ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Del Negro, Marco; Lenza, Michele; Primiceri, Giorgio E.; Tambalotti, Andrea

Working Paper

What's up with the Phillips curve?

ECB Working Paper, No. 2435

Provided in Cooperation with:
European Central Bank (ECB)

Suggested Citation: Del Negro, Marco; Lenza, Michele; Primiceri, Giorgio E.; Tambalotti, Andrea (2020) : What's up with the Phillips curve?, ECB Working Paper, No. 2435, ISBN 978-92-899-4078-8, European Central Bank (ECB), Frankfurt a. M., https://doi.org/10.2866/5407

This Version is available at:
https://hdl.handle.net/10419/229049

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Marco Del Negro, Michele Lenza,
Giorgio E. Primiceri, Andrea Tambalotti

What's up with the Phillips Curve?

No 2435 / July 2020

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

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 VARs and an estimated DSGE 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 towards 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.

Key words and phrases: Inflation, unemployment, monetary policy trade-off, VARs, DSGE models
JEL Codes: E31, E32, E37, E52

Non-Technical Summary

Why has U.S. inflation remained stable since the 1990s, even though real activity and unemployment have cycled up and down? This question has attracted considerable attention shortly after the Great Recession, when unemployment was close to 10 percent but inflation barely dipped below 1 percent, leading many economists to debate about a "missing deflation" puzzle (e.g. Hall, 2011, Ball and Mazumder, 2011, Coibion and Gorodnichenko, 2015, Del Negro et al., 2015b, Lindevand Trabandt, 2019). More recently, with unemployment below 5 percent for almost four years and inflation persistently under 2 percent, attention has instead turned to factors that may explain why inflation is not coming back (e.g. Powell, 2019, Yellen, 2019).

There is no shortage of potential explanations for this apparent disconnect between inflation and real activity over the business cycle. These explanations can be grouped in four main classes: (i) mis-measurement of either inflation or economic slack; (ii) a flatter wage Phillips curve; (iii) a flatter price Phillips curve; and (iv) a flatter aggregate demand relationship, induced by an improvement in the ability of policy to stabilize inflation.

This paper sheds light on the relative merit of these four classes of explanations, using a combination of state-of-the-art macro-econometric techniques (VARs, SVARs and estimated DSGE models). The analysis shows that changes in the structure of the labor market are unlikely to explain the inflation–real activity disconnect, since wages and, especially, unit-labor costs have displayed more stable responses than prices to rising and falling unemployment. Measurement problems can also be ruled out, since various proxies for cost pressures and

economic slack—such as the output of goods and services, the unemployment rate, hours worked, and businesses' cost per unit produced—have moved together more or less as they always have over the period under analysis.

To investigate which of the two remaining hypotheses is more accurate—a flat Phillips curve or a flat aggregate demand—we study whether the macroeconomic effects of demand shocks have changed before and after 1990. We find that the response of inflation to these shocks has drastically mitigated in the most recent sample, while real economic activity has remained reactive. This result highlights a reduced sensitivity of inflation to demand-induced cost pressures, suggesting that the slope of the Phillips curve must have fallen in the post-1990 period. The evidence that policy might have contributed to inflation stabilization by flattening the aggregate demand curve is weaker.

The paper also draws some policy implications of these findings, showing that, with a flatter Phillips curve, a central bank should recalibrate its systematic monetary policy to retain its ability to stabilize inflation.

1. INTRODUCTION

The recent history of inflation and unemployment is a puzzle. The unemployment rate has gone from below 5 percent in 2006-07 to 10 percent at the end of 2009, and back down below 4 percent over the past couple of years. These fluctuations are as *wide* as any experienced by the U.S. economy in the post-war period. In contrast, inflation has been as *stable* as 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" (e.g. Hall, 2011, Ball and Mazumder, 2011, Coibion and Gorodnichenko, 2015, Del Negro et al., 2015b, Linde 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 (e.g. 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 (e.g. Atkeson and Ohanian, 2001, Stock and Watson, 2008, Stock and Watson, 2007, Zhang et al., 2018 and Stock and Watson, 2019). The literature, which we review in more detail below, has focused on four main classes of explanations for this puzzle: (i) mis-measurement of either inflation or economic slack; (ii) a flatter wage Phillips curve; (iii) a flatter price Phillips curve; and (iv) 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 towards promising explanations for the phenomenon of interest. First, goods inflation has become much less sensitive to the business cycle after 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 et al. (2013), Coibion and Gorodnichenko (2015), Gali and Gambetti (2018), Hooper et al. (2019) 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 mis-measurement of economic slack, as well as a significant flattening of the wage Phillips curve, as the main causes 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 bi-variate relationship

[1]Some recent papers have also explored the behavior of inflation across geographic areas in the United States and across countries (e.g. Fitzgerald and Nicolini, 2014, Mavroeidis et al. 2014, McLeay and Tenreyro 2019, Hooper et al. 2019, and Geerolf, 2019). 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 et al. (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 U.S. states, they recover a flat Phillips curve once the estimates are properly re-scaled, 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.

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]

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 Zakrajsek (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

[3]In practice, we focus primarily on the relationship between inflation and unemployment, but 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.

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. 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. In 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 into 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 believes on 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 2 describes the VAR that we use for data description, whose results are then described in section 3. Section 4 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 5. Section 6 revisits the same facts presented in section 3 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 7 elaborates on some policy implications of our main findings, before offering some concluding remarks in section 8.

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 mis-measurement 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 (e.g. 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 (e.g. 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 mis-measurement, 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 et al. (2020) in this volume is 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 et al. (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 (e.g. Stock, 2011, Gordon, 2013, Hong et al., 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 (e.g. 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 et al. (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 (e.g. Fitzgerald and Nicolini, 2014, Barnichon and Mesters (2019a), Hooper et al. 2019, and McLeay and Tenreyro, 2019, echoing older work by Kareken and Solow, 1963 and 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 re-emerge 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 (e.g. Orphanides and Williams, 2005, Bernanke, 2007, Stock, 2011, Blanchard et al., 2015, Blanchard, 2016, Ball and Mazumder, 2019, Carvalho et al., 2019, Jorgensen and Lansing, 2019, and 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 (e.g. Sbordone, 2007, Auer and Fischer, 2010, Rich et al., 2013, Tallman and Zaman, 2017, Forbes, 2019a, Forbes, 2019b, Obstfeld, 2019 and Forbes et al., 2020). In a similar vein, Rubbo (2020) points to changes in the network structure of the U.S. 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

[4]Afrouzi and Yang (2019) 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.

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.

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

2.1. **Methodology.** To study macroeconomic dynamics and their change over time, we begin by adopting the following Vector-Autoregression (VAR) model:

$$y_t = c + B_1 y_{t-1} + ... + B_p y_{t-p} + u_t. \tag{2.1}$$

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 combinations 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 into more than one hypothesis on the sources of the increased disconnect between inflation and real activity. Therefore, section 5 and 6 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 has also 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 4. 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.

2.2. **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 and/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 marginal costs and inflation. Therefore, a tight labor market, 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 8 macroeconomic variables: (i) *unemployment*, measured by the civilian unemployment rate; (ii) *natural unemployment*, measured by the CBO estimate; (iii) *core inflation*, measured by the annualized quarterly growth rate of the Personal Consumption Expenditures price index, excluding food and energy; (iv) *inflation*, measured by the annualized quarterly growth rate of the GDP deflator; (v) *GDP*, measured by the logarithm of per-capita real GDP; (vi) *hours*, measured by the logarithm of per-capita hours worked in the total economy; (vii) *wage inflation*, measured by the annualized quarterly growth rate of the average hourly earnings of production and nonsupervisory employees (PNSE); and (viii) 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

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

PNSE wage inflation series. This series only covers 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 8-variable VAR over two non-overlapping samples, to investigate possible changes in the typical co-movement pattern of these variables in response to the BC shock described in section 2.1. The first sample ranges from 1964:II to 1989:IV, and the second from 1989:I to 2019:III.[6] The analysis starts in 1964:II, when the PNSE wage inflation 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 1990-91 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 shows that this choice is immaterial for the results.

The data are quarterly and the VAR includes 4 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 et al. (2015).

3. 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: (i) the attenuation of inflation fluctuations over the business cycle before and after 1990 is striking; (ii) in comparison, the co-movement of all real variables and indicators of firms' cost pressures has been remarkably stable.

3.1. 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 3.1

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

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 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.

Appendix A shows that the responses of figure 3.1 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 et al. (2019) and Angeletos et al. (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 et al. (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 3.1 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 Banbura et al. (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 3.2 plots the dynamics of all the VAR variables in this conditional forecast exercise. As in figure 3.1,

[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.

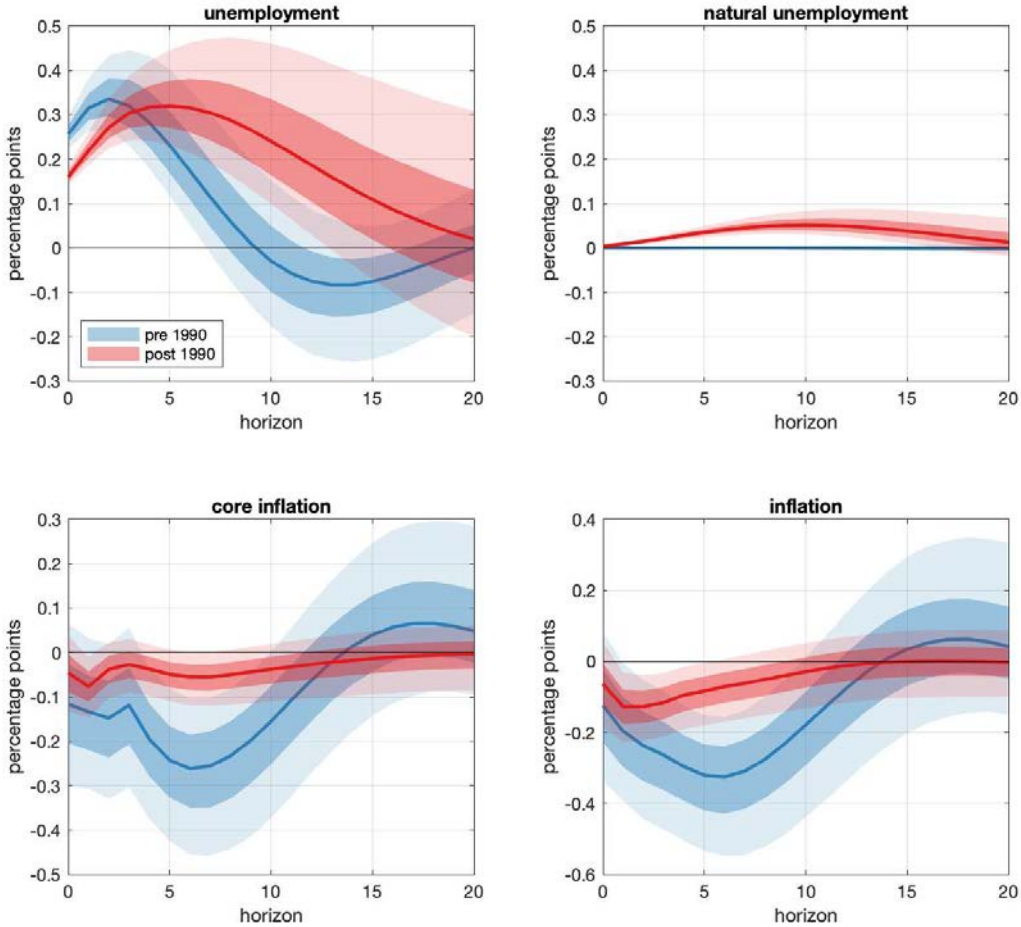

FIGURE 3.1. Impulse responses of unemployment and price inflation to a shock to the unemployment equation. The impulse responses are from the baseline VAR described in section 2.2. The shock is identified using a Cholesky strategy, with unemployment ordered first. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

the response of inflation in the second sample is much attenuated, although it now remains negative. 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.

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

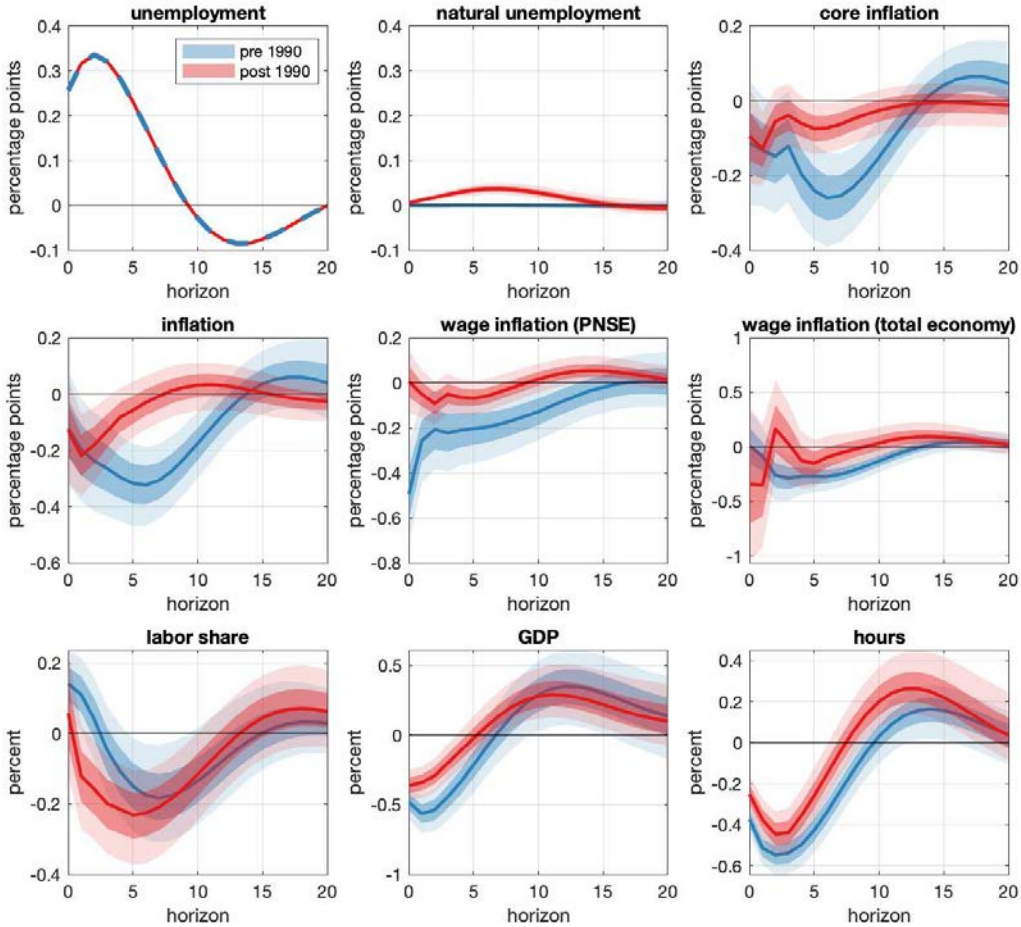

FIGURE 3.2. Response of all variables, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

regarding the stability of inflation. Interpreting this fact is the main task of the rest of the paper.

3.2. **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 (Gali and Gambetti, 2018, Hooper et al., 2019, Rognlie, 2019). This section presents VAR results broadly consistent with these findings. The second row of figure 3.2 plots the response of two measures of nominal wage inflation using the conditional

forecast approach described in the previous subsection. The first measure (PNSE) is the one used directly for the estimation of the baseline VAR. The second (total economy) 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 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 relatively high, putting upward pressure on prices and inflation. Such a story, however, would imply a substantial weakening of the comovement between unemployment and wage inflation, which seems at odds with the data. In addition, as we demonstrate in appendix A, this co-movement is approximately equally strong after 1990 regardless of whether we include or exclude the Great Recession period.

3.3. **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 is,

[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 $w = ls - h + gdp + p$.

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 3.2 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 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.

3.4. **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

[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.

[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.

relatively stable over time. The third row of figure 3.2 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.* 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 the best proxy for underlying inflationary pressures.

3.5. **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 forecasters (SPF).[12] The former was not included in the baseline VAR because it was at the zero lower bound for many years in the second sample. To avoid that period, this larger VAR is estimated excluding data after 2007. Figure 3.3 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 12 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 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

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

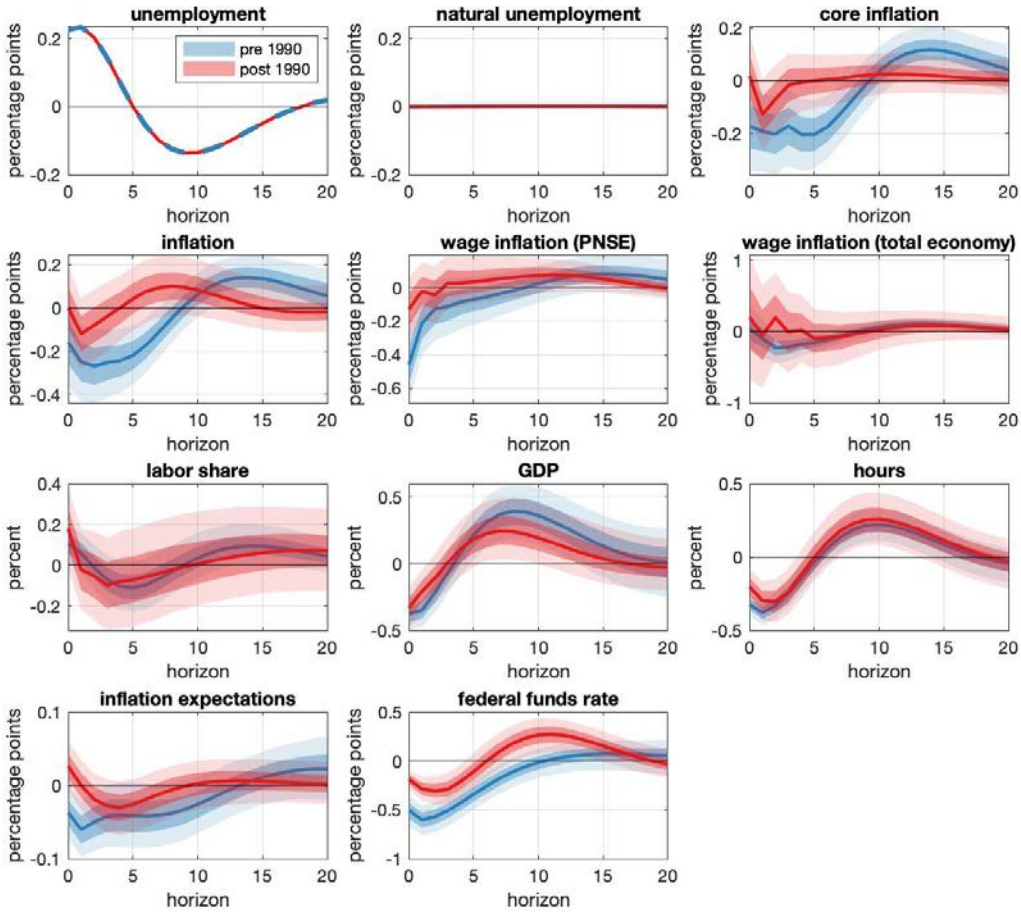

FIGURE 3.3. Responses of all variables, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2, augmented with long-term inflation expectations and the federal funds rate. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

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 5.

3.6. **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 bi-variate relationship between inflation and real activity, with no need to be more specific on its measurement. As illustrated in section 4, 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 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 5 and 6.

4. Lessons from a Stylized Model

To aid in the interpretation of the empirical facts described in section 3, 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: (i) the empirical facts of section 3 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; (ii) 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; (iii) 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, in the next sections we will introduce more information and structure to further sharpen our inference.

4.1. **A simple model of aggregate demand and supply.** The stylized model we consider consists of the following three familiar equations:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa (x_t + s_t) \tag{4.1}$$

$$x_t = E_t x_{t+1} - \sigma (i_t - E_t \pi_{t+1} - \delta_t) \tag{4.2}$$

$$i_t = E_t \pi_{t+1} + \psi_\delta \delta_t + \psi_\pi \pi_t, \tag{4.3}$$

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 (4.1) 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 3 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 (4.1) 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 (4.2) is an Euler equation, or dynamic IS equation, which connects the nominal interest

[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, in 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.

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 (4.3) 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 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

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

where $\phi \equiv (\sigma \psi_\pi)^{-1} \geq 0$ and $d_t \equiv \sigma \left(\psi_\delta - 1 \right) \delta_t$ is a simple re-scaling 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 zero lower bound. 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, 2020. 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).

4.2. **Two alternative sources of inflation stabilization.** Given the structure of the model, it is immediate to see 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.[15,16]

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,

[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.

[16]Another obvious possibility is that the volatilities of both shocks have fallen dramatically. Although the large 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.

without going to the extreme of promising a very large increase in policy rates in reaction to even small changes in inflation.[17]

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 3.

4.3. **Model solution.** This section presents the solution of the simple model described above under the assumption of i.i.d. shocks. With this simplifying assumption, expectations are zero and the model reduces to a static demand and supply framework with stochastic shocks,

$$\pi_t = \kappa \left(x_t + s_t \right)$$
$$\pi_t = -\phi \left(x_t - d_t \right),$$

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 is 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) and Barnichon and Mesters (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).

[17] 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 et al., 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.

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.

5. Interpreting the Facts with a Structural VAR

The U shock employed in section 3 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 (e.g. Smets and Wouters, 2007 or Justiniano et al., 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 4, 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 Zakrajsek (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 2 and study the impulse responses to innovations to the EBP that are orthogonal to the other variables in the system.[18] 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.[19]

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 (e.g. Caldara et al., 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 5.1 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

[18]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 et al. (2009) and Gilchrist and Zakrajsek (2012). Their dynamic systems, a 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.

[19]This is how marginal efficiency of investment shocks in Justiniano et al., 2011, risk shocks in Christiano et al., 2014 and spread shocks in Cai et al., 2019 propagate. All these shocks are identified mostly through their effect on credit spreads.

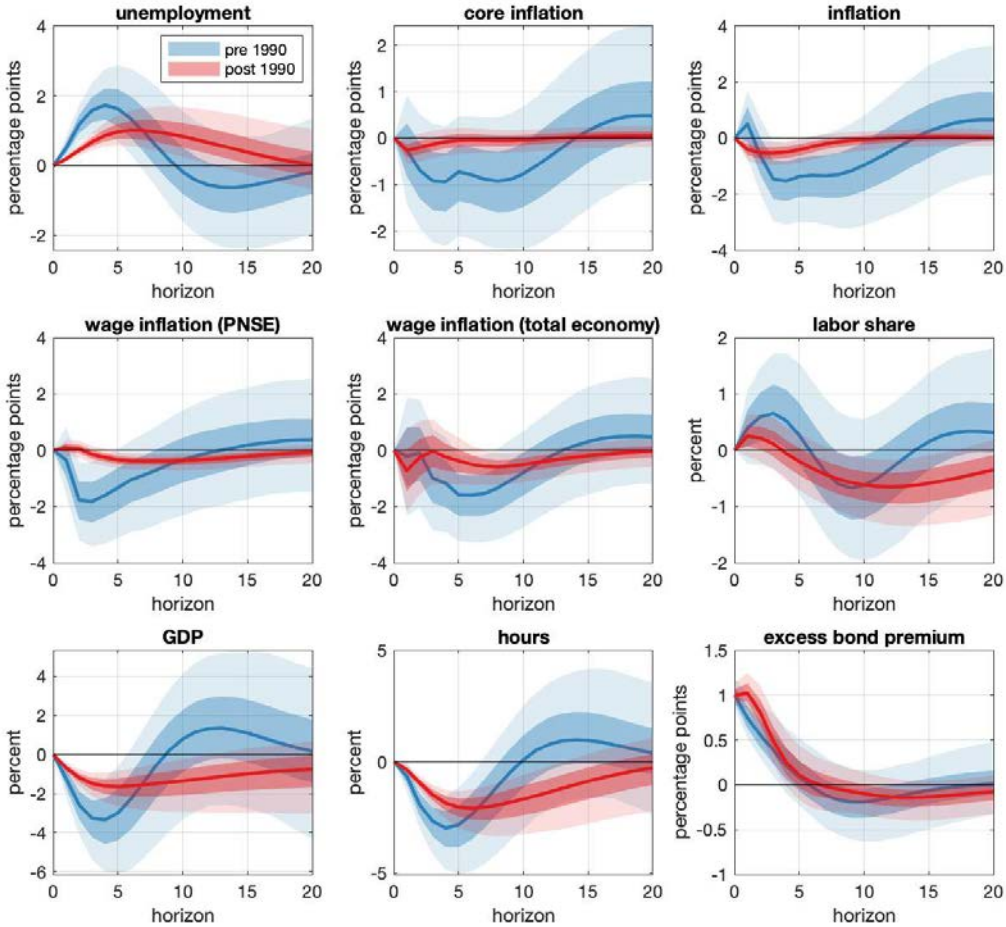

FIGURE 5.1. Impulse responses to an EBP shock, identified by assuming that it affects the excess bond premium contemporaneously, but all other variables with a lag. The impulse responses are from the baseline VAR of section 2.2, augmented with data on the excess bond premium. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1973:I to 1989:IV, and from 1989:I to 2019:III, respectively.

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 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 3. 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 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 zero lower bound on nominal interest rates.[20] 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 3. 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, 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.[21]

5.1. **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

[20] The extent to which the ZLB was binding during the Great Recession and its aftermath is much debated in the literature (e.g. Swanson and Williams, 2014, Gali and Gambetti, 2019, Eggertsson and Egiev, 2020)
[21] Christiano et al. (2015) and Gilchrist et al. (2017) develop models of this kind.

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 (e.g. Bernanke, 2007, Blanchard, 2016, Gordon, 2018 and Ball and Mazumder, 2019 among many others).

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 6.

Consider the following Phillips curve

$$\pi_t = \alpha \pi^e_{t+1|t} + (1-\alpha)\,\pi_{t-1} + \kappa\,(x_t + s_t)\,, \tag{5.1}$$

which generalizes (4.1) by allowing for inflation inertia and potentially non-rational 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

$$R^{\pi}_h = \alpha R^{\pi^e}_h + (1-\alpha)\,R^{\pi}_{h-1} + \kappa R^x_h, \tag{5.2}$$

where $R^z_h \equiv \frac{\partial E_t z_{t+h}}{\partial d_t}$ is the impulse response at horizon h of a generic variable z to the demand shock. In deriving (5.2), we used the fact that s_t is independent from the demand disturbance. This equation says that, if (5.1) holds, the impulse responses of inflation,

expected inflation and the gap to a demand shock should satisfy (5.2) 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 (5.2).

In practice, the Phillips curve (5.1) 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 (5.2). 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 5.1. 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 5.2 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 towards 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 (5.1), 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 5.2.

Another difference from Barnichon and Mesters (2019b) is that their analysis is based on 1-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 5.3 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.

6. Interpreting the Facts with an Estimated DSGE Model

As argued above, interpreting the facts of section 3 only through the lens of the VAR is challenging, as one can only speculate on the mechanisms behind them. For instance, the

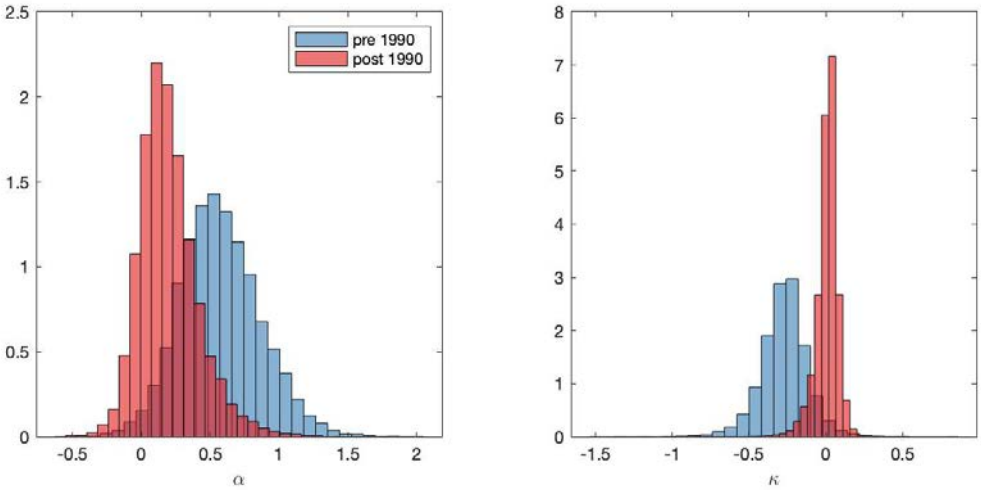

FIGURE 5.2. Histogram of inferred values of α and κ across posterior draws. These values are obtained as described in section 5.1, using the impulse responses of inflation, expected inflation and unemployment to an EBP shock, identified by assuming that it affects the excess bond premium contemporaneously, but all other variables with a lag. The impulse responses are from the baseline VAR of section 2.2, augmented with data on 1-quarter-ahead inflation expectations and the excess bond premium. The pre- and post-1990 samples consist of data from 1973:I to 1989:IV, and from 1989:I to 2019:III, respectively.

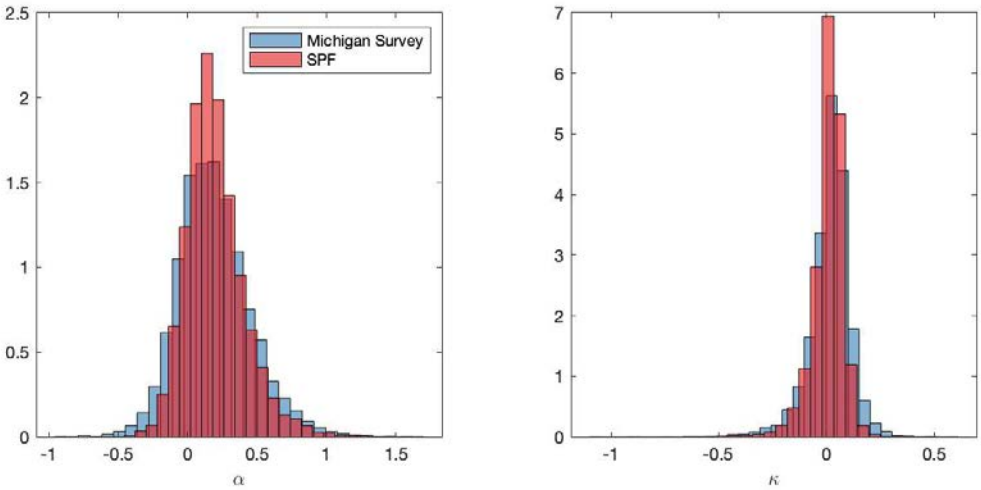

FIGURE 5.3. Histogram of inferred values of α and κ across posterior draws in the post-1990 sample (1989:I–2019:III). These values are obtained as described in section 5.1, using the impulse responses of inflation, expected inflation and unemployment to an EBP shock, identified by assuming that it affects the excess bond premium contemporaneously, but all other variables with a lag. The impulse responses are from the baseline VAR of section 2.2, augmented with data on inflation expectations and the excess bond premium. The "Michigan survey" and "SPF" histograms are obtained using data on 1-year-ahead inflation expectations from the Michigan Survey and 1-quarter ahead inflation expectations from the Survey of Professional Forecasters, respectively.

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 NY Fed 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 impacts the labor market in ways comparable to those described in section 3 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 5, 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.

6.1. **The DSGE Model.** The NY Fed DSGE model is a medium-scale New-Keynesian model. In broad strokes, the model can be described as Smets and Wouters (2007) plus financial frictions as in Bernanke et al. (1999) and Christiano et al. (2014). Variations of this model have been used in Del Negro and Schorfheide (2013), Del Negro et al. (2015b) and Del Negro et al. (2017).[22] Cai et al. (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 GDP and GDI), 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 et al. (2015a) and Laseen and Svensson (2011), to account for the zero lower bound on nominal interest rates and

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

forward guidance. Therefore, we augment the set of observables to incorporate interest rate expectations during the ZLB period. 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

$$\pi_t = \kappa_p mc_t + \frac{\iota_p}{1+\iota_p\bar{\beta}}\pi_{t-1} + \frac{\bar{\beta}}{1+\iota_p\bar{\beta}}E_t[\pi_{t+1}] + \lambda_{f,t}, \tag{6.1}$$

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

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

Its key determinant is the Calvo parameter ζ_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

$$\begin{aligned} R_t &= \rho_R R_{t-1} + (1-\rho_R)\left(\psi_\pi(\pi_t-\pi_t^*) + \psi_y(y_t-y_t^*)\right) \\ &\quad +\psi_{\triangle y}\left((y_t-y_t^*)-(y_{t-1}-y_{t-1}^*)\right) + r_t^m. \end{aligned} \tag{6.2}$$

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.[23]

6.2. **Can the DSGE replicate the VAR facts?** To assess whether the estimated DSGE model can reproduce the VAR facts of section 3, we compute the impulse responses to the

[23]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.

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.[24] The NY Fed DSGE model does not feature all the variables included in the VARs of section 3, most notably unemployment. Therefore, the results in this section are based on a DSGEVAR with slightly fewer observables. They are: hours worked per capita in log levels, a measure of price inflation, the labor share, and wage inflation.[25] This information is enough to capture the dynamics of the key economic variables involved in the connection between inflation and real activity.

Figure 6.1 shows the impulse responses of these variables to a U shock in two DSGEVARs. The blue responses are based on the DSGEVAR implied by the pre-1990 DSGE estimates; the red responses reflect the estimation on the post-1990 sample. The U shock is identified using the methodology described in section 2.1, applied to hours rather than to unemployment. As in section 3, 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.[26] The response of the labor share is also similar in the two samples and, if anything, stronger in the post-1990 estimation. Qualitatively and quantitatively, these responses are also very similar to those in section 3. 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 3.

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 into the alternative explanations of the muted

[24]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.

[25]The VARs used in section 3 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.

[26]The DSGE model, and therefore the DSGEVAR, use a slightly different measure of hours worked than the VAR of section 6.1. This is because the variable definitions in the NY Fed DSGE model are the same as in Smets and Wouters (2007).

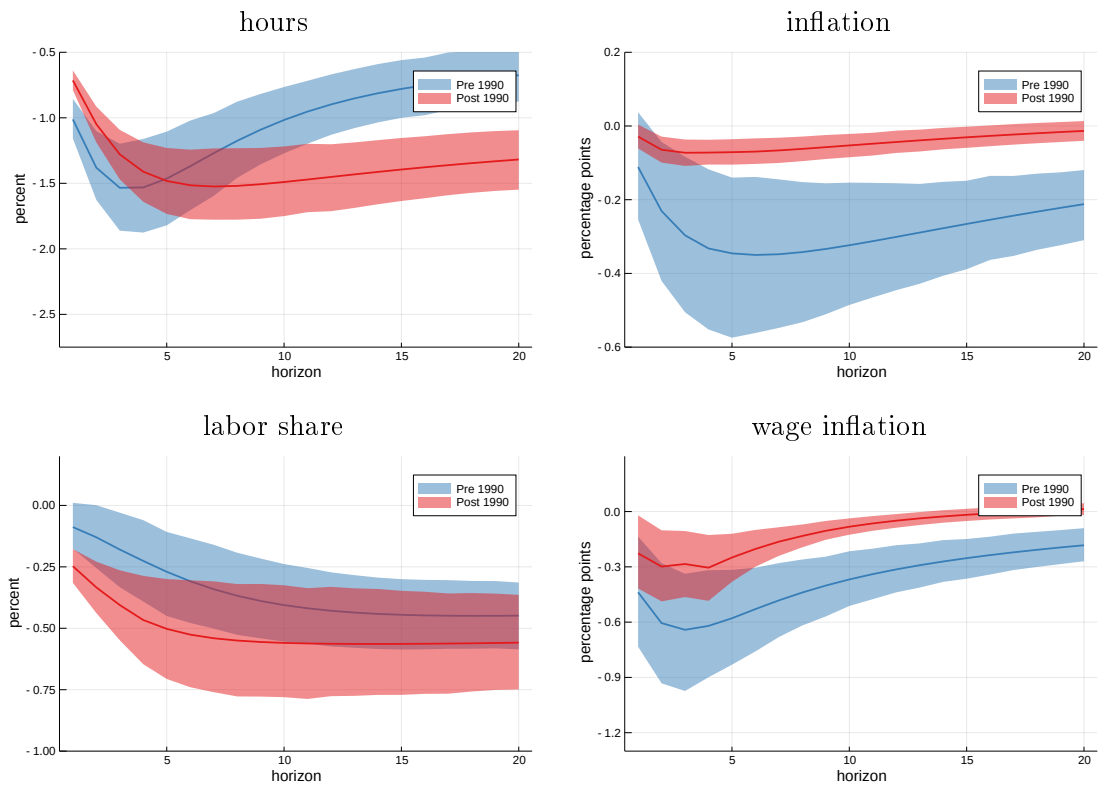

FIGURE 6.1. Impulse responses of hours worked, price inflation, the labor share, and wage inflation to a shock to the unemployment equation. The impulse responses are from the VAR approximation of the NY Fed DSGE model described in section 6 and appendix B.4. The shock is identified using a Cholesky 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:II to 1989:IV, and from 1990:I to 2019:III, respectively.

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 6.2 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.[27] In contrast, the upper right panel of figure 6.2 shows that the posterior distribution

[27]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

of the slope of the wage Phillips curve (κ_w) is quite similar across samples, although it also shifts somewhat to the right.

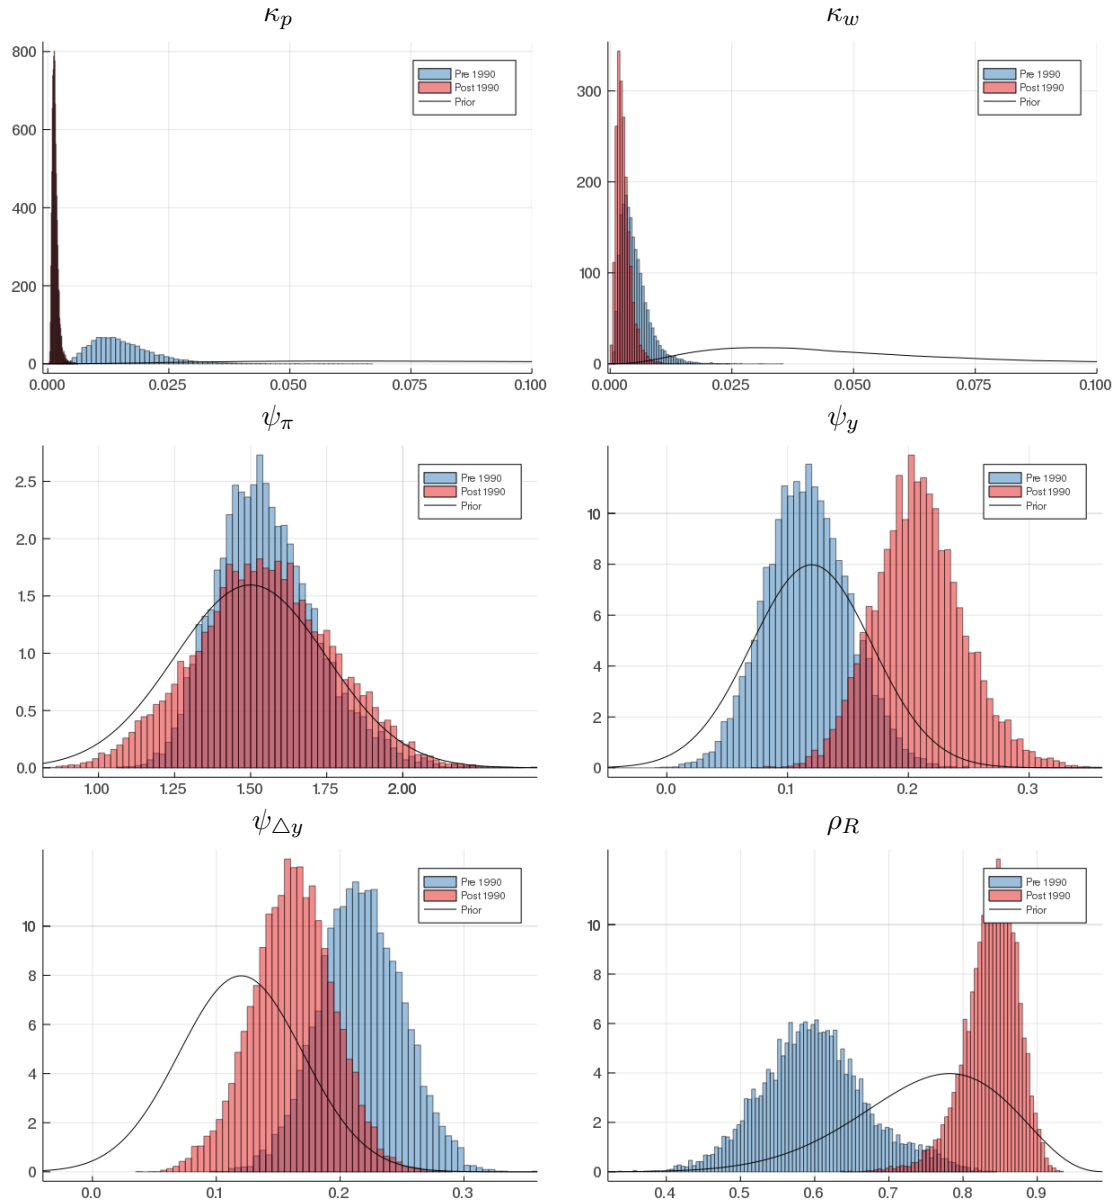

FIGURE 6.2. Prior and posterior distributions for selected parameters of the NY Fed DSGE model described in section 6 and appendix B.4. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1990:I to 2019:III, respectively.

The remaining two rows of figure 6.2 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),

caveats, figure C.1 in the appendix 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.

the output gap level (ψ_y) and its growth ($\psi_{\triangle y}$). 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.[28] 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.

6.3. **Explaining the facts.** Can the changes in the estimates of the slope of the price Phillips curve and/or the policy parameters quantitatively explain the facts of figure 6.1? Figure 6.3 answers this question. The blue and red solid 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.[29] The black lines show the counterfactual impulse responses obtained using the pre-1990 parameters, except for the slope of the Phillips curve (solid), and the policy-rule parameters (dashed). For these parameters, we use the posterior mode of the post-1990 estimates. We refer to the solid and dashed black lines as the "slope counterfactual" and the "policy counterfactual," respectively. 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, than the solid (dashed) black 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 red 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 fairing worse than the slope one.

[28] 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.

[29] These responses are essentially the same as those in figure 6.1. 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.

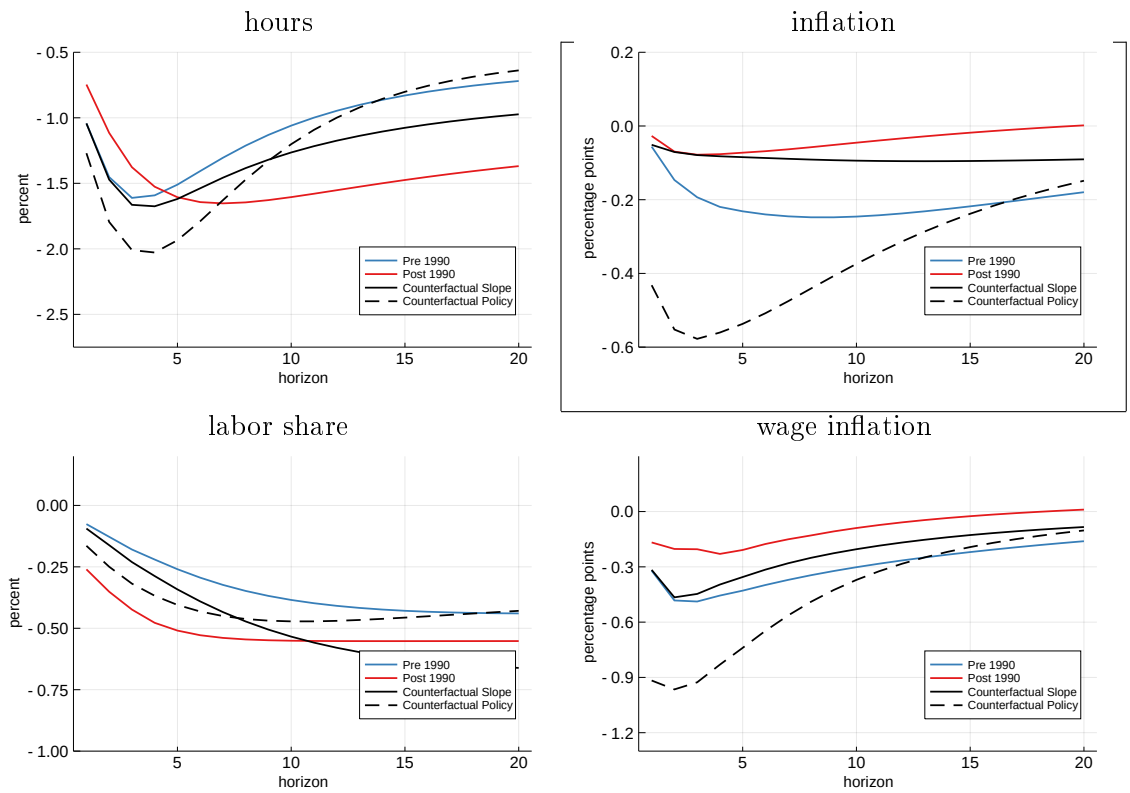

FIGURE 6.3. Impulse responses of hours worked, price inflation, the labor share, and wage inflation to a shock to the unemployment equation. The impulse responses are from the VAR approximation of the NY Fed DSGE model described in section 6 and appendix B.4. The shock is identified using a Cholesky 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 6.3 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 6.4. They are different from those in figure 6.2, 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 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 6.2. Third, the response to the output gap growth ($\psi_{\triangle y}$) 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.

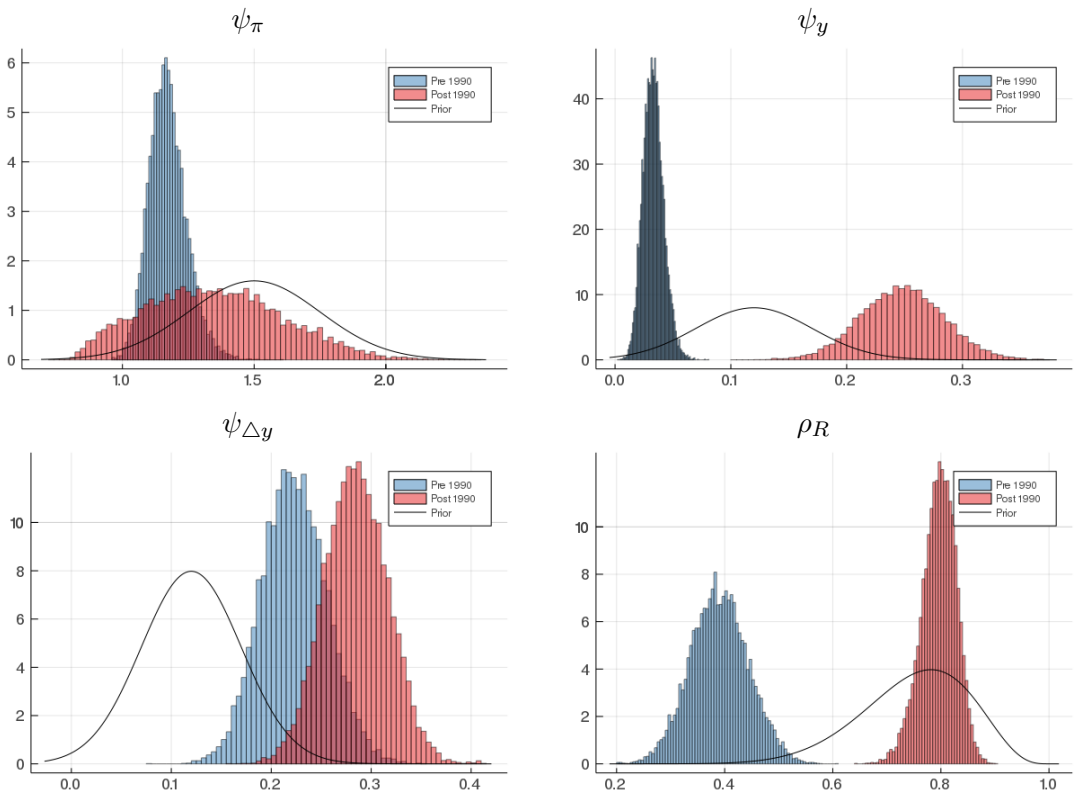

FIGURE 6.4. Prior and posterior distributions for selected parameters of the NY Fed DSGE model described in section 6 and appendix B.4. In this experiment, the model is estimated allowing for two distinct monetary policy regimes, pre- and post-1990, but assuming that all other coefficients are constant for the entire sample from 1964:II to 2019:III.

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 6.5. This is not too surprising, since this version of the model must account for the stability of inflation after 1990 only using changes in the policy parameters. The conclusion is that it is possible to find changes in 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 3 and in figure 6.1, 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 4: 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.[30]

These results stand in contrast to those of figure 6.6, 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 6.6 are in line with those in figure 6.1: the reaction of hours and the labor share is very similar across samples, but that of inflation is much more muted after 1990.[31]

7. Policy Implications

To understand the policy implications of our findings, a useful starting point is the simple model of section 4. 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 6: in response to a monetary policy shock, either a short-lived or a very persistent

[30]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 Appendix reproduce the results in figures 6.4 and 6.5 using a sample ending in 2008:III, 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.

[31]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 appendix. They are are very similar to those in figure 6.2. 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.

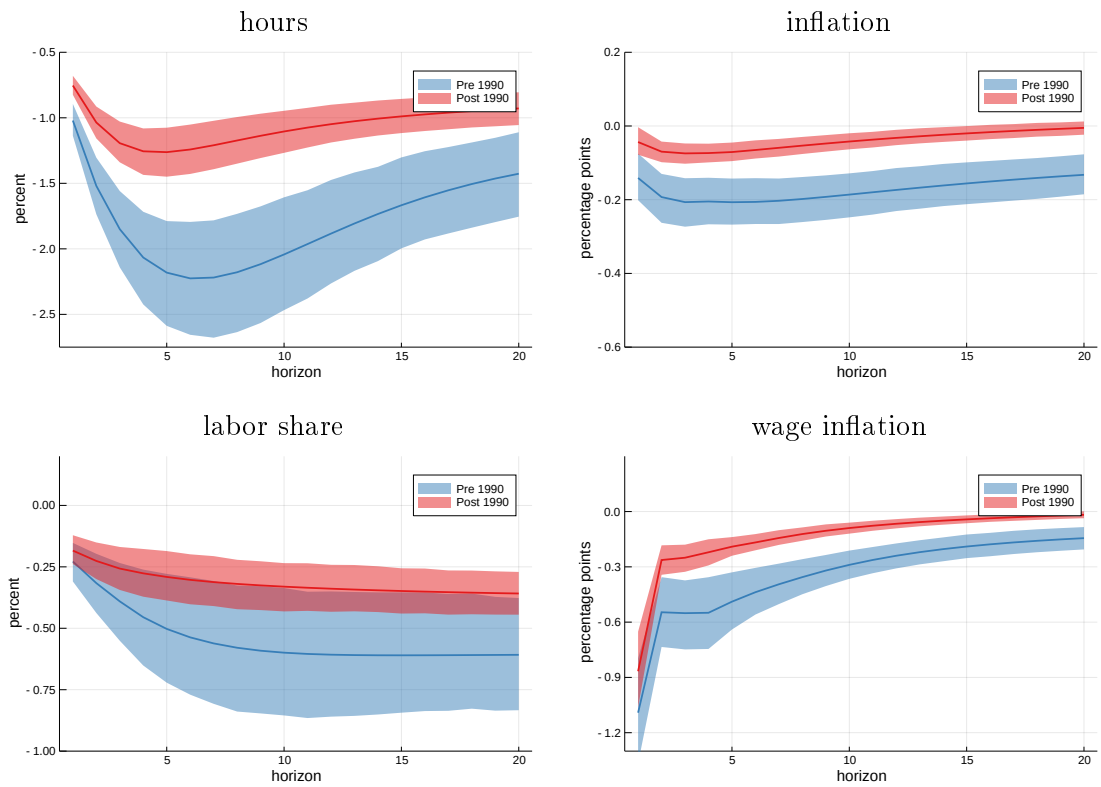

FIGURE 6.5. Impulse responses of hours worked, price inflation, the labor share, and wage inflation to a shock to the unemployment equation. The impulse responses are from the VAR approximation of the NY Fed DSGE model described in section 6 and appendix B.4. The shock is identified using a Cholesky 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 for two distinct monetary policy regimes, pre- and post-1990, but assuming that all other coefficients are constant for the entire sample from 1964:II to 2019:III.

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 re-establish its control of inflation should be to change its *systematic* 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

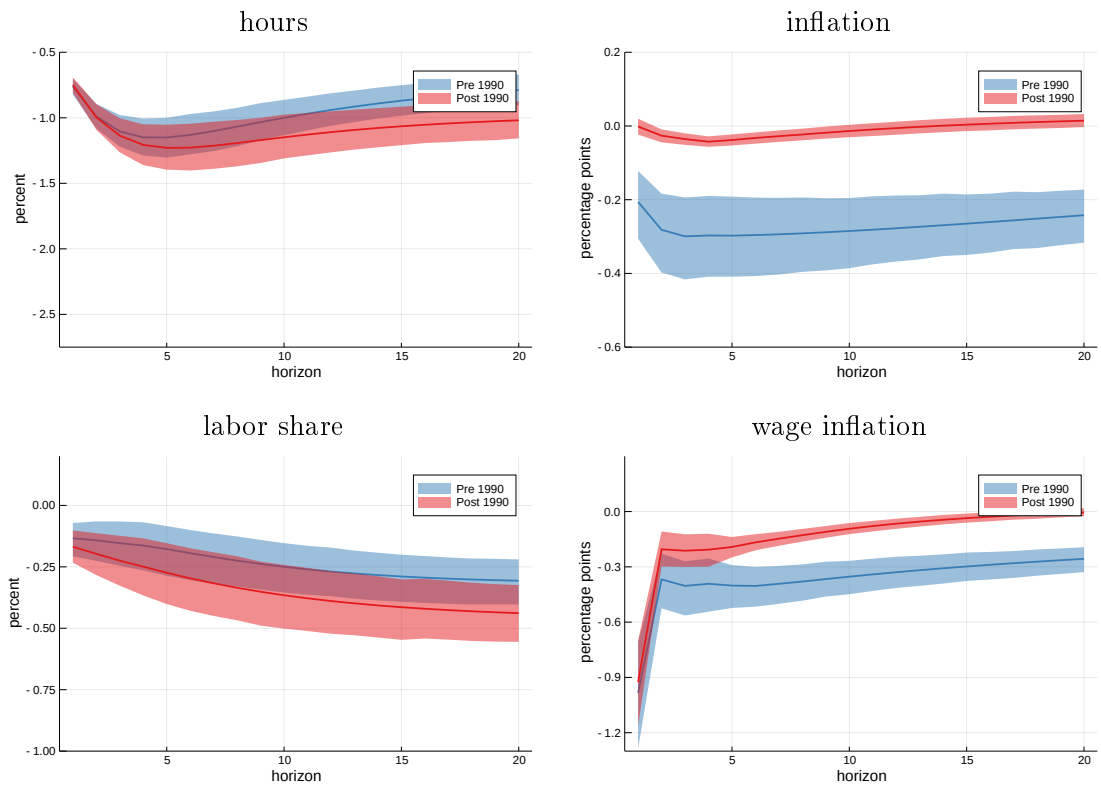

FIGURE 6.6. Impulse responses of hours worked, price inflation, the labor share, and wage inflation to a shock to the unemployment equation. The impulse responses are from the VAR approximation of the NY Fed DSGE model described in section 6 and appendix B.4. The shock is identified using a Cholesky 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 for two distinct regimes for the slope of the price Phillips curve, pre- and post-1990, but assuming that all other coefficients are constant for the entire sample from 1964:II to 2019:III.

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.

The solid lines in figure 7.1 show the model's forecasts of inflation and marginal costs starting in 2020Q1 under the estimated policy rule and two alternative settings for the

slope of the price Phillips curve κ_p: post-1990 (in red) and pre-1990 (in blue). As shown in figure 6.2 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.[32]

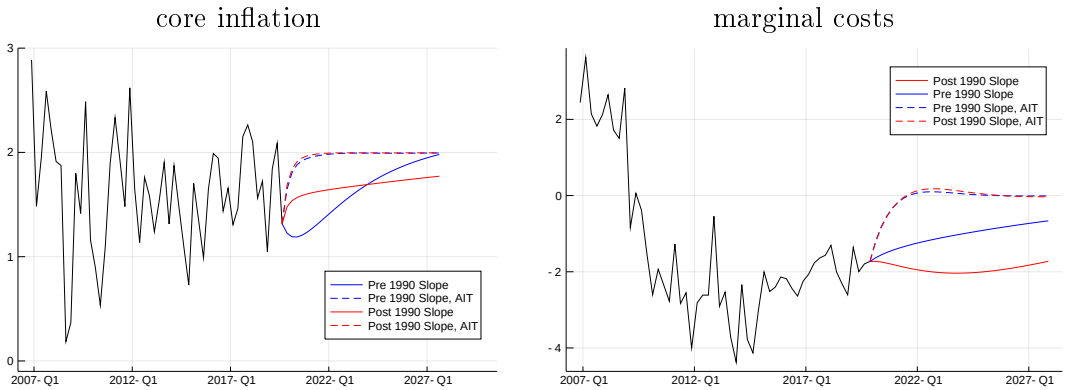

FIGURE 7.1. Forecasts of core PCE inflation and marginal costs obtained from the NY Fed DSGE model, using the modal estimate of κ_p in the pre- and post-1990 period, under the estimated policy rule and a counterfactual policy of average inflation targeting (AIT). The lines going back to 2007 represent 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.

When κ_p is high, inflation falls more initially but then it reverts relatively quickly to the FOMC'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.[33] In the

[32]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 blue forecast is computed under the assumption that κ_p returns to its pre-1990 value starting in 2020Q1.

[33]In his discussion, Olivier Blanchard makes the point that markups, which are the inverse of the real marginal costs shown in figure 7.1, are more countercyclical in recent decades than they had been before.

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 New York Fed DSGE model contributes to explain 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 debate (e.g. Mertens and Williams, 2019). The exact AIT rule used in the counterfactual simulations is

$$R_t = \frac{0.25}{1-\gamma}\pi gap_t, \tag{7.1}$$

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 FOMC's long-run goal.[34]

The dashed lines in figure 7.1 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 red lines is larger and more persistent than that for the blue lines, which represent the high κ_p economy. 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 et al., 2015b). 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

His evidence is consistent with figure 7.1: after 1990 the economy has become more "Keynesian" and real marginal costs fall more, and more persistently, following a recession.

[34]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 10 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.

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.[35] 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.

8. 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/or willing to stabilize

[35] One relevant feature of this class of models is the so-called forward guidance puzzle (e.g. Del Negro et al., 2015a) 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 et al. (2020).

inflation, making the economy more "neo-classical," 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 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 Chair Volcker that brought inflation under control in the early eighties, 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.

References

Abraham, K. G., J. C. Haltiwanger, and L. E. Rendell (2020): "How Tight is the U.S. Labor Market?" Prepared for the Spring 2020 Brookings Papers.

Afrouzi, H. and C. Yang (2019): "Dynamic Rational Inattention and the Phillips Curve," *Available at SSRN 3465793*.

Angeletos, G.-M., F. Collard, and H. Dellas (2019): "Business Cycle Anatomy," MIT, mimeo.

Aruoba, S. B. and F. Schorfheide (2010): "Sticky Prices versus Monetary Frictions: An Estimation of Policy Trade-offs," *American Economic Journal: Macroeconomics*, forthcoming.

Atkeson, A. and L. E. Ohanian (2001): "Are Phillips Curves Useful for Forecasting Inflation?" *Federal Reserve of Minneapolis Quarterly Review*, 25(1), 2–11.

Auer, R. and A. M. Fischer (2010): "The effect of low-wage import competition on U.S. inflationary pressure," *Journal of Monetary Economics*, 57, 491–503.

Ball, L. and S. Mazumder (2011): "Inflation Dynamics and the Great Recession," *Brookings Papers on Economic Activity*, 42, 337–405.

——— (2019): "A Phillips Curve with Anchored Expectations and Short-Term Unemployment," *Journal of Money, Credit and Banking*, 51, 111–137.

Banbura, M., D. Giannone, and M. Lenza (2015): "Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections," *International Journal of Forecasting*, 31, 739–756.

Barnichon, R. and G. Mesters (2019a): "Identifying Modern Macro Equations with Old Shocks," CEPR Discussion Papers 13765, C.E.P.R. Discussion Papers.

——— (2019b): "The Phillips Multiplier," CEPR Discussion Papers 13480, C.E.P.R. Discussion Papers.

Bernanke, B. (2007): "Inflation expectations and inflation forecasting," in *Speech given at the Monetary Economics Workshop of the NBER Summer Institute, Cambridge, Massachussets.*

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999): "The financial accelerator in a quantitative business cycle framework," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 21, 1341–1393.

Blanchard, O. (2016): "The Phillips Curve: Back to the '60s?" *American Economic Review*, 106, 31–34.

Blanchard, O., E. Cerutti, and L. Summers (2015): "Inflation and Activity â Two Explorations and their Monetary Policy Implications," NBER Working Papers 21726, National Bureau of Economic Research, Inc.

Boivin, J. and M. P. Giannoni (2006): "DSGE Models in a Data Rich Enviroment," *NBER Working Paper*, 12772.

Budianto, F., T. Nakata, and S. Schmidt (2020): "Average Inflation Targeting and the Interest Rate Lower Bound," BIS Working Papers no. 852.

Cai, M., M. Del Negro, M. P. Giannoni, A. Gupta, P. Li, and E. Moszkowski (2019): "DSGE Forecasts of the Lost Recovery," *International Journal of Forecasting*.

Caldara, D., C. Fuentes-Albero, S. Gilchrist, and E. Zakrajsek (2016): "The macroeconomic impact of financial and uncertainty shocks," *European Economic Review*, 88, 185–207.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–398.

Carvalho, C., S. Eusepi, E. Moench, and B. Preston (2019): "Anchored Inflation Expectations," .

Cavallo, A. (2018): "More Amazon Effects: Online Competition and Pricing Behaviors," NBER Working Papers 25138, National Bureau of Economic Research, Inc.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2015): "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 7, 110–167.

Christiano, L. J., R. Motto, and M. Rostagno (2003): "The Great Depression and the Friedman-Schwartz Hypothesis," *Journal of Money, Credit and Banking*, 35, 1119–1197.

——— (2014): "Risk Shocks," *American Economic Review*, 104, 27–65.

Clark, T. E. and T. Doh (2014): "Evaluating alternative models of trend inflation," *International Journal of Forecasting*, 30, 426–448.

Coibion, O. and Y. Gorodnichenko (2015): "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation," *American Economic Journal: Macroeconomics*, 7, 197–232.

Coibion, O., Y. Gorodnichenko, and D. Koustas (2013): "Amerisclerosis? The Puzzle of Rising U.S. Unemployment Persistence," *Brookings Papers on Economic Activity*, 44, 193–260.

Crump, R. K., S. Eusepi, M. Giannoni, and A. Sahin (2019): "A Unified Approach to Measuring u*," *Brookings Papers on Economic Activity*, 2019, 143–238.

Daly, M. C. and B. Hobijn (2014): "Downward Nominal Wage Rigidities Bend the Phillips Curve," *Journal of Money, Credit and Banking*, 46, 51–93.

Davis, S. J. (2019): "Comments on "A Unified Approach to Measuring u*"," *Brookings Papers on Economic Activity*, 223–230.

De Graeve, F. (2008): "The External Finance Premium and the Macroeconomy: US Post-WWII Evidence," *Journal of Economic Dynamics and Control*, 32, 3415 – 3440.

Del Negro, M. and S. Eusepi (2011): "Fitting Observed Inflation Expectations," *Journal of Economic Dynamics and Control*, 35, 2105–2131.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017): "Safety, liquidity, and the natural rate of interest," *Brookings Papers on Economic Activity*, Spring.

Del Negro, M., M. P. Giannoni, and C. Patterson (2015a): "The Forward Guidance Puzzle," *FRBNY Staff report*.

Del Negro, M., M. P. Giannoni, and F. Schorfheide (2015b): "Inflation in the Great Recession and New Keynesian Models," *American Economic Journal: Macroeconomics*, 7, 168–196.

Del Negro, M. and F. Schorfheide (2004): "Priors from General Equilibrium Models for VARs," *International Economic Review*, 45, 643 – 673.

——— (2013): "DSGE Model-Based Forecasting," in *Handbook of Economic Forecasting, Volume 2*, ed. by G. Elliott and A. Timmermann, Elsevier.

Eggertsson, G. and S. K. Egiev (2020): "A Unified Theory of the Great Depression and the Great Recession ," Mimeo, Brown University.

Faccini, R. and L. Melosi (2020): "Bad Jobs and Low Inflation," Mimeo, Federal Reserve Bank of Chicago.

Fernald, J. G. (2012): "A quarterly, utilization-adjusted series on total factor productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Fitzgerald, T. J. and J. P. Nicolini (2014): "Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from U.S. Cities," Working Papers 713, Federal Reserve Bank of Minneapolis.

Forbes, K. (2019a): "Inflation Dynamics: Dead, Dormant, or Determined Abroad?" NBER Working Papers 26496, National Bureau of Economic Research, Inc.

Forbes, K. J. (2019b): "How Have Shanghai, Saudi Arabia, and Supply Chains Affected U.S. Inflation Dynamics?" *Review*, 27–44.

Forbes, K. J., J. E. Gagnon, and C. G. Collins (2020): "Low Inflation Bends the Phillips Curve around the World," Peterson Institute for International Economics, Working Paper 20-06.

Fuhrer, J. and G. Moore (1995): "Inflation Persistence," *The Quarterly Journal of Economics*, 110, 127–159.

Fuhrer, J. C. (2010): "Inflation Persistence," in *Handbook of Monetary Economics*, ed. by B. M. Friedman and M. Woodford, Elsevier, vol. 3 of *Handbook of Monetary Economics*, chap. 9, 423–486.

Galí, J. (2015): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications Second edition*, no. 10495 in Economics Books, Princeton University Press.

Gali, J. and L. Gambetti (2018): "Has the U.S. Wage Phillips Curve Flattened? A Semi-Structural Exploration," Mimeo, CREI.

——— (2019): "On the Empricial (Ir)Relevance of the Zero Lower Bound Constraint," in *NBER Macroeconomics Annual*.

Galí, J. and M. Gertler (1999): "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, 44, 195–222.

Geerolf, F. (2019): "The Phillips Curve: a Relation between Real Exchange Rate Growth and Unemployment," 2018 Meeting Papers 1187, UCLA.

Giannone, D., M. Lenza, and G. E. Primiceri (2015): "Prior Selection for Vector Autoregressions," *The Review of Economics and Statistics*, 97, 436–451.

Giannone, D., M. Lenza, and L. 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, 137–173.

Gilchrist, S., R. Schoenle, J. Sim, and E. Zakrajsek (2017): "Inflation Dynamics during the Financial Crisis," *American Economic Review*, 107, 785–823.

Gilchrist, S., V. Yankov, and E. Zakrajsek (2009): "Credit market shocks and economic fluctuations: Evidence from corporate bond and stock markets," *Journal of Monetary Economics*, 56, 471–493.

Gilchrist, S. and E. Zakrajsek (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

Goldfeld, S. M. and A. S. Blinder (1972): "Some Implications of Endogenous Stabilization Policy," *Brookings Papers on Economic Activity*, 3, 585–644.

Goolsbee, A. D. and P. J. Klenow (2018): "Internet Rising, Prices Falling: Measuring Inflation in a World of E-Commerce," *AEA Papers and Proceedings*, 108, 488–92.

Gordon, R. J. (2013): "The Phillips Curve is Alive and Well: Inflation and the NAIRU During the Slow Recovery," NBER Working Papers 19390, National Bureau of Economic Research, Inc.

——— (2018): "Friedman and Phelps on the Phillips Curve Viewed from a Half Century's Perspective," NBER Working Papers 24891, National Bureau of Economic Research, Inc.

Greenwood, J., Z. Hercowitz, and P. Krusell (1998): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87, 342–36.

Hall, R. E. (2011): "The Long Slump," *American Economic Review*, 101, 431–69.

Hazell, J., J. Herreno, E. Nakamura, and J. Steinsson (2020): "The slope of the Phillips Curve: evidence from US states," Unpublished manuscript.

Hong, G. H., Z. Koczan, W. Lian, and M. Nabar (2018): "More Slack than Meets the Eye? Recent Wage Dynamics in Advanced Economies," Tech. rep., IMF Working Paper 18/50.

Hooper, P., F. S. Mishkin, and A. Sufi (2019): "Prospects for Inflation in a High Pressure Economy: Is the Phillips Curve Dead or is It Just Hibernating?" NBER Working Papers 25792, National Bureau of Economic Research, Inc.

Jones, C. I. (2020): *Macroeconomics*, W.W. Norton.

Jorgensen, P. and K. J. Lansing (2019): "Anchored Inflation Expectations and the Flatter Phillips Curve," Working Paper Series 2019-27, Federal Reserve Bank of San Francisco.

Justiniano, A., G. Primiceri, and A. Tambalotti (2011): "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 14, 101–121.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010): "Investment shocks and business cycles," *Journal of Monetary Economics*, 57, 132–145.

Kareken, J. H. and R. M. Solow (1963): *Lags in Monetary Policy*, New York: Prentice Hall, 14–96.

Laseen, S. and L. E. Svensson (2011): "Anticipated Alternative Policy-Rate Paths in Policy Simulations," *International Journal of Central Banking*, 7, 1–35.

Linde, J. and M. Trabandt (2019): "Resolving the Missing Deflation Puzzle," CEPR Discussion Papers 13690, C.E.P.R. Discussion Papers.

Mackowiak, B. and M. Wiederholt (2009): "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99, 769–803.

Mankiw, N. G. and R. Reis (2002): "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 117, 1295–1328.

Mavroeidis, S., M. Plagborg-MÃžller, and J. H. Stock (2014): "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve," *Journal of Economic Literature*, 52, 124–188.

McLeay, M. and S. Tenreyro (2019): "Optimal Inflation and the Identification of the Phillips Curve," in *NBER Macroeconomics Annual 2019, volume 34*, National Bureau of Economic Research, Inc, NBER Chapters.

Mertens, T. M. and J. C. Williams (2019): "Monetary policy frameworks and the effective lower bound on interest rates," Staff Reports 877, Federal Reserve Bank of New York.

Moulton, B. R. (2018): "The Measurement Of Output, Prices, And Productivity: Whatâs Changed Since The Boskin Commission?" Tech. rep., The Brookings Institution.

Obstfeld, M. (2019): "Global Dimensions of U.S. Monetary Policy," NBER Working Papers 26039, National Bureau of Economic Research, Inc.

Orphanides, A. and J. C. Williams (2005): "Inflation scares and forecast-based monetary policy," *Review of Economic Dynamics*, 8, 498–527.

Powell, J. H. (2019): "Challenges for Monetary Policy," Speech delivered at the "Challenges for Monetary Policy" symposium, sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming.

Primiceri, G. E. (2006): "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy," *Quarterly Journal of Economics*, 121, 867–901.

——— (2019): "Comments on "A Unified Approach to Measuring u*"," *Brookings Papers on Economic Activity*, 215–223.

Rich, R. W., M. H. Linder, and R. Peach (2013): "The parts are more than the whole: separating goods and services to predict core inflation," *Current Issues in Economics and Finance*.

Rognlie, M. (2019): "Comment on "Optimal Inflation and the Identification of the Phillips Curve," in *NBER Macroeconomics Annual 2019, volume 34*, National Bureau of Economic Research, Inc, NBER Chapters.

Rubbo, E. (2020): "Networks, Phillips Curves and Monetary Policy," Unpublished manuscript.

Sargent, T. J. (1999): *The Conquest of American Inflation*, Princeton University Press.

Sbordone, A. M. (2002): "Prices and unit labor costs: a new test of price stickiness," *Journal of Monetary Economics*, 49, 265–292.

——— (2007): "Globalization and Inflation Dynamics: The Impact of Increased Competition," in *International Dimensions of Monetary Policy*, National Bureau of Economic Research, Inc, NBER Chapters, 547–579.

Sims, C. A. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

Smets, F. and R. Wouters (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97, 586 – 606.

Solon, G., R. Barsky, and J. A. Parker (1994): "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *The Quarterly Journal of Economics*, 109, 1–25.

Stansbury, A. and L. H. Summers (2020): "More Rents, Or Less Rent-sharing? The Role Of Declining Worker Power In Recent Macro Trends," Manuscript prepared for the Spring 2020 Brookings Papers.

Stock, J. (2011): "Discussion of Ball and Mazumder, âInflation Dynamics and the Great Recessionâ." *Brookings Papers on Economic Activity*, 42, 387–402.

Stock, J. H. and M. W. Watson (1999): "Business cycle fluctuations in US macroeconomic time series," Elsevier, vol. 1 of *Handbook of Macroeconomics*, chap. 1, 3 – 64.

——— (2007): "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39, 3–33.

——— (2008): "Phillips Curve Inflation Forecasts," NBER Working Papers 14322, National Bureau of Economic Research, Inc.

——— (2019): "Slack and Cyclically Sensitive Inflation," NBER Working Papers 25987, National Bureau of Economic Research, Inc.

Swanson, E. T. and J. C. Williams (2014): "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates," *American Economic Review*, 104, 3154–3185.

Tallman, E. W. and S. Zaman (2017): "Forecasting inflation: Phillips curve effects on services price measures," *International Journal of Forecasting*, 33, 442–457.

Taylor, J. B. (1980): "Aggregate Dynamics and Staggered Contracts," *Journal of Political Economy*, 88, 1–23.

——— (1993): "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.

Woodford, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press.

Yellen, J. (2019): "Remark delivered at a public event, "What's (not) up with inflation?" hosted by the Hutchins Center on Fiscal & Monetary Policy at Brookings," available online at https://www.brookings.edu/research/former-fed-chair-janet-yellen-on-why-the-answer-to-the-inflation-puzzle-matters/.

Zhang, B., J. C. Chan, and J. L. Cross (2018): "Stochastic volatility models with ARMA innovations: An application to G7 inflation forecasts," CAMA Working Papers 2018-32, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.

Appendix A. Facts: Robustness and Extensions

This appendix shows that the key facts described in section 3 are robust to many obvious deviations from the baseline setup. To conduct these robustness checks and extensions we add variables to the baseline VAR, which in some case also results in shorter samples. Therefore, we perform these checks using the conditional forecasts introduced in section 3.1, since this technique typically leads to narrower posterior credible sets and sharper inference.

More measures of price inflation. The first robustness check consists of extending the baseline VAR specification to include three additional inflation measures, based on the Consumer Price Index (CPI), core CPI and the headline PCE price index. Figure A.1 shows that the attenuation in the response of these measures of inflation in the second sample is similar to that of the two inflation variables included in the baseline model. We conclude

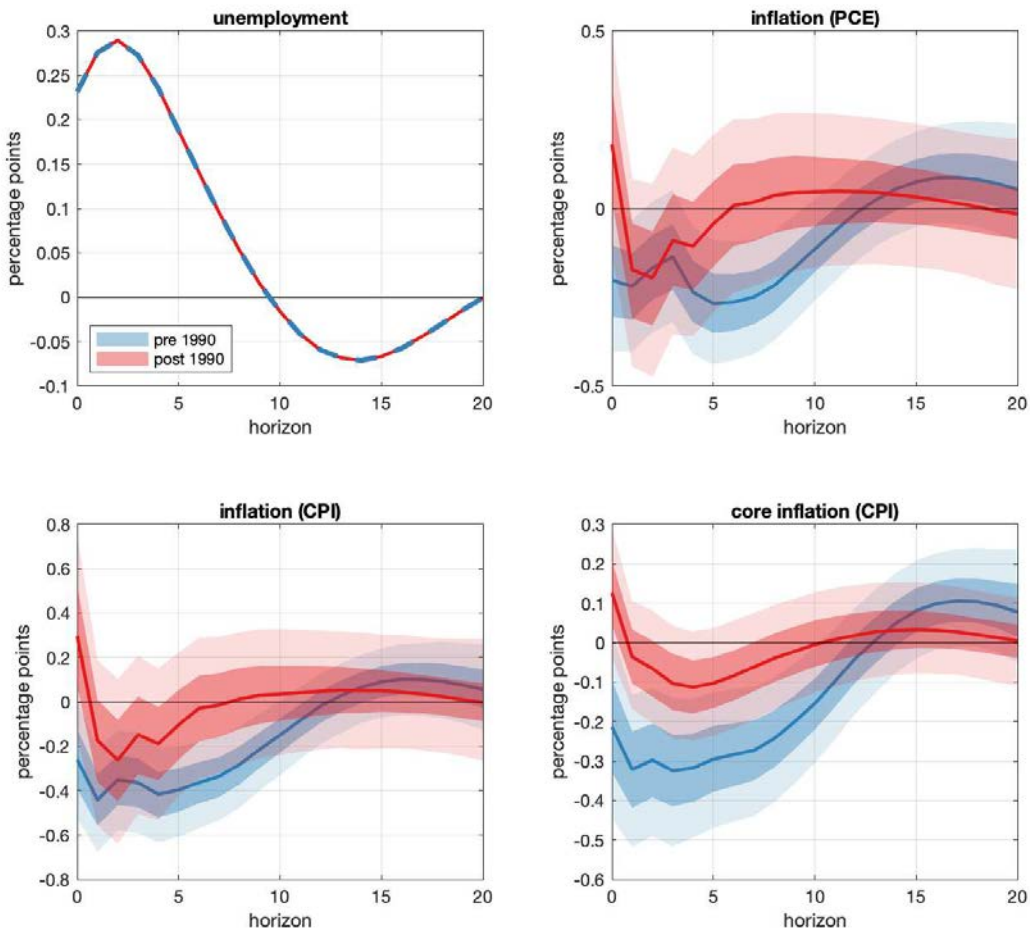

FIGURE A.1. Responses of alternative measures of price inflation, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2, augmented with data on PCE, CPI and core CPI inflation. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

that the marked decrease in the association of price inflation to labor market slack does not depend on how we measure goods prices.

More wage measures. Measuring wages and especially their cyclicality is notoriously difficult. For example, standard wage series like those used in the baseline VAR do not account for changes in the mix of workers over the business cycle, which makes them appear less cyclical (e.g. Solon et al., 1994). A better wage indicator in this respect is the employment cost index (ECI), which focuses on pay rates for specific types of work, thus

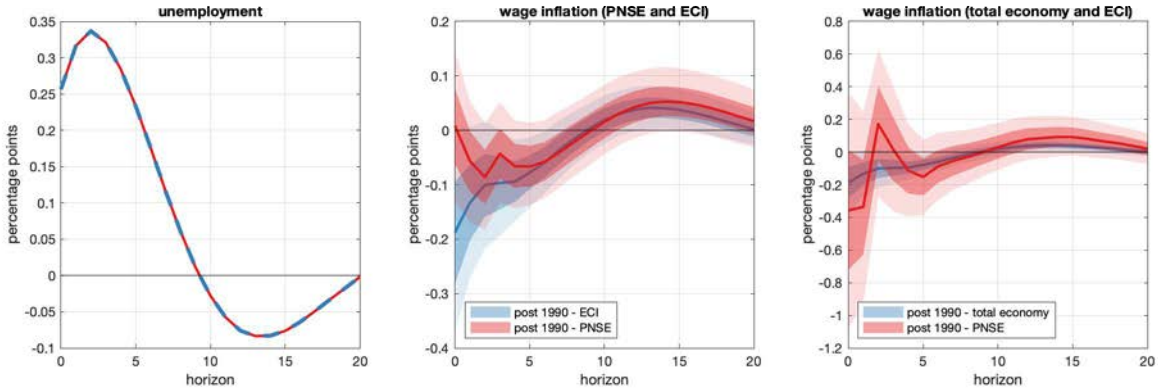

FIGURE A.2. Responses of alternative measures of wage inflation, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2, augmented with the annualized quarterly growth rate of the employment cost index (wages and salaries of all workers in private industries). The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

controlling for these composition effects. These data, however, are only available since 1980 (or 1975, but without seasonal adjustment). Therefore, using them in the first sample is problematic. But we can add a wage inflation measure from the ECI (wages and salaries of all workers in private industries) to the VAR in the second sample. Figure A.2 compares the response of the ECI wage measure to the ones in the baseline VAR. As expected, the ECI is more clearly cyclical than the other wage measures, falling by a larger amount than the PNSE on impact. Moreover, its response is much more sharply estimated than that of the wage measure for the total economy, confirming that nominal wages remain sensitive to business cycle shocks after 1990.

Figure A.3 presents another interesting piece of evidence on the behavior of wages. It compares the conditional forecasts of wage inflation, as measured by the ECI, when the VAR is estimated with and without data after 2007. The purpose of this exercise is to verify the extent to which the Great Recession and the subsequent period of slow recovery might have impacted wage dynamics. In particular, one hypothesis often mentioned to explain the missing disinflation over the Great Recession is that downward nominal wage rigidity might have limited the fall in wages, and hence in costs and inflation. The conditional forecasts in the figure suggest that this factor is unlikely to be a major driver of the reduced cyclical

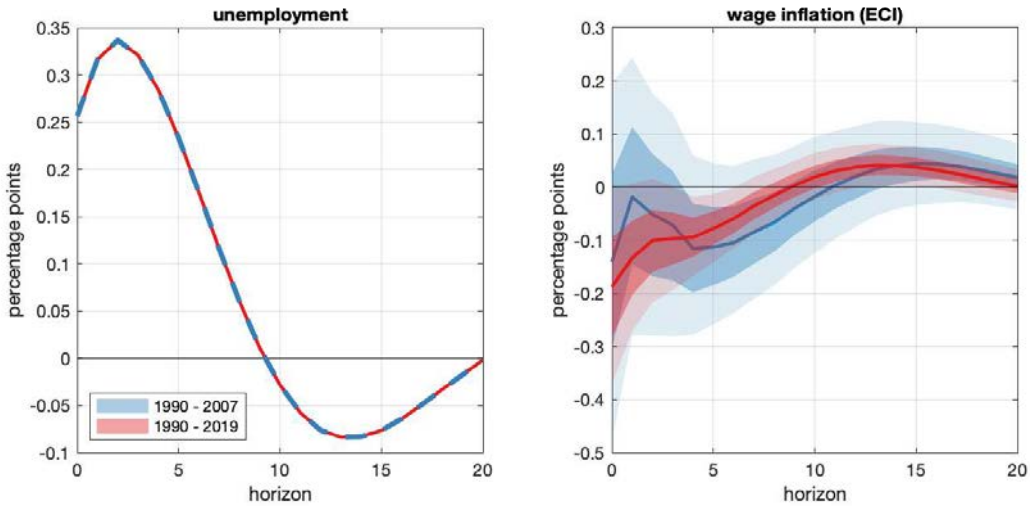

FIGURE A.3. Responses of alternative measures of wage inflation, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2, augmented with the annualized quarterly growth rate of the employment cost index (wages and salaries of all workers in private industries). The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The 1990-2007 and 1990-2019 samples consist of data from 1989:I to 2007:IV and from from 1989:I to 2019:III, respectively.

sensitivity of inflation since 1990, given that the dynamics of wage inflation are similar regardless of whether we include post-2007 data.

More measures of the labor share. To further probe the stability of the business cycle correlation between unemployment and unit labor costs, we also extend the baseline VAR to include two additional commonly used measures of the labor share. These are the logarithm of the labor share in the nonfarm business sector (NFBS) and in the nonfinancial corporate sector (NFCS), from the productivity and costs release. Figure A.4 plots the behavior of these two variables, which is similar to that of the labor share in the total economy included in the baseline VAR.

More measures of employment and real activity. We now augment the baseline VAR with the logarithm of the employment-to-population ratio and of the ratio between GDP and the CBO estimate of potential GDP. Figure A.5 illustrates that the responses of these additional variables are almost exactly the same in the two samples, confirming the results based on hours and GDP of the previous subsections.

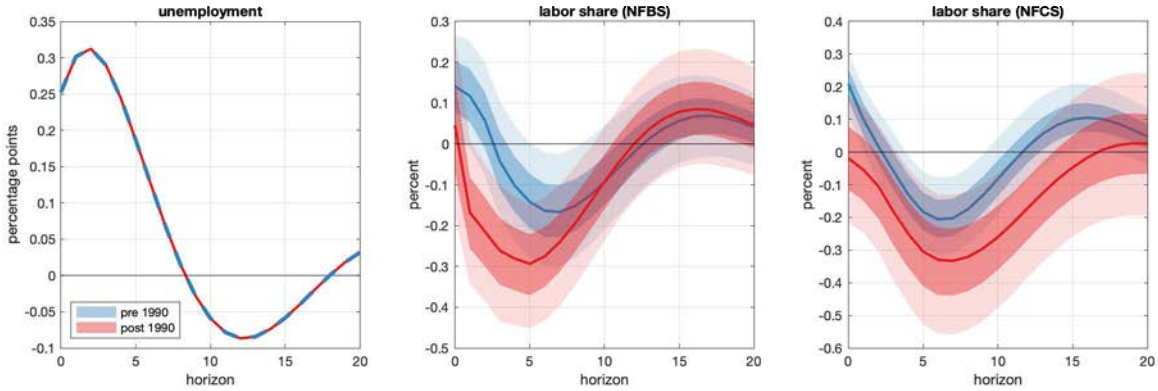

FIGURE A.4. Responses of alternative measures of the labor share, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2, augmented with the logarithm of the labor share in the nonfarm business sector and in the nonfinancial corporate sector. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

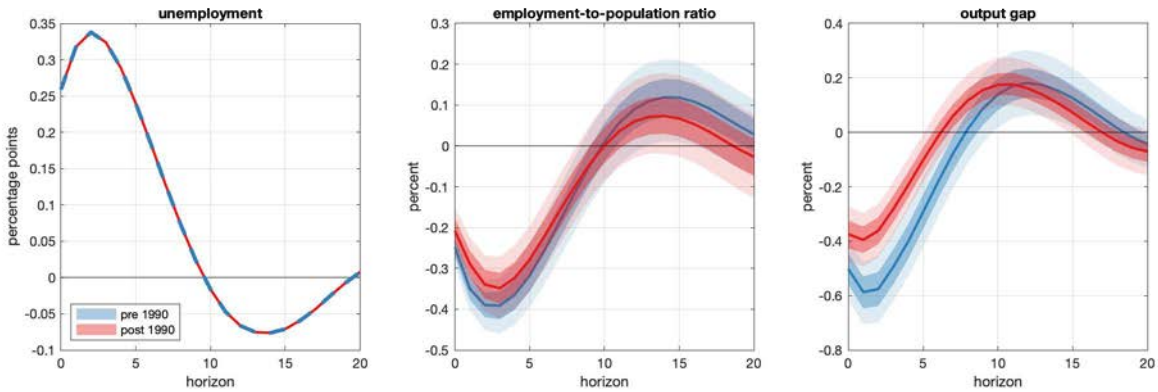

FIGURE A.5. Responses of the employment-to-population ratio and the output gap, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2, augmented with the logarithm of the employment-to-population ratio and of the ratio between GDP and potential GDP from the Congressional Budget Office (CBO). The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

Alternative ways of splitting the sample. As mentioned in section 2.2, there are at least two other dates in which we could have reasonably split the sample: 1984—right after the Volcker disinflation—and 1994—when inflation was more clearly stabilized around 2%. Figures A.6 and A.7 show that splitting the sample at these two alternative dates has little

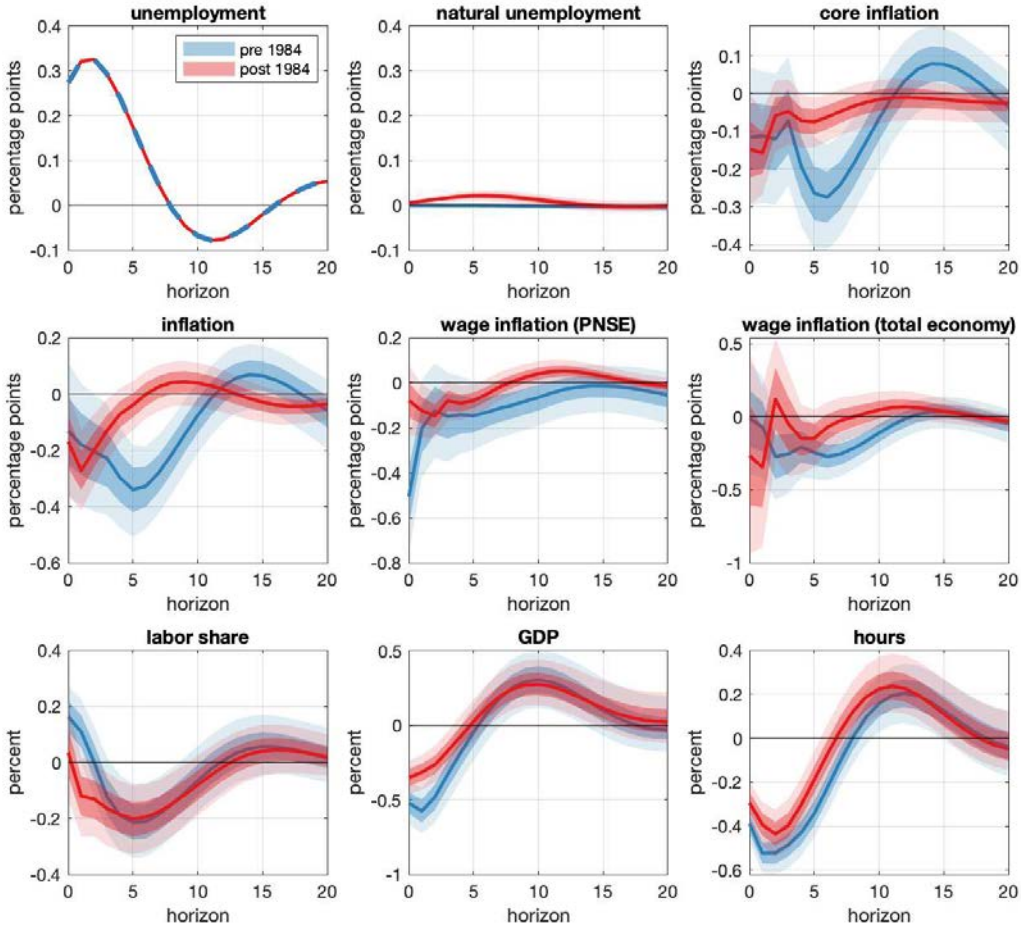

FIGURE A.6. Responses of all variables, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1984 samples consist of data from 1964:II to 1983:IV, and from 1983:I to 2019:III, respectively.

impact on the main results. Inflation is much more stable in the second sample, while the dynamics of the real variables change little.

Impulse responses to a typical business cycle shock. Finally, as a complementary exercise, figure A.8 presents the impulse responses of unemployment and inflation to a typical business cycle shock. We define this shock as the linear combination of structural

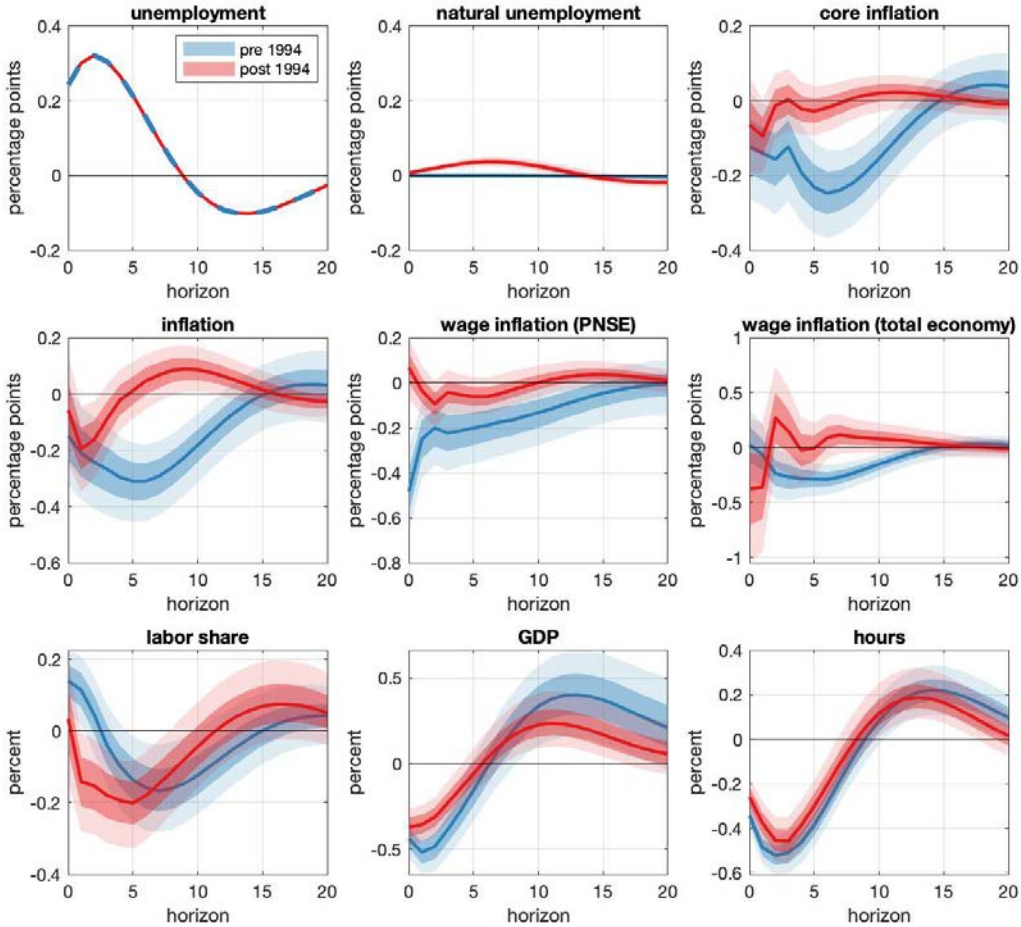

FIGURE A.7. Responses of all variables, conditional on unemployment following the path in the first subplot. These responses are computed by applying the methodology described in section 3.1 to the baseline VAR of section 2.2. The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1994 samples consist of data from 1964:II to 1993:IV, and from 1993:I to 2019:III, respectively.

disturbances that drives the largest share of unemployment variation at business cycle frequencies, as in Giannone et al. (2019) and Angeletos et al. (2019).[36] The impulse responses

[36]More precisely, we compute the spectral density of y_t implied by the VAR representation (2.1). We use this spectrum to evaluate the variance of unemployment at business cycle frequencies (those related to periods of length between 6 and 32 quarters, as in Stock and Watson, 1999), as a function of the columns of all possible $\Sigma^{\frac{1}{2}}$ matrices. The BC shock is pinned down by the column associated with the maximal variance of unemployment at these frequencies. For technical details and the specific implementation of this idea, see Giannone et al. (2019) and Angeletos et al. (2019).

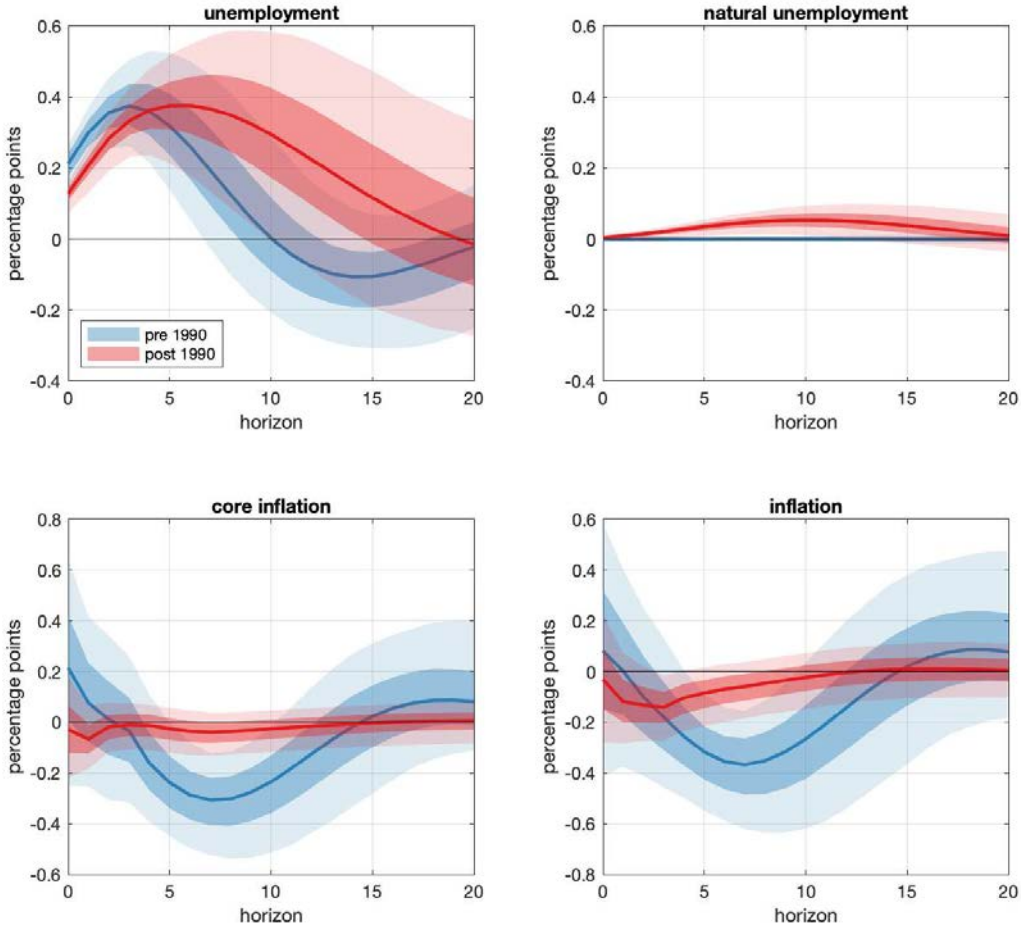

FIGURE A.8. Impulse responses of unemployment and price inflation to a business cycle shock. The impulse responses are from the baseline VAR described in section 2.2. The business cycle shock is the linear combination of structural disturbances that drives the highest share of unemployment variation at business cycle frequencies, and it is identified using the methodology of Giannone et al. (2019) and Angeletos et al. (2019). The solid lines are posterior medians, while the shaded areas correspond to 68- and 95-percent posterior credible regions. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1989:I to 2019:III, respectively.

computed following this approach are essentially identical to those of figure 3.1, suggesting that the two combinations of structural shocks—the one responsible for the bulk of business cycle fluctuations and the one associated with the one-step-ahead forecast error in unemployment—are virtually the same.

Appendix B. Details on the DSGE Model

This appendix provides some details of the DSGE model used in section 6.

B.1. **Equilibrium conditions.** The model economy is populated by eight classes of agents: 1) a continuum of households, who consume and supply differentiated labor; 2) competitive labor aggregators that combine labor supplied by individual households; 3) competitive final good-producing firms that aggregate the intermediate goods into a final product; 4) a continuum of monopolistically competitive intermediate good producing firms; 5) competitive capital producers that convert final goods into capital; 6) a continuum of entrepreneurs who purchase capital using both internal and borrowed funds and rent it to intermediate good producing firms; 7) a representative bank collecting deposits from the households and lending funds to the entrepreneurs; and finally 8) a government, composed of a monetary authority that sets short-term interest rates and a fiscal authority that sets public spending and collects taxes. We solve each agent's problem, and derive the resulting equilibrium conditions, which we approximate around the non-stochastic steady state. Since the derivation follows closely the literature (e.g.Christiano et al., 2005), we describe here the log-linearized conditions.

Growth in the economy is driven by technological progress, $Z_t^* = e^{\frac{1}{1-\alpha}\tilde{z}_t} Z_t^p e^{\gamma t}$, which is assumed to include a deterministic trend ($e^{\gamma t}$), a stochastic trend (Z_t^p), and a stationary component ($\tilde{z}_t$), where α is the income share of capital (after paying mark-ups and fixed costs in production). Trending variables are divided by Z_t^* to express the model's equilibrium conditions in terms of the stationary variables. In what follows, all variables are expressed in log deviations from their steady state, and steady-state values are denoted by $*$-subscripts.

The stationary component of productivity $\tilde{z}_t$ and the growth rate of the stochastic trend $z_t^p = log(Z_t^p/Z_{t-1}^p)$ are assumed to follow AR(1) processes:

$$\tilde{z}_t = \rho_z \tilde{z}_{t-1} + \sigma_z \varepsilon_{z,t},\ \varepsilon_{z,t} \sim N(0,1). \tag{B.1}$$

$$z_t^p = \rho_{z^p} z_{t-1}^p + \sigma_{z^p} \epsilon_{z^p,t},\ \epsilon_{z^p,t} \sim N(0,1). \tag{B.2}$$

The growth rate of technology evolves thus according to

$$z_t \equiv log(Z_t^*/Z_{t-1}^*) - \gamma = \frac{1}{1-\alpha}(\tilde{z}_t - \tilde{z}_{t-1}) + z_t^p, \tag{B.3}$$

where γ is the steady-state growth rate of the economy.

The *optimal allocation of consumption* satisfies the following Euler equation:

$$\begin{aligned} c_t = &-\frac{1-\bar{h}}{\sigma_c(1+\bar{h})}\left(R_t - E_t[\pi_{t+1}] + cy_t\right) + \frac{\bar{h}}{1+\bar{h}}\left(c_{t-1} - z_t\right) \\ &+ \frac{1}{1+\bar{h}}E_t\left[c_{t+1} + z_{t+1}\right] + \frac{(\sigma_c - 1)}{\sigma_c(1+\bar{h})}\frac{w_* L_*}{c_*}\left(L_t - E_t[L_{t+1}]\right), \end{aligned} \tag{B.4}$$

where c_t is consumption, L_t denotes hours worked, R_t is the nominal interest rate, and π_t is inflation. The parameter σ_c captures the degree of relative risk aversion while $\bar{h} \equiv he^{-\gamma}$ depends on the degree of habit persistence in consumption, h, and steady-state growth. This equation includes hours worked because utility is non-separable in consumption and leisure.

The convenience yield cy_t contains both a liquidity component cy_t^l and a safety component cy_t^s

$$cy_t = cy_t^l + cy_t^s, \tag{B.5}$$

where we let each premium be given by the sum of two AR(1) processes, one that captures highly persistent movements ($cy_t^{P,l}$ and $cy_t^{P,s}$) with autoregressive coefficients fixed at .99, and one that captures transitory fluctuations ($\tilde{cy}_t^{P,l}$ and $\tilde{cy}_t^{P,s}$).

The *optimal investment decision* satisfies the following relationship between the level of investment i_t, measured in terms of consumption goods, and the value of capital in terms of consumption q_t^k:

$$i_t = \frac{q_t^k}{S'' e^{2\gamma}(1+\bar{\beta})} + \frac{1}{1+\bar{\beta}}\left(i_{t-1} - z_t\right) + \frac{\bar{\beta}}{1+\bar{\beta}}E_t\left[i_{t+1} + z_{t+1}\right] + \mu_t. \tag{B.6}$$

This relationship shows that investment is affected by investment adjustment costs (S'' is the second derivative of the adjustment cost function) and by an exogenous process μ_t, which we call "marginal efficiency of investment", that alters the rate of transformation between consumption and installed capital (see Greenwood et al. (1998)). The shock μ_t follows an AR(1) process with parameters ρ_μ and σ_μ. The parameter $\bar{\beta} \equiv \beta e^{(1-\sigma_c)\gamma}$ depends

on the intertemporal discount rate in the household utility function, β, on the degree of relative risk aversion σ_c, and on the steady-state growth rate γ.

The *capital stock*, $\bar{k}_t$, which we refer to as "installed capital", evolves as

$$\bar{k}_t = \left(1 - \frac{i_*}{\bar{k}_*}\right)\left(\bar{k}_{t-1} - z_t\right) + \frac{i_*}{\bar{k}_*} i_t + \frac{i_*}{\bar{k}_*} S'' e^{2\gamma}(1+\bar{\beta})\mu_t, \tag{B.7}$$

where $i_*/\bar{k}_*$ is the steady state investment to capital ratio. Capital is subject to variable capacity utilization u_t; *effective capital* rented out to firms, k_t, is related to $\bar{k}_t$ by:

$$k_t \;=\; u_t - z_t + \bar{k}_{t-1}. \tag{B.8}$$

The optimality condition determining the *rate of capital utilization* is given by

$$\frac{1 -}{} r_t^k = u_t, \tag{B.9}$$

where r_t^k is the rental rate of capital and ψ captures the utilization costs in terms of foregone consumption.

Real marginal costs for firms are given by

$$mc_t \;=\; w_t + \alpha L_t - \alpha k_t, \tag{B.10}$$

where w_t is the real wage. From the optimality conditions of goods producers it follows that all firms have the same *capital-labor ratio*:

$$k_t \;=\; w_t - r_t^k + L_t. \tag{B.11}$$

We include financial frictions in the model, building on the work of Bernanke et al. (1999), Christiano et al. (2003), De Graeve (2008), and Christiano et al. (2014). We assume that banks collect deposits from households and lend to entrepreneurs who use these funds as well as their own wealth to acquire physical capital, which is rented to intermediate goods producers. Entrepreneurs are subject to idiosyncratic disturbances that affect their ability to manage capital. Their revenue may thus turn out to be too low to pay back the loans received by the banks. The banks therefore protect themselves against default risk by pooling all loans and charging a spread over the deposit rate. This spread may vary as a function of entrepreneurs' leverage and riskiness.

The *realized return on capital* is given by

$$\tilde{R}_t^k - \pi_t = \frac{r_*^k}{r_*^k + (1-\delta)} r_t^k + \frac{(1-\delta)}{r_*^k + (1-\delta)} q_t^k - q_{t-1}^k, \tag{B.12}$$

where $\tilde{R}_t^k$ is the gross nominal return on capital for entrepreneurs, r_*^k is the steady state value of the rental rate of capital r_t^k, and δ is the depreciation rate.

The *excess return on capital* (the spread between the expected return on capital and the riskless rate) can be expressed as a function of the convenience yield cy_t, the entrepreneurs' leverage (i.e. the ratio of the value of capital to net worth), and "risk shocks" $\tilde{\sigma}_{\omega,t}$ capturing mean-preserving changes in the cross-sectional dispersion of ability across entrepreneurs (see Christiano et al. (2014)):

$$E_t \left[\tilde{R}_{t+1}^k - R_t \right] = cy_t + \zeta_{sp,b} \left(q_t^k + \bar{k}_t - n_t \right) + \tilde{\sigma}_{\omega,t}, \tag{B.13}$$

where n_t is entrepreneurs' net worth, $\zeta_{sp,b}$ is the elasticity of the credit spread to the entrepreneurs' leverage $(q_t^k + \bar{k}_t - n_t)$. $\tilde{\sigma}_{\omega,t}$ follows an AR(1) process with parameters ρ_{σ_ω} and σ_{σ_ω}. *Entrepreneurs' net worth* n_t evolves in turn according to

(B.14)

$$\begin{aligned} n_t \quad = \quad & \zeta_{n,\tilde{R}^k} \left(\tilde{R}_t^k - \pi_t \right) - \zeta_{n,R} \left(R_{t-1} - \pi_t + cy_{t-1} \right) + \zeta_{n,qK} \left(q_{t-1}^k + \bar{k}_{t-1} \right) + \zeta_{n,n} n_{t-1} \\ & - \gamma_* \tfrac{v_*}{n_*} z_t - \tfrac{\zeta_{n,\sigma_\omega}}{\zeta_{sp,\sigma_\omega}} \tilde{\sigma}_{\omega,t-1}, \end{aligned}$$

where the ζ's denote elasticities, that depend among others on the entrepreneurs' steady-state default probability $F(\bar{\omega})$, where γ_* is the fraction of entrepreneurs that survive and continue operating for another period, and where v_* is the entrepreneurs' real equity divided by Z_t^*, in steady state.

The *production function* is

$$y_t = \Phi_p \left(\alpha k_t + (1-\alpha) L_t \right), \tag{B.15}$$

where $\Phi_p = 1 + \Phi / y_*$, and Φ measures the size of fixed costs in production. The *resource constraint* is:

$$y_t = g_* g_t + \frac{c_*}{y_*} c_t + \frac{i_*}{y_*} i_t + \frac{r_*^k k_*}{y_*} u_t. \tag{B.16}$$

where $g_t = log(\frac{G_t}{Z_t^* y_* g_*})$ and $g_* = 1 - \frac{c_* + i_*}{y_*}$. *Government spending* g_t is assumed to follow the exogenous process:

$$g_t = \rho_g g_{t-1} + \sigma_g \varepsilon_{g,t} + \eta_{gz} \sigma_z \varepsilon_{z,t}.$$

Optimal decisions for price and wage setting deliver the *price and wage Phillips curves*, which are respectively:

$$\pi_t = \kappa_p mc_t + \frac{\iota_p}{1 + \iota_p \bar{\beta}} \pi_{t-1} + \frac{\bar{\beta}}{1 + \iota_p \bar{\beta}} E_t[\pi_{t+1}] + \lambda_{f,t}, \tag{B.17}$$

and

(B.18)

$$w_t = \frac{(1 - \zeta_w \bar{\beta})(1 - \zeta_w)}{(1 + \bar{\beta})\zeta_w((\lambda_w - 1)\epsilon_w + 1)} \left(w_t^h - w_t \right) - \frac{1 + \iota_w \bar{\beta}}{1 + \bar{\beta}} \pi_t + \frac{1}{1 + \bar{\beta}} \left(w_{t-1} - z_t + \iota_w \pi_{t-1} \right) + \frac{\bar{\beta}}{1 + \bar{\beta}} E_t \left[w_{t+1} + z_{t+1} + \pi_{t+1} \right] + \lambda_{w,t},$$

where $\kappa = \frac{(1-\zeta_p \bar{\beta})(1-\zeta_p)}{(1+\iota_p \bar{\beta})\zeta_p((\Phi_p - 1)\epsilon_p + 1)}$, the parameters ζ_p, ι_p, and ϵ_p are the Calvo parameter, the degree of indexation, and the curvature parameter in the Kimball aggregator for prices, and ζ_w, ι_w, and ϵ_w are the corresponding parameters for wages. w_t^h measures the household's marginal rate of substitution between consumption and labor, and is given by:

$$w_t^h = \frac{1}{1 - \bar{h}} \left(c_t - \bar{h} c_{t-1} + \bar{h} z_t \right) + \nu_l L_t, \tag{B.19}$$

where ν_l characterizes the curvature of the disutility of labor (and would equal the inverse of the Frisch elasticity in the absence of wage rigidities). The mark-ups $\lambda_{f,t}$ and $\lambda_{w,t}$ follow the exogenous ARMA(1,1) processes:

$$\lambda_{f,t} = \rho_{\lambda_f} \lambda_{f,t-1} + \sigma_{\lambda_f} \varepsilon_{\lambda_f,t} - \eta_{\lambda_f} \sigma_{\lambda_f} \varepsilon_{\lambda_f,t-1},$$

and

$$\lambda_{w,t} = \rho_{\lambda_w} \lambda_{w,t-1} + \sigma_{\lambda_w} \varepsilon_{\lambda_w,t} - \eta_{\lambda_w} \sigma_{\lambda_w} \varepsilon_{\lambda_w,t-1}.$$

Finally, the monetary authority follows a *policy feedback rule*:

$$\begin{aligned} R_t &= \rho_R R_{t-1} + (1 - \rho_R) \left(\psi_1 (\pi_t - \pi_t^*) + \psi_2 (y_t - y_t^*) \right) \\ &+ {}_3 \left((y_t - y_t^*) - (y_{t-1} - y_{t-1}^*) \right) + r_t^m. \end{aligned} \tag{B.20}$$

where π_t^* is a time-varying inflation target, y_t^* is a measure of the "full-employment level of output," and r_t^m captures exogenous departures from the policy rule.

The time-varying inflation target π_t^* is meant to capture the rise and fall of inflation and interest rates in the estimation sample.[37] As in Aruoba and Schorfheide (2010) and Del Negro and Eusepi (2011), we use data on long-run inflation expectations in the estimation of the model. This allows us to pin down the target inflation rate to the extent that long-run inflation expectations contain information about the central bank's objective. The time-varying *inflation target* evolves according to

$$\pi_t^* = \rho_{\pi^*}\pi_{t-1}^* + \sigma_{\pi^*}\epsilon_{\pi^*,t}, \tag{B.21}$$

where $0 < \rho_{\pi^*} < 1$ and $\epsilon_{\pi^*,t}$ is an iid shock. We model π_t^* as a stationary process, although our prior for ρ_{π^*} will force this process to be highly persistent.

The "full-employment level of output" y_t^* represents the level of output that would obtain if prices and wages were fully flexible and if there were no markup shocks. This variable along with the natural rate of interest r_t^* are obtained by solving the model without nominal rigidities and markup shocks (that is, equations (B.4) through (B.19) with $\zeta_p = \zeta_w = 0$, and $\lambda_{f,t} = \lambda_{w,t} = 0$).

The exogenous component of the policy rule r_t^m evolves according to the following process

$$r_t^m = \rho_{r^m} r_{t-1}^m + \epsilon_t^R + \sum_{k=1}^{K} \epsilon_{k,t-k}^R, \tag{B.22}$$

where ϵ_t^R is the usual contemporaneous policy shock, and $\epsilon_{k,t-k}^R$ is a policy shock that is known to agents at time $t-k$, but affects the policy rule k periods later, that is, at time t. We assume that $\epsilon_{k,t-k}^R \sim N(0, \sigma_{k,r}^2),\ i.i.d.$ As argued in Laseen and Svensson (2011), such anticipated policy shocks allow us to capture the effects of the zero lower bound on nominal interest rates, as well as the effects of forward guidance in monetary policy.

B.2. **State Space Representation and Data.** We use the method in Sims (2002) to solve the system of log-linear approximate equilibrium conditions and obtain the transition

[37] The assumption that the inflation target moves exogenously is of course a simplification. A more realistic model would for instance relate movements in trend inflation to the evolution of the policy makers' understanding of the output-inflation trade-of, as in Sargent (1999) or Primiceri (2006).

equation, which summarizes the evolution of the vector of state variables s_t:

$$s_t = \mathcal{T}(\theta)s_{t-1} + \mathcal{R}(\theta)\epsilon_t. \tag{B.23}$$

where θ is a vector collecting all the DSGE model parameters and ϵ_t is a vector of all structural shocks. The state-space representation of our model is composed of the transition equation (B.23), and a system of measurement equations:

$$Y_t = \mathcal{D}(\theta) + \mathcal{Z}(\theta)s_t, \tag{B.24}$$

mapping the states into the observable variables Y_t, which we describe in detail next.

The estimation of the model is based on data on real output growth (including both GDP and GDI measures), consumption growth, investment growth, real wage growth, hours worked, inflation (measured by core PCE and GDP deflators), short- and long- term interest rates, 10-year inflation expectations, market expectations for the federal funds rate up to 6 quarters ahead, Aaa and Baa credit spreads, and total factor productivity growth unadjusted for variable utilization. Measurement equations (B.24) relate these observables to the model variables as follows:

(B.25)

$$\begin{aligned}
\text{GDP growth} &= 100\gamma + (y_t - y_{t-1} + z_t) + e_t^{gdp} - e_{t-1}^{gdp} \\
\text{GDI growth} &= 100\gamma + (y_t - y_{t-1} + z_t) + e_t^{gdi} - e_{t-1}^{gdi} \\
\text{Consumption growth} &= 100\gamma + (c_t - c_{t-1} + z_t) \\
\text{Investment growth} &= 100\gamma + (i_t - i_{t-1} + z_t) \\
\text{Real Wage growth} &= 100\gamma + (w_t - w_{t-1} + z_t) \\
\text{Hours} &= \bar{L} + L_t \\
\text{Core PCE Inflation} &= \pi_* + \pi_t + e_t^{pce} \\
\text{GDP Deflator Inflation} &= \pi_* + \delta_{gdpdef} + \gamma_{gdpdef} * \pi_t + e_t^{gdpdef} \\
\text{FFR} &= R_* + R_t \\
\text{FFR}_{t,t+j}^{e} &= R_* + E_t\left[R_{t+j}\right],\ j = 1, ..., 6 \\
\text{10y Nominal Bond Yield} &= R_* + E_t\left[\tfrac{1}{40}\textstyle\sum_{j=0}^{39} R_{t+j}\right] + e_t^{10y} \\
\text{10y Infl. Expectations} &= \pi_* + E_t\left[\tfrac{1}{40}\textstyle\sum_{j=0}^{39} \pi_{t+j}\right] \\
\text{Baa - 10-year Treasury Spread} &= cy_*^l + cy_*^s + SP_* + E_t\tfrac{1}{80}\textstyle\sum_{j=0}^{79}\left[\tilde{R}_{t+j+1}^k - R_{t+j}\right] + e_t^{Baa} \\
\text{TFP growth, demeaned} &= z_t + \tfrac{\alpha}{1-\alpha}\left(u_t - u_{t-1}\right) + e_t^{tfp}.
\end{aligned}$$

All variables are measured in percent. The terms π_* and R_* measure respectively the net steady-state inflation rate and short-term nominal interest rate, expressed in percentage terms, and $\bar{L}$ captures the mean of hours (this variable is measured as an index). We assume that some of the variables are measured with "error," that is, the observed value equals the model implied value plus an AR(1) exogenous process e_t^* that can be thought of either measurement errors or some other unmodeled source of discrepancy between the model and the data, as in Boivin and Giannoni (2006). For instance, the terms e_t^{gdp} and e_t^{gdi} capture measurement error of total output.[38] Alternatively, for the long-term nominal interest rate, the term e_t^{10y} captures fluctuations in term premia not captured by the model.

B.3. **Inference, Prior and Posterior Parameter Estimates.** We estimate the model using Bayesian techniques. This requires the specification of a prior distribution for the model parameters. For most parameters common with Smets and Wouters (2007), we use the same marginal prior distributions. As an exception, we favor a looser prior than Smets and Wouters (2007) for the quarterly steady state inflation rate π_*; it is centered at 0.75% and has a standard deviation of 0.4%. Regarding the financial frictions, we specify priors for the parameters SP_*, $\zeta_{sp,b}$, ρ_{σ_ω}, and σ_{σ_ω}, while we fix the parameters corresponding to the steady state default probability and the survival rate of entrepreneurs, respectively. In turn, these parameters imply values for the parameters of (B.14). Information on the priors is provided in Table 2.

B.4. **Data Construction.** Data on real GDP (GDPC), the GDP deflator (GDPDEF), core PCE inflation (PCEPILFE), nominal personal consumption expenditures (PCEC), and nominal fixed private investment (FPI) are produced at a quarterly frequency by the Bureau of Economic Analysis, and are included in the National Income and Product Accounts (NIPA). Average weekly hours of production and nonsupervisory employees for total private industries (AWHNONAG), civilian employment (CE16OV), and the civilian non-institutional population (CNP16OV) are produced by the Bureau of Labor Statistics (BLS) at a monthly frequency. The first of these series is obtained from the Establishment Survey,

[38] We introduce correlation in the measurement errors for GDP and GDI, which evolve as follows:

$$\begin{aligned} e_t^{gdp} &= \rho_{gdp} \cdot e_{t-1}^{gdp} + \sigma_{gdp}\epsilon_t^{gdp},\ \epsilon_t^{gdp} \sim i.i.d.N(0,1) \\ e_t^{gdi} &= \rho_{gdi} \cdot e_{t-1}^{gdi} + \varrho_{gdp} \cdot \sigma_{gdp}\epsilon_t^{gdp} + \sigma_{gdi}\epsilon_t^{gdi},\ \epsilon_t^{gdi} \sim i.i.d.N(0,1). \end{aligned}$$

The measurement errors for GDP and GDI are thus stationary in *levels*, and enter the observation equation in first differences (e.g. $e_t^{gdp} - e_{t-1}^{gdp}$ and $e_t^{gdp} - e_{t-1}^{gdp}$). GDP and GDI are also cointegrated as they are driven by a comment stochastic trend.

and the remaining from the Household Survey. Both surveys are released in the BLS Employment Situation Summary. Since our models are estimated on quarterly data, we take averages of the monthly data. Compensation per hour for the non-farm business sector (COMPNFB) is obtained from the Labor Productivity and Costs release, and produced by the BLS at a quarterly frequency. The data are transformed following Smets and Wouters (2007), with the exception of the civilian population data, which are filtered using the Hodrick-Prescott filter to remove jumps around census dates. The federal funds rate is obtained from the Federal Reserve Board's H.15 release at a business day frequency. We take quarterly averages of the annualized daily data and divide by four. Let Δ denote the temporal difference operator. Then:

$$
\begin{aligned}
\text{Output growth} &= 100 * \Delta LN((GDPC)/CNP16OV) \\
\text{Consumption growth} &= 100 * \Delta LN((PCEC/GDPDEF)/CNP16OV) \\
\text{Investment growth} &= 100 * \Delta LN((FPI/GDPDEF)/CNP16OV) \\
\text{Real wage growth} &= 100 * \Delta LN(COMPNFB/GDPDEF) \\
\text{Hours worked} &= 100 * LN((AWHNONAG * CE16OV/100)/CNP16OV) \\
\text{GDP Deflator Inflation} &= 100 * \Delta LN(GDPDEF) \\
\text{Core PCE Inflation} &= 100 * \Delta LN(PCEPILFE) \\
\text{FFR} &= (1/4) * FEDERAL\ FUNDS\ RATE
\end{aligned}
$$

Long-run inflation expectations are obtained from the Blue Chip Economic Indicators survey and the Survey of Professional Forecasters available from the FRB Philadelphia's Real-Time Data Research Center. Long-run inflation expectations (average CPI inflation over the next 10 years) are available from 1991Q4 onward. Prior to 1991Q4, we use the 10-year expectations data from the Blue Chip survey to construct a long time series that begins in 1979Q4. Since the Blue Chip survey reports long-run inflation expectations only twice a year, we treat these expectations in the remaining quarters as missing observations and adjust the measurement equation of the Kalman filter accordingly. Long-run inflation expectations $\pi_t^{O,40}$ are therefore measured as

$$\text{10y Infl Exp} = (\text{10-year average CPI inflation forecast} - 0.50)/4.$$

where 0.50 is the average difference between CPI and GDP annualized inflation from the beginning of the sample to 1992. We divide by 4 to express the data in quarterly terms.

We measure *Spread* as the annualized Moody's Seasoned Baa Corporate Bond Yield spread over the 10-Year Treasury Note Yield at Constant Maturity. Both series are available from the Federal Reserve Board's H.15 release. Like the federal funds rate, the spread data are also averaged over each quarter and measured at a quarterly frequency. This leads to:

$$\text{Spread} = (1/4) * (Baa\ Corporate\ -\ 10\ year\ Treasury).$$

Similarly,

$$\text{10y Bond yield} = (1/4) * (10\ year\ Treasury).$$

Lastly, TFP growth is measured using John Fernald's TFP growth series, unadjusted for changes in utilization. That series is demeaned, divided by 4 to express it in quarterly growth rates, and divided by Fernald's estimate of $(1-\alpha)$ to convert it in labor augmenting terms:

$$\text{TFP growth, demeaned } = (1/4)*(\text{Fernald's TFP growth, unadjusted, demeaned})\ /(1-\alpha).$$

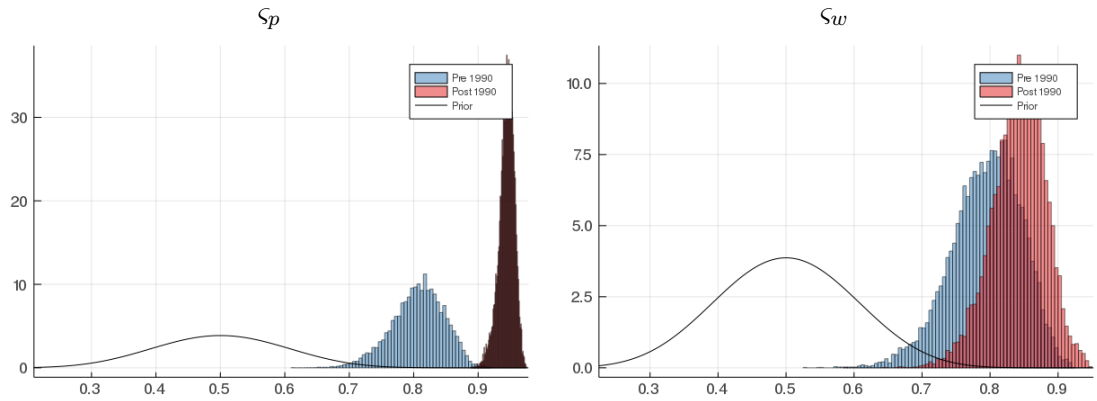

Figure C.1. Prior and posterior distributions for some additional parameters of the NY Fed DSGE model described in section 6 and appendix B.4. The pre- and post-1990 samples consist of data from 1964:II to 1989:IV, and from 1990:I to 2019:III, respectively.

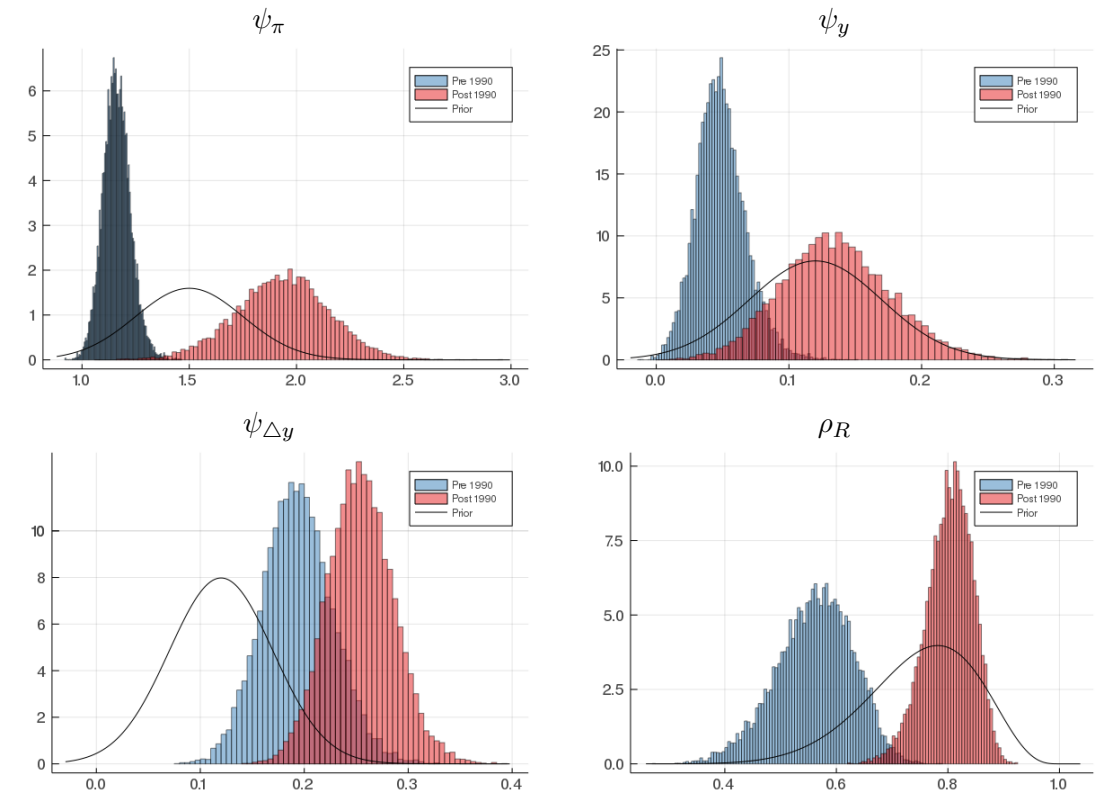

Figure C.2. Prior and posterior distributions for selected parameters of the NY Fed DSGE model described in section 6 and appendix B.4. In this experiment, the model is estimated allowing for two distinct monetary policy regimes, pre- and post-1990, but assuming that all other coefficients are constant for the entire sample that now excludes the ZLB period, from 1964:II to 2008:III.

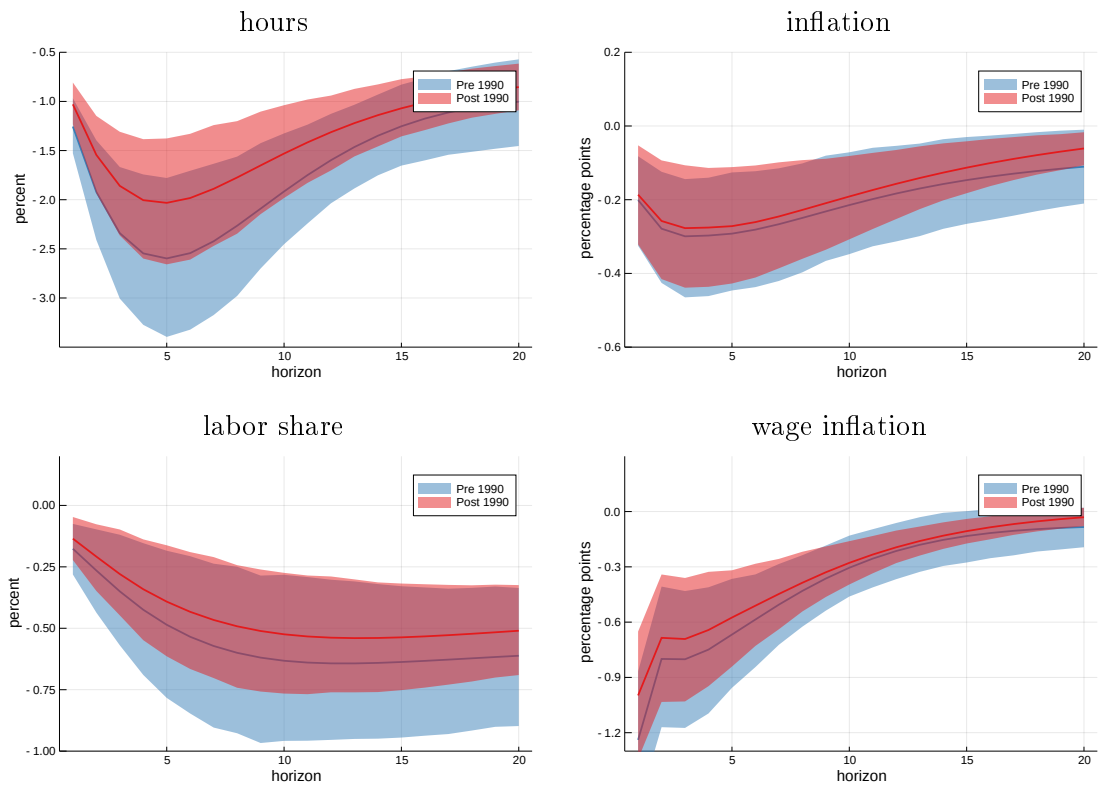

FIGURE C.3. Impulse responses of hours worked, price inflation, the labor share, and wage inflation to a shock to the unemployment equation. The impulse responses are from the VAR approximation of the NY Fed DSGE model described in section 6 and appendix B.4. The shock is identified using a Cholesky 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 for two distinct monetary policy regimes, pre- and post-1990, but assuming that all other coefficients are constant for the entire sample that now excludes the ZLB period, from 1964:II to 2008:III.

	Dist	Mean	Std Dev		Dist	Mean	Std Dev
Policy Parameters							
$_1$	Normal	1.50	0.25	ρ_{r^m}	Beta	0.50	0.20
$_2$	Normal	0.12	0.05	σ_{r^m}	InvG	0.10	2.00
$_3$	Normal	0.12	0.05	σ_{ant1}	InvG	0.20	4.00
ρ_R	Beta	0.75	0.10				
Nominal Rigidities Parameters							
ζ_p	Beta	0.50	0.10	ζ_w	Beta	0.50	0.10
ι_p	Beta	0.50	0.15	ι_w	Beta	0.50	0.15
ϵ_p	-	10.00	fixed	ϵ_w	-	10.00	fixed
Other Endogenous Propagation and Steady State Parameters							
100γ	Normal	0.40	0.10	S''	Normal	4.00	1.50
α	Normal	0.30	0.05		Beta	0.50	0.15
$100(\beta^{-1}-1)$	Gamma	0.25	0.10	π_*	-	0.50	fixed
σ_c	Normal	1.50	0.37	γ_{gdpdef}	Normal	1.00	2.00
h	Beta	0.70	0.10	δ_{gdpdef}	Normal	0.00	2.00
ν_l	Normal	2.00	0.75	$\bar{L}$	Normal	-45.00	5.00
δ	-	0.03	fixed	λ_w	-	1.50	fixed
Φ_p	Normal	1.25	0.12	g_*	-	0.18	fixed
Financial Frictions Parameters							
$F(\bar{\omega})$	-	0.03	fixed	$\zeta_{sp,b}$	Beta	0.05	0.00
SP_*	Gamma	2.00	0.10	γ_*	-	0.99	fixed
Exogenous Process Parameters							
ρ_g	Beta	0.50	0.20	σ_g	InvG	0.10	2.00
ρ_b	Beta	0.50	0.20	σ_b	InvG	0.10	2.00
ρ_μ	Beta	0.50	0.20	σ_μ	InvG	0.10	2.00
$\rho_z til$	Beta	0.50	0.20	$\sigma_{\tilde{z}}$	InvG	0.10	2.00
ρ_{σ_ω}	Beta	0.75	0.15	σ_{σ_ω}	InvG	0.05	4.00
ρ_{π_*}	-	0.99	fixed	σ_{π_*}	InvG	0.03	6.00
ρ_{z^p}	Beta	0.50	0.20	σ_{z^p}	InvG	0.10	2.00
ρ_{λ_f}	Beta	0.50	0.20	σ_{λ_f}	InvG	0.10	2.00
ρ_{λ_w}	Beta	0.50	0.20	σ_{λ_w}	InvG	0.10	2.00
η_{λ_f}	Beta	0.50	0.20	η_{gz}	Beta	0.50	0.20
η_{λ_w}	Beta	0.50	0.20				
Measurement Error Parameters							
$\mathcal{C}_{me}$	-	1.00	fixed	ϱ_{gdp}	Normal	0.00	0.40
ρ_{gdp}	Normal	0.00	0.20	σ_{gdp}	InvG	0.10	2.00
ρ_{gdi}	Normal	0.00	0.20	σ_{gdi}	InvG	0.10	2.00
ρ_{10y}	Beta	0.50	0.20	σ_{10y}	InvG	0.75	2.00
ρ_{tfp}	Beta	0.50	0.20	σ_{tfp}	InvG	0.10	2.00
ρ_{gdpdef}	Beta	0.50	0.20	σ_{gdpdef}	InvG	0.10	2.00
ρ_{pce}	Beta	0.50	0.20	σ_{pce}	InvG	0.10	2.00

Note: For Inverse Gamma prior mean and SD, τ and ν reported.

TABLE 2. Prior Distribution

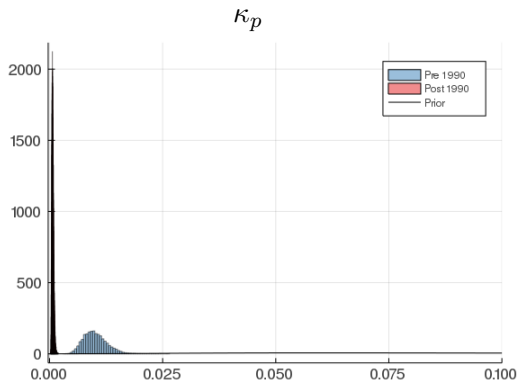

FIGURE C.4. Prior and posterior distributions for the slope of the price Phillips curve in the NY Fed DSGE model described in section 6 and appendix B.4. In this experiment, the model is estimated allowing for two distinct regimes for the slope of the price Phillips curve, pre- and post-1990, but assuming that all other coefficients are constant for the entire sample from 1964:II to 2019:III.

Acknowledgements

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.

Marco Del Negro
Federal Reserve Bank of New York, New York, United States; email: marco.delnegro@ny.frb.org

Michele Lenza
European Central Bank, Frankfurt am Main, Germany; ECARES; email: michele.lenza@ecb.europa.eu

Giorgio E. Primiceri
Northwestern University, Evanston, United States; CEPR; NBER; email: g-primiceri@northwestern.edu

Andrea Tambalotti
Federal Reserve Bank of New York, New York, United States; email: a.tambalotti@gmail.com

© European Central Bank, 2020

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the Social Science Research Network electronic library or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB's website.

PDF ISBN 978-92-899-4078-8 ISSN 1725-2806 doi:10.2866/5407 QB-AR-20-087-EN-N