his frustration in explaining the Phillips curve to undergraduate economics students. Despite its centrality to macroeconomics, there are scant empirical illustrations of its existence. He joked that Paul Krugman's economic textbook solves this problem by only showing data on the Phillips curve from 1955 to 1968; N. Gregory Mankiw's textbook only shows data from 1961 to 1968, and Ben Bernanke's textbook shows no data at all (Bernanke agreed and laughed in response).

Wolfers also inquired about the confidence intervals around u^* in the paper. He interpreted the paper to estimate a very small probability that u^* was below 5½ percent in 1973. At Brookings in 1973, Arthur Okun presented a paper estimating the natural rate of unemployment to be between 4 and 5 percent, which the current paper views as statistically unlikely, indicating the uncertainty about estimates of u^* through time.[4]

Gerald Cohen argued that despite the centrality of inflation expectations in many modern versions of the Phillips curve, most economists are often overly precise about estimates of inflation expectations. He referenced comments from Primiceri, for example, noting that many survey respondents expect inflation to be about 2.5 percent rather than the Federal Reserve's 2 percent target over time. Cohen referenced a 2015 *Brookings Paper* studying inflation expectations in New Zealand, which showed that they were well above the reality of realized inflation as well as of the central bank's inflation target.[5] Therefore, estimating the Phillips curve with some precise estimate of inflation expectations, possibly based on forecasts from the Blue-Chip Economic Indicators survey filled out by informed professionals, may be detached from inflation expectations in reality.

Laurence Meyer gave his interpretation of the paper, notably that the Phillips curve is operative but not particularly relevant, or, as he put it, "The Phillips curve is alive, but who the hell cares?" In particular, the Phillips curve may exist, but it takes massive declines in unemployment to get small increases in inflation because the curve is so flat, meaning that in practice there is very little trade-off. Meyer noted that the paper goes on to argue that inflation expectations are instead central to the inflation process, and to get inflation, central bankers have to change the public's expectations of future inflation. Meyer argued, however, that this is an incredibly difficult task and represents a major crisis in central banking; central banks cannot run tight labor markets or change expectations to get inflation up to target. Both Japan and Europe, for example, are still a long way from

4. Arthur M. Okun, "Upward Mobility in a High-Pressure Economy," *Brookings Papers on Economic Activity*, no. 1 (1973): 207–52, https://www.brookings.edu/wp-content/uploads/1973/01/1973a_bpea_okun_fellner_greenspan.pdf. In this paper, Okun actually says that "at the present time, the controversial range for the target unemployment rate extends from 4 to 5 percent."

5. Saten Kumar, Hassan Afrouzi, Olivier Coibion, and Yuriy Gorodnichenko, "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand," *Brookings Papers on Economic Activity*, Fall 2015, 151–208.

achieving their inflation targets on a sustained basis, despite significant forward guidance designed to increase inflation expectations as well as accommodative monetary policy driving tight labor markets. Under this interpretation, estimates of the natural rate of unemployment essentially seem irrelevant, given the flatness of the Phillips curve and how immovable inflation expectations are. The United States, however, appears to be in an advantageous position in this regard, given that inflation is close to target and inflation expectations seem well anchored. As Meyer put it, "We are just beautiful. But everybody else is screwed."

John Haltiwanger commented on the paper's dual approach to measure u^* using labor market entry and exit from unemployment and estimates of the Phillips curve. The former in particular involves measures of labor market tightness, which is difficult to measure. Early versions of labor market tightness involved vacancies over unemployment, but more recent versions involve modeling a broader job matching process. New models might measure slack as effective vacancies over effective searches, where effective searches include people out of the labor force as well those in unemployment. He cited Robert Hall and Sam Schulhofer-Wohl's recent paper showing that the probability of getting a job for someone who wants one but has fallen out of the labor force is relatively high and is also highly cyclical.[6] Measures such as these, which better measure labor market slack, might be better than u^* or $\overline{u}$ and might be used to provide a better fit of the Phillips curve.

Haltiwanger also commented on the suggested relationship in the paper between declining inflows into unemployment and declining labor market dynamism and reallocation. Sympathetic with this idea, he remarked that there might be a decline in the volatility of idiosyncratic labor market shocks that would be consistent with this trend, perhaps driven by changes in the composition of firms in the economy. However, the evidence has been pushing against this idea. Notably, new labor productivity data at the firm level in the United States shows an increasing dispersion of productivity growth across firms, a trend consistent with *rising* labor market frictions or wedges—a worrying sign.

Richard Cooper, echoing similar comments from Justin Wolfers, remarked on how different the paper's estimates of u^* were from historical estimates. Specifically, the *1962 Economic Report of the President*, which was written

6. Robert E. Hall and Sam Schulhofer-Wohl, "Measuring Job-Finding Rates and Matching Efficiency with Heterogeneous Job-Seekers," *American Economic Journal: Macroeconomics* 10, no. 1 (2018): 1–32, https://web.stanford.edu/~rehall/JobFindingRates.

by economic luminaries such as James Tobin and Arthur Okun, targeted an interim unemployment rate of 4 percent, an estimate below the 95 percent confidence interval in the paper.[7] He asked where historical estimates went wrong, questioning the paper's estimates of u^* for the early 1960s. Robert Gordon countered that the evidence was in favor of the paper's estimates because inflation was high in the 1960s as a result of underestimating u^* at the time.

Martin Baily reviewed generally the concept of u^* as originally conceived by Milton Friedman, noting that the paper treated u^* in the same spirit.[8] Namely, under Friedman's original framework, the unemployment rate falls below its natural rate accidentally. This is due to the fact that monetary policy affects inflation, which in turn drives down real wages. Because of downward nominal wage rigidity, firms tend to lay off workers instead of cutting their wages, hence leading to temporary surges in unemployment where the unemployment rate deviates from its natural rate absent wage rigidity. However, this basic framework may have shifted, primarily due to structural changes in the labor market institutions governing employment specifically related to weaker worker bargaining power and the decline of unions; making the wage setting process different could change the degree to which nominal wage rigidity binds. In addition, Baily reinforced the point made by David Romer: that by saying the Phillips curve has been "rescued" by including inflation expectations without a real theory behind how inflation expectations form, the authors' argument in favor of the Phillips curve leaves something to be desired.

George Perry echoed the point made by Martin Baily about labor unions. In the first 30 years after World War II, strong labor unions that could negotiate wages and salaries made wages more reactive to labor market conditions. This could give way to wage price spirals and inflation, making inflation more sensitive to labor market conditions and economic shocks more generally. During this period, the Phillips curve could be estimated well. The decline of labor unions since then, however, has instead led to a structural change in this relationship, Perry argued, and trying to find alternative specifications for the Phillips curve ignores these structural changes. Namely, the change came relatively abruptly after the 1980s, driven

7. Council of Economic Advisers, *1962 Economic Report of the President* (Washington: U.S. Government Printing Office, 1962), 8, https://fraser.stlouisfed.org/title/45/item/8133.

8. Milton Friedman, "The Role of Monetary Policy," speech given at Eightieth Annual Meeting of the American Economic Association and then published in the *American Economic Review* 58, no. 1 (1968): 1–17.

by foreign competition that weakened the bargaining power of domestic labor unions. Foreign competition also put general downward pressure on wages and prices. Trying to fit the same model to these two distinct periods invites policy mistakes. For example, current monetary policy—namely, inflation targeting—assumes that some modest amount of inflation can maximize employment regardless of the institutional setup of wage and price setting.

Stefano Eusepi thanked the discussants for their observations, promising to incorporate them into the paper. First, he addressed critiques of the Phillips curve by making a distinction about how exactly it has flattened in recent years. Namely, in their model, with the same size of the output gap, they can explain both the low inflation since 2008 and the high inflation of the 1970s. The reason for this is that the Phillips curve has only flattened in one dimension: its relationship with the current unemployment gap. It has *not* flattened in relation to the present discounted value of future unemployment gaps. This creates a situation where even if the unemployment gap is large, but it is expected to revert back to zero, then the Phillips curve is quite flat and inflation does not react. However, if the gap is expected to be permanent, then the slope of the Phillips curve can be quite large and inflation does materialize. In this way, the Phillips curve is not actually flat.

Eusepi also commented on the formation of inflation expectations. Although their current paper does not model inflation expectations, he noted other papers the authors had written on the behavior of long-term inflation expectations that vary endogenously with monetary policy.[9] These models can explain inflation expectations based on survey data, both in the 1970s and today.

James Stock referenced comments from Frederic Mishkin and Laurence Meyer on the flat Phillips curve. He asked about how inflation expectations are developed and their relationship with monetary policy. In this context, if the unemployment rate were expected to increase by a large amount, then monetary policy should respond by a large amount. If this is the case, then expectations should not move as much.

Ayşegül Şahin discussed the difference between u^* and $\bar{u}$, noting the usefulness of both measures. She referred to $\bar{u}$ as an anchor of u^*, and she cited the Great Recession as a prominent example. During the Great

9. Carlos Carvalho, Stefano Eusepi, Emanuel Moench, and Bruce Preston, "Anchored Inflation Expectations," working paper, https://centridiricerca.unicatt.it/complexity-Eusepi_AnchoredDraft_Final_Final.pdf.

Recession, when unemployment was high, $\bar{u}$ was trending down, due to secular factors. Therefore, recognizing this downward trend in $\bar{u}$ would have increased policymakers' confidence that there was substantial slack in the labor market.

In terms of bargaining power, Şahin referenced work by Andrew Glover and Jacob Short showing that aging is affecting the labor share of income because older workers extract less of the profit generated by firms.[10] The aging of the population is consistent with the decline in bargaining power, which also explains the decline in the inflow rate to unemployment discussed in the paper.

10. Andrew Glover and Jacob Short, "Demographic Origins of the Decline in Labor's Share," working paper, 2018, https://www.cemfi.es/ftp/pdf/papers/Seminar/shortj-gs2018_lsage.pdf.