MARCO DEL NEGRO
Federal Reserve Bank of New York

GIORGIO E. PRIMICERI
Northwestern University

MICHELE LENZA
European Central Bank

ANDREA TAMBALOTTI
Federal Reserve Bank of New York

What's Up with the Phillips Curve?

ABSTRACT The business cycle is alive and well, and real variables respond to it more or less as they always did. Witness the Great Recession. Inflation, in contrast, has gone quiescent. This paper studies the sources of this disconnect using vector autoregressions and an estimated dynamic stochastic general equilibrium model. It finds that the disconnect is due primarily to the muted reaction of inflation to cost pressures, regardless of how they are measured—a flat aggregate supply curve. A shift in policy toward more forceful inflation stabilization also appears to have played some role by reducing the impact of demand shocks on the real economy. The evidence rules out stories centered around changes in the structure of the labor market or in how we should measure its tightness.

The recent history of inflation and unemployment is a puzzle. The unemployment rate has gone from below 5 percent in 2006–2007 to 10 percent at the end of 2009, and back down below 4 percent in 2018–2019. These fluctuations are as wide as any experienced by the US economy in the postwar period. In contrast, inflation has been as stable as

Conflict of Interest Disclosure: Marco Del Negro is vice president in the Macroeconomic and Monetary Studies Function of the Research and Statistics Group of the Federal Reserve Bank of New York and coeditor of the *Journal of Applied Econometrics*; Michele Lenza is head of section for the Monetary Policy Research Division of the European Central Bank, the Forecasting and Business Cycle group leader and research coordinator for the European Central Bank, and a member of the Forecast Steering Committee for the European Central Bank; Giorgio E. Primiceri is a professor of economics at Northwestern University and coeditor of the *American Economic Journal: Macroeconomics*; Andrea Tambalotti is vice president of the Federal Reserve Bank of New York. Beyond these affiliations, the authors did not receive financial support from any firm or person for this paper or from any firm or person with a financial or political interest in this paper. The views expressed in this paper are those of the authors and do not necessarily reflect those of Northwestern University, the *American Economic Journal*, the *Journal of Applied Econometrics*, the Federal Reserve Bank of New York, or the European Central Bank.

ever, with core inflation almost always between 1 and 2.5 percent, except for short bouts below 1 percent in the darkest hours of the Great Recession.

Much has been written about this disconnect between inflation and unemployment. In the early phase of the expansion, when unemployment was close to 10 percent and inflation barely dipped below 1 percent, the search was for the "missing deflation" (Hall 2011; Ball and Mazumder 2011; Coibion and Gorodnichenko 2015; Del Negro, Giannoni, and Schorfheide 2015; Lindé and Trabandt 2019). More recently, with unemployment below 5 percent for almost four years and inflation persistently under 2 percent, attention has turned to factors that may explain why inflation is not coming back (Powell 2019; Yellen 2019). Beyond this recent episode, a reduction in the cyclical correlation between inflation and real activity has been evident at least since the 1990s (Atkeson and Ohanian 2001; Stock and Watson 2007, 2008, 2019; Zhang, Chan, and Cross forthcoming). The literature, which we review in more detail below, has focused on four main classes of explanations for this puzzle: (1) mismeasurement of either inflation or economic slack; (2) a flatter wage Phillips curve; (3) a flatter price Phillips curve; and (4) a flatter aggregate demand relationship, induced by an improvement in the ability of policy to stabilize inflation.

This paper tries to distinguish among these four competing hypotheses using a variety of time series methods. We find overwhelming evidence in favor of a flatter price Phillips curve. Some of the evidence is also consistent with a change in policy that has led to a flatter aggregate demand relationship.

The analysis starts by illustrating a set of empirical facts regarding the dynamics of inflation in relation to other macroeconomic variables, using a vector autoregression (VAR). Many of these facts are already known, but the dynamic, multivariate perspective offered by the VAR makes it easier to consider them jointly, enhancing our ability to point toward promising explanations for the phenomenon of interest. First, goods inflation has become much less sensitive to the business cycle since 1990, consistent with most of the literature on the severe illness of the Phillips curve. Second, this is true to a lesser extent for nominal wage inflation: the wage Phillips curve is in better health than its price counterpart, as also found by Coibion, Gorodnichenko, and Koustas (2013), Coibion and Gorodnichenko (2015), Galí and Gambetti (2019), Hooper, Mishkin, and Sufi (2020), and Rognlie (2019). Third, there is little change before and after 1990 in the business cycle dynamics of the most popular indicators of inflationary pressures relative to each other, especially when compared to the pronounced reduction in the responsiveness of inflation.

These indicators include measures of labor market activity, such as the unemployment rate and its deviations from the natural rate, hours, the employment-to-population ratio, and unit labor costs, as well as broader notions of resource utilization, such as GDP and its deviation from measures of potential. Fourth, the decline in the sensitivity of inflation to the business cycle is heterogenous across goods and services. In particular, Stock and Watson (2019) document that the relationship between cyclical unemployment and inflation has changed very little over time for certain categories of goods and services that are better measured and less exposed to international competition. Our VAR analysis produces results that are consistent with these findings, but we do not report them here since they are not necessary to draw our main conclusions.[1]

Together, the first three facts listed above lead us to reject mismeasurement of economic slack, as well as a significant flattening of the wage Phillips curve, as the main cause of the emergence of the inflation–real activity disconnect since 1990. We draw this conclusion because those two explanations are inconsistent with the small change in the dynamic relationship between the most common indicators of cost pressures before and after 1990, at the same time as inflation became much more stable.[2] A further implication of this finding is that we can focus the rest of the investigation on the bivariate relationship between inflation and real activity, without having to take a stance on the most appropriate measure of the latter. Any indicator commonly used in the literature will do.[3]

1. Some recent papers have also explored the behavior of inflation across geographic areas in the United States and across countries (Fitzgerald and Nicolini 2014; Mavroeidis, Plagborg-Møller, and Stock 2014; McLeay and Tenreyro 2019; Hooper, Mishkin, and Sufi 2020; Geerolf 2020). They generally find that the correlation between inflation and unemployment in the cross section is stronger and more stable than in the time series. Hazell and others (2020) provide a guide to translate this cross-sectional evidence into the time series elasticity that is of interest to most of the macroeconomics literature. Using data on US states, they recover a flat Phillips curve once the estimates are properly rescaled, with some evidence of a further reduction in the slope coefficient after 1990. Fully reconciling this evidence across geographies and exchange rate regimes with our conclusion requires more work.

2. We cannot rule out that all the indicators of cost pressures that we include in our analysis have become a poorer proxy for the "true" real marginal costs that drive firms' pricing decisions after 1990. However, it is unlikely that an unobserved change in the dynamics of those costs could have left almost no trace on the joint dynamics of all those indicators.

3. In practice, we focus primarily on the relationship between inflation and unemployment, but we continue to do so in the context of a VAR that also includes other macroeconomic variables. We focus on unemployment because it is arguably the most straightforward and widely discussed measure of the health of the real economy, as well as the most commonly used independent variable in Phillips curve regressions.

This conclusion marks the boundary to which we can push the VAR for purely descriptive purposes. As illustrated in a recent influential paper by McLeay and Tenreyro (2019), the observed relationship between inflation and real activity is the result of the interaction between aggregate demand and supply. The latter captures the *positive* relationship between inflation and real activity, usually associated with the price Phillips curve. Higher inflation is connected with higher marginal costs, which in turn tend to rise in expansions, when demand is high, the labor market is tight, and wages are under pressure. On the contrary, the economy's aggregate demand captures a *negative* relationship between inflation and real activity, which reflects the endogenous response of monetary policy to inflationary pressures. When inflation is high, the central bank tightens monetary policy, thus slowing the real economy. Therefore, the observed cyclical disconnect between inflation and real activity might be the result of either a flat Phillips curve—the slope hypothesis—or a flat aggregate demand, generated by a forceful response of monetary policy to inflation. In the limit in which the central bank pursues perfect inflation stability, inflation is observed to be insensitive to the cycle, regardless of the slope of the Phillips curve. We refer to this second possible explanation for the stability of inflation as the policy hypothesis.

Distinguishing between these two hypotheses is a classic identification problem that requires economic assumptions that were not needed for the data description exercise in the first part of the paper. We impose those restrictions through two complementary approaches, a structural VAR (SVAR) and an estimated dynamic stochastic general equilibrium (DSGE) model. In the SVAR, we identify cyclical fluctuations that can be plausibly attributed to a demand disturbance. To do so, we follow Gilchrist and Zakrajšek (2012) and use their data on the excess bond premium (EBP) to identify a financial shock that propagates through the economy like a typical demand shock, by depressing both real activity and inflation. We choose this shock as a proxy for demand disturbances because it accounts for a significant fraction of the business cycle fluctuations behind the facts described in the first part of the paper. In response to this demand shock, inflation barely reacts in the post-1990 sample, while it used to fall significantly before 1990. This result indicates that the slope of the aggregate supply relationship must have fallen since 1990. Intuitively, the demand shock acts as an instrument for cost pressures in the Phillips curve, identifying its slope. If real activity declines in response to an EBP shock, as it clearly does in both samples, and this lowers cost pressures (i.e., if the instrument is not weak), a muted response of inflation implies a flat Phillips curve.

Although this evidence clearly points in the direction of a very flat aggregate supply curve after 1990, it does not rule out the possibility that monetary policy might have also contributed to the observed stability of inflation. In fact, the main implication of this hypothesis is that monetary policy should lean more heavily against inflation by limiting the impact of demand shocks on the real fluctuations. At the limit in which inflation is perfectly stable, demand shocks should leave no footprint on the real variables. The impulse responses to the EBP shock are far from implying no reaction of the real variables to the demand disturbance, as we would expect if the stability of inflation were due to monetary policy, although they do point to some stabilization, at least in the short run.

The SVAR evidence that we just described helps narrow down the relative contribution of the slope and policy hypotheses for the stability of inflation. To provide an even sharper quantification of their respective roles, we turn to an estimated DSGE model. This exercise is subject to the typical trade-off associated with imposing tighter economic restrictions on the data. On the one hand, we can map the slope and policy hypotheses directly onto parameters of the model that we can estimate on data before and after 1990. On the other hand, the results of this exercise hinge on the entire structure of the model, rather than on a looser set of identifying assumptions as in the SVAR. Therefore, they stand or fall together with the observer's beliefs about the validity of that structure as a representation of the data. To support the case in favor of the model's validity for our purposes, we show that it reproduces the facts generated by the reduced-form VAR used for data description in the first part of the paper. In terms of the two hypotheses, the DSGE estimates point in the direction of a much flatter Phillips curve in the second sample. If we assume that the slope of the Phillips curve is the only parameter that changes after 1990, the estimated model still broadly reproduces the empirical facts. If we only allow policy to change, the estimated model falls short.

Together, the results of the SVAR and DSGE produce two conclusions. First, there is strong support for the slope hypothesis: the slope of the Phillips curve has fallen very substantially after 1990, although it has not gone all the way to zero. Second, there is also some evidence that the policy hypothesis—or other structural changes contributing to a flatter aggregate demand curve—might have contributed to reduce the cyclical sensitivity of inflation, but this evidence is weaker.

The rest of the paper proceeds as follows. The remainder of this section reviews the literature. Section I describes the VAR that we use for data description, the results of which are then described in section II. Section III introduces a stylized aggregate demand and supply framework inspired by McLeay and Tenreyro (2019), which illustrates the fundamental identification problem underlying the interpretation of the observed relationship between inflation and real activity. This model also guides the interpretation of the impulse responses to the EBP shock presented in section IV. Section V revisits the same facts presented in section II from the perspective of an estimated DSGE model and uses that model to further explore the relative contribution of the slope and policy hypothesis to the observed stability of inflation. Section VI elaborates on some policy implications of our main findings, before offering some concluding remarks in section VII.

THE LITERATURE The literature has explored four main classes of explanations for the reduction in the observed correlation between inflation and real activity.

The first set of explanations is related to mismeasurement of either inflation or economic slack. In the inflation dimension, much of the debate has focused on the role of new products and quality adjustment in the construction of price indexes and in the measurement of output and productivity, especially following the introduction of technologies with a very visible impact on everyday life, such as the internet and smart phones (Moulton 2018). This branch of the literature has also explored the recent emergence of online retailing as a source of transformation in firms' pricing practices (Cavallo 2018; Goolsbee and Klenow 2018). By focusing on cyclical fluctuations, our analysis mostly bypasses these considerations, since they primarily pertain to the level of measured inflation. In addition, the inflation–real activity disconnect predates the potential effect of information technology on price mismeasurement, further reducing the potential explanatory power of this hypothesis for our phenomenon of interest.

On the real activity front, the definition and measurement of economic slack have been the subject of a vast literature. Abraham, Haltiwanger, and Rendell (2020) in this volume offer a very recent example. Much of this work has focused on the estimation of potential output and the natural rate of unemployment as reference points to assess the cyclical position of the economy and its influence on inflation. Crump and others (2019) provide a comprehensive discussion of this literature, which features many prominent contributions in the Brookings papers. Our results on the stability of the co-movement of various measures of cost pressures should reduce the weight put on explanations of the inflation disconnect based on the idea that any one measure of slack might be less representative of underlying inflationary pressures after 1990, for instance, due to changes in the relationship between measured unemployment and the overall health of the labor market (Stock 2011; Gordon 2013; Hong and others 2018; Ball and Mazumder 2019).

A second set of explanations for the emergence of the inflation–real activity disconnect focuses on a flatter wage Phillips curve and, more in general, on structural transformations in the labor market and its connection with the goods market (Daly and Hobijn 2014; Stansbury and Summers 2020; Faccini and Melosi 2020). Taken together, our results suggest that whatever structural change might have occurred in the labor market, it is unlikely to be the leading cause of inflation stability. In a recent Brookings paper, Crump and others (2019) capture some of these structural transformations in their Phillips curve estimates by anchoring the inference on the natural rate of unemployment to disaggregated data on labor market flows. This procedure produces a model of inflation that accounts for its dynamics throughout the sample. However, doing so requires a low slope coefficient, as stressed by Davis (2019) and Primiceri (2019) in their discussions.

A third set of explanations focuses on the role of policy in delivering stable inflation. The idea is that a stronger response of monetary policy to inflation flattens the aggregate demand curve, weakening the connection between inflation and real fluctuations, even if the aggregate supply relationship is unchanged (Fitzgerald and Nicolini 2014; Barnichon and Mesters 2019a; Hooper, Mishkin, and Sufi 2020; McLeay and Tenreyro 2019; Kareken and Solow 1963; Goldfeld and Blinder 1972). An implication of this hypothesis is that the Phillips curve is hibernating: a stronger correlation between inflation and business cycles would reemerge if monetary policy reacted less to inflation, as it probably did before the 1990s. Consistent with this view, we also find that monetary policy played some role in stabilizing inflation over the cycle. However, our evidence suggests that policy did not entirely succeed in eliminating demand-driven real fluctuations, implying that it cannot be the dominant driver of the inflation–real activity disconnect. This result, however, leaves open the possibility that changes in monetary (and perhaps fiscal) policy were behind the low-frequency fluctuations in inflation related to its slow rise between the mid-1960s and 1979 and its return to 2 percent over the subsequent two decades, as argued for instance by Primiceri (2006).

Related to this policy hypothesis is the large literature on the role of inflation expectations and their anchoring (Orphanides and Williams 2005; Bernanke 2007; Stock 2011; Blanchard, Cerutti, and Summers 2015; Blanchard 2016; Ball and Mazumder 2019; Carvalho and others 2019; Jorgensen and Lansing 2019; Barnichon and Mesters 2019b). Empirically, expectations are now less volatile than they were before 1990, as we also find in our VAR. However, this observation does not establish that changes in their formation, perhaps in response to shifts in the conduct or communication of policy, represent an autonomous source of inflation stability. Rather, our evidence suggests that the behavior of inflation expectations mostly reflects the inflation stability induced by the flattening of the aggregate supply curve, instead of being its primary source.

In conclusion, our results support a fourth set of explanations that attribute the inflation–real activity disconnect to forces that reduce the response of goods prices to the cost pressures faced by firms, lowering the slope of the structural price Phillips curve. This is the slope hypothesis, which takes several variants. The most prominent is the one that attributes a reduction in the response of prices to marginal costs to the increased relevance of global supply chains, heightened international competition, and other effects of globalization (Sbordone 2007; Auer and Fischer 2010; Peach, Rich, and Linder 2013; Tallman and Zaman 2017; Forbes 2019a, 2019b; Obstfeld 2020; Forbes, Gagnon, and Collins 2020). In a similar vein, Rubbo (2020) points to changes in the network structure of the US production sector.[4]

Compared to this literature that concentrates on estimating the slope of the Phillips curve as a summary statistic of the connection between inflation and real activity, our VAR approach explores more broadly the dynamic relationships among real and nominal variables to draw conclusions on the mechanisms that drive them and how they have changed since 1990. Another advantage of our approach is that it focuses on business cycle dynamics, abstracting from lower-frequency trends and other developments that might be less informative about the Phillips curve relationship. As a result, we do not address the reasons why inflation has been stubbornly below most central banks' targets for the better part of the last decade.

I. Methodology and Data

The objective of this paper is to shed light on the possible causes of the widely acknowledged attenuation of the response of inflation to labor market slack over the past three decades. This section illustrates the methodology and the data that we use to document this fact and its

4. Afrouzi and Yang (2020) connect changes in the conduct of monetary policy directly to the slope of the Phillips curve, straddling the two strands of the literature that we just discussed. They present a model in which rationally inattentive price setters respond less to aggregate shocks when monetary policy is committed to inflation stabilization.

relationship with the behavior of a broad set of macroeconomic variables, whose joint dynamics might help to discriminate among alternative explanations.

I.A. Methodology

To study macroeconomic dynamics and their change over time, we begin by adopting the following vector autoregression (VAR) model:

$$y_t = c + B_1 y_{t-1} + \ldots + B_p y_{t-p} + u_t . \tag{1}$$

In this expression, y_t is an $n \times 1$ vector of macroeconomic variables, which is modeled as a function of its own past values, a constant term, and an $n \times 1$ vector of forecast errors (u_t) with covariance matrix Σ. The reduced-form shocks (u_t) are a linear combination of n orthogonal structural disturbances (ε_t), which we write as $u_t = \Gamma\varepsilon_t$.

VARs are flexible multivariate time series models, which provide a rich account of the complex forms of autocorrelation and cross-correlation that are typical of macroeconomic variables. To synthesize and illustrate these relationships, we study the dynamic response of the variables of interest to a typical unemployment shock. We identify this "U shock" using a Cholesky scheme with unemployment ordered first. This shock corresponds to the linear combination of structural disturbances that drives the one-step-ahead forecast error in unemployment. The impulse responses to this shock tell us how the system evolves in the future if next quarter's unemployment rate turns out to be higher than expected.

The specific combination of shocks responsible for the one-step-ahead forecast error in unemployment accounts for the bulk of business cycle fluctuations in real activity, but it ignores other sources of macroeconomic variation. As a consequence, this part of the analysis has little to say about the substantial share of inflation variability that is independent of the U shocks. Instead, it focuses on the component of inflation that responds to real business cycle impulses, which is the essence of the Phillips curve. Moreover, this approach does not attempt to pin down the precise identity of the structural disturbances driving fluctuations, as in more typical structural VARs. Doing so would require additional economic assumptions, which are not necessary to document many of the empirical facts regarding the attenuated response of inflation to labor market slack that have been discussed in the literature. The advantage of this methodology is that we can illustrate these facts in the context of a unified, dynamic, multivariate statistical framework, without imposing any theoretical restriction.

A limitation of our method is that, without economic restrictions, the facts that it uncovers can be mapped onto more than one hypothesis on the sources of the increased disconnect between inflation and real activity. Therefore, sections IV and V will also explore more economically demanding approaches—SVARs and DSGE models—in order to further pinpoint the source of the empirical observations illustrated below.

The VAR approach to data description that we pursue in this section also has some advantages over the much more popular direct estimation of a Phillips curve, defined as the relationship between inflation (and its lags) and some measure of slack. First, such an inflation equation is embedded in the VAR, which therefore encompasses the single-equation approaches as long as the same variables used in them are included in y_t. Second, embedding such an equation into a multivariate framework explicitly recognizes the challenging identification problem of distinguishing a Phillips curve, which represents the economy's aggregate supply, from its aggregate demand. We illustrate this challenge in the context of a stylized New Keynesian model in section III. Third, looking at the response of inflation and the other variables to the combination of shocks responsible for the one-step-ahead forecast error of unemployment produces more flexible measures of economic slack than those based on specific indicators of potential output or natural unemployment—two notoriously elusive concepts. Fourth, we do not need explicit measures, or a model, of inflation expectations, as long as the variables included in the VAR approximately span the information set used by agents to form those expectations. This aspect of the analysis is especially important, since inflation expectations are a key ingredient in most formulations of the Phillips curve. At the same time, given their relevance, we also consider VAR specifications that include a direct measure of expectations in the vector y_t.

I.B. Data

What variables should the VAR include to provide a comprehensive view of the forces shaping the connection between inflation and the labor market? To answer this question, we refer to an intuitive description of that connection, which is embedded in most formal and informal frameworks built around a price or wage Phillips curve.[5] When firms try to hire more workers to satisfy higher demand for their output, wages tend to rise. Given labor productivity, this increase in wages is associated with higher

5. A simple example of such a framework is the New Keynesian model with sticky wages and prices in chapter 6 of Galí (2015).

marginal costs and inflation. Therefore, a tight labor market and rising wages, costs, and inflation tend to occur together in response to demand shocks.

To characterize these channels, data on inflation and unemployment are not enough. In addition, we need measures of wages, labor productivity, and firms' costs to capture the intermediate steps of the transmission. Therefore, we propose a baseline VAR that includes eight macroeconomic variables: (1) *unemployment*, measured by the civilian unemployment rate; (2) *natural unemployment*, measured by the CBO estimate; (3) *core inflation*, measured by the annualized quarterly growth rate of the personal consumption expenditures (PCE) price index, excluding food and energy; (4) *inflation*, measured by the annualized quarterly growth rate of the GDP deflator; (5) *GDP*, measured by the logarithm of per capita real GDP; (6) *hours*, measured by the logarithm of per capita hours worked in the total economy; (7) *wage inflation*, measured by the annualized quarterly growth rate of the average hourly earnings of production and nonsupervisory employees (PNSE); and (8) the *labor share*, measured by the logarithm of the share of labor compensation in GDP.

Besides unemployment, this VAR includes a block of variables referring to the total economy: GDP, hours, the labor share, and the GDP deflator. These variables can be combined to compute a measure of hourly nominal compensation in the total economy. The growth rate of this measure of nominal wages closely tracks compensation per hour in the nonfarm business sector, a commonly used indicator of labor costs. The problem with both these series is that they are extremely volatile at high frequencies, obscuring the underlying wage dynamics over the business cycle. For this reason, the baseline VAR also includes the PNSE wage inflation series. This series only covers about 80 percent of private industries, but it is substantially less noisy than the more comprehensive ones that we mentioned. Finally, in addition to GDP inflation, our model also includes core PCE inflation, given its importance as a gauge of underlying inflationary pressures.

We estimate this eight-variable VAR over two nonoverlapping samples, to investigate possible changes in the typical co-movement pattern of these variables in response to the U shock described in section I.A. The first sample ranges from 1964:Q2 to 1989:Q4, and the second from 1989:Q1 to 2019:Q3.[6] The analysis starts in 1964:Q2, when the PNSE wage inflation

6. The first four observations are used as initial conditions, since our VAR has four lags. Therefore, effectively, the estimation starts in 1965:Q2 in the first sample and in 1990:Q1 in the second one.

series first becomes available. The date at which we split the sample is the result of a compromise. On the one hand, there is some evidence that macroeconomic dynamics might have changed around 1984, after the first phase of the so-called Volcker disinflation. On the other hand, a simple inspection of the data suggests that inflation has been most stable starting around the mid-1990s, after the opportunistic disinflation engineered by the Federal Reserve under Alan Greenspan following the 1990–1991 recession. As a compromise between these two alternatives, we split the sample in 1990. This choice also has the advantage of creating two samples of fairly similar lengths. Section A of the online appendix shows that this choice is immaterial for the results.

The data are quarterly and the VAR includes four lags. It is estimated with Bayesian methods and a standard Minnesota prior, given the relatively high number of variables and short sample sizes. The tightness of the prior is chosen based on the data-driven method described in Giannone, Lenza, and Primiceri (2015).

II. Facts

This section documents a number of known and new facts concerning the relationship between unemployment, inflation, and some other key macroeconomic variables. These empirical findings lead us to two main conclusions: the attenuation of inflation fluctuations over the business cycle before and after 1990 is striking, and, in comparison, the co-movement of all real variables and indicators of firms' cost pressures has been remarkably stable.

II.A. Fact 1: Unemployment and Price Inflation

We begin by presenting the impulse responses of unemployment and inflation to a U shock in the two samples. Figure 1 shows that inflation falls significantly as unemployment rises in the first sample. This finding suggests that the U shock is characterized by a strong demand component, which explains why traditional Phillips curves estimated over this sample have a negative slope. In the second sample, instead, unemployment increases by a roughly similar amount, but the responses of both inflation measures are muted. In fact, the response of core inflation is statistically indistinguishable from zero throughout the horizon, while that of GDP inflation is a bit more negative and borderline significant after about one year. In addition, the very flat response of natural unemployment

Figure 1. Impulse Responses to an Unemployment Shock in the Baseline VAR

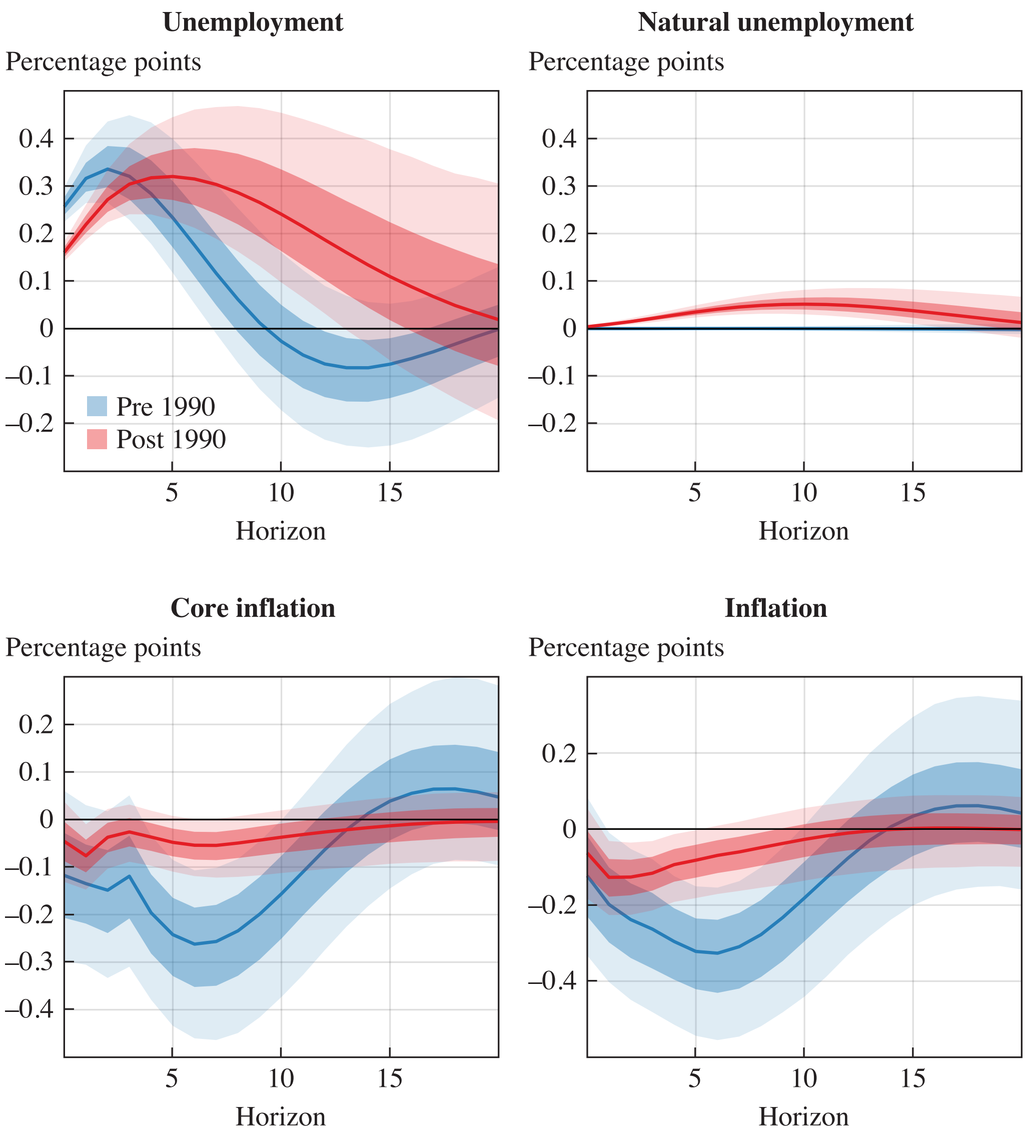

Source: Authors' calculations.

Note: The impulse responses are from the baseline VAR described in section I.B. The shock is identified using a Cholesky strategy, with unemployment ordered first. The solid lines are posterior medians, while the shaded areas correspond to 68 percent and 95 percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:Q2 to 1989:Q4 and from 1989:Q1 to 2019:Q3, respectively.

indicates that the shock only captures business cycle variation. Therefore, looking at the reaction of unemployment or the unemployment gap to this shock would produce identical results.

Online appendix A shows that the responses of figure 1 are nearly identical to those to a typical business cycle shock, obtained as the linear combination of structural disturbances that drives the largest share of unemployment variation at business cycle frequencies, as in Giannone, Lenza, and Reichlin (2019) and Angeletos, Collard, and Dellas (2019). This result suggests that the combination of shocks associated with the one-step-ahead forecast error in unemployment and the one responsible for the bulk of business cycle fluctuations are virtually the same. Our finding also casts some doubt on the interpretation of the muted response of inflation to business cycle shocks proposed by Angeletos, Collard, and Dellas (2019), since they do not explain why such response was much more vigorous before 1990.

The response of unemployment to a U shock in figure 1 is more persistent in the second sample. This feature of economic fluctuations is evident even from the raw data, and it is consistent with the lengthening of expansions in the last thirty years. However, this change in the profile of unemployment fluctuations does not play much of a role in accounting for the attenuated response of inflation in the second sample. We illustrate this point with an exercise that forces the response of unemployment to be identical in the two samples. Specifically, we compute the responses of all variables as the difference between their forecast conditional on a specific path of unemployment and their unconditional forecast, following the methodology of Bańbura, Giannone, and Lenza (2015). As this common path in both samples, we choose the median response of unemployment to a U shock in the first sample.[7] Figure 2 plots the dynamics of all the VAR variables in this conditional forecast exercise. As in figure 1, the response of inflation in the second sample is much attenuated, although it now remains negative. Online appendix A shows that this change in inflation dynamics over the two samples is not limited to the two measures of inflation included in the baseline VAR, but it extends to a number of other commonly used inflation series.

7. This conditional forecast approach recovers the most likely sequence of shocks to guarantee that unemployment follows a given path. In this respect, it has a slightly different interpretation relative to the impulse responses, because the latter are based on a single shock perturbing the economy at horizon zero.

Figure 2. Impulse Responses Conditional on a Path for Unemployment in the Baseline VAR

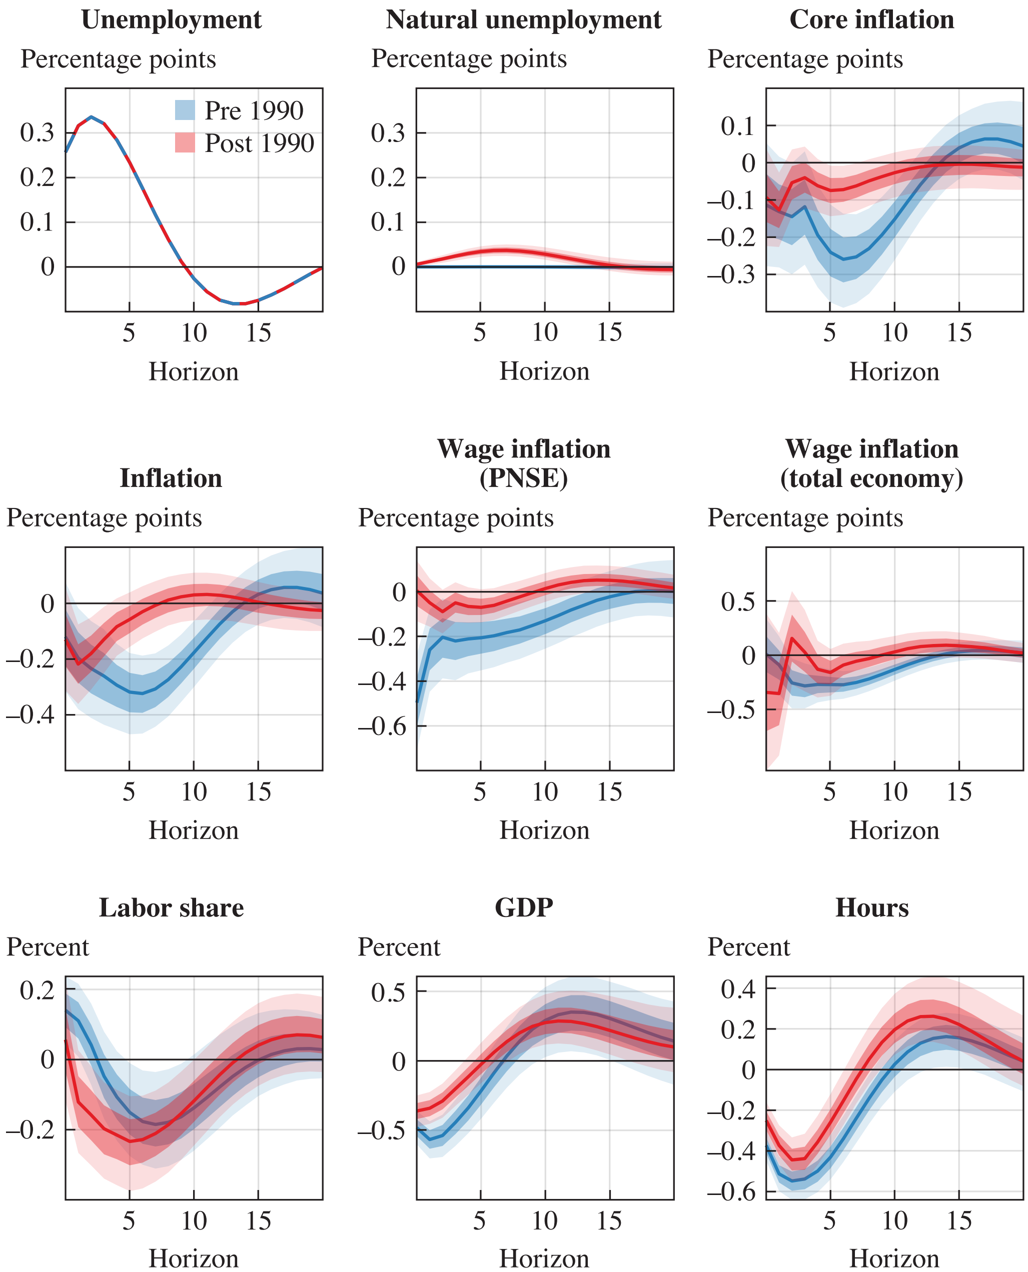

Source: Authors' calculations.

Note: These responses are computed by applying the methodology described in section II.A to the baseline VAR of section I.B. The solid lines are posterior medians, while the shaded areas correspond to 68 percent and 95 percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:Q2 to 1989:Q4 and from 1989:Q1 to 2019:Q3, respectively.

These findings can be summarized into a first key stylized fact: the sensitivity of goods price inflation to labor market slack has decreased dramatically after 1990. This fact provides a complementary, more dynamic, characterization of many findings in the literature regarding the stability of inflation. Interpreting this fact is the main task of the rest of the paper.

II.B. Fact 2: Unemployment and Wage Inflation

A substantial body of recent work finds that the connection of wage inflation to labor market slack remains stronger than that of goods inflation (Galí and Gambetti 2019; Hooper, Mishkin, and Sufi 2020; Rognlie 2019). This section presents VAR results broadly consistent with these findings. The second row of figure 2 plots the response of two measures of nominal wage inflation using the conditional forecast approach described in the previous subsection. The first measure (PNSE, middle graph) is the one used directly for the estimation of the baseline VAR. The second (total economy, right graph) is that implied by the data on the labor share, hours, output, and GDP inflation.[8] The reaction of the PNSE series is attenuated in the post-1990 period, while the response of the total economy measure shows more similarities in the two samples. Therefore, we take the balance of the evidence as consistent with the view that the connection between wage inflation and unemployment remains alive, although it is weaker in the more recent period. As shown in online appendix A the sensitivity of wage inflation to unemployment after 1990 is even stronger when wages are measured with the employment cost index (ECI), which is arguably a better measure of the cyclicality of wages than the ones used in this section. Unfortunately, the ECI is only available starting in 1980, preventing a full comparison of its behavior pre- and post-1990. We summarize these findings in the form of a second stylized fact: the sensitivity of nominal wage inflation to labor market slack has diminished after 1990, but less than that of price inflation.

One implication of this fact is that explanations of the unemployment-inflation disconnect involving a much reduced responsiveness of wage inflation to labor market slack are not very plausible. For example, a popular narrative attributes the stability of inflation during the Great Recession to the existence of downward nominal wage rigidities: if firms are reluctant or unable to lower nominal wages, their marginal costs should remain

8. In logs, the labor share (ls) is defined as the sum of the nominal wage (w) and hours (h), minus real GDP (gdp) and the price level (p), or $ls \equiv w + h - gdp - p$. Therefore, this measure of the (log) nominal wage is constructed as $w = ls - h + gdp + p$.

relatively high, putting upward pressure on prices and inflation. Such a story, however, would imply a substantial weakening of the co-movement between unemployment and wage inflation, which seems at odds with the data. In addition, as we demonstrate in online appendix A, this co-movement is approximately equally strong after 1990 regardless of whether we include or exclude the Great Recession period.

II.C. Fact 3: Unemployment and Unit Labor Costs

One obvious difficulty in interpreting the evidence on the connection between nominal wage inflation and unemployment presented in the previous section is that it also partly reflects a weaker response of goods inflation. Mechanically, nominal wage inflation is the sum of real wage inflation and goods inflation. Therefore, the former will appear less responsive to the cycle if the latter reacts less given the dynamics of the real wage. A more helpful approach to evaluate the implications of wage dynamics for inflation, therefore, is to study more direct measures of how wages contribute to firms' marginal costs.[9] The most popular proxy for aggregate real marginal costs are unit labor costs (or, equivalently, the labor share). With constant returns in production, (log) unit labor costs are proportional to (log) marginal costs. Under more general assumptions, this proportionality no longer holds, but unit labor costs are likely to remain a more accurate gauge of the cost pressures faced by firms than nominal wage inflation.[10]

Figure 2 shows that the forecast of the labor share conditional on the usual path of unemployment is very similar in the two samples. This observation leads to the third stylized fact: the co-movement of unemployment and the labor share over the business cycle is stable over time. This fact supports and further refines the view according to which labor market developments are unlikely to be the main source of the change in inflation dynamics over the past thirty years. The claim is not that labor market

9. In pricing problems based on cost minimization, firms' marginal costs are the key driver of their pricing decisions. As a result, the evolution of aggregate marginal costs is the fundamental source of inflation in a large class of models with nominal rigidities. These models include those with staggered price setting, as in Calvo (1983) and Taylor (1980), as well as those with sticky information or rational inattention, as in Mankiw and Reis (2002) or Mackowiak and Wiederholt (2009).

10. With constant returns to scale production, a firm's log marginal cost is proportional to its log unit labor cost, defined as $ulc = w - (gdp - h)$. With homogeneous factor markets, marginal cost is equalized across firms, so that the *aggregate* log labor share ($ls = w + h - gdp - p$) is proportional to the average real marginal cost.

dynamics have not changed since 1990. More narrowly, the statement is that, whatever those changes might have been, they did not have a significant impact on the dynamics of firms' marginal costs, at least as seen through the lens of a proxy such as the labor share. The next section adds one further dimension to this claim by showing that the same can be said of other well-known aggregate proxies for firms' cost pressures.

II.D. Fact 4: Unemployment and Other Measures of Real Activity

The previous subsection argued that unit labor costs are likely to be the most informative variable on the extent to which cost pressures originating in the labor market are transmitted to goods prices.[11] Next, we show that the dynamics of many other variables used in the literature to capture real sources of inflationary pressure, from the labor market or otherwise, are also relatively stable over time. The third row of figure 2 reports the conditional forecasts of hours and output. These responses are essentially identical over time, implying a fourth stylized fact: the business cycle correlation among several indicators of real activity has not changed in the two samples. Online appendix A further shows that these empirical patterns also hold for the output gap and the employment-to-population ratio, when we add these variables to the baseline VAR.

The important conclusion that we draw from these results is that the severe illness of the reduced-form relationship between inflation and real activity cannot be cured by picking a different indicator of either labor or goods market slack among those commonly used in the literature. In fact, the remarkable stability in the dynamic relationships between all the real variables that we have considered suggests that the diagnosis of what ails inflation should be independent of one's view on the best proxy for underlying inflationary pressures.

II.E. Adding Interest Rates and Expected Inflation

In this subsection, we augment the model with data on the federal funds rate and on long-term inflation expectations from the survey of professional

11. A key implication of firms' cost minimization is that marginal cost is equalized across all inputs. As a result, marginal cost pressures—measured by comparing wages to labor productivity—provide a comprehensive view of the cost pressures faced by firms, even if the input whose direct cost is rising is not labor. The main difficulty in operationalizing this observation is the measurement of the marginal cost and marginal benefit of labor, i.e., the wage and the marginal product of labor. Available measures of wages and (average) labor productivity capture those marginal concepts only under restrictive assumptions.

forecasters (SPF).[12] The former was not included in the baseline VAR because it was at the zero lower bound (ZLB) for many years in the second sample. To avoid that period, this larger VAR is estimated excluding data after 2007. Figure 3 plots the conditional forecast of all the variables in the model. Compared to the baseline, the conditioning path of unemployment, which is, as usual, the median response of unemployment to the U shock in the pre-1990 sample, returns to zero faster, although its inverted *S* shape is otherwise similar. Moreover, the estimated responses are more uncertain in the second sample, since it is shorter by about twelve years. However, the main empirical facts documented so far are robust to these changes.

Focusing on the newly added variables, the response of the federal funds rate in the two samples has a similar shape, but it is less persistent after 1990. That of inflation expectations is more muted in the second sample, similar to inflation. At the same time, the gap between the two variables falls significantly in the first sample, while it is more stable in the second, just as is the case with inflation itself. This observation suggests that the reduction in the sensitivity of inflation to business cycles goes beyond what can be explained through the increased stability of long-run inflation expectations. The extent to which more-anchored expectations simply reflect the increased stability of inflation, as opposed to being one of its independent sources, remains an open question. We will return to this issue in section IV.

II.F. Summary of the Key Facts

The four stylized facts documented above lead us to two important conclusions, which crucially inform the rest of the analysis. First, the change in the business cycle dynamics of inflation before and after 1990 stands out compared to that of all the real variables that we have considered. Second, the co-movement of these real variables is remarkably similar before and after 1990. Together, these two observations suggest that we can focus the rest of the analysis on the bivariate relationship between inflation and real activity, with no need to be more specific on its measurement. As illustrated in section III, however, this significant narrowing of the scope of the inquiry is not sufficient to conclude that the anemic response of inflation to the cycle is due to a flattening of the structural Phillips curve. The reason is that a flattening of the aggregate demand relationship, perhaps

12. The data on long-term inflation expectations are constructed as in Clark and Doh (2014) and Del Negro and others (2017).

Figure 3. Impulse Responses Conditional on a Path for Unemployment in a Larger VAR

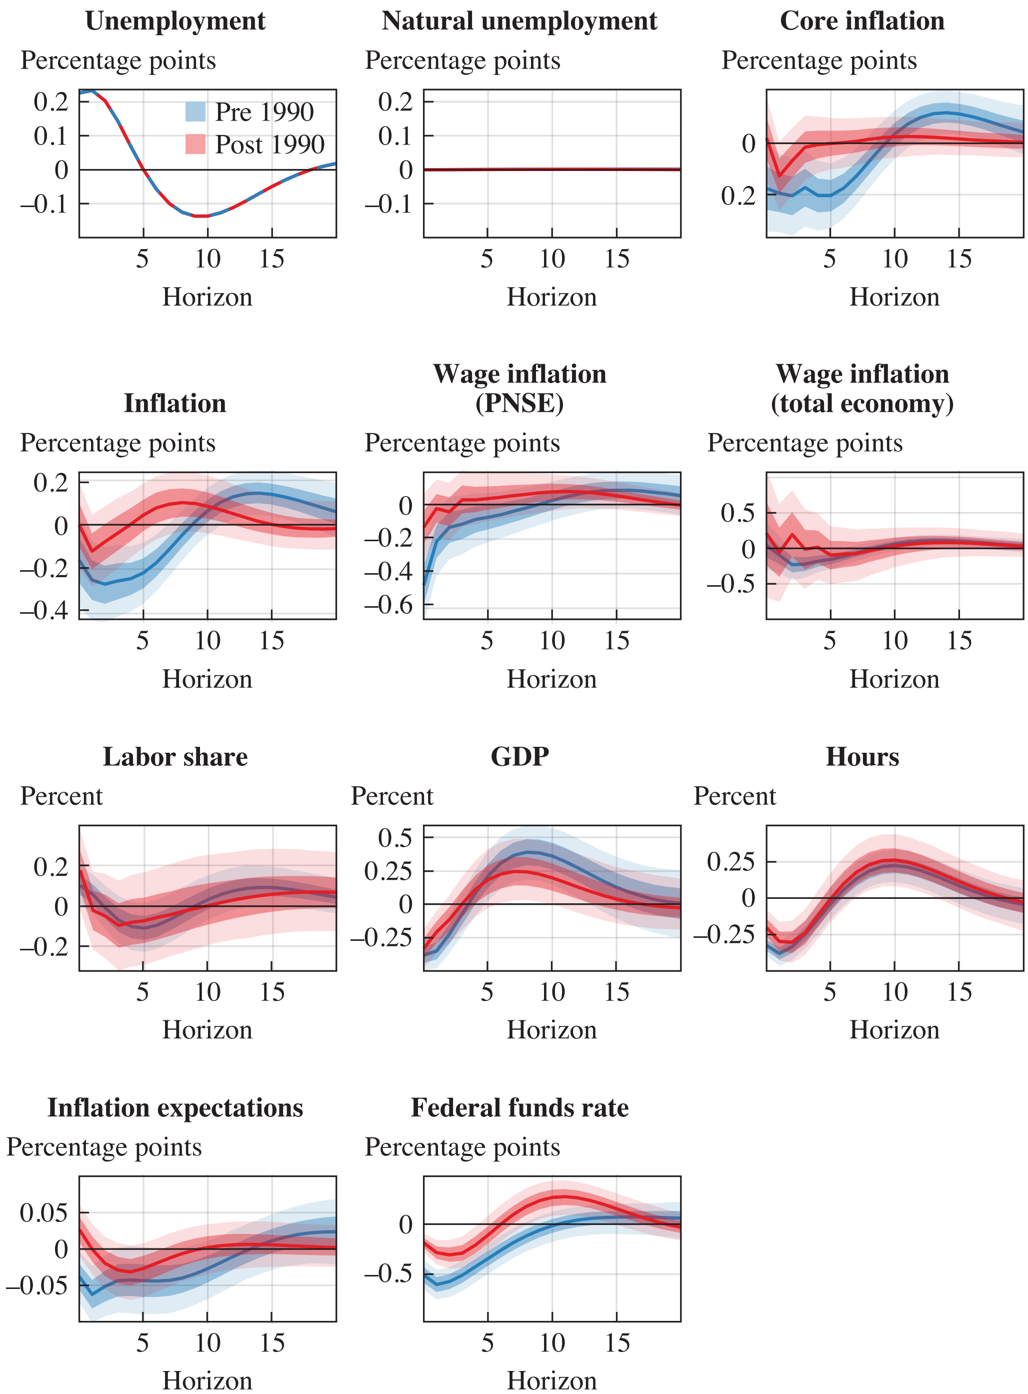

Source: Authors' calculations.

Note: These responses are computed by applying the methodology described in section II.A to the baseline VAR of section I.B, augmented with long-term inflation expectations and the federal funds rate. The solid lines are posterior medians, while the shaded areas correspond to 68 percent and 95 percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:Q2 to 1989:Q4 and from 1989:Q1 to 2019:Q3, respectively.

induced by a more forceful reaction of monetary policy to inflation, could in principle result in more stable inflation. Further distinguishing between these two possibilities requires putting more structure on the problem, as we will then do in sections IV and V.

III. Lessons from a Stylized Model

To aid in the interpretation of the empirical facts described in section II, we now introduce a stylized model of the joint determination of inflation and real activity. This model, which is directly inspired by McLeay and Tenreyro (2019), is based on the textbook New Keynesian framework of Woodford (2003) and Galí (2015). However, its implications for the nature of business cycles under alternative hypotheses regarding the possible sources of inflation stability are quite general, as we argue below. We use this simple model to make three essential points: (1) the empirical facts of section II are consistent with two possible explanations of the stability of inflation after 1990: either a reduction in the sensitivity of pricing decisions to marginal cost pressures or a change in the conduct of monetary policy; (2) the key implication that differs across these two hypotheses is that real activity is driven predominantly by demand-type shocks in the first case and supply-type (or cost-push) disturbances in the second; and (3) unfortunately, it is difficult to empirically verify which shocks—demand or supply—are prevalent in the post-1990 period based on the co-movement pattern between inflation and real activity, because the variation of inflation is minimal. Therefore, next we will introduce more information and structure to further sharpen our inference.

III.A. A Simple Model of Aggregate Demand and Supply

The stylized model we consider consists of the following three familiar equations:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa(x_t + s_t), \tag{2}$$

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - \delta_t), \text{ and} \tag{3}$$

$$i_t = E_t \pi_{t+1} + \psi_\delta \delta_t + \psi_\pi \pi_t, \tag{4}$$

where π_t represents price inflation, $x_t \equiv y_t - y_t^n$ is the output gap (defined as the log deviation of output from a measure of potential), i_t is the nominal interest rate, and s_t and δ_t are exogenous disturbances. In this formulation, $E_t\pi_{t+1}$ and $E_t x_{t+1}$ denote rational expectations of next period's inflation and output gap.[13]

Equation (2) is the model's structural Phillips curve, an aggregate supply relationship that maps higher output gaps into higher inflation. It is based on the dependence of inflation on firms' marginal costs—a fairly general feature of optimal pricing problems (Sbordone 2002)—in combination with some simplifying assumptions that make marginal costs proportional to the output gap. These simplifying assumptions, however, are not restrictive for our analysis, given that facts 3 and 4 in section II document a stable dynamic relationship between many real activity variables usually employed to measure slack, such as unemployment, hours, GDP, and unit labor costs. Therefore, given this stability, we do not need to take a strong stance on what x_t precisely represents in our stylized model and will simply refer to it generically as real activity, or the output gap. Finally, the supply, or cost-push shock (s_t) in equation (2) stems from fluctuations in desired markups, which explains why it is scaled by the slope κ.[14]

The other two equations constitute the demand block of the model. In particular, equation (3) is an Euler equation, or dynamic investment-savings (IS) equation, which connects the nominal interest rate to real activity. The strength of this negative relationship is governed by the parameter $\sigma > 0$. In addition, the equation is perturbed by the shock δ_t, which can be interpreted as capturing fluctuations in the Wicksellian natural rate of interest, due to technology or demand disturbances. We will refer to it as a demand shock, for short. Equation (4) is a simple interest rate rule that represents the response of the monetary policy authority to economic developments. This specification allows for a direct response of the policy rate to the IS shock—for reasons that will become clear shortly—and to inflation, where the Taylor principle requires $\psi_\pi > 0$. Adding a term in the

13. Lowercase letters denote logs, so that, for instance, $y_t \equiv \log Y_t$, where Y_t is the level of output. Y_t^n is natural output, the level of output that would be observed in the absence of nominal rigidities.

14. Under Calvo pricing, fluctuations in desired markups have the same effect on inflation as those in real marginal costs. Therefore, at the limit in which prices never change, the sensitivity of inflation to both real activity and desired markups captured by the parameter κ goes to zero and inflation becomes perfectly stable. More generally, we could allow for other sources of exogenous supply shocks, which might include, for instance, exogenous shifts in inflation expectations that are not fully captured by the rational expectations term $E_t\pi_{t+1}$. In the presence of such shocks, the variability of inflation is not zero even when $\kappa = 0$, but the qualitative implications of the model described below do not change.

output gap, as in Taylor (1993), or a monetary policy shock would not change the model's key qualitative implications.

Plugging the policy rule into the IS equation produces a negative relationship between inflation and real activity of the form

$$\pi_t = -\phi(x_t - E_t x_{t+1} - d_t),$$

where $\phi \equiv (\sigma\psi_\pi)^{-1} \geq 0$ and $d_t \equiv \sigma(\psi_\delta - 1)\delta_t$ is a simple rescaling of the demand shock. If the demand disturbance is observable, monetary policy can perfectly offset it by setting $\psi_\delta = 1$. More generally, demand shocks are likely to be transmitted to the economy at least partially, either because the monetary authority observes them with noise, or because it chooses not to react to them fully, or because it is prevented from doing so by the ZLB. All of these scenarios are captured in this model by $\psi_\delta < 1$, which implies some pass-through of these shocks into inflation.

The negative slope of this aggregate demand equation reflects the fact that monetary policy leans against inflation by raising the real interest rate, which in turns lowers real activity. This feature of aggregate demand does not depend on the exact specification of the interest rate rule, as long as the real interest rate responds positively to inflation. As shown by McLeay and Tenreyro (2019), this is also a feature of aggregate demand under optimal monetary policy. In this respect, our approach to modeling monetary policy and that of McLeay and Tenreyro (2019) are isomorphic, even if they derive the aggregate demand equation directly from an optimal policy problem, without relying on the IS equation. In comparison, our setup with an IS equation and a policy rule is more explicit about some potential sources of demand shocks, but its key implications are the same.

In sum, at a high level of generality, the model is just an aggregate supply (AS) and aggregate demand (AD) framework, similar to those typically found in intermediate macroeconomics textbooks, such as Jones (2018). In fact, most of the intuition that we derive from this framework stems exactly from this underlying demand and supply structure, as in McLeay and Tenreyro (2019).

III.B. Two Alternative Sources of Inflation Stabilization

Given the structure of the model, it is immediately apparent that stable inflation can be the result of at least two changes in the economy: first, a flat structural Phillips curve, which corresponds to $\kappa \to 0$; second, a very elastic aggregate demand curve, which corresponds to $\phi \to 0$ or, in terms of the interest rate rule, to $\psi_\pi \to \infty$. In what follows, we will take these two extreme parametric restrictions as stylized representations of the two alternative hypotheses on the ultimate source of the observed inflation stability that have been most discussed in the literature.[15]

The first hypotheses—the (Phillips curve) slope hypothesis, for short—is that inflation is stable because changes in the structure of goods markets or in firms' pricing practices have produced a structural disconnect between inflation and marginal cost pressures. The literature has explored many mechanisms that might lead to such a disconnect, as reviewed in the introduction. Distinguishing among them is beyond the scope of this paper.

The second hypothesis that we focus on—the policy hypothesis, for short—is that inflation is stable because monetary policy now leans more heavily against inflation than it did in the first part of the sample, thus reducing its variability in equilibrium. This is the hypothesis favored by McLeay and Tenreyro (2019). In our stylized model, this hypothesis amounts to assuming that ψ_π has increased in the second part of the sample. In the limit with $\psi_\pi \to \infty$, inflation becomes perfectly stable. In practice, there are many channels through which a change in the actual conduct of monetary policy and in the communication and public understanding of its objectives can affect inflation dynamics and inflation expectations, without going to the extreme of promising a very large increase in policy rates in reaction to even small changes in inflation.[16]

15. In our stylized model, a flat aggregate demand curve could also result from $\sigma \to \infty$, although its intercept (and thus inflation) would still depend on the demand shock in this case. We do not focus on this possibility because there is not much evidence that the responsiveness of the real economy to interest rates has increased since 1990. If anything, there is some discussion of a reduced pass-through of interest rates, and of financial conditions more generally, onto real variables, especially since the financial crisis.

Another obvious possibility is that the volatilities of both shocks have fallen dramatically. Although the large body of literature on the Great Moderation indeed suggests that the volatility of (at least some) shocks did fall in the mid-1990s, we do not focus on this possible explanation because the volatility of real variables has not fallen nearly as much as that of inflation, at least in response to business cycle shocks. In this respect, our main object of inquiry is the reduction in the volatility of inflation relative to that of its plausible real drivers, conditional on business cycle shocks.

16. Inflation expectations do not play a crucial independent role in our model because under rational expectations they are a function of the same shocks that drive inflation. In this respect, stable inflation and inflation expectations are two manifestations of the same phenomenon. Carvalho and others (2019) discuss the notion of expectation anchoring theoretically and empirically in the context of a model with learning, in which expectations are not as tightly linked to actual inflation as under rational expectations. Jorgensen and Lansing (2019) also study the implications of a learning model for the observed connection between inflation and real activity.

The rest of this section derives some basic implications of these two alternative hypotheses in the context of our stylized model, and discusses the extent to which they are consistent with the evidence in section II.

III.C. Model Solution

This section presents the solution of the simple model described above under the assumption of independent and identically distributed shocks. With this simplifying assumption, expectations are zero and the model reduces to a static demand and supply framework with stochastic shocks,

$$\pi_t = \kappa(x_t + s_t)$$

$$\pi_t = -\phi(x_t + d_t),$$

whose solution is

$$x_t = \frac{\phi}{\phi + \kappa} d_t - \frac{\kappa}{\phi + \kappa} s_t$$

$$\pi_t = \frac{\phi\kappa}{\phi + \kappa} d_t + \frac{\phi\kappa}{\phi + \kappa} s_t.$$

The particular form of this solution is, of course, model specific, but its economics are simple and quite general. Demand shocks induce a positive correlation between inflation and the output gap. With direct observations on d_t, or an instrument for it, it would be possible to estimate the slope of the Phillips curve κ by comparing the response of x_t and π_t to the shock d_t, as in Barnichon and Mesters (2019a, 2019b). On the contrary, supply shocks induce a negative correlation between inflation and the output gap, from whose strength we could infer the demand parameter ϕ. When demand and supply shocks cannot be directly observed, the correlation between inflation and the output gap is not informative on either ϕ or κ, as in the classic identification problem. This is the basic point nicely illustrated by McLeay and Tenreyro (2019).

To further clarify this identification challenge, and to shed light on how to potentially overcome it, consider the solution of the model under the two alternative sources of inflation stabilization that we discussed above, $\kappa \to 0$ or $\phi \to 0$. Inflation is zero in both cases, but $x_t = d_t$ under the slope hypothesis, while $x_t = -s_t$ under the policy hypothesis. In other words, when the slope of the structural Phillips curve is zero—for example, because prices are insensitive to marginal cost pressures—the economy becomes more Keynesian, and demand shocks are the predominant drivers of output fluctuations. On the contrary, when policy leans very heavily against inflation, the economy tracks the flexible price equilibrium, it becomes more neoclassical, and economic fluctuations are driven by supply or cost-push disturbances. In sum, which hypothesis—slope or policy—is a better explanation of post-1990 inflation stability simply depends on whether post-1990 business cycles were mainly driven by supply or demand disturbances.

A popular approach to distinguish between demand- and supply-driven fluctuations is to exploit the co-movement pattern of real activity and inflation. A strongly positive correlation would signal the prevalence of demand shocks, while a negative one would favor the predominance of supply innovations. Unfortunately, this strategy is not effective in our case, given the observed stability of inflation: if inflation varies very little over the business cycle—as it has since the 1990s—it also carries limited information to help us separate demand from supply shocks based on the sign of its co-movement with real activity. This is why, in the next sections, we will attempt to tackle this identification challenge by bringing either more information to the table or more theoretical restrictions on the impact matrix, of the form provided, for instance, by full-blown DSGE models.

IV. Interpreting the Facts with a Structural VAR

The U shock employed in section II is a useful descriptive tool, which helps to focus the empirical analysis on the dynamics of inflation and real activity occurring over the cycle. This exercise focuses on the frequencies at which the connection between the nominal and real side of the economy is usually thought to be most evident, as well as those at which monetary policy might have the most significant impact on these dynamics. But most business cycle models with multiple shocks suggest that these dynamics reflect the responses of the economy to a mixture of structural shocks, even if one of them might be preponderant (Smets and Wouters 2007; Justiniano, Primiceri, and Tambalotti 2010). In terms of the stylized model presented above, the U shock would be a combination of demand and supply disturbances, with weights that depend on the relative variance of those shocks, as well as the structural parameters of the economy, including the slope of the aggregate demand and supply equations. As argued in section III, it is therefore impossible to determine the main source of inflation stability—

a flat aggregate supply or demand—unless we can distinguish the two kinds of shocks more precisely. This is the task of this section.

More specifically, we use data on the excess bond premium (EBP) constructed by Gilchrist and Zakrajšek (2012) to identify a credit market disturbance, which we interpret as a proxy for demand shocks. To do so, we add the EBP to the baseline VAR of section I and study the impulse responses to innovations to the EBP that are orthogonal to the other variables in the system.[17] The idea is that innovations to the EBP capture disruptions in credit markets that propagate through the rest of the economy largely as demand shocks. When credit is tight, as signaled by a high EBP, investment falls, reducing aggregate demand and generating further reactions in the economy that also lead to lower labor demand, lower wages, lower income, and ultimately lower inflation.[18]

This strategy does not hinge on the identification of genuinely exogenous credit supply shocks, which can be hard to disentangle from other disturbances affecting financial markets, such as uncertainty shocks or even monetary policy shocks (Caldara and others 2016). All we need is that these innovations to the EBP propagate through the economy by shifting primarily the demand for labor and goods, regardless of their ultimate origin, and that this is true to roughly the same extent before and after 1990. This is the maintained assumption in the rest of this analysis.

Figure 4 presents the impulse responses to the EBP innovation described above (to save space, we omit the response of natural unemployment, given that it is flat). The EBP shocks are more volatile in the second sample, mostly reflecting the sharp spike in credit spreads during the financial crisis. Their standard deviation is approximately 60 percent higher after 1990 than before. Therefore, we normalize the size of the shock in both samples so that it increases the EBP by 1 percentage point on impact at the

17. In practice, we order the EBP last in the VAR and use a recursive scheme to identify the shock to the EBP equation. The results are similar if we order the EBP first in the Cholesky ordering, or if we compute impulse responses to the combination of shocks with the highest contribution to the EBP's variance at business cycle frequencies. This identification strategy is similar to that pursued by Gilchrist, Yankov, and Zakrajšek (2009) and Gilchrist and Zakrajšek (2012). Their dynamic systems, a factor-augmented vector autoregression (FAVAR) and a VAR, respectively, also include "fast moving" variables, such as asset prices, which they place below the EBP in the Cholesky ordering. We do not have such variables in our system, so we order the EBP last.

18. This is how marginal efficiency of investment shocks in Justiniano, Primiceri, and Tambalotti (2011), risk shocks in Christiano, Motto, and Rostagno (2014), and spread shocks in Cai and others (2019) propagate. All these shocks are identified mostly through their effect on credit spreads.

Figure 4. Impulse Responses to an EBP Shock

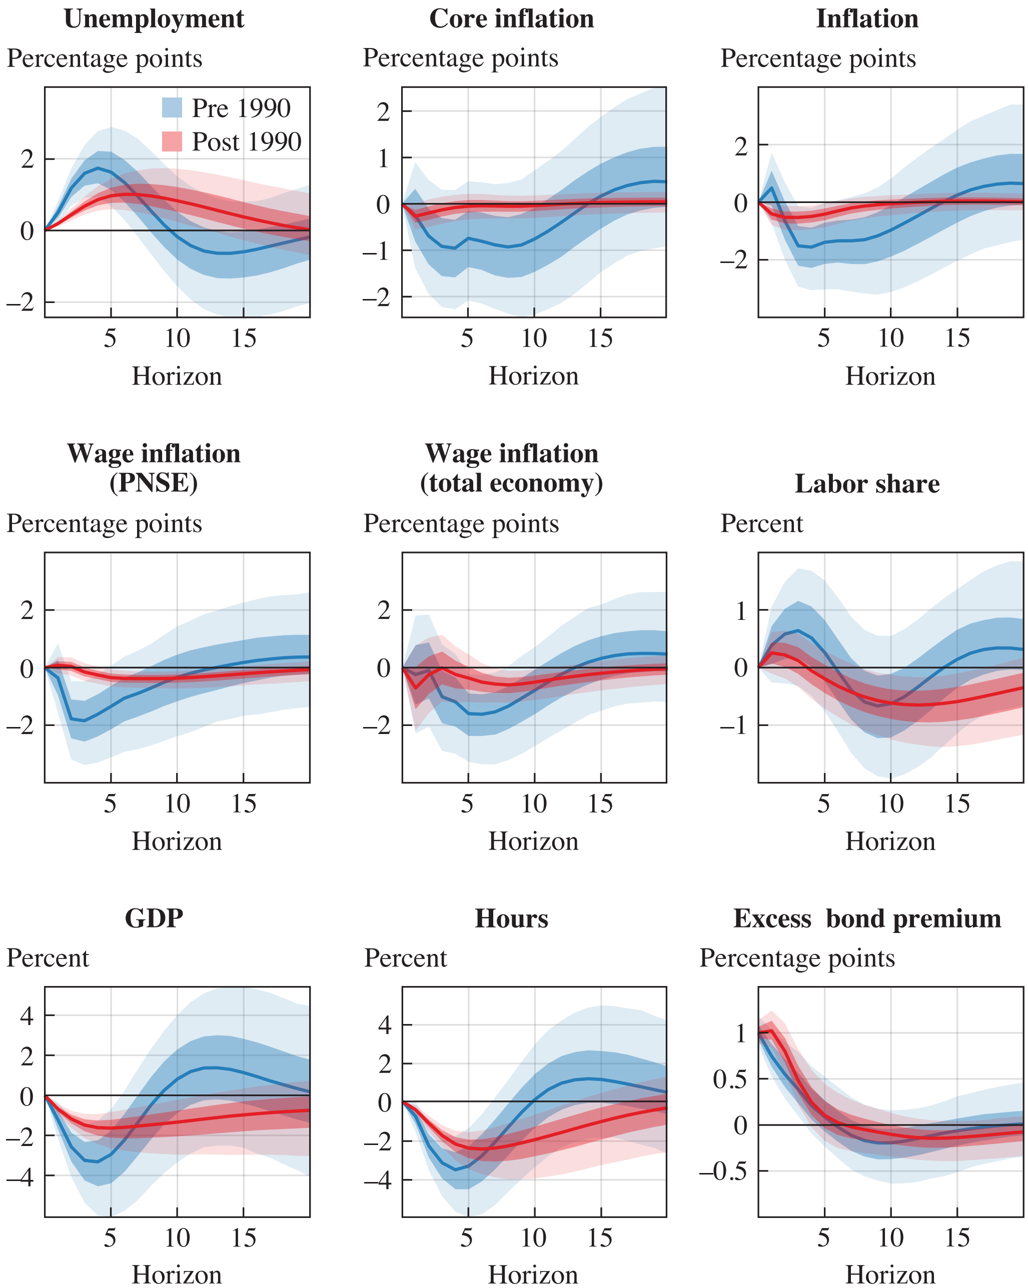

Source: Authors' calculations.

Note: The EBP shock is identified by assuming that it affects the excess bond premium contemporaneously, but all other variables with a lag. The impulse responses are from the baseline VAR of section I.B, augmented with data on the EBP. The solid lines are posterior medians, while the shaded areas correspond to 68 percent and 95 percent posterior credible regions. The pre- and post-1990 samples consist of data from 1973:Q1 to 1989:Q4 and from 1989:Q1 to 2019:Q3, respectively.

median draw. After this normalization, the response of the EBP to its innovation has the same shape in the two samples, which simplifies the evaluation of the changes in the reaction of the other variables.

Among these variables, the inflation rates barely react in the second sample, consistent with the findings in section II. According to the intuition developed within the stylized model presented above, most of the information to distinguish between the slope and policy hypotheses should come from the responses of the real variables. Under the policy hypothesis, the economy should become more neoclassical, with demand shocks having smaller effects on the real variables. On the contrary, under the slope hypothesis the economy should become more Keynesian, with demand shocks becoming more destabilizing. Comparing the second sample to the first, there is evidence of an attenuated response of unemployment, GDP, and hours in the first few quarters after the shock. This piece of evidence is consistent with the policy hypothesis, especially considering that this is the horizon at which monetary policy has arguably the most bite on the real economy. However, the response of all the real variables is much more persistent in the second sample. For instance, the post-1990 response of unemployment remains statistically and economically positive for a substantially longer period of time. As a consequence, the effect of the EBP shock on unemployment, cumulated over a five-year horizon, is actually overall larger in the second sample than in the first. Similar considerations hold for GDP, hours, and the labor share.

On balance, this exercise provides fairly strong evidence in favor of the slope hypothesis, given that the response of inflation has become much more muted than that of the real variables. However, the experiment does not completely rule out an important contribution of monetary policy in better insulating the economy from demand shocks, and hence delivering more stable inflation. Parsing this evidence into sharper conclusions on the relative contribution of these two developments to the stability of inflation since 1990 is very difficult to do without putting more structure on the identification problem. This is what we do in the next section.

An alternative, perhaps more intuitive, way of presenting these results is to retell them through the perspective of the Great Recession. Most observers agree that the Great Recession originated from the financial crisis that preceded it. Although that shock had a complex origin and it affected the economy through many channels, it had the hallmarks of a typical demand shock. In our VAR, one of the main manifestations of this shock is a massive increase in the EBP. If this shock was indeed primarily a demand shock, the fact that inflation fell by a very limited amount is a strong indication that the slope of the Phillips curve must be very low, at least relative to what it used to be before the 1990s.

At the same time, the fact that the economy weathered the storm without a collapse more akin to that experienced during the Great Depression is consistent with monetary policy (arguably with some fiscal help) having been able to limit the impact of the shock to the real economy. And perhaps the real effects of the shock could have been counteracted even more effectively had it not been for the limits imposed by the ZLB on nominal interest rates.[19] Therefore, the evidence might also be consistent with an improvement in the ability of policymakers to limit the damage caused by demand shocks on the real economy. This second conclusion, however, requires taking a stance on the size of the shock that hit the economy, in comparison, for instance, to the one that occurred in the early 1930s. Although many commentators have compared the extent of the financial disruption during the financial crisis to that associated with the Great Depression, it is difficult to make such a comparison formally. Therefore, we consider this second conclusion more tentative than the one regarding the reduction in the slope of the Phillips curve, which is more directly supported by the evidence.

An important caveat to the line of reasoning pursued in this section is that it is predicated on the assumption that EBP shocks, and the Great Recession, were both primarily demand disturbances. More precisely, the requirement is that they should affect inflation through their impact on the conventional measures of cost pressures that we have analyzed in section II. If, on the contrary, these disturbances reflect an important cost-push component, the evidence shown above is harder to interpret. Suppose that an increase in the EBP shock, caused by raising financing costs, induces firms to charge higher prices through channels that are not manifested in changes in the labor share or other standard measures of slack, and that this occurs to a greater extent in the second part of the sample than in the first. In this case, inflation might end up being flat in response to that shock, as in the post-1990 evidence above, since the impulse would act as a positive cost-push shock in the Phillips curve at the same time as it depresses other sources of cost pressure through the conventional demand channels.[20]

19. The extent to which the ZLB was binding during the Great Recession and its aftermath is much debated in the literature (Swanson and Williams 2014; Galí and Gambetti 2020; Eggertsson and Egiev 2020).

20. Christiano, Eichenbaum, and Trabandt (2015) and Gilchrist and others (2017) develop models of this kind.

IV.A. The Role of Inflation Expectations

Our SVAR analysis so far has focused on the direct relationship between inflation and unemployment, conditional on demand shocks. This conditional correlation is informative on the slope of the aggregate supply curve, helping to distinguish it from that of aggregate demand. However, to translate our findings on this conditional correlation into statements about the sensitivity of prices to costs, or equivalently into the slope of the expectations-augmented Phillips curve of Friedman and Phelps, we must take a stance on the dynamics of inflation expectations. The concern, as articulated by Stock (2011) and Barnichon and Mesters (2019b) for instance, is that anchored expectations produce a flatter aggregate supply curve, potentially confounding changes in the coefficient on real activity in the structural Phillips curve. In fact, many observers refer to the expectations-augmented Phillips curve in attributing the transition to stable inflation since the 1990s to a change in the dynamics of inflation expectations, rather than to a change in the responsiveness of prices to costs (Bernanke 2007; Blanchard 2016; Gordon 2018; Ball and Mazumder 2019).

One approach to addressing this concern is to be specific about the microfoundations of firms' pricing decision, as we do in the DSGE model of the next section. In this section, we take an intermediate step in evaluating the potential of a change in the dynamics of expectations to be behind the finding of a flatter Phillips curve after 1990. This step is based on the impulse responses to the EBP shock derived above. This approach is very similar to that pursued by Barnichon and Mesters (2019b), but in contrast to their claim, it confirms our main finding of a large decline in the slope of the structural Phillips curve over the past three decades. This result is also consistent with the DSGE analysis of section IV.

Consider the following Phillips curve,

$$\pi_t = \alpha \pi^e_{t+1|t} + (1-\alpha)\pi_{t-1} + \kappa(x_t + s_t), \tag{5}$$

which generalizes equation (2) by allowing for inflation inertia and potentially nonrational inflation expectations. This expression also imposes that inflation has no permanent effect on real activity, consistent with many empirical versions of the Phillips curve, which are vertical in the long run. By replacing the time index t with $t + h$, taking an expectation conditional on time-t information and a derivative with respect to a demand shock at time t on both sides, we obtain

$$R^{\pi}_h = \alpha R^{\pi^e}_h + (1-\alpha) R^{\pi}_{h-1} + \kappa R^x_h, \tag{6}$$

where $R_h^z = \dfrac{\partial E_t z_{t+h}}{\partial d_t}$ is the impulse response at horizon h of a generic variable z to the demand shock. In deriving equation (6), we used the fact that s_t is independent from the demand disturbance. Equation (6) says that, if equation (5) holds, the impulse responses of inflation, expected inflation, and the gap to a demand shock should satisfy equation (6) at every horizon h. Therefore, if we observed the true impulse responses of these variables, we would be able to infer the values of α and κ that satisfy equation (6).

In practice, the Phillips curve—equation (5)—does not hold exactly in the data. In addition, we can only estimate R_h^π, $R_h^{\pi^e}$, and R_h^x, as opposed to observing their population value. Despite these difficulties, we can still compute a value for α and κ by minimizing the distance between our estimates of the right- and the left-hand side of equation (6). To operationalize this strategy, we augment the structural VAR of the previous section with a measure of one-quarter-ahead inflation expectations from the Survey of Professional Forecasters. This VAR produces impulse responses that are nearly identical to those of figure 4. In addition, the reaction of expected inflation is similar to that of actual inflation, although more muted. For each posterior draw of the impulse responses of inflation, expected inflation, and unemployment (which proxies for the gap) to the EBP shock, we compute a value of α and κ using a linear projection. Figure 5 summarizes the outcome of this procedure. The right panel displays the histogram of the inferred value of κ in the pre- and post-1990 sample. This value clearly shifts toward zero in the second sample, consistent with the slope hypothesis. The estimate of α also declines, indicating a more inertial inflation process.

Accounting for inflation inertia in equation (5), which is important to fit the behavior of inflation as demonstrated by Fuhrer and Moore (1995), Galí and Gertler (1999), and Fuhrer (2010), is also crucial to the inference on the slope of the Phillips curve. If we restrict $\alpha = 1$, as in Barnichon and Mesters (2019b), our results are more consistent with theirs: κ still declines after 1990, but less than in figure 5.

Another difference from Barnichon and Mesters (2019b) is that their analysis is based on one-year-ahead inflation expectations from the Michigan Survey, as in Coibion and Gorodnichenko (2015). Unfortunately, this series is only available starting in 1978, which does not give us enough observations for a reliable estimation before 1990. However, figure 6 shows that the results for the post-1990 sample with the Michigan data are similar to our baseline with SPF data, confirming the robustness of our findings.

Figure 5. Posterior Distribution of Inferred Values of α and κ

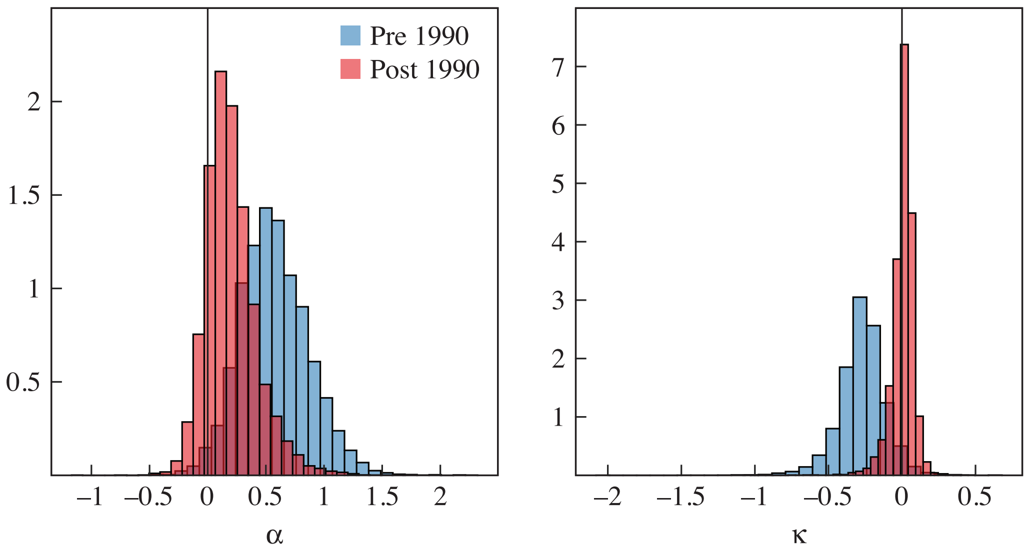

Source: Authors' calculations.

Note:These values are obtained as described in section IV.A, using the impulse responses of inflation, expected inflation, and unemployment to an EBP shock, identified by assuming that it affects the EBP contemporaneously, but all other variables with a lag. The impulse responses are from the baseline VAR of section I.B, augmented with data on one-quarter-ahead inflation expectations and the EBP. The pre- and post-1990 samples consist of data from 1973:Q1 to 1989:Q4 and from 1989:Q1 to 2019:Q3, respectively.

Figure 6. Posterior Distribution of Inferred Values of α and κ in the Post-1990 Sample

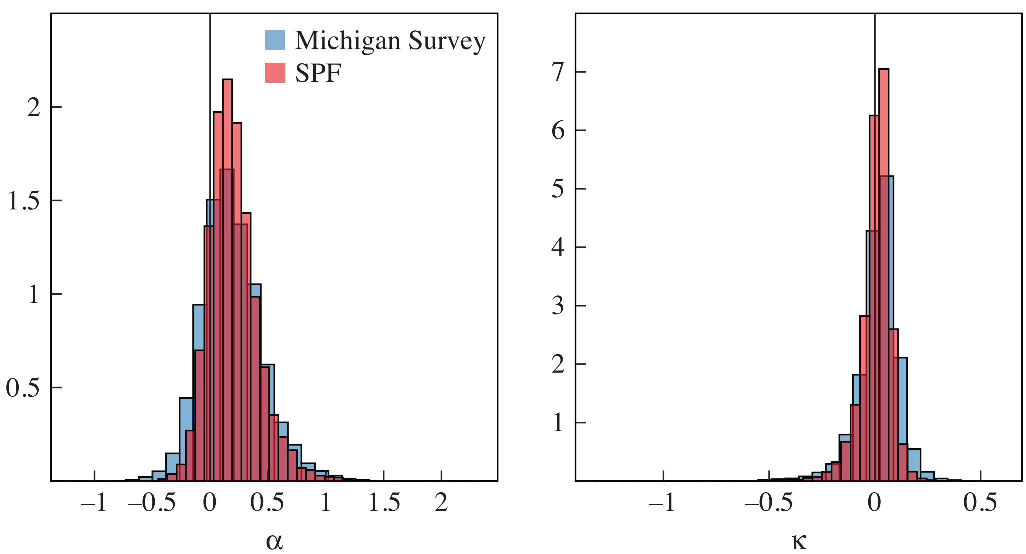

Source: Authors' calculations.

Note: These values are obtained as described in section IV.A, using the impulse responses of inflation, expected inflation, and unemployment to an EBP shock, identified by assuming that it affects the EBP contemporaneously, but all other variables with a lag. The impulse responses are from the baseline VAR of section I.B, augmented with data on inflation expectations and the EBP. The Michigan Survey and SPF histograms are obtained using data on one-year-ahead inflation expectations from the Michigan Survey and one-quarter ahead inflation expectations from the Survey of Professional Forecasters, respectively.

V. Interpreting the Facts with an Estimated DSGE Model

As argued above, interpreting the facts of section II only through the lens of the VAR is challenging, as one can only speculate on the mechanisms behind them. For instance, the reduced sensitivity of inflation to labor market conditions in the second part of the sample could be due to differences in monetary policy or in other aspects of the economy. In the previous section, we made progress on this issue by identifying a demand disturbance in a structural VAR. In this section, instead, we interpret these facts through the lens of a fully specified DSGE model.

In such a structural model, interpreting the evidence is straightforward, since the mechanism is identified as part of the estimation. Of course, using structural models comes with its own challenges, most notably the fact that the model may be misspecified. From the perspective of the question at hand, such misspecification may imply that the DSGE cannot reproduce the same facts as the VAR. The first part of this section addresses this concern. It shows that the Federal Reserve Bank of New York's DSGE model, when estimated over the same subsamples as the VAR, reproduces the VAR facts both qualitatively and by and large quantitatively. In particular, a shock that has an impact on the labor market in ways comparable to those described in section II produces a muted response of inflation after 1990. The second part of the section investigates what changes in the model's structural parameters deliver these much attenuated responses. As in section IV, we focus on two explanations: changes in the conduct of monetary policy and changes in the slope of the structural price Phillips curve. We begin with a very brief description of the DSGE model.

V.A. The DSGE Model

The NY Fed DSGE model is a medium-scale New Keynesian model. In broad strokes, the model can be as Smets and Wouters (2007) plus financial frictions as in Bernanke, Gertler, and Gilchrist (1999) and Christiano, Motto, and Rostagno (2014). Variations of this model have been used in Del Negro and Schorfheide (2013), Del Negro, Giannoni, and Schorfheide (2015), and Del Negro and others (2017).[21] Cai and others (2019) document the model's real-time forecasting performance over the past ten years. The model is estimated using multiple measures of output growth (both

21. More specifically, the model is the same as that in Del Negro and others (2017), except that the processes for productivity growth and the safety premium were assumed to be more persistent in that paper.

GDP and gross domestic income), the growth in consumption, investment, the real wage, hours worked, two measures of inflation (both core PCE and the GDP deflator), long-run inflation expectations, the federal funds rate, the ten-year Treasury yield, the Baa-Treasury spread, and the series of total factor productivity growth constructed by Fernald (2012). We also allow for anticipated policy shocks as in Del Negro, Giannoni, and Patterson (2015) and Laséen and Svensson (2011), to account for the ZLB on nominal interest rates and forward guidance. Therefore, we augment the set of observables to incorporate interest rate expectations during the ZLB period. Online appendix B provides all the equilibrium conditions, the definition of the observables, and the specification of the priors, which are the same as in Smets and Wouters (2007) for all parameters common to the two models.

Here we only report two of the model's equations because their parameters are mentioned in the analysis that follows. The first equation is the price Phillips curve,

$$\pi_t = \kappa_p mc_t + \frac{\iota_p}{1+\iota_p\bar{\beta}}\pi_{t-1} + \frac{\bar{\beta}}{1+\iota_p\bar{\beta}}E_t[\pi_{t+1}] + \lambda_{f,t}, \tag{7}$$

where π_t is inflation, mc_t is marginal cost, and $\lambda_{f,t}$ captures exogenous fluctuations in desired markups. The slope of the Phillips curve is

$$\kappa_p = \frac{\left(1-\zeta_p\bar{\beta}\right)\left(1-\zeta_p\right)}{\left(1+\iota_p\bar{\beta}\right)\zeta_p\left(\left(\Phi_p-1\right)\epsilon_p+1\right)}.$$

Its key determinant is the Calvo parameter $\zeta_{p,}$ which represents the fraction of firms that do not adjust their price in every period. Therefore, a higher ζ_p means that prices are stickier, which makes the Phillips curve flatter. The other parameters entering the slope are the degree of indexation ι_p, the curvature parameter in the Kimball aggregator for prices ϵ_p, the size of fixed costs in production Φ_p, and $\bar{\beta}$, the discount rate adjusted for steady-state growth.

The second important equation is the interest rate rule followed by the monetary authority,

$$\begin{aligned} R_t = {} & \rho_R R_{t-1} + (1-\rho_R)\left(\psi_\pi(\pi_t - \pi_t^*) + \psi_y(y_t - y_t^*)\right) \\ & + \psi_{\Delta y}\left((y_t - y_t^*) - (y_{t-1} - y_{t-1}^*)\right) + r_t^m, \end{aligned} \tag{8}$$

where R_t is inflation, y_t and y_t^* are actual and natural output, so that their difference is a measure of the output gap, π_t^* is a time-varying inflation target, and r_t^m captures exogenous departures from the policy rule.[22]

V.B. Can the DSGE Replicate the VAR Facts?

To assess whether the estimated DSGE model can reproduce the VAR facts of section II, we compute the impulse responses to the U shock in the VAR implied by the estimated DSGE model. Del Negro and Schorfheide (2004) show how to construct the VAR approximation of the DSGE model, which they call DSGEVAR(∞). We refer to this approximation as DSGEVAR in the remainder of the paper.[23] The NY Fed DSGE model does not feature all the variables included in the VARs of section II, most notably unemployment. Therefore, the results in this section are based on a DSGEVAR with slightly fewer observables. These are hours worked per capita in log levels, a measure of price inflation, the labor share, and wage inflation.[24] This information is enough to capture the dynamics of the key economic variables involved in the connection between inflation and real activity.

Figure 7 shows the impulse responses of these variables to a U shock in two DSGEVARs, one based on the DSGEVAR implied by the pre-1990 DSGE estimates and the other reflecting the estimation on the post-1990 sample. The U shock is identified using the methodology described in section I.A., applied to hours rather than to unemployment. As in section II, the response of the labor market to a U shock is similar in the two samples, although it is slightly smaller on impact and more persistent in the second sample.[25] The response of the labor share is also similar in the two samples and, if anything, stronger in the post-1990 estimation. Qualitatively and

22. The time-varying inflation target π_t^*, which was not present in the original Smets and Wouters (2007) specification, is a very persistent process that captures the secular rise and fall of inflation and nominal interest rates in the estimation sample. As in Del Negro and Eusepi (2011), we use data on long-run inflation expectations as an observable in the estimation, which help the model to account for those low-frequency movements.

23. For any given vector of DSGE model parameters, the VAR approximation of the DSGE model is what would be obtained by generating artificial data from the DSGE model and estimating a VAR on such generated data. Del Negro and Schorfheide (2004) show how to compute the matrices of this VAR in population, without actually generating any artificial data.

24. The VARs used in section II also include the log level of per capita GDP. This variable is not stationary in the NY Fed DSGE model. Therefore we cannot include it in the DSGEVAR.

25. The DSGE model, and therefore the DSGEVAR, uses a slightly different measure of hours worked than the VAR in section IV.A. This is because the variable definitions in the NY Fed DSGE model are the same as in Smets and Wouters (2007).

Figure 7. Impulse Responses to an Unemployment Shock in the VAR Approximation of the DSGE

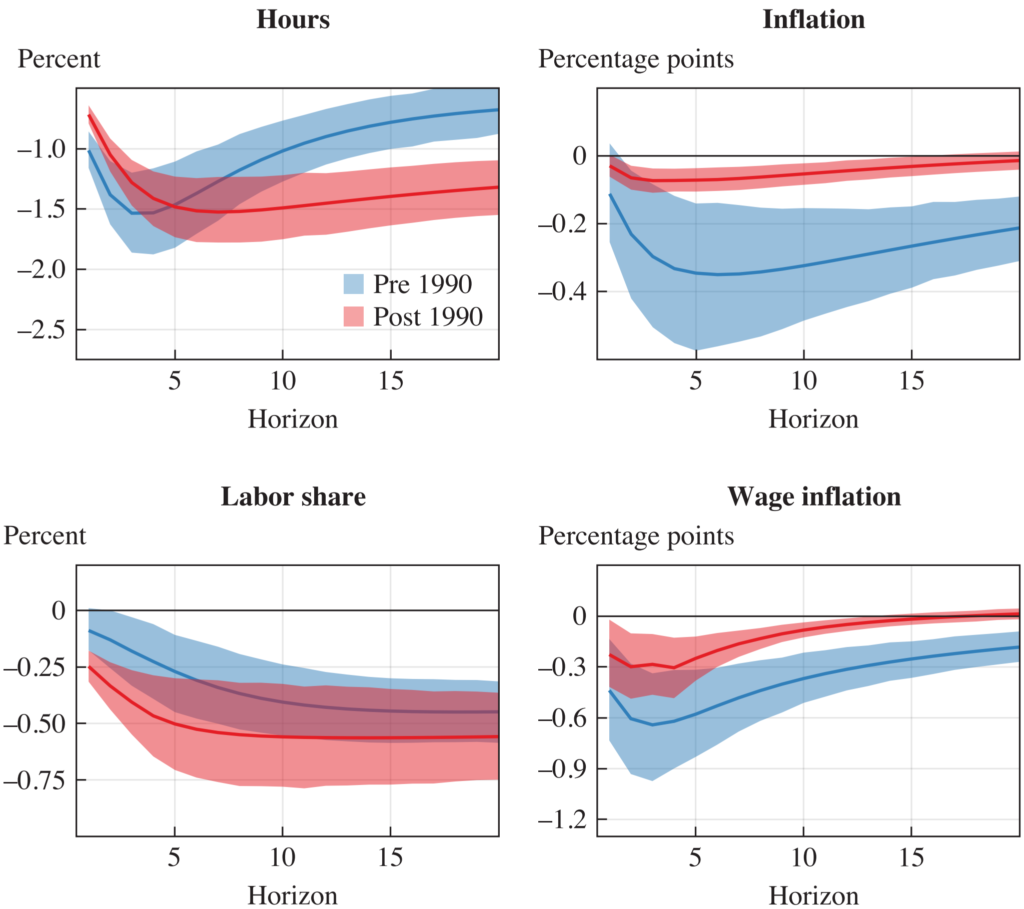

Source: Authors' calculations.

Note: The impulse responses are from the VAR approximation of the NY Fed DSGE model (DSGEVAR) described in section V. The shock is identified using a Cholesky strategy, with unemployment ordered first. The solid lines are posterior medians, while the shaded areas correspond to 95 percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:Q2 to 1989:Q4 and from 1990:Q1 to 2019:Q3, respectively.

quantitatively, these responses are also very similar to those in section II. The DSGE model therefore confirms that the transmission of U shocks to marginal costs is very similar across the two samples. On the contrary, the response of inflation is notably different: it is very muted in the second sample, as in the VAR. Finally, the responses of nominal wage inflation are somewhat weaker in the second sample, but not as weak as that of inflation, also consistent with the results of section II.

What changes in the estimated DSGE parameters are responsible for the differences in the DSGEVAR responses before and after 1990? To answer this question, we concentrate on the parameters that map more directly onto the alternative explanations of the muted response of inflation to the business cycle that we are comparing: the slope of the price and wage Phillips curve and the parameters describing monetary policy. The first row of figure 8 focuses on the former. The upper left panel shows the posterior distribution of the slope of the price Phillips curve κ_p in the two samples, which has declined substantially over time.[26] In contrast, the upper right panel of figure 8 shows that the posterior distribution of the slope of the wage Phillips curve (κ_w) is quite similar across samples, although it also shifts somewhat to the right.

The remaining two rows of figure 8 focus on the parameters of the monetary policy rule, which capture the inertia of the interest rate (ρ_R) and its response to inflation (ψ_π), the output gap level (ψ_y) and its growth ($\psi_{\Delta y}$). The posterior distribution of the response to inflation does not change much across samples. The response to the output gap increases after 1990, while that to its growth rate falls. Interest rate persistence is notably higher after 1990.[27] The implications of the estimated decline in the slope of the price Phillips curve κ_p for the dynamics of inflation are clear, at least qualitatively: they make inflation less responsive to real activity, as in the data. In contrast, the combined implications of the changes in the policy rule parameters are less obvious. Therefore, we now move to investigate which of these changes in estimated parameters might explain the observed changes in the response of inflation to U shocks.

V.C. Explaining the Facts

Can the changes in the estimates of the slope of the price Phillips curve or the policy parameters quantitatively explain the facts of figure 7?

26. We focus on the slope of the price Phillips curve κ_p rather than on the underlying structural parameters for two reasons. First, the aggregate data used in the estimation of the DSGE model identify the slope, not the underlying structural parameters. For example, the price stickiness parameter ζ_p and Kimball aggregator parameter ϵ_p are not separately identified, as they only affect the slope and do not enter anywhere else in the model. We follow Smets and Wouters (2007) in estimating the former and calibrating the latter, but this is an arbitrary choice. Therefore, we cannot interpret the change in the slope reliably in terms of the underlying structural parameters. Second, only the slope, and not the underlying parameters, matters for the dynamics of the system and for the outcome of the counterfactuals presented below. With these caveats, online appendix figure C.1 reports the posterior estimates of the price stickiness parameter ζ_p, which, according to our estimation, is the key driver of the change in the slope.

27. During the recent ZLB episode, the model accounts for the forced deviations of the interest rate from the estimated policy rule, as well as for forward guidance, using anticipated monetary policy shocks.

Figure 8. Prior and Posterior Distributions for Selected Parameters of the DSGE Model

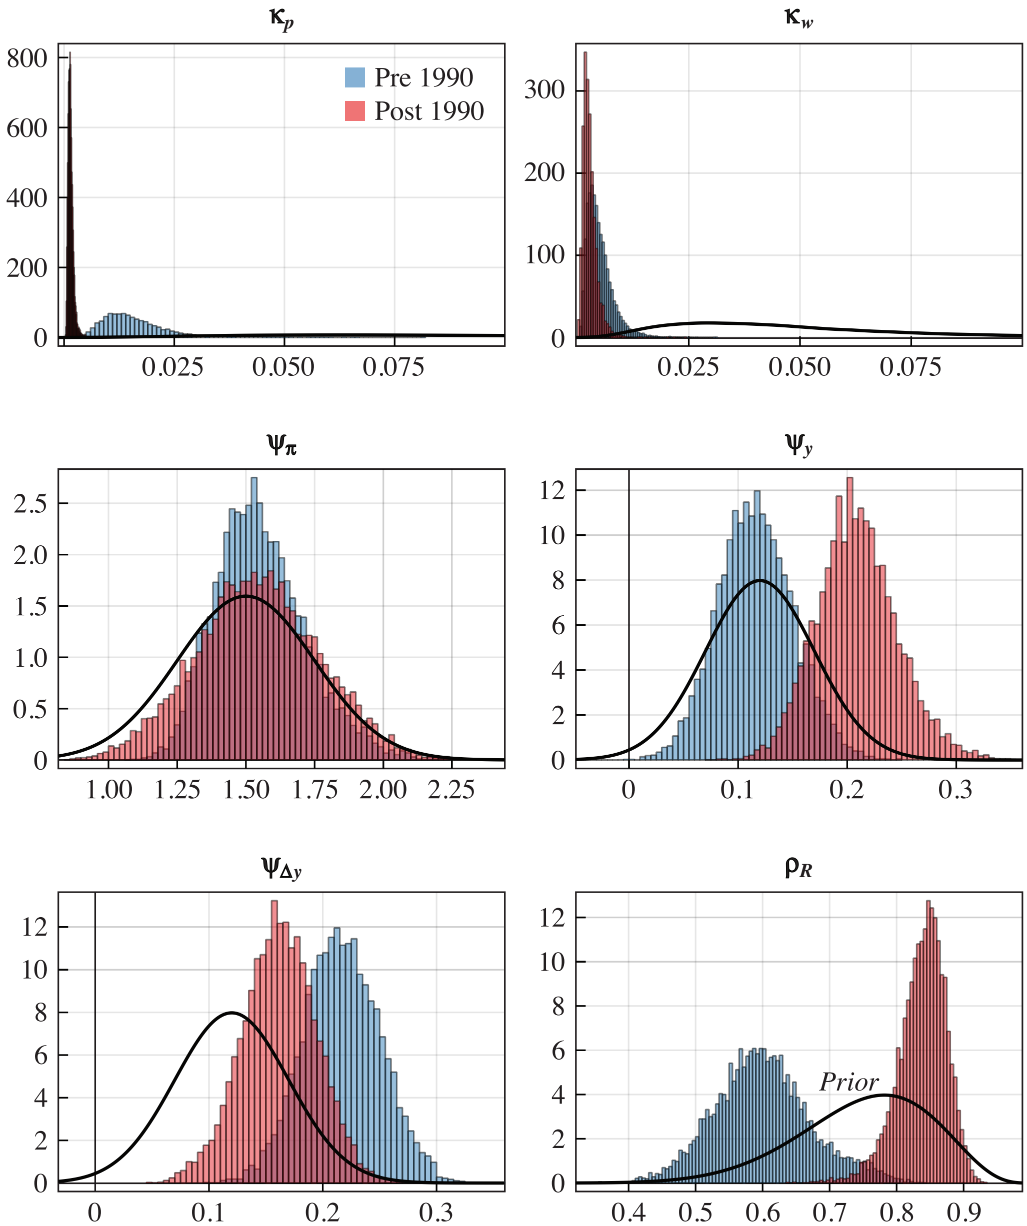

Source: Authors' calculations.

Note: The pre- and post-1990 samples consist of data from 1964:Q2 to 1989:Q4 and from 1990:Q1 to 2019:Q3, respectively.

Figure 9 answers this question. The pre-1990 and post-1990 lines show the impulse responses to a business cycle shock in the DSGEVAR implied by the posterior mode of the DSGE parameters before and after 1990.[28] The other two lines show the counterfactual impulse responses obtained using the pre-1990 parameters, except for the slope of the Phillips curve and the policy-rule parameters. For these parameters, we use the posterior mode of the post-1990 estimates. The point of this exercise is the following: if changes in the slope (policy) fully account for the differences in the impulse responses across samples, then the counterfactual policy line should be as close as possible to the solid red line. In terms of the responses of hours and the labor share, both counterfactuals are close to the post-1990 line, which represents the actual post-1990 responses. However, this is not the case for the responses of price and wage inflation. For inflation, the slope counterfactual is essentially on top of the actual post-1990 response. In contrast, the policy counterfactual produces an even stronger reaction than before 1990. Finally, both counterfactuals tend to overestimate the response of wage inflation, with the policy counterfactual again faring worse than the slope one.

Figure 9 indicates that changes in the slope of the price Phillips curve alone can explain the muted response of inflation to U shocks. Together with the fact that the policy rule is not very different before and after 1990, except for the estimate of ρ_R, this result suggests that a change in the policy rule is not necessary to account for the behavior of inflation after 1990. But could there still be a role for policy if we took a change in the slope of the Phillips curve off the table?

To answer this question, we estimate the DSGE model before and after 1990 allowing only the policy rule parameters to change. The other parameters are assumed to be constant over the entire sample. The purpose of this exercise is to give the policy hypothesis the best shot at explaining the facts. The policy rule parameters estimated as part of this exercise are reported in figure 10. They are different from those in figure 8, which are based on an estimation in which all parameters can change before and after 1990 to best fit the data. First, the distribution of the response to inflation (ψ_π) is skewed to the right after 1990, relative

28. These responses are essentially the same as those in figure 7. The difference is that, there, the solid lines are the medians of the posterior distribution of the impulse responses. Here, the solid lines are the impulse responses at the mode of the posterior distribution of the parameters. We use the posterior mode of the parameters here because this approach makes it easier to conduct the counterfactual described below.

Figure 9. Impulse Responses to an Unemployment Shock in the DSGEVAR, with Counterfactuals

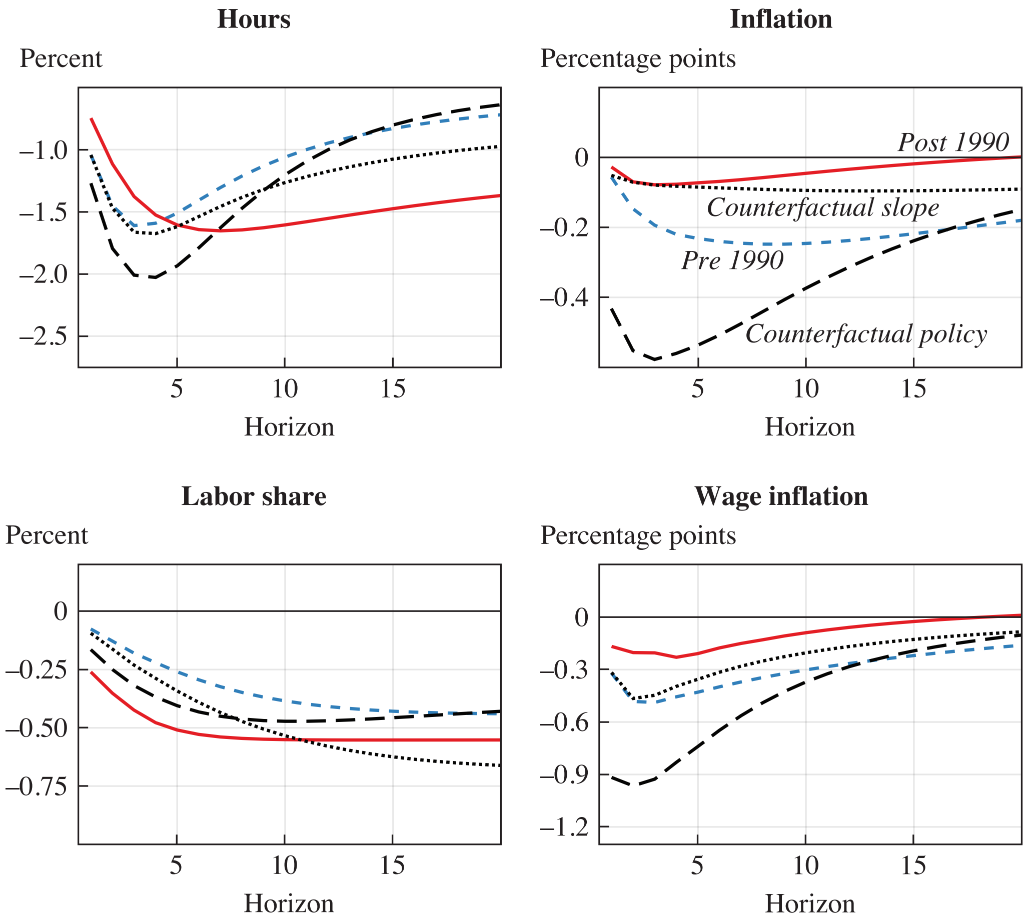

Source: Authors' calculations.

Note: The impulse responses are from the VAR approximation of the NY Fed DSGE model (DSGEVAR) described in section V. The shock is identified using a Cholesky strategy, with unemployment ordered first. The pre-1990 and post-1990 lines are impulse responses constructed using the modal posterior estimates in the pre- and post-1990 samples, respectively. The counterfactual slope responses are obtained using the pre-1990 modal posterior estimates for all parameters except the slope of the Phillips curve, for which we use the post-1990 posterior mode. The counterfactual policy responses are obtained using the pre-1990 modal posterior estimates for all parameters except the policy rule ones, for which we use the post-1990 posterior mode.

Figure 10. Prior and Posterior Distributions for Selected Parameters of the DSGE Model, Estimated with Different Monetary Policy before and after 1990

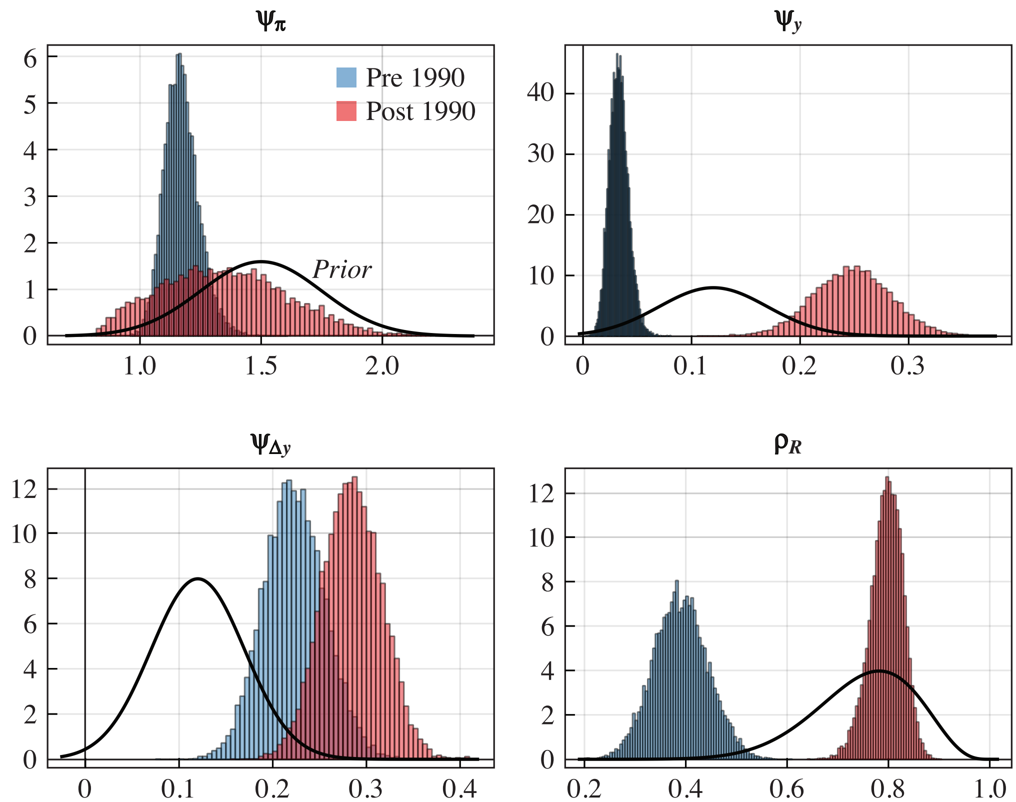

Source: Authors' calculations.
Note: In this experiment, the model is estimated allowing the coefficients of the monetary policy rule to be different before and after 1990, but assuming that all other coefficients are constant for the entire sample from 1964:Q2 to 2019:Q3.

to the pre-1990 estimates. Second, the differences in the response to the output gap (ψ_y) before and after 1990 are larger than in figure 8. Third, the response to the output gap growth ($\psi_{\Delta y}$) is now also stronger in the second sample, rather than weaker. Finally, the persistence in interest rates (ρ_R) remains much higher post 1990, which implies that monetary policy has stronger control of inflation, as shown for instance by Woodford (2003).

Together, these changes in the policy rule parameters do result in a more muted response of inflation to the U shock in the second sample, as shown in figure 11. This is not too surprising, since this version of the model must account for the stability of inflation after 1990 using only changes in the policy parameters. The conclusion is that it is possible to find changes in

Figure 11. Impulse Responses to an Unemployment Shock in the DSGEVAR, Estimated with Different Monetary Policy before and after 1990

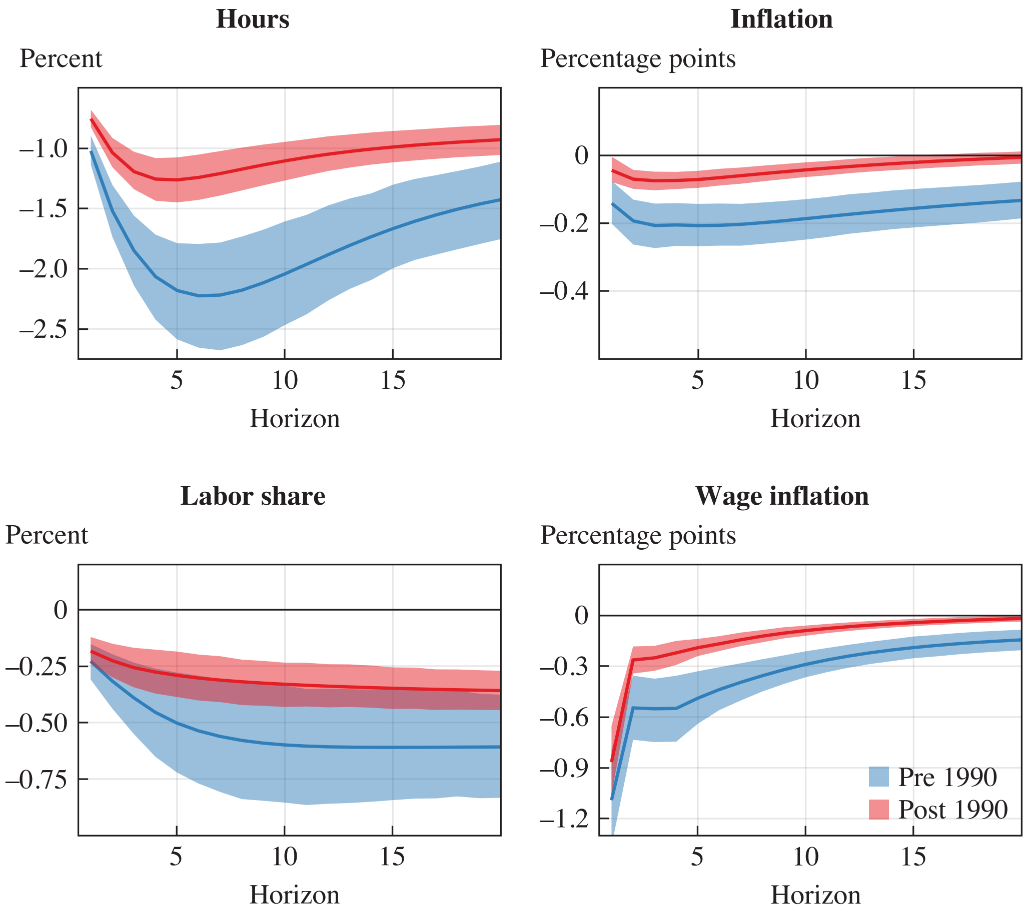

Source: Authors' calculations.

Note: The impulse responses are from the VAR approximation of the NY Fed DSGE model (DSGEVAR) described in section V. The shock is identified using a Cholesky strategy, with unemployment ordered first. The solid lines are posterior medians, while the shaded areas correspond to 95 percent posterior credible regions. In this experiment, the model is estimated allowing the coefficients of the monetary policy rule to be different before and after 1990, but assuming that all other coefficients are constant for the entire sample from 1964:Q2 to 2019:Q3.

the policy rule parameters that stabilize inflation as in the data. However, this success in the inflation dimension comes at the cost of also muting the responses of hours and the labor share after 1990. This is in contrast with the empirical facts highlighted in section II and in figure 7, in which hours and the labor share are, if anything, more responsive after 1990. These results can be interpreted in light of the discussion in section III: to the extent that the main shocks hitting the economy are demand shocks, policy can effectively control inflation by neutralizing their real effects. But the evidence that this has occurred is weak.[29]

These results stand in contrast to those of figure 12, where we repeat the same exercise imposing that only the slope of the price Phillips curve has changed between the two samples. The responses in figure 12 are in line with those in figure 7: the reaction of hours and the labor share is very similar across samples, but that of inflation is much more muted after 1990.[30]

VI. Policy Implications

To understand the policy implications of our findings, a useful starting point is the simple model of section III. In that framework, both real activity and inflation increase in response to a positive shift in aggregate demand. With a flatter aggregate supply curve, however, any given shift in demand has a larger impact on real activity and a smaller effect on inflation. As it turns out, this simple "static" intuition also holds in the quantitative DSGE model of section V: in response to a monetary policy shock, either a short-lived or a very persistent one, inflation reacts less and real activity reacts more when the Phillips curve slope is low. In essence, with a flatter Phillips curve, it is more difficult for monetary authorities to steer inflation in any particular direction using unsystematic policy measures.

The fact that monetary policy cannot as easily control inflation by engineering isolated, unexpected shifts in aggregate demand should not be surprising. After all, a flatter Phillips curve corresponds to a structural change in the economy. As a consequence, the best way for monetary policy to reestablish its control of inflation should be to change its systematic

29. The exercise that we just described might not in fact be the best shot for the policy hypothesis because the post-1990 sample includes a long period at the ZLB, when monetary policy could not react as much to U shocks. Figures C.2 and C.3 in the online appendix reproduce the results in figures 10 and 11 using a sample ending in 2008:Q3, before the onset of the ZLB period and the Great Recession. These parameter estimates indicate an even stronger response to inflation in the post-1990 period, but the impulse responses in figure C.3 do not suggest any attenuation in the response of inflation to a U shock after 1990, contrary to the evidence.

30. The corresponding posterior distributions of the slope of the price Phillips curve κ_p pre- and post-1990 are shown in figure C.4 in the online appendix. They are very similar to those in figure 8. In this exercise the slope of the price Phillips curve changes across samples due to changes in the price stickiness parameter ζ_p. Results are similar when both the slope of the price Phillips curve κ_p and the indexation parameter ι_p are allowed to change.

Figure 12. Impulse Responses to an Unemployment Shock in the DSGEVAR, Estimated with Different Slopes of the Price Phillips Curve before and after 1990

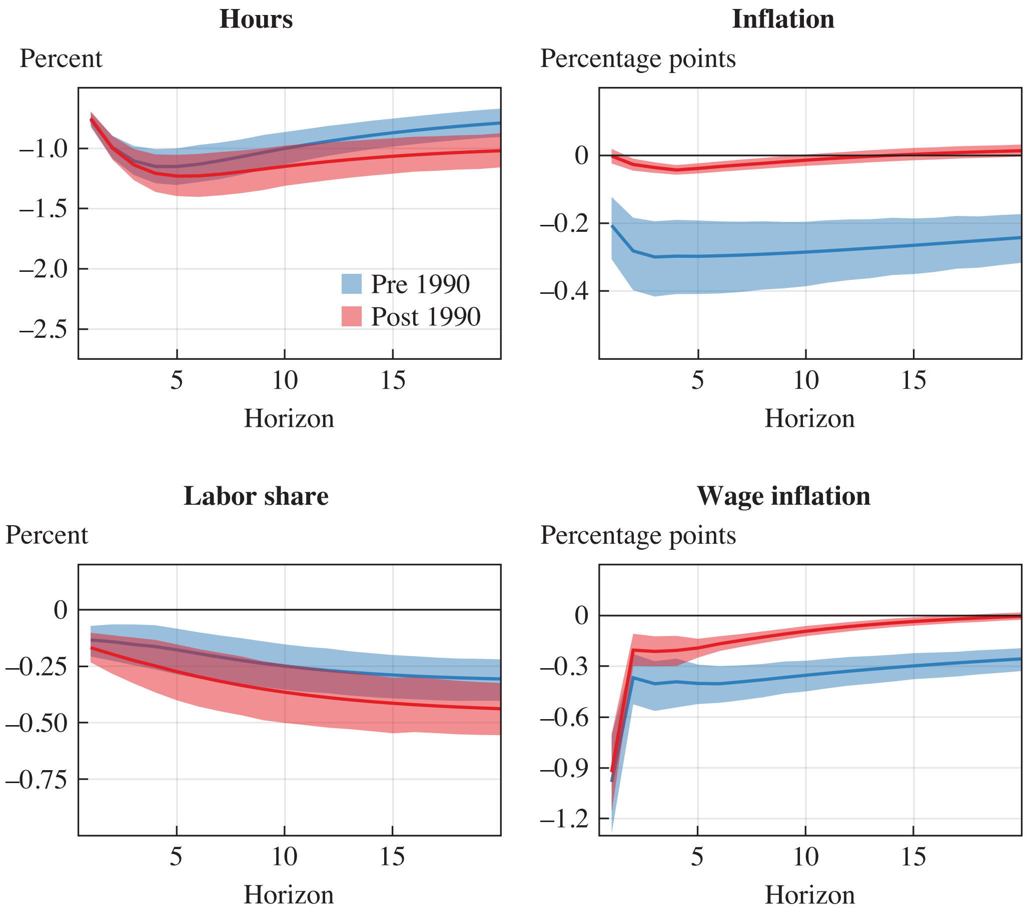

Source: Authors' calculations.
Note: The impulse responses are from the VAR approximation of the NY Fed DSGE model (DSGEVAR) described in section V. The shock is identified using a Cholesky strategy, with unemployment ordered first. The solid lines are posterior medians, while the shaded areas correspond to 95 percent posterior credible regions. In this experiment, the model is estimated allowing the slope of the price Phillips curve to be different before and after 1990, but assuming that all other coefficients are constant for the entire sample from 1964:Q2 to 2019:Q3.

reaction to the state of the economy, as opposed to through a series of policy shocks. In addition, intuitively, the value of adopting such a policy rule should be higher exactly when the slope of the Phillips curve is low and the trade-off implied by a business-as-usual policy is unfavorable.

We illustrate this principle quantitatively using the DSGE model discussed in the previous section. We begin by showing that our main explanation for the muted response of inflation to real activity—a lower slope of the Phillips curve—is also consistent with the persistent weakness of inflation since the Great Recession. We then demonstrate that alternative policy rules—such as average inflation targeting—can be as effective at stabilizing inflation with a flat Phillips curve as it would have been with a steeper one.

Figure 13 shows the model's forecasts of inflation and marginal costs starting in 2020:Q1 under the estimated policy rule and two alternative settings for the slope of the price Phillips curve κ_p: post-1990 and pre-1990. As shown in figure 8 both the pre- and post-1990 values of κ_p are quite low in absolute terms. For ease of exposition, and to emphasize the relative differences, we will refer to the pre-1990 κ_p as "high" and to the post-1990 κ_p as "low" in the remainder of the section.[31]

When κ_p is high, inflation falls more initially, but then it reverts relatively quickly to the Federal Open Market Committee's long-run goal. When κ_p is low, the shortfall of inflation is much more persistent, to the point that it is not closed even by the end of the forecast horizon (left panel). The differences in the inflation forecasts can be explained in terms of those for real marginal costs (right panel). When κ_p is high, marginal costs revert more quickly to the steady state, while they remain persistently depressed with a low κ_p. Intuitively, a high κ_p brings the economy closer to one with flexible prices, in which real marginal costs never deviate from their desired level. With a low κ_p, on the contrary, the deviations of marginal costs from their desired level is more persistent, due to a stronger endogenous propagation.[32] In the current situation, this means that the drag on marginal costs brought about by the Great Recession and by the other negative shocks that followed continues to exert a negative impact on inflation. Therefore, a flat estimated Phillips curve in the NY Fed DSGE model contributes toward explaining why inflation has been persistently below target over the past decade.

We now repeat the same exercise, but under the alternative policy rule of average inflation targeting (AIT). We focus on this specific rule because it is a much discussed monetary policy strategy in the current

31. The two values for κ_p, as well as the values of the other DSGE model parameters used to generate the forecasts, are the modal estimates of the two-regime estimation where only the slope of the price Phillips curve is allowed to be different across the two regimes. The forecast is computed under the assumption that κ_p returns to its pre-1990 value starting in 2020:Q1.

32. In his discussion, Olivier Blanchard makes the point that markups, which are the inverse of the real marginal costs shown in figure 13, are more countercyclical in recent decades than they had been before. His evidence is consistent with figure 13: after 1990 the economy has become more Keynesian and real marginal costs fall more, and more persistently, following a recession.

Figure 13. DSGE Forecasts of Core PCE Inflation and Marginal Costs under Different Monetary Policies

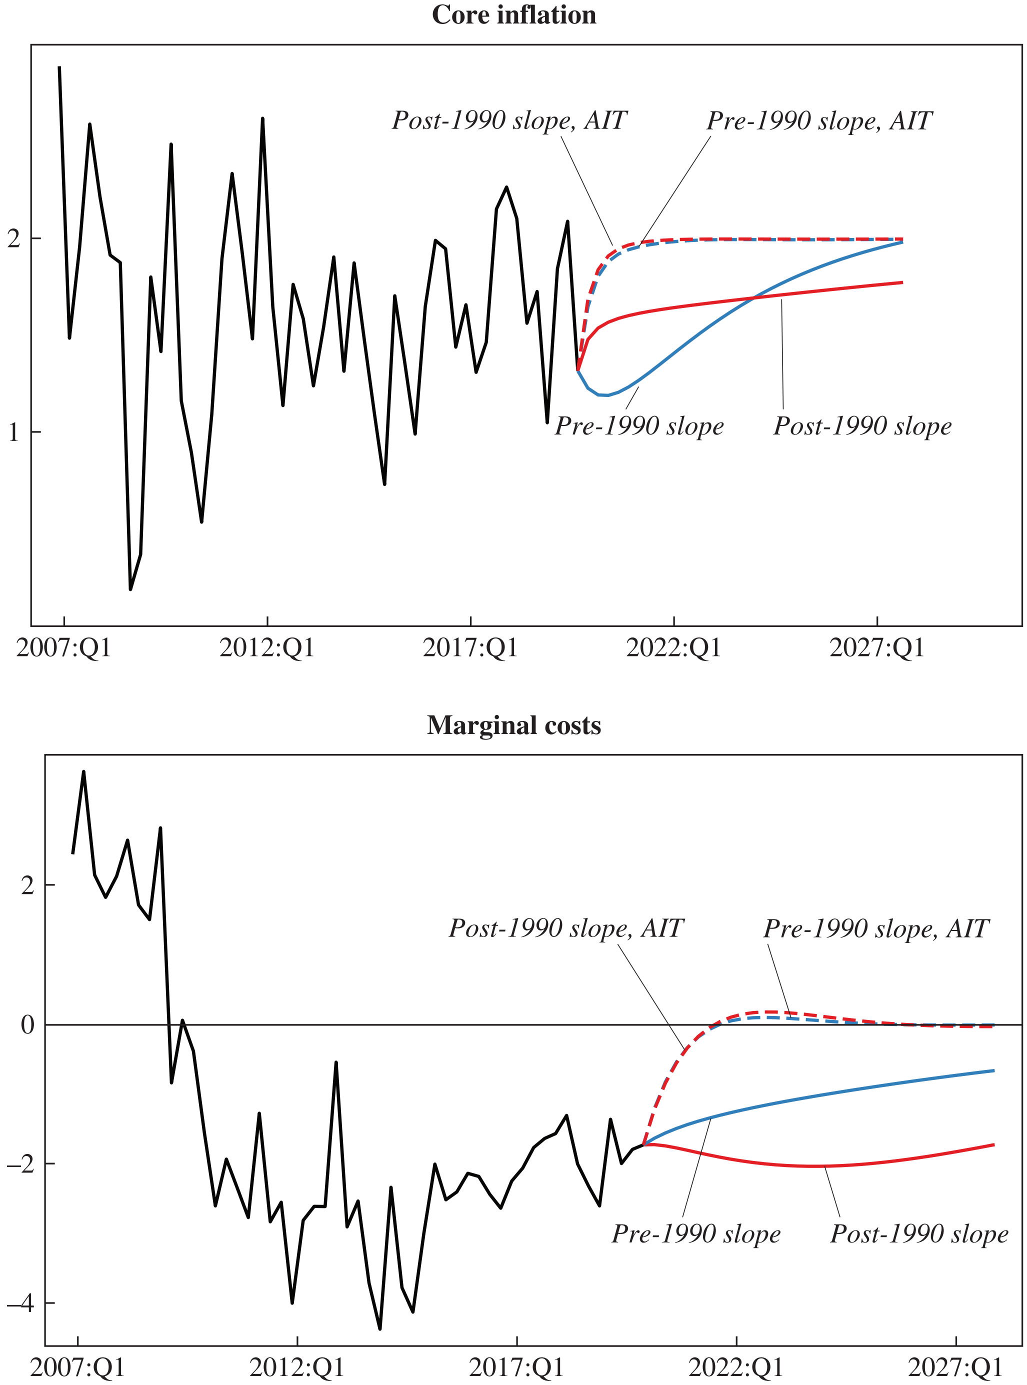

Source: Authors' calculations.

Note: These forecasts use the modal estimate of κ_p in the pre- and post-1990 period, under the estimated policy rule, and a counterfactual policy of average inflation targeting. The lines going back to 2007 represent data for inflation and the smoothed values of marginal costs, which are not observed. The two values for κ_p, as well as the values of the other DSGE parameters used to generate the forecasts, are the modal estimates of the two-regime estimation where only the slope of the price Phillips curve is allowed to be different across the pre- and post-1990 regimes.

debate (Mertens and Williams 2019). The exact AIT rule used in the counterfactual simulations is

$$R_t = \frac{0.25}{1-\gamma}\pi gap_t, \tag{9}$$

where the inflation gap is computed as

$$\pi gap_t = \pi_t + \gamma \pi gap_{t-1}$$

and all the variables are in deviation from their steady state. For inflation, this is 2 percent—the Federal Open Market Committee's long-run goal.[33]

The dashed lines in figure 13 show that switching to AIT has a more significant effect on the dynamics of marginal costs in the low κ_p economy: the gap between the solid and dashed lines is larger and more persistent with the low post-1990 slope than with the high pre-1990 one. As a consequence, marginal costs quickly revert to steady state, leading to nearly identical paths of inflation in the high and low κ_p economies.

In sum, a flat Phillips curve requires the monetary authority to work harder to stabilize inflation: unemployment needs to get lower to bring inflation back to target after a recession, everything else being equal. In equilibrium, however, a flat Phillips curve also makes the economy more Keynesian, implying that systematic monetary policy can persistently affect the dynamics of marginal costs (Del Negro, Giannoni, and Schorfheide 2015). As a result, the ability of policy to achieve its objectives is not compromised. A corollary of this general principle is that systematic policies like average inflation targeting could be especially effective in bringing inflation back to 2 percent in the current environment, as shown by the simulations above. An important caveat to this sanguine conclusion is that we obtained it in a rational expectations New Keynesian model where private agents perfectly understand the monetary policy strategy in place.[34]

33. The AIT gap πgap_t is the discounted average of past inflation rates, where we set $\gamma = 0.93$ to produce a half-life of ten quarters. In this example, we set $\pi gap_t = 0$ at the beginning of the forecast, resulting in no overshooting. Lower initial values of the gap, reflecting the accumulated shortfall of inflation from target over the past few years, would generate an overshoot of inflation over 2 percent.

34. One relevant feature of this class of models is the so-called forward guidance puzzle (Del Negro, Giannoni, and Patterson 2015), according to which policies based on promises of future actions are too successful in stabilizing the economy. This is one of the channels through which AIT works well to stabilize inflation in our simulations. For an assessment of the effectiveness of AIT under bounded rationality, see Budianto, Nakata, and Schmidt (2020).

However, the idea that monetary policy has a tighter grip on the real economy when aggregate supply is flatter goes beyond the confines of this specific class of models. We conclude that, even in the current environment, monetary policy can be as effective as it ever was in achieving price stability, as long as it pursues the appropriate strategy.

VII. Concluding Remarks

How do inflation and unemployment co-move over the business cycle? This paper explored how the answer to this question changed over the past few decades, and why. Unlike much of the existing Phillips curve literature, we address this question in a multivariate, dynamic context—studying impulse responses to typical unemployment shocks in a VAR, rather than regression coefficients in a Phillips curve—because the persistence of the business cycle impulses matters, and because looking jointly at several variables provides more clues on possible mechanisms and explanations.

We find that the persistence of unemployment fluctuations has increased somewhat after 1990, consistent with longer but shallower recoveries. The same is true of many other measures of labor and goods market activity that are commonly employed as proxies for cost pressures, including real wages and unit labor costs. The exception to this picture of relative stability over time is inflation, which has become far less sensitive to business cycle shocks.

Who is the culprit? Together, our findings rule out explanations that hinge on unemployment having become less relevant as an indicator of wage and price pressures or on the demise of the wage Phillips curve. We are left to consider two main possibilities, which have potentially very different implications for monetary policy. One is a lower slope of the price Phillips curve, leading to a more Keynesian economy, in which demand shocks dominate business cycles. The other is that policy is better able and willing to stabilize inflation, making the economy more neoclassical, with fluctuations dominated by supply shocks instead.

The casual observation that inflation and unemployment were on the same side of the Federal Reserve's dual mandate for many years after the Great Recession casts some doubt on the policy hypothesis. To support this observation more formally, we use a structural VAR and the NY Fed DSGE model to identify shocks and mechanisms. The SVAR suggests that inflation has become less responsive to shocks to credit spreads that have a large impact on real activity. If these shocks mostly shift aggregate demand, this evidence supports the slope hypothesis. This is the same conclusion reached by the DSGE model. Although it imposes more restrictions than the VAR, the DSGE can replicate the VAR impulse responses, and it attributes the reduction in the sensitivity of inflation to business cycle shocks to a reduction in the slope of the structural Phillips curve. Changes in the policy parameters, alone, are less successful at explaining the facts.

Although our analysis points more decisively in the direction of the slope hypothesis than of the policy hypothesis, it does not imply that monetary policy did not play a role in stabilizing inflation. By most accounts, it was the Federal Reserve under Volcker that brought inflation under control in the early 1980s, when our estimates find that the price Phillips curve was still alive and well. Moreover, our study leaves a number of important questions unanswered. First among them are what structural forces underlie the reduced sensitivity of inflation to cost pressures. We leave this question for future research.

ACKNOWLEDGMENTS Prepared for the 2020 Spring Brookings Papers on Economic Activity. We thank Ethan Matlin and especially William Chen for excellent research assistance, our discussants Olivier Blanchard and Chris Sims, as well as Richard Crump, Keshav Dogra, Simon Gilchrist, Pierre-Olivier Gourinchas, Rick Mishkin, Argia Sbordone, and the conference participants for useful comments and suggestions. Jim Stock provided much appreciated guidance during the research and writing process. The views expressed in this paper are those of the authors and do not necessarily represent those of the European Central Bank, the Eurosystem, the Federal Reserve Bank of New York, or the Federal Reserve System.

References

Abraham, Katherine G., John C. Haltiwanger, and Lea E. Rendell. 2020. "How Tight Is the US Labor Market?" In the present volume of *Brookings Papers on Economic Activity*.

Afrouzi, Hassan, and Choongryul Yang. 2020. "Dynamic Rational Inattention and the Phillips Curve." Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3465793#.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2019. "Business Cycle Anatomy." Working Paper. https://economics.mit.edu/files/15230.

Atkeson, Andrew, and Lee E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve of Minneapolis Quarterly Review* 25, no. 1: 2–11.

Auer, Raphael, and Andreas M. Fischer. 2010. "The Effect of Low-Wage Import Competition on U.S. Inflationary Pressure." *Journal of Monetary Economics* 57, no. 4: 491–503.

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

Ball, Laurence, and Sandeep Mazumder. 2019. "A Phillips Curve with Anchored Expectations and Short-Term Unemployment." *Journal of Money, Credit and Banking* 51, no.1: 111–37.

Bańbura, Marta, Domenico Giannone, and Michele Lenza. 2015. "Conditional Forecasts and Scenario Analysis with Vector Autoregressions for Large Cross-Sections." *International Journal of Forecasting* 31, no. 3: 739–56.

Barnichon, Régis, and Geert Mesters. 2019a. "Identifying Modern Macro Equations with Old Shocks," Discussion Paper 13765. London: Centre for Economic Policy Research. https://cepr.org/active/publications/discussion_papers/dp.php?dpno=13765.

Barnichon, Régis, and Geert Mesters. 2019b. "The Phillips Multiplier." Discussion Paper 13480. London: Centre for Economic Policy Research. https://cepr.org/active/publications/discussion_papers/dp.php?dpno=13480.

Bernanke, Ben S. 2007. "Inflation Expectations and Inflation Forecasting." Speech given at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Mass., July 10.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist. 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics*, vol. 1C, edited by John B. Taylor and Michael Woodford. Amsterdam: North-Holland.

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

Blanchard, Olivier, Eugenio Cerutti, and Lawrence Summers. 2015. "Inflation and Activity—Two Explorations and Their Monetary Policy Implications." Working Paper 21726. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w21726.

Budianto, Flora, Taisuke Nakata, and Sebastian Schmidt. 2020. "Average Inflation Targeting and the Interest Rate Lower Bound." Working Paper 852. Basel: Bank for International Settlements. https://www.bis.org/publ/work852.htm.

Cai, Michael, Marco Del Negro, Marc P. Giannoni, Abhi Gupta, Pearl Li, and Erica Moszkowski. 2019. "DSGE Forecasts of the Lost Recovery." *International Journal of Forecasting* 35, no. 4: 1770–89.

Caldara, Dario, Cristina Fuentes-Albero, Simon Gilchrist, and Egon Zakrajšek. 2016. "The Macroeconomic Impact of Financial and Uncertainty Shocks." *European Economic Review* 88, 185–207.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12, no. 1: 383–98.

Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston. 2019. "Anchored Inflation Expectations." Working Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3018198.

Cavallo, Alberto. 2018. "More Amazon Effects: Online Competition and Pricing Behaviors." Working Paper 25138. Cambridge, Mass.: National Bureau of Economic Research.

Christiano, Lawrence J., Martin S. Eichenbaum, and Mathias Trabandt. 2015. "Understanding the Great Recession." *American Economic Journal: Macroeconomics* 7, no. 1: 110–67.

Christiano, Lawrence J., R. Motto, and M. Rostagno. 2014. "Risk Shocks." *American Economic Review* 104, no. 1: 27–65.

Clark, Todd E., and Taeyoung Doh. 2014. "Evaluating Alternative Models of Trend Inflation." *International Journal of Forecasting* 30, no. 3: 426–48.

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.

Coibion, Olivier, Yuriy Gorodnichenko, and Dmitri Koustas. 2013. "Amerisclerosis? The Puzzle of Rising U.S. Unemployment Persistence." *Brookings Papers on Economic Activity*, Fall, 193–260.

Crump, Richard K., Stefano Eusepi, Marc P. Giannoni, and Ayşegül Şahin. 2019. "A Unified Approach to Measuring *u**." *Brookings Papers on Economic Activity*, Spring, 143–238.

Daly, Mary C., and Bart Hobijn. 2014. "Downward Nominal Wage Rigidities Bend the Phillips Curve." *Journal of Money, Credit and Banking* 46, no. S2: 51–93.

Davis, Steven J. 2019. Comment on "A Unified Approach to Measuring *u**." *Brookings Papers on Economic Activity*, Spring, 223–30.

Del Negro, Marco, and Stefano Eusepi. 2011. "Fitting Observed Inflation Expectations." *Journal of Economic Dynamics and Control* 35, no. 12: 2105–31.

Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti. 2017. "Safety, Liquidity, and the Natural Rate of Interest." *Brookings Papers on Economic Activity*, Spring, 235–316.

Del Negro, Marco, Marc P. Giannoni, and Christina Patterson. 2015. "The Forward Guidance Puzzle." Staff Report 574. Federal Reserve Bank of

New York. https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr574.pdf.

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.

Del Negro, Marco, and Frank Schorfheide. 2004. "Priors from General Equilibrium Models for VARs." *International Economic Review* 45, no. 2: 643–73.

Del Negro, Marco, and Frank Schorfheide. 2013. "DSGE Model-Based Forecasting." In *Handbook of Economic Forecasting*, vol. 2, edited by Graham Elliott and Allan Timmermann. Amsterdam: North-Holland.

Eggertsson, Gauti B., and Sergey K. Egiev. 2020. "Fundamental Driven Liquidity Traps: A Unified Theory of the Great Depression and the Great Recession." Working Paper. https://drive.google.com/file/d/1vKfVvF9RwsMiMoqdEdOwfbrdLbuM345/view.

Faccini, Renato, and Leonardo Melosi. 2020. "Bad Jobs and Low Inflation." Working Paper 2020-09. Federal Reserve Bank of Chicago. https://www.chicagofed.org/publications/working-papers/2020/2020-09.

Fernald, John. 2012. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Working Paper 2012-19. Federal Reserve Bank of San Francisco. http://www.frbsf.org/economic-research/publications/working-papers/2012/19/.

Fitzgerald, Terry J., and Juan Pablo Nicolini. 2014. "Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from U.S. Cities." Working Paper 713. Federal Reserve Bank of Minneapolis. https://www.minneapolisfed.org/research/working-papers/is-there-a-stable-relationship-between-unemployment-and-future-inflation-evidence-from-us-cities.

Forbes, Kristin J. 2019a. "Inflation Dynamics: Dead, Dormant, or Determined Abroad?" *Brookings Papers on Economic Activity*, Fall, 257–319.

Forbes, Kristin J. 2019b. "How Have Shanghai, Saudi Arabia, and Supply Chains Affected U.S. Inflation Dynamics?" *Review* 101, no. 1: 27–44.

Forbes, Kristin J., Joseph E. Gagnon, and Christopher G. Collins. 2020. "Low Inflation Bends the Phillips Curve around the World." Working Paper 20-06. Washington: Peterson Institute for International Economics. https://www.piie.com/publications/working-papers/low-inflation-bends-phillips-curve-around-world.

Fuhrer, Jeffrey C. 2010. "Inflation Persistence." In *Handbook of Monetary Economics*, edited by Benjamin M. Friedman and Michael Woodford. Amsterdam: North-Holland.

Fuhrer, Jeffrey C., and George Moore. 1995. "Inflation Persistence." *Quarterly Journal of Economics* 110, no. 1: 127–59.

Galí, Jordi. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, 2nd ed. Princeton, N.J.: Princeton University Press.

Galí, Jordi, and Luca Gambetti. 2019. "Has the U.S. Wage Phillips Curve Flattened? A Semi-Structural Exploration." Working Paper 25476. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w25476.

Galí, Jordi, and Luca Gambetti. 2020. "On the Empirical (Ir)relevance of the Zero Lower Bound Constraint." in *NBER Macroeconomics Annual* 34:141–70.

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

Geerolf, François. 2020. "The Phillips Curve: A Relation between Real Exchange Rate Growth and Unemployment." Working Paper. https://fgeerolf.com/phillips.pdf.

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri. 2015. "Prior Selection for Vector Autoregressions." *Review of Economics and Statistics* 97, no. 2: 436–51.

Giannone, Domenico, Michele Lenza, and Lucrezia Reichlin. 2019. "Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed since the Crisis?" *International Journal of Central Banking* 15, no. 5: 137–73.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek. 2017. "Inflation Dynamics during the Financial Crisis." *American Economic Review* 107, no. 3: 785–823.

Gilchrist, Simon, Vladimir Yankov, and Egon Zakrajšek. 2009. "Credit Market Shocks and Economic Fluctuations: Evidence from Corporate Bond and Stock Markets." *Journal of Monetary Economics* 56, no. 4: 471–93.

Gilchrist, Simon, and Egon Zakrajšek. 2012. "Credit Spreads and Business Cycle Fluctuations." *American Economic Review* 102, no. 4: 1692–720.

Goldfeld, Stephen M., and Alan S. Blinder. 1972. "Some Implications of Endogenous Stabilization Policy." *Brookings Papers on Economic Activity*, no. 3: 585–644.

Goolsbee, Austan D., and Peter J. Klenow. 2018. "Internet Rising, Prices Falling: Measuring Inflation in a World of E-Commerce." In *American Economic Association Papers and Proceedings*, vol. 108. Pittsburgh, Pa.: AEA Publications.

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

Gordon, Robert J. 2018. "Friedman and Phelps on the Phillips Curve Viewed from a Half Century's Perspective." *Review of Keynesian Economics* 6, no. 4: 425–36.

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

Hazell, Jonathan, Juan Herreno, Emi Nakamura, and Jón Steinsson. 2020. "The Slope of the Phillips Curve: Evidence from US States." Working Paper.

Hong, Gee Hee, Zsóka Kóczán, Weicheng Lian, and Malhar Nabar. 2018. "More Slack than Meets the Eye? Recent Wage Dynamics in Advanced Economies." Working Paper 18/50. Washington: International Monetary Fund. https://www.imf.org/en/Publications/WP/Issues/2018/03/09/More-Slack-than-Meets-the-Eye-Recent-Wage-Dynamics-in-Advanced-Economies-45692.

Hooper, Peter, Frederic S. Mishkin, and Amir Sufi. 2020. "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?" *Research in Economics* 74, no. 1: 26–62.

Jones, Charles I. 2018. *Macroeconomics*, 4th ed. New York: W.W. Norton.

Jorgensen, Peter Lihn, and Kevin J. Lansing. 2019. "Anchored Inflation Expectations and the Flatter Phillips Curve." Working Paper 2019-27. Federal Reserve Bank of San Francisco. https://www.frbsf.org/economic-research/publications/working-papers/2019/27/.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2010. "Investment Shocks and Business Cycles." *Journal of Monetary Economics* 57, no. 2: 132–45.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2011. "Investment Shocks and the Relative Price of Investment." *Review of Economic Dynamics* 14, no. 1: 101–21.

Kareken, John, and Robert M. Solow. 1963. "Lags in Fiscal and Monetary Policy." In *Stabilization Policies: A Series of Research Studies Prepared for the Commission on Money and Credit*, edited by E. Cary Brown. Englewood Cliffs, N.J.: Prentice Hall.

Laséen, Stefan, and Lars E. O. Svensson. 2011. "Anticipated Alternative Policy-Rate Paths in Policy Simulations." *International Journal of Central Banking* 7, no. 3: 1–35.

Lindé, Jesper, and Mathias Trabandt. 2019. "Resolving the Missing Deflation Puzzle." Discussion Paper 13690. London: Centre for Economic Policy Research. https://cepr.org/active/publications/discussion_papers/dp.php?dpno=13690.

Mackowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99, no. 3: 769–803.

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117, no. 4: 1295–328.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock. 2014. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature* 52, no. 1: 124–88.

McLeay, Michael, and Silvana Tenreyro. 2019. "Optimal Inflation and the Identification of the Phillips Curve." *NBER Macroeconomics Annual* 34:199–255.

Mertens, Thomas M., and John C. Williams. 2019. "Monetary Policy Frameworks and the Effective Lower Bound on Interest Rates." In *American Economic Association Papers and Proceedings*, vol. 109. Pittsburgh, Pa.: AEA Publications.

Moulton, Brent R. 2018. *The Measurement of Output, Prices, and Productivity: What's Changed since the Boskin Commission?* Washington: Brookings Institution. https://www.brookings.edu/research/the-measurement-of-output-prices-and-productivity/.

Obstfeld, Maurice. 2020. "Global Dimensions of U.S. Monetary Policy." *International Journal of Central Banking* 16, no. 1: 73–132.

Orphanides, Athanasios, and John C. Williams. 2005. "Inflation Scares and Forecast-Based Monetary Policy." *Review of Economic Dynamics* 8, no. 2: 498–527.

Peach, Richard, Robert Rich, and M. Henry Linder. 2013. "The Parts Are More than the Whole: Separating Goods and Services to Predict Core Inflation." *Current Issues in Economics and Finance* 19, no. 7: 1–8.

Powell, Jerome H. 2019. "Challenges for Monetary Policy." Speech given at the Economic Policy Symposium: Challenges for Monetary Policy, Federal Reserve Bank of Kansas City, Jackson Hole, Wyo., August 23.

Primiceri, Giorgio E. 2006. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy." *Quarterly Journal of Economics* 121, no. 3: 867–901.

Primiceri, Giorgio E. 2019. Comment on "A Unified Approach to Measuring u^*," by Richard K. Crump, Stefano Eusepi, Marc Giannoni, and Ayşegül Şahin. *Brookings Papers on Economic Activity*, Spring, 223–230.

Rognlie, Matthew. 2019. "Comment on Optimal Inflation and the Identification of the Phillips Curve." *NBER Macroeconomics Annual* 34:267–79.

Rubbo, Elisa. 2020. "Networks, Phillips Curves and Monetary Policy." Working Paper. https://scholar.harvard.edu/files/elisarubbo/files/rubbo_jmp.pdf.

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

Sbordone, Argia M. 2007. "Globalization and Inflation Dynamics: The Impact of Increased Competition." In *International Dimensions of Monetary Policy*, edited by Jordi Galí and Mark J. Gertler. Chicago: University of Chicago Press.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97, no. 3: 586–606.

Stansbury, Anna, and Lawrence H. Summers. 2020. "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy." In the present volume of *Brookings Papers on Economic Activity*.

Stock, James H. 2011. Comment on "Inflation Dynamics and the Great Recession," by Laurence Ball and Sandeep Mazumder, *Brookings Papers on Economic Activity*, Spring, 387–402.

Stock, James H., and M. W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39, no. s1: 3–33.

Stock, James H., and M. W. Watson. 2008. "Phillips Curve Inflation Forecasts." Working Paper 14322. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w14322.

Stock, James H., and M. W. Watson. 2019. "Slack and Cyclically Sensitive Inflation." Working Paper 25987. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w25987.

Swanson, Eric T., and John C. Williams. 2014. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." *American Economic Review* 104, no. 10: 3154–85.

Tallman, Ellis W., and Saeed Zaman. 2017. "Forecasting Inflation: Phillips Curve Effects on Services Price Measures." *International Journal of Forecasting* 33, no. 2: 442–57.

Taylor, John B. 1980. "Aggregate Dynamics and Staggered Contracts." *Journal of Political Economy* 88, no. 1: 1–23.

Taylor, John B. 1993. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy* 39:195–214.

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

Yellen, Janet. 2019. "Former Fed Chair Janet Yellen on Why the Answer to the Inflation Puzzle Matters." Remark, Hutchins Center on Fiscal and Monetary Policy at Brookings, October 3. https://www.brookings.edu/research/former-fed-chair-janet-yellen-on-why-the-answer-to-the-inflation-puzzle-matters/.

Zhang, Bo, Joshua C. C. Chan, and Jamie L. Cross. Forthcoming. "Stochastic Volatility Models with ARMA Innovations: An Application to G7 Inflation Forecasts." *International Journal of Forecasting*.

Comments and Discussion

COMMENT BY

OLIVIER BLANCHARD[1] This is a very nice paper. On the methodological front, it moves away from the traditional single equation estimation of the Phillips curve relation and shows the many useful ways in which one can go back and forth between vector autogression (VAR) representations and dynamic stochastic general equilibrium (DSGE) models to think about the relation between inflation and unemployment. On the conclusions front, the paper largely confirms current wisdom, which is reassuring. It reaches three main conclusions: The relation between inflation and unemployment reflects primarily a causal effect of unemployment on inflation. The effect of unemployment on price inflation has become weaker. The effect of unemployment on wage inflation has also become weaker but less so.

In my comments, I shall focus mostly on one aspect of their results, the adjustment of prices and wages to movements in unemployment. I shall argue that we should focus more on the movements in markups—the ratio of prices to nominal unit labor costs—and think about their policy implications.

First, however, I shall briefly comment on two aspects of their methodology.

The first is, rightly, a focus on potential reverse causality from inflation to unemployment, leading to a correlation between the error term and unemployment. A recent paper by McLeay and Tenreyro (2019) has argued that a change in monetary policy could indeed be the explanation behind the change in the Phillips curve relation. To explore the issue, the authors use the VAR equivalent of an instrumental variable single equation approach,

1. Thanks to Mark Bils for discussions and to Chris Erceg for simulations.

by looking at the co-movements in inflation and unemployment conditional on a demand shock. They take as demand shocks movements in the excess bond premium, an unusual choice and one which strikes me as somewhat unconvincing. The factors behind the movements in the premium, namely, variations in market risk aversion, may well have both strong supply and demand effects. I would have preferred the use of a more conventional instrument, such as movements in cyclically adjusted fiscal balances or Romer-Romer fiscal series, which seem more likely to affect demand than supply in the short run. The authors conclude that reverse causality does not appear to play a major role in explaining the change in the relation between unemployment and inflation. I suspect the use of a more standard instrument would lead to the same conclusion.

The second is a master class in how to go from reduced form VARs to structural VARs to estimated DSGEs and back. The idea of using VARs to derive impulse price responses conditional on similar impulse responses for unemployment across two samples is clever—if probably subject to a Lucas critique. Comparing implied VAR representations from DSGE models to actual VARs under different assumptions about parameters to get a sense of how different sets of assumptions fit various aspects of the data is also clever and well done. I learned a lot from the paper.

Turning now to the conclusions. The various approaches used in the paper all suggest that the price response to unemployment has significantly decreased over time. The evidence on the response of wages to unemployment is less conclusive. Their VAR approach suggests that the wage response has also decreased, but less than that of prices. The DSGE approach gives a stronger, more statistically significant, answer, namely, that the wage response has decreased by less than the price response.

I want to focus in my comments on the behavior of markups. To come to a conclusion about the behavior of markups, however, one needs to know the response of nominal unit costs—and thus not only the response of wages but also the response of productivity—to unemployment. In private conversation, the authors have kindly given me the response of productivity, which looks largely similar in both samples. I conclude that the authors provide some evidence of countercyclical markups: when unemployment increases, wages decrease, but markups increase, more so than they used to, leading to only a muted response of prices.

I shall take this as license to think about the cyclical behavior of markups and its policy implications. I do so because there is substantial evidence that the behavior of markups is more complex and more important than we used to believe. In the Dixit-Stiglitz model of monopolistic

Figure 1. Breakdown of the Year-on-Year Change in Core Inflation in the Euro Area

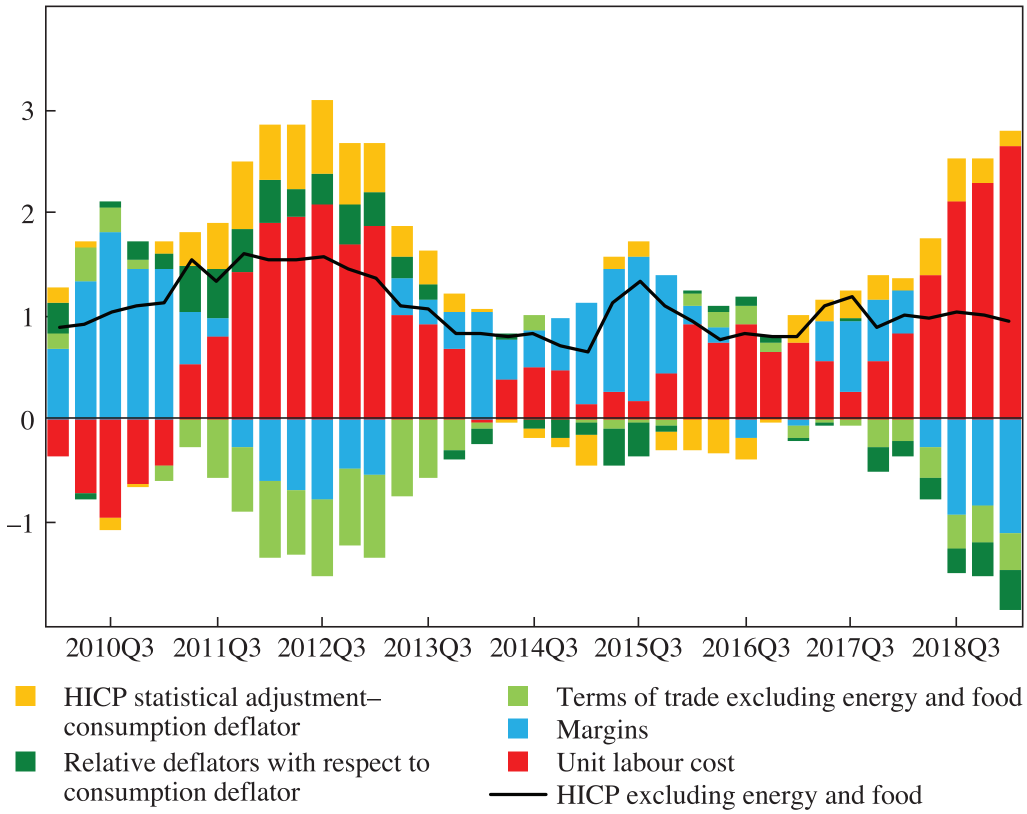

Source: Diev, Kalantzis, and Lalliard (2019).

competition—still a standard building block of most DSGE models—the elasticity of substitution between goods is constant, and so is the markup. Reality suggests a more complex picture.

Let me go through two additional pieces of evidence.

The first is the evolution of markups in the euro zone. Figure 1, taken from Diev, Kalantzis, and Lalliard (2019), shows the evolution of core inflation since 2010 in the euro zone and its decomposition between terms of trade effects, nominal unit labor costs, and markups. Concentrating on the end of the sample, two aspects are striking. The first is the increase in unit labor costs inflation since 2015, reflecting (behind the scene) the pressure of lower unemployment on wage inflation. The second is how this increase in cost inflation has not translated in price inflation: as costs have increased, markups have decreased nearly in unison. In other words, markups have been strongly countercyclical, decreasing as output increased and unemployment decreased.

The second piece of evidence builds on the work of Bils, Klenow, and Malin (2018), who argued that markups were strongly countercyclical. Following the approach of their paper, I construct series for markups, real output, and nominal unit labor costs in the US nonfarm business sector. Markups are constructed as the ratio of the price deflator to the nominal unit labor costs, thus as the inverse of the labor share. The results of a regression of log markups on log real output and log nominal unit costs are given for two subsamples, 1980–1999 and 2000–2019. All three series are hp-filtered.

$$\text{1980–1999} \qquad \mu = -0.10y - 0.08ulc + \epsilon$$

$$R^2 = 0.09, \sigma(\mu) = 1.2 \text{ percent}$$

$$\text{2000–2019} \qquad \mu = -0.23^{**}y - 0.60^{**}\,ulc + \epsilon$$

$$R^2 = 0.69, \sigma(\mu) = 1.9 \text{ percent}$$

**significant at the 95 percent confidence level

In both samples, the coefficients on output and nominal unit labor costs are negative. Markups indeed appear to be countercyclical, and an increase in nominal costs leads to a decrease in markups, suggesting real or nominal rigidities. The results are, however, significantly different across the two samples. In the earlier sample, the coefficients are small and insignificant. In the later sample, both coefficients are large and strongly significant. Also, the variation in markups, measured by the standard deviation, is substantially higher in the later sample.

This difference across the two samples is robust to other ways of detrending, such as first differencing, and other ways of measuring marginal cost, not using wages, which Bils, Klenow, and Malin (2018) insisted might not represent the marginal cost of labor but rather the price of materials. It suggests that, for various reasons, prices have become more inertial, and markups more variable than they used to be. I shall not explore what the reasons might be, but take these two facts as stylized facts and draw potential policy implications. Let me state the bottom line: both increased price rigidity and larger markup shocks strengthen the case for wage inflation targeting over price inflation targeting.

Consider the simplistic nominal wage and price equations:

$$w = ap + (1 - a)Ep - cu + \eta$$

$$p = bw + (1 - b)Ew + \epsilon.$$

The coefficients a and b proxy for the stickiness of wages and prices. The lower a or b, the more predetermined prices or wages. Assume ϵ and η are white noise and capture distortions, so that the optimal level of unemployment is invariant to any of the two shocks and, given the white noise assumptions, equal to zero.

Consider now the effect of a markup shock, $\epsilon > 0$. Given the white noise assumption, Ep and Ew are both equal to zero, so the two equations become:

$$w = ap - cu$$

$$p = bw + \epsilon.$$

Now consider the effects of the shock under two alternative monetary regimes, price inflation targeting, $p = 0$, and wage inflation targeting, $w = 0$.

Under price inflation targeting, unemployment is given by:

$$u = \frac{1}{bc}\epsilon.$$

Under wage inflation targeting, unemployment is given by:

$$u = \frac{a}{c}\epsilon.$$

Given that $a \leq 1$, the implication is that, in response to markup shocks, wage inflation targeting always dominates price inflation targeting. And that the higher the price rigidity (the lower b) and the higher the variance of markup shocks, the more wage inflation targeting dominates.

The intuition is straightforward. In response to a positive markup shock, maintaining a stable price level requires a decrease in the nominal wage, and this decrease requires in turn an increase in unemployment. The less the price responds to the wage, the more the wage has to decrease and the more the unemployment rate has to increase. In contrast, maintaining a stable wage level, and thus allowing the price to increase with the markup, requires a smaller increase in unemployment.

The same conclusions follow from simulations of a larger model. The simulations below—courtesy of Chris Erceg, with whom I am doing joint work on this issue—show the effects of a markup shock in a midsize New Keynesian model with nominal wage and nominal price rigidities. The simulations show the effect of an AR(1) markup shock, with AR

Figure 2. Output Response to a Markup Shock

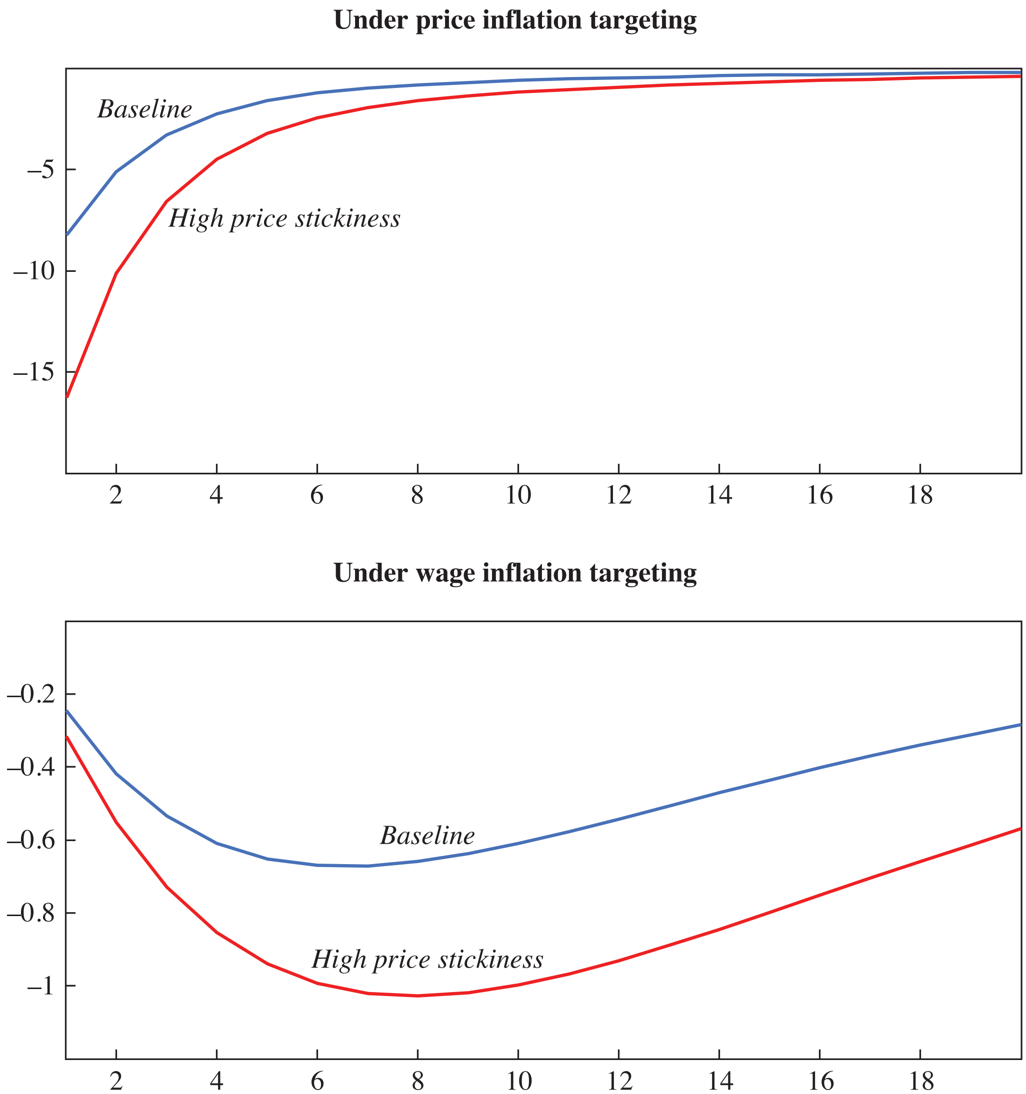

Source: Author's calculations.

coefficient 0.9, under either strict price inflation targeting or strict wage inflation targeting, and in each case, with either baseline or high price rigidity (captured by the Calvo coefficient on the expected length of price setting).

Figure 2 yields two conclusions. First, under the baseline, the output cost is much lower under wage inflation than price inflation targeting, –0.2 percent on impact and –0.6 percent at the trough under wage inflation targeting, versus –8.0 percent on impact under price inflation targeting. Second, the output cost under price inflation targeting is much larger when there is high price rigidity, –16 percent versus –8 percent under the baseline.

Let me finally turn from the effects of markup shocks to the effects of demand shocks. Extend the simplistic model above to add an aggregate demand equation and an interest rate rule:

$$w = ap + (1 - a)Ep - cu + \eta$$

$$p = bw + (1 - b)Ew + \epsilon$$

$$u = r + \epsilon_u$$

$$r = \alpha\big(dp + (1 - d)w\big).$$

Unemployment is an increasing function of the policy rate and a demand shock, ϵ_u. The policy rate responds to both the price and the wage. One can think of price inflation targeting as $d = 1$ and wage inflation targeting as $d = 0$. Consider an unexpected demand shock, ϵ_u. Solving for unemployment under price inflation targeting gives:

$$u = \frac{1 - ab}{1 - ab + cb\alpha}\epsilon_u.$$

Solving for unemployment under wage inflation targeting gives instead:

$$u = \frac{1 - ab}{1 - ab + c\alpha}\epsilon_u.$$

So long as $b < 1$, the effect of the demand shock on unemployment is smaller under wage than price inflation. The intuition is again straightforward. If prices respond little to wages and thus respond little to unemployment, and if the interest rate responds to prices, the monetary response will be small, and demand shocks will have a large effect on unemployment. This is particularly clear if one considers the case where $b = 0$. Under wage targeting, $u = (1/(1 + c\alpha))\epsilon_u$, but under price targeting, $u = \epsilon_u$; in that case, as prices do not react, monetary policy does not react either, and the demand shock has a full effect on unemployment.

Again, it is useful to look at the question in a larger model. Based on the same model as above, figure 3 shows the effect of an AR(1) demand shock, with AR coefficient 0.9, under price inflation targeting and under wage inflation targeting and, in each case, under baseline and high price rigidity. It yields two conclusions. The effect of the shock on output is substantially

Figure 3. Output Response to a Demand Shock

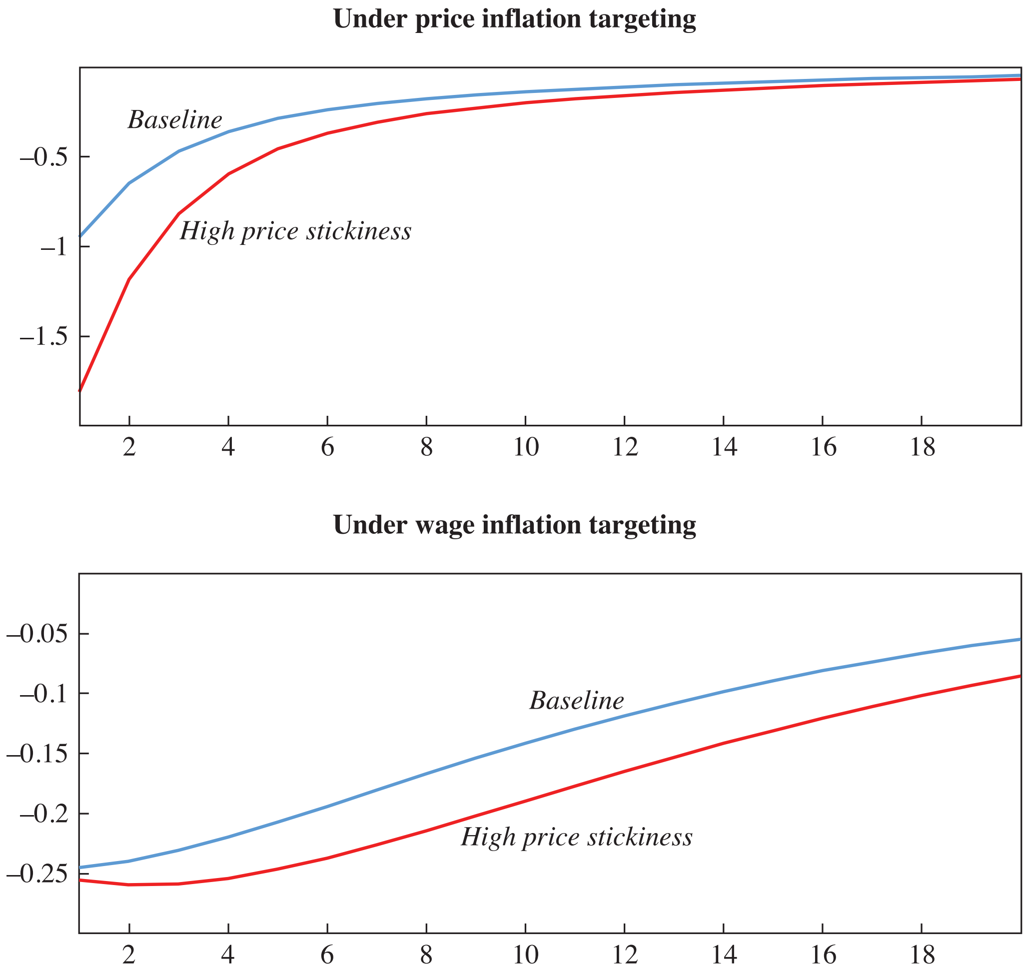

Source: Author's calculations.

smaller under wage targeting. And the higher the price rigidity, the larger the comparative advantage of wage over price inflation targeting.

The implications of nominal price rigidity for the effects of demand shocks are less robust than for the effects of markup shocks: in principle, a smaller response of prices to unemployment can be offset by a more aggressive interest rate rule (a larger α in the model above). But if the signal from prices is noisy because of movements in markups, a more aggressive rule will lead to too strong a reaction to a markup shock, an undesirable outcome.

These comments are not the place to go further in exploring wage inflation targeting. But I see the conclusions of the paper, as well as a body of other evidence, as suggesting that the nature of the Phillips curve has changed, that the behavior of prices given unit labor costs has changed, and that this is ground enough to explore wage inflation targeting.

REFERENCES FOR THE BLANCHARD COMMENT

Bils, Mark, Peter Klenow, and Benjamin Malin. 2018. "Resurrecting the Role of the Product Market Wedge in Recessions." *American Economic Review* 108, nos. 4-5: 1118–46.

Diev, Pavel, Yannick Kalantzis, and Antoine Lalliard. 2019. "Why Have Strong Wage Dynamics Not Pushed Up Inflation in the Euro Area?" Bulletin de la Banque de France 225/6, September-October. https://publications.banque-france.fr/en/why-have-strong-wage-dynamics-not-pushed-inflation-euro-area.

McLeay, Michael, and Silvana Tenreyro. 2019. "Optimal Inflation and the Identification of the Phillips Curve." *NBER Macroeconomics Annual* 34:199–255.

COMMENT BY

CHRISTOPHER A. SIMS The main fact the paper documents is that the relation between a one-dimensional measure of real activity and the rate of wage or price inflation has weakened between pre-1990 and post-1990 periods in the United States. This is a robust result, emerging with a variety of measures of real activity and a variety of time series modeling approaches. But despite using tools capable of unraveling a richer story, the paper maintains a one-dimensional conceptual framework, with a single driving force for a "business cycle" that moves all variables in a repeating pattern. The paper acknowledges that this does not account for all variation in inflation but suggests that longer-run variation in inflation can be set aside when focusing on its "cyclical" variation. This leads the paper into a discussion of why inflation has become less variable over a long span of time, while making no effort to account for the fact that inflation's level has come steadily down over this span of time and has recently been persistently below Federal Reserve targets.

Most variation in inflation is not movement along a Phillips curve. Since the paper does not display impulse responses to disturbances other than its composite unemployment-shifting shock, I have estimated a simple three-variable vector autoregression (VAR), with monthly data on industrial production, the personal consumption expenditures (PCE) deflator, and average hourly earnings of production and nonsupervisory employees. The impulse responses, pre- and post-1990, are displayed in figures 1 and 2.

Figure 1. Pre-1990 Impulse Responses

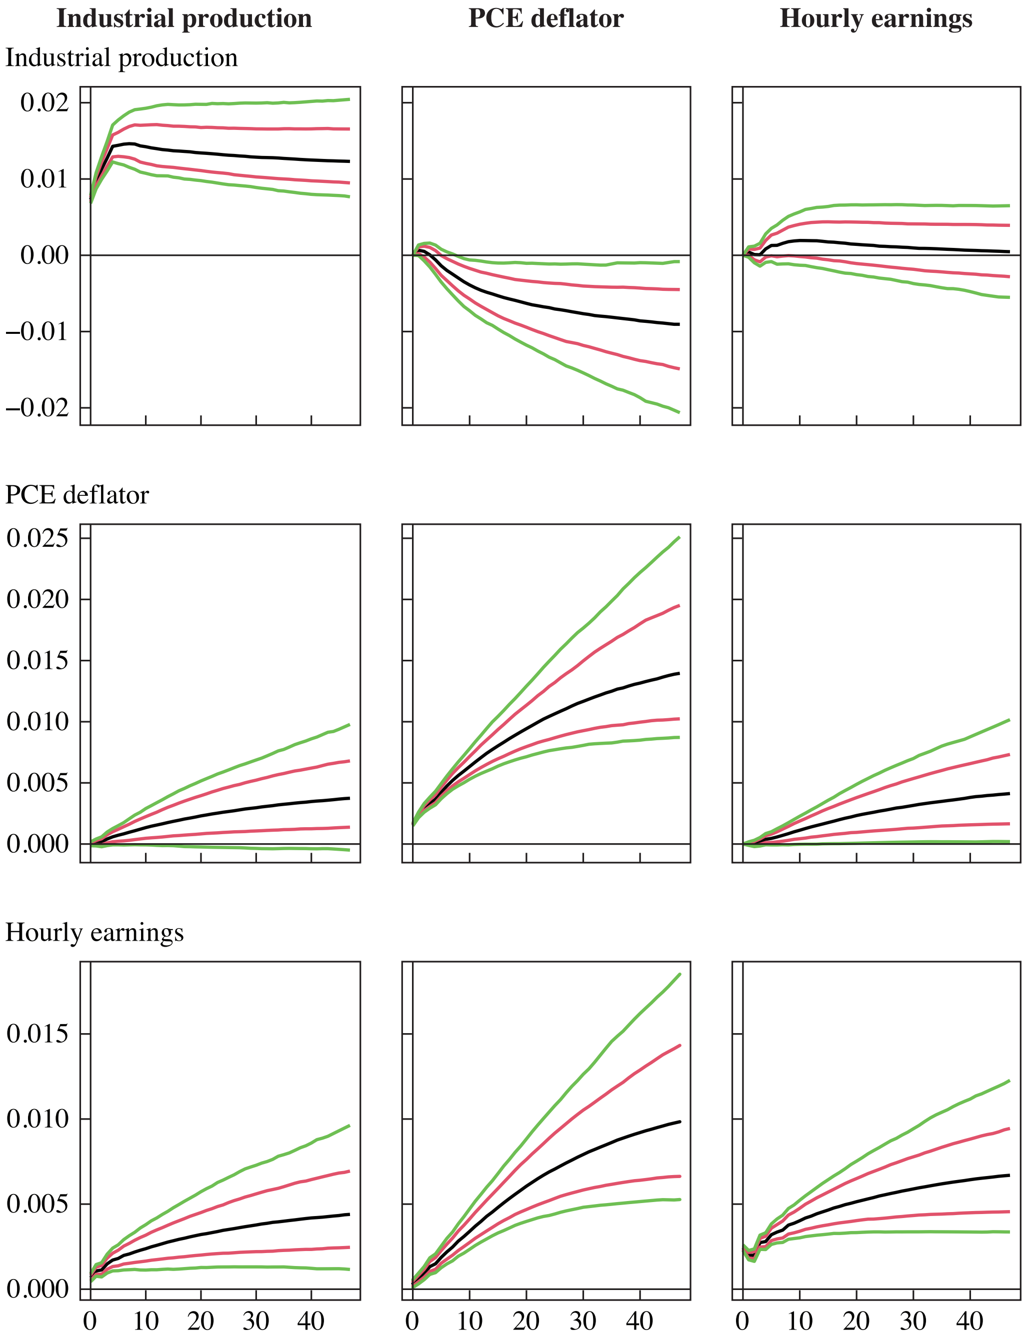

Source: Author's calculations.

Notes: Impulse responses from a VAR estimated with monthly data for industrial production, PCE deflator, and average hourly earnings for January 1965 through December 1989. The impulse responses have been orthogonalized by Cholesky decomposition, with industrial production first in the ordering. Estimation used a Minnesota prior. The two error bands are 68 percent and 90 percent bands.

Figure 2. Post-1990 Impulse Responses

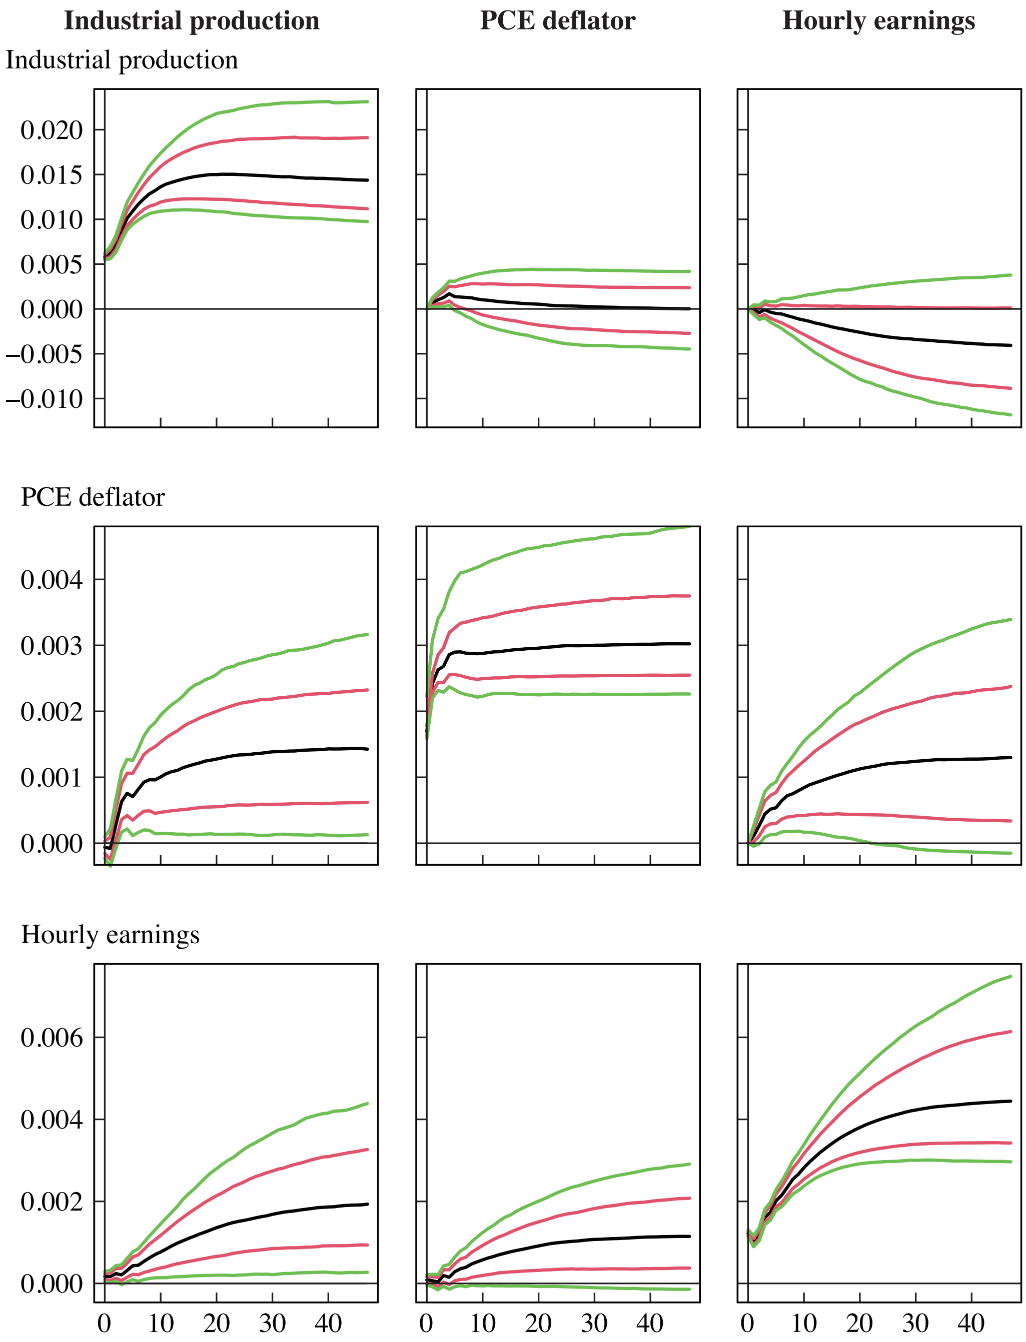

Source: Author's calculations.

Notes: Impulse responses from a VAR estimated with monthly data for industrial production, PCE deflator, and average hourly earnings for January 1990 through May 2016. The impulse responses have been orthogonalized by Cholesky decomposition, with industrial production first in the ordering. Estimation used a Minnesota prior. The two error bands are 68 percent and 90 percent bands.

The first column of each of these figures captures the main message of the paper: there is a disturbance that moves output, prices, and wages in the same direction. Scanning across the first row, we can see that this first-column shock accounts for most of the variation in output. (All plots in each row have the same scale, so that visually smaller responses account for less variation in the row variable.) The size of the first-column disturbance to industrial production is about the same in both plots, with the median estimate of long-run response of industrial production (at the end of the five years shown) around 1.4 percent in both periods. The scale of the price and wage rows of the plots is quite different in the two periods. The median estimate of the long-run response of the PCE deflator to this first-column shock is around 0.12 percent in the post-1990 period and around 0.40 percent in the earlier period. For the wage response the corresponding responses are 0.20 percent and 0.5 percent. While these differences are in the same direction as found in the paper, it is worth noting that the 68 percent error bands overlap.

What is omitted from the paper is the message of the bottom two rows of the plot arrays. For both prices and wages, in both time periods, most of the variation is being generated by disturbances that do not move output in the same direction as wages and prices. In the pre-1990 case, the second column, which accounts for most of the variance in prices and wages, is a shock that moves output down and both prices and wages up. This probably reflects the oil shock stagflation of the 1970s. In the post-1990 case the corresponding shock is in the third column. It also moves output down and inflation and wages up, though the output decline is only marginally statistically significant. The second column post-1990 and the third column pre-1990 are similar in showing almost no response of industrial output and substantial price and wage responses, though the relative importance of these shocks for explaining price inflation is greater after 1990.

The responses in the second two rows are all at a smaller scale after 1990, but the relative importance for explaining wages and prices of the first column shock, which behaves like a movement along a Phillips curve, is about the same in both periods. These responses do fit the story that the slope of the Phillips curve—the response of wages and prices to the level of business activity—has declined. They do not fit the additional claim in the paper that this is the main source of the decline in variability of inflation.

The United States, Japan, and the euro area all moved steadily toward zero interest rates and below-target inflation after 1990. Benhabib, Schmitt-Grohé, and Uribe (2001) showed us that, once the zero bound on interest rates is recognized, the standard modeling of inflation dynamics with a Taylor rule policy leads to a drift toward near-zero interest rates and low or negative inflation. Their argument depends on there not being a reliable fiscal expansion response to the occurrence of the low interest rates, but in light of recent experience this seems realistic.

The Benhabib, Schmitt-Grohé, and Uribe (2001) model can explain persistent low inflation, persistent low interest rates, inability of the Federal Reserve to affect inflation, and insensitivity of inflation to real disturbances. It does not explain the high levels of real government debt in the United States or the low real interest rates. We don't have a model that pulls all these facts together. But Benhabib, Schmitt-Grohé, and Uribe (2001) seem to offer a better starting point for understanding these facts than a narrow focus on the Phillips curve.

I also have one narrower criticism of the paper. I'm not convinced there is any substantial difference between wage and price inflation in the changes between the pre-1990 and post-1990 periods. Both in figure 3 and in figure 4 in the paper, the sizes of the responses of wage and price inflation are quite similar in the post-1990 period. In figure 2 in the paper the responses of wage inflation seem weaker, not stronger, than those of prices. So the claim that the wage Phillips curve is still important after 1990 does not seem supported by the reduced form statistical analysis. It depends on the results in section IV, which invokes the theoretical structure of the New York Federal Reserve model, and in particular the notion of separate wage and price Phillips curves with mutually uncorrelated shocks and driven by distinct real variables. The idea that workers have pricing power and control the quantity of their labor to maintain a wage markup, implicit in this specification of a wage Phillips curve, is in my view at best a modeling convenience. So this aspect of the paper's conclusions might need several grains of salt.

REFERENCE FOR THE SIMS COMMENT

Benhabib, Jess, Stephanie Schmitt-Grohé, and Martín Uribe. 2001. "The Perils of Taylor Rules." *Journal of Economic Theory* 96, no. 1–2: 40–69.

GENERAL DISCUSSION Rick Mishkin began by observing that the authors' conclusions about the slope of the Phillips curve are not consistent with evidence from analyses of Phillips curves using state- and city-level data in, for example, papers by McLeay and Tenreyro and by Hooper, Mishkin, and Sufi.[1] Importantly, disaggregated data allow these researchers to control for endogenous monetary policy because states and cities within a monetary union experience the same monetary policy setting. These analyses find that the Phillips curve slope did not change significantly after 1990, in contrast to the current paper's conclusions. Mishkin asked whether the authors predict a difference between the disaggregated and aggregate data that would help reconcile this micro evidence with the authors' finding that the Phillips curve slope declined after 1990.

Giovanni Ricco observed that the estimation of trend inflation and unemployment is a critical step in fitting a Phillips curve, which is primarily concerned with the relationship between the cyclical component of inflation and the cyclical component of slack. He commented that the vector autoregressions (VARs) used in the authors' main exercise would implicitly fit a deterministic trend to the data, potentially leading to misestimation of the cyclical components of inflation and unemployment. He asked whether the authors had examined whether this feature affects their results.

Robert Gordon made two comments. First, he reemphasized Sims's reference to supply shocks as crucial to understanding inflation before 1990. He noted, for example, that most of the variation in inflation between 1973 and 1985 can be explained by unfavorable supply shocks that occurred during the 1975 and 1980 oil crisis episodes. Similarly, the very sharp decline of inflation between 1981 and 1985 was due not just to high unemployment but also to the dollar's appreciation during that period. Gordon said controlling for these supply disturbances is important for understanding the shape of the Phillips curve and observed that they are not explicitly considered in the paper.

Second, Gordon observed that two broad hypotheses have been offered to explain why inflation has been subdued in the last ten years despite a dramatic decline in unemployment. One is that the Phillips curve has flattened, as the authors argue. The other is that the rate of unemployment consistent with stable inflation (NAIRU) has declined. The NAIRU hypothesis has been discussed extensively in the context of the low inflation of the 1990s, but Gordon observed that some evidence suggests the NAIRU has declined even further in the last decade. Globalization, the decline in worker bargaining power, the decline in computer prices, and the increased

1. Michael McLeay and Silvana Tenreyro, "Optimal Inflation and the Identification of the Phillips Curve," *NBER Macroeconomics Annual* 34 (2019): 199–255; Peter Hooper, Frederic S. Mishkin, and Amir Sufi, "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or Is It Just Hibernating?," *Research in Economics* 74, no. 1 (2020): 26–62.

importance of computer and IT-related goods are all examples of factors that could have reduced the NAIRU and subdued inflation.

However, he observed that a decline in the NAIRU cannot account for why inflation declined very modestly after the sharp increase in unemployment after the 2008 financial crisis or during the period of weak activity around 2015. Gordon noted that in a contest between the NAIRU and the flat Phillips curve hypotheses, he would likely fall in favor of the flat Phillips curve explanation, consistent with the authors' conclusions.

Giorgio Primiceri responded to comments from Ricco by noting that the authors' VAR approach allows them to measure the inflationary response specifically to an identified cyclical shock. More specifically, the authors verify that estimated trend variables like the NAIRU do not respond to the business cycle shock used in the VAR exercise, suggesting that the identifying disturbance is in fact a cyclical one.

Primiceri acknowledged comments from Sims that the evidence that the business cycle's correlation with wage inflation has declined less than its correlation with price inflation is somewhat tenuous and noted that this finding depends partly on the measure of wage inflation used. For example, the authors find that the correlation between unemployment and wage inflation measured by the employment cost index remains relatively strong before and after 1990; in contrast, the correlation between unemployment and average hourly earnings declined about in line with the unemployment-price inflation correlation after 1990.

Primiceri acknowledged that the business cycle shock does not capture all the relevant historical business cycle variation and that it may not capture important movements in lower-frequency variables that affect inflation. However, the authors' focus in this paper was to examine why inflation does not move over the business cycle as much as it did in the past; the shock was chosen to help address this question.

In addition, while Primiceri acknowledged the excess bond premium shock used in the VAR exercise is a mix of a demand and supply shock, as Sims and Blanchard both suggested, he also noted that were one to replace it with a more traditional demand shock, like an unemployment shock, the main results would hold. The authors' interpretation of the response of inflation to the shock remains valid if the shock contains a prevalent demand component, which the literature surrounding the excess bond premium suggests it does.

Andrea Tambalotti responded to comments from Blanchard, acknowledging that the dynamics of markups are very important for explaining inflation's behavior and perhaps more so in the most recent period. He also noted that targeting wage inflation, which appears to have a stronger correlation with the business cycle, would be an interesting policy in the environment of flatter aggregate supply. However, he referenced work by himself, Justiniano, and Primiceri that examines the quantitative losses associated with stabilizing wage inflation and stabilizing price inflation and finds that the two strategies perform similarly.[2] That result might hint that the wedges between wages and prices are not terribly important for explaining the disconnect between inflation and the business cycle.

Marco Del Negro first addressed comments by both discussants that the business cycle shock is somewhat simplistic. He stressed that the authors test the response of inflation to multiple shocks throughout the paper. In particular, the exercise using the estimated dynamic stochastic general equilibrium (DSGE) model features a large suite of shocks to many variables in the model, and the results from that exercise are very much consistent with the results from the VAR exercise with one business cycle shock.

Del Negro also replied to comments from Mishkin, referencing work by Hazell, Herreno, Nakamura, and Steinsson, who point out that regressions of state- or city-level inflation data on unemployment may be confounded by the fact that different locations feature different unemployment and adjustment dynamics.[3] Controlling for these differences across locations could be possible with data on, for example, location-specific inflation expectations or structural unemployment, but these data are not readily available.

He also responded to Blanchard's comment that markups have become much more countercyclical in the last twenty years than they were prior. Consistent with this observation, Del Negro affirmed that the authors' DSGE model estimated on post-1990 data suggests that real marginal costs (the inverse of aggregate markups) move quite procyclically. In addition, consistent with Blanchard's comment about targeting wage inflation, he noted that systematic policy in the DSGE model is very effective at moving marginal costs precisely because of their stronger relation with the business cycle.

2. Alejandro Justiniano, Giorgio E. Primiceri, and Andrea Tambalotti, "Is There a Trade-Off between Inflation and Output Stabilization?," *American Economic Journal: Macroeconomics* 5, no. 2 (2013) 1–31.

3. Jonathon Hazell, Juan Herreno, Emi Nakamura, and Jón Steinsson, "The Slope of the Phillips Curve: Evidence from US States," Working Paper (2020).