

The Phillips Multiplier*

Regis Barnichon^(a) and *Geert Mesters*^(b)

^(a) Federal Reserve Bank of San Francisco and CEPR

^(b) Universitat Pompeu Fabra, Barcelona GSE

April 30, 2019

Abstract

We propose a methodology for determining the inflation-unemployment trade-off faced by a central bank, i.e., the ability of a central bank to transform unemployment into inflation (and vice versa) via its interest rate policy. We introduce the Phillips multiplier as a statistic to non-parametrically characterize the trade-off and its dynamic nature. Inference on the Phillips multiplier is based on a simple instrumental variable regression of cumulative inflation on cumulative unemployment using monetary shocks as instruments and weak instrument robust methods. We compute the Phillips multiplier for the US and the UK and document that the trade-off went from being very large in the pre-1990 sample period to being small (but significant) post-1990. In contrast to earlier evidence of a substantial flattening of the slope of Phillips curve, we find that the decline in the trade-off is mostly due to the anchoring of inflation expectations.

JEL classification: C14, C32, E32, E52.

Keywords: Inflation-Unemployment trade-off, Dynamic Multiplier, Instrumental variables, Phillips curve.

*First draft: February 2018. We thank Oscar Jorda, Sylvain Leduc, Nicolas Petrosky-Nadeau, Frank Schorfheide, Adam Shapiro, Jim Stock, Paolo Surico, Mark Watson, Dan Wilson and participants at the 2018 Annual SNDE Symposium, the 2018 IAAE Annual Meeting, the 2018 NBER Summer Institute and the CEPR “Low Inflation and Wage Dynamics: Implications for Monetary Policy and Financial Stability conference at Banca Italia for helpful comments. The views expressed in this paper are the sole responsibility of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System. Mesters acknowledge support from the Spanish Ministry of Science and Technology: Grant ECO2015-68136-P, Marie Skłodowska-Curie Actions: FP7-PEOPLE-2012-COFUND Action grant agreement no: 600387, the Spanish Ministry of Economy and Competitiveness through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2011-0075) and Fundación BBVA scientific research grant (PR16-DAT-0043) on Analysis of Big Data in Economics and Empirical Applications.

1 Introduction

When I say there is a trade-off between inflation and unemployment, I do not mean [that there is] a stable downward-sloping Phillips curve. [...] At its heart, the inflation-unemployment trade-off is [...] a claim that changes in policy push inflation and unemployment in opposite directions. Mankiw (2001)

The existence of an inflation-unemployment trade-off is at the core of monetary policy making, because central banks rely on this trade-off to “transform” unemployment into inflation (and vice-versa) through their interest rate policy. The ability of a central bank to control inflation thus depends on the magnitude of this trade-off, or more colloquially put, on the “unemployment cost” of lowering inflation.

The inflation-unemployment trade-off faced by policy makers is traditionally inferred from a Phillips curve linking inflation to real activity, or more generally from a multivariate structural model involving a Phillips curve. Unfortunately, such approach suffers from two empirical challenges:¹ (i) specification uncertainty, and (ii) endogeneity issues. First, there is uncertainty about the model’s relevant set of explanatory variables and about the appropriate dynamic lag structure (e.g., Gordon (2011)). Second, endogeneity issues are pervasive. For the Phillips curve, confounding from supply shocks, unobserved inflation expectations and measurement error in the natural unemployment rate all lead to highly uncertain and potentially biased coefficient estimates (e.g., Mavroeidis et al. (2014)).

In this paper, we avoid these empirical issues by directly characterizing the inflation-unemployment as the object of interest. We introduce a statistic –*the Phillips multiplier*– which is defined as the expected cumulative change in inflation caused by a monetary shock that lowers expected unemployment by 1ppt. To this extent, the Phillips multiplier directly captures the central bank’s inflation-unemployment trade-off across different horizons in line with the definition of Mankiw (2001). Moreover, we show that the Phillips multiplier

¹To be specific we imply the following two-step methods, which consist in (i) specifying and estimating a Phillips curve as in the limited-information approach (e.g., Mavroeidis, Plagborg-Møller and Stock (2014)) or a full system of structural equations as in a DSGE model (full information approach, e.g., Schorfheide (2011)), and (ii) inferring the policy-relevant trade-off from the estimated coefficients.

can be estimated by a simple instrumental variable regression where we regress cumulative inflation on cumulative unemployment using monetary shocks as instruments.² Importantly, we provide weak instrument robust confidence sets which are constructed by inverting a serial correlation corrected Anderson and Rubin (1949) statistic.

Compared to characterizations of the central bank’s inflation-unemployment trade-off based on a structural model, the Phillips multiplier offers a number of benefits. First, the Phillips multiplier is a non-parametric characterization of the trade-off: it captures how policy-induced changes in unemployment affect inflation over different time scales, from short time scales (less than a year), to medium term time scales (biennial), to longer time scales. As such, the Phillips multiplier does not rely on any specific structural model and is largely robust to model mis-specification.³ Second, thanks to the use of instruments, the Phillips multiplier avoids the identification issues that have plagued the Phillips curve literature; confounding factors and measurement error. Finally, by adopting weak instrument robust inference methods our findings do not depend on the strength of the instruments.

Equipped with our Phillips multiplier, we revisit important lessons from the Phillips curve literature in a *robust* and *identified* setting: (i) how large is the inflation-unemployment trade-off? (ii) what are the dynamics of the trade-off, i.e. how fast and for how long can the central bank affect inflation with a policy-induced change in unemployment? is the long-run trade-off infinite, as implied by a vertical long-run Phillips curve? (iii) has the trade-off changed over time? Is the trade-off smaller in the more recent period, as suggested by the flattening of the Phillips curve (e.g., Ball and Mazumder (2011), Coibion and Gorodnichenko (2015)), or is that flattening spurious and instead due to mis-specification, confounding from supply factors or to a mis-measured natural rate of unemployment? And if the trade-off is

²Alternative estimation methods such as LP-IV and SVAR-IV can also be considered, see Jordà (2005), Stock and Watson (2018), Mertens and Ravn (2013) and Olea, Stock and Watson (2018). However, these methods do require a two-step implementation where first the impulse responses to the monetary policy shocks are recovered and second the Phillips multiplier is computed as the cumulative ratio of the impulse responses. Such two-step approaches are less efficient and require more efforts when constructing confidence bounds for the Phillips multiplier. We provide a detailed discussion in Section 2.

³Note that similar as for LP-IV and SVAR-IV we do require specifying a set of control variables. In this sense our approach does require some modeling decisions, but we view these as mild compared to specifying a set of structural equations.

indeed smaller, is it because of better anchored inflation expectations or because of a lower sensitivity of inflation to economic slack?

We compute the US Phillips multiplier using as external instruments the Romer and Romer (2004) narratively identified monetary policy changes as well as the high frequency identified monetary policy shocks pioneered by Kuttner (2001) and Gürkaynak, Sack and Swanson (2005). Both identification schemes have their own advantages: the Romer and Romer (2004) shocks cover a long sampling period (1969-2007), while the high frequency shocks only cover the recent 1990-2008 period but are arguably more convincing in terms of the exogeneity assumption. For the UK Phillips multiplier we rely on the narratively identified monetary shocks of Cloyne and Hurtgen (2016).

For the full sample estimates the Phillips multiplier is very similar for the US and the UK. The multiplier is initially indeterminate due to the lagged reaction of inflation and unemployment to monetary policy, but the multiplier is significantly negative over the medium terms (2 years) and in the long run (4-5 years). In contrast, the unconditional Phillips multiplier (i.e., estimated by OLS) is close to zero across all horizons. This result highlights the importance of using instrument to address endogeneity issues and cautions against popular OLS-based estimates of the trade-off (e.g., Simon, Matheson and Sandri (2013)).

The full sample results mask the fact that the Phillips multiplier, and thus the inflation-unemployment trade-off, has changed considerably between the pre- and the post-1990 periods. Indeed, for the US we find that in the pre-1990 period, the Phillips multiplier strictly increases (in absolute value) with the horizon and diverges at longer horizons, pointing to an infinite trade-off. This is due to the inertial nature of inflation during this period: inflation reacts with a one year lag to changes in unemployment, and the response of inflation is much more persistent than that of unemployment. As a result, the Fed's ability to steer inflation with policy-induced unemployment movements is large.⁴

In the post-1990 period the situation is very different. While the Phillips multiplier is

⁴In the language of the Phillips curve literature, our non-parametric result corresponds to a close to vertical long-run Phillips curve, i.e., the absence of any long-run trade-off between inflation and unemployment: The Fed can only lower unemployment at the cost of an ever increasing rate of inflation.

still significantly different from zero in the medium run, it is much smaller and there is no indication of divergence in the long run. The implication is that the Fed’s inflation-unemployment trade-off is substantially smaller, or in other words, that the Fed’s ability to steer inflation with policy-induced unemployment movements has declined considerably.

This change in the trade-off in the post-1990 period could be due to the anchoring of inflation expectations or to a change in the sensitivity of inflation to economic slack (holding expectations constant). To separate these mechanisms, we show that the Phillips multiplier can be written as the ratio of two separate dynamic multipliers: (i) a multiplier capturing the trade-off holding inflation expectations constant, and (ii) a multiplier capturing the degree of anchoring of inflation expectations.

We find no compelling evidence of a decline in the sensitivity of inflation to economic slack (holding expectations constant). In contrast to earlier evidence of a flattening of the Phillips curve based on OLS estimates (e.g., Ball and Mazumder (2011), Blanchard (2016)), we find that the sensitivity of inflation to economic slack (holding expectations constant) is relatively stable between the two sample periods.⁵ Instead, the anchoring of inflation expectations is responsible for most (if not all) of the change in the trade-off. Inflation expectations went from being completely unanchored to almost perfectly anchored. In the pre-1990 period, inflation expectations responded one-for-one with inflation following monetary shocks, and this “zero-anchoring” implies a large long-run trade-off. Intuitively, any policy-induced transitory change in inflation feeds into inflation expectations, which then feeds into future inflation and so on, making the transitory change in inflation persistent or even permanent. In the post-1990 period however, inflation expectations respond little to policy-induced changes in inflation, and this implies a much lower trade-off. Intuitively, with well-anchored inflation expectations, inflation movements have no second-round effects through inflation expectations, and the overall inflation response to a change in policy is smaller and less persistent.

⁵Overall, our findings caution against popular OLS estimates of the trade-off or Phillips curve coefficients, since those estimates can be severely affected by confounding factors.

We see our approach to characterize the trade-off with the Phillips multiplier as paralleling the literature on the fiscal multiplier, see the review of Ramey (2016). The fiscal literature features two prominent ways to estimate the fiscal multiplier: (i) a model-based approach based on DSGE models (where the multiplier is indirectly inferred from the estimated parameters of the model) and (ii) an instrumental variable approach where the multiplier can be obtained directly from IV-type regressions involving the cumulative sum of output and government spending, see Ramey and Zubairy (2018) for a prominent example. While the Phillips curve literature has so far relied on (i) by taking as a starting point the existence of a well-specified Phillips curve relationship, our paper instead follows the non-parametric route pioneered by Ramey and Zubairy (2018) in the fiscal literature.

While we characterize the inflation-unemployment trade-off following the definition of Mankiw (2001), it is worth noting that other characterizations of the inflation-unemployment trade-off are also possible within our framework. For instance, the sacrifice ratio is typically defined, see Ball (1994), as the cumulative increase in unemployment from a one percentage point *permanent* reduction in inflation. This definition of the inflation-unemployment trade-off relies on the assumption that a change in policy has a permanent effect on inflation which might not hold uniformly across time, see Benati (2015).⁶ Importantly, our approach—using instruments to directly identify the inflation-unemployment trade-off—can be equally well adopted to identify the sacrifice ratio.

More broadly, in the context of the vast Phillips curve literature (e.g., Gordon (2011)), our effort to characterize the trade-off in an identified setting relates most closely to a few papers estimating the trade-off or the sacrifice ratio *conditional* on identified demand shocks (King and Watson (1994), Cecchetti and Rich (2001) and Benati (2015)). Different however from our approach, all these papers were based on structural VAR models, where demand shocks are identified from exclusion or sign restrictions. Empirically such approaches have led to large estimation uncertainty and strong dependence on the model specification and

⁶Indeed, the sacrifice ratio relies on the existence of a unit-root in inflation, which is most credible only in the pre-Volker sampling period (Benati (2015)).

the identification scheme, most likely because “the identifying restrictions are tenuous and generate weak or invalid instruments” (Cecchetti and Rich (2001)). In contrast, our proposed approach is robust to weak instruments. In spirit but not in substance, our approach shares some similarities with Ball (1994)’s non-parametric estimate of the sacrifice ratio. Like us, Ball (1994) did not rely on any Phillips curve model, and like our external instruments, his identification of dis-inflationary episodes was based on a narrative approach. Our method nests Ball (1994)’s approach in an instrumental regression setting.

The remainder of this paper is organized as follows. In Section 2 we define the central bank’s define the inflation-unemployment trade-off following Mankiw (2001) and introduce the Phillips multiplier as a statistic for characterizing the trade-off. We discuss the estimation of the Phillips multiplier and discuss its relation to existing methods for characterizing the trade-off between inflation and unemployment. Section 4 presents our full sample estimates of the Phillips multiplier for the US and the UK. Section 5 discussed the changes in the trade off for different sampling periods and Section 6 explains the changes by studying the anchoring of inflation expectations. Section 7 concludes.

2 The trade-off and the Phillips multiplier

In this section we introduce the Phillips multiplier as a statistic for characterizing the central bank’s trade-off between unemployment and inflation: the ability of a central bank to use its instrument (the policy rate) to transform unemployment into inflation (or vice-versa).

In line with Mankiw (2001)’s definition of the inflation-unemployment trade-off, we define the central bank’s average trade-off between inflation and unemployment as the average change in inflation caused by a change in policy that lowers the unemployment rate by 1ppt over the next h periods:

$$T_h = \frac{\partial \bar{\pi}_{t:t+h}}{\partial i_t} \Big|_{\varepsilon_t^i=1} / \frac{\partial \bar{u}_{t:t+h}}{\partial i_t} \Big|_{\varepsilon_t^i=1}, \quad h \geq 0, \quad (1)$$

where $\bar{y}_{t:t+h} = \frac{1}{h} \sum_{j=0}^h y_{t+j}$ denotes the average value of some variable y over $[t, t+h]$ and

$\left. \frac{\partial \bar{y}_{t:t+h}}{\partial i_t} \right|_{\varepsilon_t^i=1}$ denotes the marginal effect of an exogenous unit change ε_t^i in the central bank instrument i_t on some average variable $\bar{y}_{t:t+h}$.

Since the average responses of π and u are computed over increasingly larger windows $[t, t+h]$ as h increases, T_h captures the average trade-off between inflation and unemployment at different levels of time aggregation: from short time scales (less than a year, $h < 4$ using quarterly data), to medium term time scales (biennial, $h = 8$), to longer time scales (e.g., $h = 20$). In this way, T_h is a non-parametric characterization of the dynamic nature of the trade-off faced by a central bank.

In this work we measure the average trade-off directly using a statistic called the Phillips multiplier. We define the Phillips multiplier as

$$\mathcal{P}_h = \mathcal{R}_h^{\bar{\pi}} / \mathcal{R}_h^{\bar{u}}, \quad h = 0, 1, 2, \dots, \quad (2)$$

where $\mathcal{R}_j^{\bar{\pi}}$ and $\mathcal{R}_j^{\bar{u}}$ in (2) are the impulse responses of average inflation and unemployment to a one-unit policy shock ε_t^i . Specifically, $\mathcal{R}_h^{\bar{\pi}} = \frac{1}{h} \sum_{j=0}^h \mathcal{R}_j^{\pi}$ and $\mathcal{R}_h^{\bar{u}} = \frac{1}{h} \sum_{j=0}^h \mathcal{R}_j^u$, with \mathcal{R}_j^y the causal impulse response at horizon j of a variable y to a one-unit shock ε_t^i :

$$\mathcal{R}_h^y = E(y_{t+h} | \varepsilon_t^i = 1, w_t) - E(y_{t+h} | \varepsilon_t^i = 0, w_t), \quad (3)$$

where w_t is a vector of control variable.⁷

The Phillips multiplier \mathcal{P}_h is the natural statistical counterpart to our definition of the average trade-off: it measures how the forecast for inflation changes when a change in policy lowers unemployment by 1ppt over the next h periods. The Phillips multiplier parallels the concept of the government spending multiplier in the fiscal literature, see Ramey and Zubairy (2018).⁸

A key advantage of the Phillips multiplier is that it can be estimated by an IV regres-

⁷While the impulse response of the average value of a variable is little unusual in the macroeconomic literature, this is simply the cumulative impulse response of the variable discounted by the horizon.

⁸The government spending multiplier is often measured as the ratio of the cumulative impulse responses of output and government spending, Ramey and Zubairy (2018).

sion of cumulative inflation on cumulative unemployment using monetary policy shocks as instruments. Importantly, one does not need to compute the impulse responses of inflation and unemployment to a monetary shock.

Denote by ξ_t an instrumental variable for the policy shock ε_t^i . The following simple proposition summarizes the result.

Proposition 1. *The multiplier \mathcal{P}_h can be estimated from the cumulative regression*

$$\sum_{j=0}^h \pi_{t+j} = \mathcal{P}_h \sum_{j=0}^h u_{t+j} + w_t' \gamma_h + e_{t+h} \quad (4)$$

where $\sum_{j=0}^h u_{t+j}$ is instrumented by ξ_t and where w_t is a vector of control variables.

Proof. See appendix A. □

The empirical characterization of the multiplier, implied by (4), leads to a simple and robust way for conducting inference on the multiplier. We adopt weak instrument robust methods for constructing confidence sets for \mathcal{P}_h , see Andrews, Stock and Sun (2019). In particular, since the multiplier is computed using a single instrument ξ_t we rely on the Anderson and Rubin (1949) (AR) statistic for inference. Note that the model is just identified which implies that in the case of strong instruments the AR statistic does not sacrifice power when compared to the standard t -test based on the normal limiting distribution of the 2SLS estimator, e.g. Andrews et al. (2019). More importantly, inference based on the AR statistic does not depend on the strength of the instruments, which is of crucial importance as recent contributions have shown that monetary shocks do not necessarily explain a large portion of the variance in macroeconomic variables, see Gorodnichenko and Lee (2017) and Plagborg-Møller and Wolf (2018). We provide a detailed account for the implementation of the AR statistic in Appendix B.

We emphasize that alternative estimation methods for the Phillips multiplier can also be considered. For instance, it is possible to use LP-IV or SVAR-IV methods, as discussed in Jordà (2005), Mertens and Ravn (2013), Stock and Watson (2018) and Olea et al. (2018),

which would first compute the impulse responses of inflation and unemployment to monetary policy shocks and subsequently the Phillips multiplier according to equation (2). There are two difficulties with such approach. First, such two-step method is less efficient when compared to the one-step approach implied by proposition 1. Second, the construction of confidence bounds, based on either standard or weak instrument robust methods, is more cumbersome for the two-step method.

3 Comparison with a Phillips curve implied trade off

In this section we compare the characterization of the inflation-unemployment trade-off using the Phillips multiplier with the more traditional approach. The latter consists of first estimating a Phillips curve which may or may not be embedded in a larger structural model, and then using the estimated coefficients to back out the policy-relevant trade-off. When compared to this approach, the Phillips multiplier offers a number of benefits.

First, by being defined directly from the impulse response functions of inflation and unemployment to a monetary shock, the Phillips multiplier avoids the endogeneity issues that have plagued the estimation of the Phillips curve, see Mavroeidis et al. (2014). In particular, since monetary shocks are orthogonal to supply shocks and to the natural unemployment rate u_t^* ,⁹ the Phillips multiplier will allow us to estimate the trade-off (i) without bias from confounding supply shocks and (ii) without any need for a measure of the natural unemployment rate (because $\mathcal{R}_h^{u-u^*} = \mathcal{R}_h^u$). Note that these confounding factors have led to large estimation uncertainty in the Phillips curve literature, see Mavroeidis et al. (2014).

Second, unlike the Phillips curve, the Phillips multiplier characterizes the trade-off in a *non-parametric* fashion. This provides robustness to mis-specification, which is particularly attractive given the large uncertainty regarding the set of relevant explanatory variables to include in the Phillips curve as well as the appropriate dynamic lag specification. See Gordon (2011), King and Watson (1994) and the discussion in the appendix.

⁹Under the common assumption that monetary policy is neutral under flexible prices (e.g., Galí (2015)).

More generally, the coefficients of the Phillips curve alone may not be enough to characterize the inflation-unemployment trade-off faced by the central bank. Indeed, since the central bank influences real activity through its interest rate policy, the specification of the (IS) curve may also matter. In fact while the trade-off reduces to the slope of the Phillips curve in the basic New-Keynesian model (Galí (2015)), this is not a general result, and the trade-off typically depends on both the Phillips *and* the (IS) curve coefficients.¹⁰ We provide such an example in Appendix C. Since the (IS) curve suffers from similar specification and endogeneity issues as the Phillips curve (Fuhrer and Rudebusch (2004)), the estimation issues get compounded. The non-parametric and identified nature of the Phillips multiplier, a statistic specifically designed to characterize the trade-off, allows us to side-step these issues.

4 Empirical estimates of the Phillips multiplier

In this section we compute the Phillips multiplier for the US and the UK. Since our direct estimation of the trade-off relies on the existence of external instruments for exogenous changes in policy, the sample period we consider is dictated by the availability of relevant instruments.

Our baseline estimates for the US are based on the Romer and Romer (2004) narrative measure of exogenous monetary policy changes. As an alternative (considered in the next section), we will also rely on the recent high-frequency identification (HFI) approach pioneered by Kuttner (2001) and use surprises in fed funds futures prices around FOMC announcement as instruments for monetary shocks.¹¹ Inflation is measured as the (annualized) quarter-to-quarter change in the PCE price level. The sample covers 1969-2007 thanks to Tenreyro and Thwaites (2016)'s extension of the Romer and Romer series.

For the UK we rely on the Cloyne and Hurtgen (2016) narratively identified measures

¹⁰Intuitively, the central bank's trade-off between inflation and unemployment depends on the relative speed with which monetary policy affects unemployment versus inflation, and this relative speed depends on the dynamic specifications of *both* the Phillips curve *and* the (IS) curve.

¹¹See Gertler and Karadi (2015), Miranda-Agrippino (2016), Nakamura and Steinsson (2018) for more recent examples of the use of HFI instruments to identify the effects of monetary policy.

of exogenous monetary policy changes. Following Cloyne and Hurtgen (2016) inflation is measured as the (annualized) quarter-to-quarter change in the retail prices index (excluding mortgage interest payments). The sample period covers 1975-2007.

For all specifications throughout this paper, we use quarterly data and include four lags of inflation and unemployment as control variables w_t in regression (4). Finally, as implied by equation (3) we conduct all our analyses under a stationarity assumption.¹²

4.1 US Results

Figure 1 shows our baseline estimate for the US Phillips multiplier for horizon $h = 0$ until $h = 20$ and also reports the Olea and Pflueger (2013) F -stats from the first stage of the instrumental variable regression (4). The F -statistics document to what extent the monetary policy shocks are correlated with cumulative unemployment at horizon h . The F -stats statistics are reasonable (around 10) for h between 8 and 14 quarters, but they do not exceed the serial correlation robust threshold of Olea and Pflueger (2013) which requires F -statistics above 23. Therefore we rely on weak instrument robust methods for computing the confidence sets for the Phillips multiplier. In particular, in Figure 1 the shaded area shows the 95% confidence region obtained from inverting the serial correlation robust Anderson and Rubin (1949) statistic, see Appendix B for details. Importantly, these bands do not depend on the strength of the instruments. The thick line shows the 2SLS point estimate which coincides nearly perfectly with the minimum AR statistic. For comparison purposes we also show the standard 95% confidence bands obtained from the possibly invalid normal limiting distribution of the 2SLS estimator. We note that these bands are symmetric by construction but they do not uniformly dominate the AR-based confidence sets.¹³

The Phillips multiplier is initially indeterminate, but as time goes by \mathcal{P}_h becomes negative, reaches about -2ppt after $h = 12$ quarters and diverges thereafter. Note that a large (or infinite) long-run trade-off implies that a transitory policy-induced change in unemployment

¹²We note that the main ideas of our methodology can be applied when a unit root is present.

¹³Notably, the upper bound of the confidence region is over estimated by the standard bounds, which is due to the imposed symmetry.

has a very persistent (or permanent) effect on inflation. With a large trade-off, the Fed’s ability to steer inflation with policy-induced unemployment movements is large. Equivalently, a large trade-off implies that the Fed could lower unemployment permanently only at the cost of ever increasing inflation.¹⁴

To highlight the importance of proper identification, i.e. conditioning on the monetary shocks, we show in Figure 2 the Phillips multiplier (top panel) together with its unconditional counterpart (bottom panel). The unconditional multiplier is obtained by the OLS regression of average inflation on average unemployment, i.e. computing \mathcal{P}_h using OLS instead of IV. The difference in the multipliers is striking, as the identified Phillips multiplier substantially larger (in absolute value) than suggested by the OLS estimate. This result is in line with what one would expect from confounding from supply shocks (since supply shocks lead to a positive correlation between inflation and unemployment) and measurement error in the unemployment gap (leading to an attenuation bias). A popular approach (especially in policy institutions) consists in inferring the inflation–unemployment trade-off from Phillips curve regressions estimated by OLS (e.g., Simon et al. (2013)). Our result highlights the importance of properly identifying the trade-off and cautions against trade-off estimates based on OLS regressions.

To better understand the dynamic of the multiplier, we can separately study its individual components, namely the impulse responses of average inflation and unemployment. Indeed, the numerator $\mathcal{R}_h^{\bar{u}}$ and denominator $\mathcal{R}_h^{\bar{\pi}}$ of \mathcal{P}_h are the impulse responses of the average levels of inflation and unemployment over the next h period following a monetary policy shock. Using local projections with monetary shocks (Jordà (2005), Stock and Watson (2018)), we

¹⁴An infinite long-run trade-off is reminiscent of the accelerationist Phillips curve of the 1970s where unemployment was related to *changes* in inflation. We will come back to this point in the next section, where we study the multiplier over different sample periods.

can estimate the impulse responses of average inflation and average unemployment from

$$\bar{\pi}_{t:t+h} = \xi_t \beta_h^\pi + w_t' \gamma_h^\pi + e_{h,t+h}^\pi \tag{5}$$

$$\bar{u}_{t:t+h} = \xi_t \beta_h^u + w_t' \gamma_h^u + e_{h,t+h}^u$$

where ξ_t is the Romer and Romer monetary shock, and the parameters β_h^π and β_h^u capture the effect of a one unit change in the monetary shocks on average inflation and unemployment. Since, the multiplier is the ratio of the impulse responses we normalize the impulse responses such that the monetary shock leads to an unemployment response that peaks at 1ppt.¹⁵

As shown in the right-column of Figure 1, the multiplier is initially indeterminate because both unemployment and inflation respond with a lag to monetary policy, the well known transmission lags of monetary policy (e.g., Svensson (1997)).¹⁶ Then, the multiplier builds up (in absolute value) over time, because the response of inflation is delayed relative to that of unemployment and because the response of inflation is more persistent than that of unemployment: as average unemployment starts to mean-revert to its unconditional mean at $h \simeq 14$ (which ceteris paribus increases $\mathcal{P}_h = \mathcal{R}_h^\pi / \mathcal{R}_h^u$ in absolute value), inflation remains elevated and shows no sign of mean reversion. As a result, the multiplier keeps increasing and diverges. The central bank can trigger a persistent (possibly permanent) change in inflation at a finite unemployment cost, i.e., the trade-off is large.

4.2 UK Results

Figure 3 shows the Phillips multiplier for the UK. The F -statistics are smaller, notably in the first 10 quarters, indicating that identification is poor for these horizons. As discussed above, this explains the erratic behavior of the multiplier in the these first periods and stresses the

¹⁵Since we only use one instrument, the ratio of the β 's corresponds with the Phillips multiplier, see Appendix A for details. From this it is also easy to see that rescaling both β^π and β^u bears no influence on the results and is only for explanatory purposes.

¹⁶The fact that the initial responses of $\bar{\pi}_{t:t+h}$ and $\bar{u}_{t:t+h}$ are close to zero implies that the multiplier can behave erratic during the first quarters and the confidence bands are typically large or even infinite as we shall see below.

importance of using weak instrument robust confidence bounds.

Overall, the UK and US Phillips multipliers are quite similar in their dynamics. The UK multiplier is initially indeterminate due to the lagged response of inflation and unemployment. After $H = 10$ the multiplier becomes negative and decreases as the horizon increases. We do find modest differences between the US and the UK in the levels of the multiplier at large horizons, but overall a clear negative trade-off between inflation and unemployment exists 3-4 years after a policy shock.

5 The US Phillips multiplier over time

Our results based on the full sample mix very different policy regimes that blur the interpretability of our results. In fact, a number of Phillips curve-based studies, mainly for the US, have suggested substantial changes in the persistence of inflation as well as in the magnitude of the inflation-unemployment trade-off; from the close to unit-root behavior of inflation in the 1970s (e.g., King and Watson (1994)) to the flattening of the Phillips curve in the post-1990 period (e.g., Ball and Mazumder (2011) and Blanchard (2016)).

In this section, we parallel the Phillips curve literature and study the evolution of the US Phillips multiplier over time, using our non-parametric and identified approach to reassess the evolution of the trade-off over time.¹⁷ Importantly, unlike these earlier Phillips-curve studies typically based on OLS regressions, our approach will allow us to discard confounding effects that could give the illusion of a change in the trade-off, notably a change in the variance contribution of supply shocks (which varies the magnitude of the endogeneity bias), or mis-measured changes in the natural rate of unemployment.¹⁸

To investigate a possible change in the Phillips multiplier since 1990, we focus on the US and draw on HFI monetary surprises.¹⁹ Specifically, we follow Gertler and Karadi (2015) and

¹⁷For the UK we found that the instruments are unfortunately not informative enough to speak to possible changes over time.

¹⁸For instance, if the natural rate of unemployment was over-estimated during the late 1990s, the unemployment gap would have been under-estimated (not small enough), leading to a downward-biased estimate of the slope of the Phillips curve.

¹⁹Unfortunately, the Romer and Romer narrative instruments cannot be used to study the trade-off in

consider changes in Federal Funds futures rates (FF) around FOMC announcement dates as external instruments.²⁰ This measure is plausibly uncorrelated with other shocks because they are changes across a short announcement window. To estimate the Phillips multiplier using HFI instruments, we proceed as with the Romer-Romer shocks except that we add as control variables the Greenbook forecasts for inflation and unemployment. This is to capture any changes in the Federal Funds futures rates driven by the part of information set of the Fed that differs from that of the market.

Since HFI monetary surprises are only available in the more recent period (1990-2008),²¹ we will report two sets of results to assess the evolution of the Phillips multiplier over time: first, the Phillips multiplier estimated using the Romer-Romer narrative shocks over 1969-1989, and second the Phillips multiplier estimated using the HFI monetary surprises over 1990-2008.

The results are shown in Figure 4. While the estimated multipliers are similarly indeterminate at short horizons, the top panels of Figure 4 show that their behaviors differ markedly at longer horizons. The pre-1990 multiplier clearly diverges with the horizon, while the post-1990 multiplier is roughly stable at some small (but non-zero) long-run value (see the bottom panel of Figure 4). In other words, in the pre-1990 period, the trade-off is large: monetary policy can have a very persistent, perhaps permanent, effect on inflation at a finite unemployment cost. In contrast, in the post-1990 period the trade-off is small and does not diverge. Consequently, the multiplier is substantially smaller in the post-1990 period with a multiplier at about -0.15 after 3-4 years, an order of magnitude smaller than in the pre-1990 period.²²

the more recent period (such as post-1990), because they have very low relevance (low first-stage F-stats) leading to infinite AR-based confidence sets and thus cannot be used to estimate the Phillips multiplier. See Ramey (2016) for a forceful account of the difficulty of using the Romer-Romer shocks to learn about the effects of monetary policy in the post-1990 period.

²⁰As Gertler and Karadi (2015), we found that “FF4”, the three month ahead monthly Fed Funds futures, provided the best first-stage, and we will present results based on these instruments.

²¹We intentionally exclude the post-2008 period (the zero lower-bound period) during which forward guidance played a more prominent role and the Fed started employing unconventional monetary tools that could have different effects than conventional monetary tools, see Swanson (2017).

²²Interestingly, notice that the multiplier is much better identified at longer horizons for the post-1990 sampling period. The weak instrument robust confidence sets are narrow indicating the informativeness of

To help understand this flattening of the multiplier, Figure 5 plots the impulse responses of average inflation and unemployment in the two sample periods; pre-1990 and post-1990. The reasons for the change in the multiplier are twofolds: for a given change in unemployment, the response of inflation (i) is a lot more muted (one order of magnitude smaller) in the more recent period, and (ii) no longer displays inertia relative to unemployment and in fact appears to mirror the impulse response of unemployment. In the next section, we will explore to what extent the anchoring of inflation expectations can be behind this change in the behavior of inflation and the trade-off.

6 The trade-off and anchoring of inflation expectations

Two mechanisms could explain a change in the trade-off in the post-1990 period: (i) a change in the sensitivity of inflation to economic slack (holding expectations constant), and/or (ii) a change in the anchoring of inflation expectations.

In this section, we separate these two channels by decomposing the Phillips multiplier into the ratio of two separate dynamic multipliers: (i) a multiplier capturing the trade-off holding inflation expectations constant, and (ii) a multiplier capturing the degree of anchoring of inflation expectation.

6.1 Decomposing the Phillips multiplier

With some simple algebra, we can re-write the Phillips multiplier as

$$\mathcal{P}_h = \frac{\mathcal{K}_h}{1 - \mathcal{A}_h}, \quad h = 0, 1, 2, \dots \quad (6)$$

the high frequency instruments.

where \mathcal{K}_h is the “expectation-augmented Phillips multiplier” and \mathcal{A}_h is the “anchoring multiplier” which are defined as follows

$$\begin{aligned}\mathcal{K}_h &= \mathcal{R}_h^{\bar{\pi}-\bar{\pi}^e} / \mathcal{R}_h^{\bar{u}} \\ & \quad h = 0, 1, 2, \dots, \\ \mathcal{A}_h &= \mathcal{R}_h^{\bar{\pi}^e} / \mathcal{R}_h^{\bar{\pi}}\end{aligned}\tag{7}$$

where $\mathcal{R}_h^{\bar{\pi}-\bar{\pi}^e}$ is the average impulse response of $\pi_{t+h} - \pi_{t+h+1|t}^e$ to a monetary shock over h periods, with $\pi_{t+1|t}^e$ denoting expected future inflation as of time t , and $\mathcal{R}_h^{\bar{\pi}^e}$ is the average impulse response of expected inflation.

The multiplier \mathcal{K}_h is the dynamic non-parametric analog of κ , the slope of the expectation-augmented Phillips curve (e.g., Coibion and Gorodnichenko (2015))

$$\pi_t - \pi_{t+1|t}^e = -\kappa(u_t - u_t^*) + e_t.\tag{8}$$

\mathcal{A}_h is a dynamic measure of the degree of anchoring of inflation $\left. \frac{\partial \pi_{t+1|t}^e}{\partial i_t} \right|_{\varepsilon_t^i=1} / \left. \frac{\partial \pi_t}{\partial i_t} \right|_{\varepsilon_t^i=1}$, that is it captures how policy-induced inflation movements pass through to inflation expectations.²³ Full anchoring of inflation expectations corresponds to $\mathcal{A}_h = 0$, while full pass-through of inflation to inflation expectations implies $\mathcal{A}_h = 1$. Note how the \mathcal{A}_h parameter shares close similarities with the “ α ” parameter in the Phillips curve literature and Solow-Tobin tests of the natural rate hypothesis, e.g., Gordon, Solow, Perry and Gordon (1970). In earlier Phillips curve regressions, α was traditionally the loading on past inflation meant to proxy for expected inflation ($\pi_{t+1|t}^e = \alpha\pi_{t-1}$). In that context, a value of α close to 1 implied the existence of a (close to) unit-root in inflation and no long-run trade-off between inflation and unemployment.

Expression (6) makes clear that two factors can lead to a decline in the Phillips multiplier:

- (i) a decline in the sensitivity of inflation to economic slack holding inflation expectations constant (a decrease in \mathcal{K}_h), and
- (ii) an increase in the anchoring of inflation expectations

²³By focusing on the elasticity of inflation expectations to a shock, this measure is similar in spirit to other measures found in the literature, see e.g., Gürkaynak, Levin and Swanson (2010).

(an increase in \mathcal{A}_h).

As $h \rightarrow \infty$, \mathcal{K}_∞ can be seen as capturing the cumulative direct effect of economic slack on inflation, while \mathcal{A}_∞ captures the cumulative second-round effects arising from the adjustment of inflation expectations to inflation movements that then feed into inflation and in turn affect inflation expectations further, etc... Given our interest in the decline in the long-run value of the Fed's trade-off, we will let $h \rightarrow \infty$ to discuss two cases:

(i) With no anchoring of inflation expectations ($\mathcal{A}_\infty = 1$), the trade-off diverges as

$\mathcal{P}_\infty = \frac{\mathcal{K}_\infty}{1-\mathcal{A}_\infty} \xrightarrow{\mathcal{A}_\infty \rightarrow 1} \infty$. Intuitively, any policy-induced transitory change in inflation will feed into inflation expectations, which will then feed into future inflation; making the transitory change in inflation permanent (in the limit where $\mathcal{A}_\infty \rightarrow 1$).²⁴ The central bank is able to engineer large (i.e., permanent) movements in inflation with small (i.e., transitory) movements in unemployment: the trade-off is large as the central bank can only lower unemployment permanently at the cost of ever increasing inflation.

(ii) With some anchoring of inflation expectations ($\mathcal{A}_\infty < 1$), there is incomplete pass-through between inflation and inflation expectations, which implies that the Phillips multiplier remains finite with $\mathcal{P}_\infty \geq \mathcal{K}_\infty$.²⁵ Intuitively, if after a policy change, inflation expectations do not fully adjust to the overall change in inflation (incomplete pass-through), movements in inflation have limited second-round effects and the cumulative change in inflation is finite, i.e., policy-induced unemployment movements will generate only transitory inflation movements. The stronger the anchoring of inflation expectations (the smaller \mathcal{A}_∞), the smaller (in absolute value) the long-run Phillips multiplier, i.e., the smaller the long-run trade-off. In the limit case where $\mathcal{A}_h = 0$, we have $\mathcal{P}_h = \mathcal{K}_h$.

²⁴In the language of the Phillips curve, the slope of the Phillips curve depends on inflation expectations, so that while the Phillips curve may be downward slopping in the short-run because current inflation deviate from inflation expectations, in the long-run, the slope of the Phillips curve tends to infinity if inflation expectations full adjust to movements in inflation.

²⁵An imperfect pass-through is clearly at odds with the full-information rational expectation. See Coibion, Gorodnichenko and Kamdar (2017) for consistent supporting evidence of departures from full-information rational expectations, and Akerlof, Dickens and Perry (2000) for some implications of near-rational expectations on the long-run slope of the Phillips curve.

6.2 Anchoring versus Flattening

To assess whether the large decline in \mathcal{P}_h after 1990 is due to a change in the sensitivity of inflation to economic slack (holding expectations constant), and/or a change in the anchoring of inflation expectations, we estimate the dynamic multipliers \mathcal{K}_h pre- and post-1990. If we do not detect any change in \mathcal{K}_h over time, this means that the anchoring of inflation expectations accounts for the decline in the Phillips multiplier.

To measure inflation expectations, we rely on the inflation forecasts of households from the University of Michigan Survey of Consumers, and we use households' forecast of price changes over the next 12 months, consistent with Coibion and Gorodnichenko (2015)'s finding that household forecasts appear to be a more relevant measure of inflation forecasts for the Phillips curve than professional forecasts.²⁶

Figure 6 plots the multipliers \mathcal{P}_h and \mathcal{K}_h , estimated for the pre-1990 and post-1990 samples using the Romer-Romer and HFI identification strategies.

The most striking result is the difference between \mathcal{P}_h^{pre90} and \mathcal{K}_h^{pre90} . Once we control for inflation expectations, the multiplier between inflation and unemployment drops by an order of magnitude, from about -3.0 to -0.3 at $H = 20$ (left column). In contrast, controlling for inflation expectations has little effect on the post-1990 estimates. Although uncertainty is substantial, we find $\mathcal{P}_{20}^{post90} \simeq \mathcal{K}_{20}^{post90} \simeq -.15$.

Interestingly, \mathcal{K}_h , the dynamic multiplier between inflation and unemployment holding inflation expectations constant is pretty stable pre- and post-1990 (at least within sampling uncertainty), with $\mathcal{K}_{20}^{pre90} \simeq -.30$ and $\mathcal{K}_{20}^{post90} \simeq -.15$. This absence of compelling evidence of a decline in \mathcal{K}_h stands in contrast with Phillips curve OLS estimates that point to a severe flattening of κ , the slope of the Phillips curve. Again, the presence of confounding factors could affect these earlier results, and our approach allows us to side-step these endogeneity issues.

The difference between the Phillips multiplier and the expectation-augmented multiplier

²⁶We obtain similar results with median one-year ahead forecasts from the Survey of Professional Forecasters or with longer-term 10-year ahead inflation expectations as estimated by the Board of Governors.

can be understood more clearly by studying the underlying impulse responses. In Figure 7 we show the impulse responses of inflation and expected inflation to the monetary shocks pre- and post-1990. Before 1990, the impulse responses of inflation expectations tracks that of inflation expectations, and as a result $\mathcal{A}_{20}^{pre90} \simeq 1$, which implies that $\mathcal{P}_{20} = \frac{\mathcal{K}_{20}}{1-\mathcal{A}_{20}}$ is large and diverging.²⁷ In contrast, after 1990 there is no significant response of inflation expectations to a monetary shock and $\mathcal{A}_{20}^{post90} \simeq 0$, which implies a small and finite trade-off with $\mathcal{P}_{20} \simeq \mathcal{K}_{20} \simeq -.15$. We conclude that the anchoring of inflation expectations is responsible for the decline in the Phillips multiplier over time.

7 Conclusion

The inflation-unemployment trade-off faced by policy makers is traditionally inferred from the coefficients of an estimated Phillips curve. However, such approach is fraught with specification and endogeneity issues. In this paper we propose an instrumental variable based approach to directly characterize the central bank’s trade-off with the *Phillips multiplier*; the cumulative change in inflation caused by a policy shock that raises unemployment by one percentage point.

Using instruments for the US and the UK, we revisit the main lessons of the Phillips curve literature in a robust and identified setting. We find that (i) over short time scales (less than a year) the trade-off is indeterminate because of transmission lags in policy, (ii) over medium time scales (biennial) the trade-off is significantly negative, and (iii) for the US the trade-off went from being very large in the pre-1990 sample period to being small (but still significant) in the post-1990 period, i.e., during the onset of inflation targeting. Using inflation expectation data for the US, we find that most of the change in the US trade-off owes to the anchoring of inflation expectations. In contrast, the overall sensitivity of inflation to economic slack (holding inflation expectations constant) is not markedly lower between the two periods.

²⁷In the language of the Phillips curve, the long-run slope of the Phillips curve is vertical and there is no long-run trade-off between inflation and unemployment

References

- Akerlof, G. A., Dickens, W. T. and Perry, G. L.: 2000, Near-Rational Wage and Price Setting and the Long-Run Phillips Curve, *Brookings Papers on Economic Activity* **31**(1), 1–60.
- Anderson, T. W. and Rubin, H.: 1949, Estimation of the parameters of a single equation in a complete system of stochastic equations, *Ann. Math. Statist.* **20**(1), 46–63.
- Andrews, I., Stock, J. and Sun, L.: 2019, Weak instruments in iv regression: Theory and practice, *Annual Review of Economics* . forthcoming.
- Ball, L.: 1994, What determines the sacrifice ratio?, *Monetary policy*, The University of Chicago Press, pp. 155–193.
- Ball, L. and Mazumder, S.: 2011, Inflation dynamics and the great recession, *Brookings Papers on Economic Activity* pp. 337–405.
- Barth, M. and Ramey, V.: 2001, The cost channel of monetary transmission, *NBER macroeconomics annual* **16**, 199–240.
- Benati, L.: 2015, The long-run Phillips curve: A structural VAR investigation, *Journal of Monetary Economics* **76**, 15–28.
- Blanchard, O.: 2016, The phillips curve: Back to the '60s?, *American Economic Review* **106**(5), 31–34.
- Cecchetti, S. G. and Rich, R. W.: 2001, Structural estimates of the u.s. sacrifice ratio, *Journal of Business & Economic Statistics* **19**(4), 416–427.
- Cloyne, J. and Hurtgen, P.: 2016, The macroeconomic effects of monetary policy: A new measure for the united kingdom, *American Economic Journal: Macroeconomics* **8**(4), 75–102.
- Coibion, O. and Gorodnichenko, Y.: 2015, Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation, *American Economic Journal: Macroeconomics* **7**(1), 197–232.
- Coibion, O., Gorodnichenko, Y. and Kamdar, R.: 2017, The Formation of Expectations, Inflation and the Phillips Curve, *NBER Working Papers 23304*, National Bureau of Economic Research, Inc.
- Fuhrer, J. C. and Rudebusch, G. D.: 2004, Estimating the Euler equation for output, *Journal of Monetary Economics* **51**(6), 1133–1153.
- Gali, J.: 2011, The Return Of The Wage Phillips Curve, *Journal of the European Economic Association* **9**, 436–461.
- Galí, J.: 2015, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press.

- Gali, J. and Gertler, M.: 1999, Inflation dynamics: A structural econometric analysis, *Journal of Monetary Economics* **44**, 195–222.
- Gertler, M. and Karadi, P.: 2015, Monetary Policy Surprises, Credit Costs, and Economic Activity, *A EJ: Macroeconomics* **7**, 44–76.
- Gordon, R. J.: 2011, The history of the phillips curve: Consensus and bifurcation, *Economica* **78**, 10–50.
- Gordon, R. J., Solow, R., Perry, G. and Gordon, R.: 1970, The recent acceleration of inflation and its lessons for the future, *Brookings Papers on Economic Activity* **1970**(1), 8–47.
- Gorodnichenko, Y. and Lee, B.: 2017, A Note on Variance Decomposition with Local Projections. Working paper.
- Gürkaynak, R. S., Levin, A. and Swanson, E.: 2010, Does inflation targeting anchor long-run inflation expectations? evidence from the us, uk and sweden, *Journal of the European Economic Association* **8**(6), 1208–1242.
- Gürkaynak, Refet, S., Sack, B. and Swanson, E.: 2005, The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models, *American Economic Review* **95**, 425–436.
- Jordà, O.: 2005, Estimation and Inference of Impulse Responses by Local Projections, *The American Economic Review* **95**, 161–182.
- King, R. G. and Watson, M. W.: 1994, The post-war U.S. phillips curve: a revisionist econometric history, *Carnegie-Rochester Conference Series on Public Policy* **41**, 157–219.
- Kuttner, K. N.: 2001, Monetary policy surprises and interest rates: Evidence from the Fed funds futures market, *Journal of Monetary Economics* **47**(3), 523–544.
- Mankiw, N. G.: 2001, The Inexorable and Mysterious Tradeoff between Inflation and Unemployment, *Economic Journal* **111**(471), 45–61.
- Mavroeidis, S., Plagborg-Møller, M. and Stock, J. H.: 2014, Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve, *Journal of Economic Literature* **52**, 124–188.
- Mertens, K. and Ravn, M. O.: 2013, The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States, *American Economic Review* **103**(4), 1212–1247.
- Miranda-Agrippino, S.: 2016, Unsurprising shocks: information, premia, and the monetary transmission, *Bank of England working papers 626*, Bank of England.
- Nakamura, E. and Steinsson, J.: 2018, High frequency identification of monetary non-neutrality: The information effect, *Quarterly Journal of Economics* .

- Newey, W. K. and West, K. D.: 1994, Automatic lag selection in covariance matrix estimation, *The Review of Economic Studies* **61**(4), 631–653.
- Olea, J. L. M. and Pflueger, C.: 2013, A robust test for weak instruments, *Journal of Business & Economic Statistics* **31**(3), 358–369.
- Olea, J. L., Stock, J. H. and Watson, M. W.: 2018, Inference in structural vector autoregressions identified with an external instrument, *Technical report*. Working paper.
- Plagborg-Møller, M. and Wolf, C. K.: 2018, Instrumental variable identification of dynamic variance decompositions. working paper.
- Ramey, V.: 2016, Macroeconomic Shocks and Their Propagation, in J. B. Taylor and H. Uhlig (eds), *Handbook of Macroeconomics*, Elsevier, Amsterdam, North Holland.
- Ramey, V. A. and Zubairy, S.: 2018, Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data, *Journal of Political Economy* **126**.
- Romer, C. D. and Romer, D. H.: 2004, A new measure of monetary shocks: Derivation and implications, *American Economic Review* **94**, 1055–1084.
- Schorfheide, F.: 2011, Estimation and evaluation of dsge models: progress and challenges, *Technical report*, National Bureau of Economic Research.
- Simon, J., Matheson, T. and Sandri, D.: 2013, The dog that didnt bark: Has inflation been muzzled or was it just sleeping?, *World Economic Outlook* pp. 79–95.
- Stock, J. H. and Watson, M. W.: 2018, Identification and estimation of dynamic causal effects in macroeconomics using external instruments, *The Economic Journal* **128**(610), 917–948.
- Svensson, L. E. O.: 1997, Inflation forecast targeting: Implementing and monitoring inflation targets, *European Economic Review* **41**(6), 1111–1146.
- Swanson, E. T.: 2017, Measuring the effects of federal reserve forward guidance and asset purchases on financial markets, *Working Paper 23311*, National Bureau of Economic Research.
- Tenreyro, S. and Thwaites, G.: 2016, Pushing on a string: Us monetary policy is less powerful in recessions, *American Economic Journal: Macroeconomics* **8**, 43–74.

Appendix A: Proof of Proposition 1

If ξ_t is an instrument for the policy shock ε_t^i so that $\xi_t = \alpha\varepsilon_t^i + \nu_t$ with ν_t classical measurement error, the horizon h impulse responses for inflation and unemployment – under a linearity assumption – are given by the population coefficients in the regressions

$$\pi_{t+h} = \xi_t \beta_h^\pi + w_t' \gamma_h^\pi + e_{h,t+h}^\pi, \quad \text{and} \quad u_{t+h} = \xi_t \beta_h^u + w_t' \gamma_h^u + e_{h,t+h}^u. \quad (9)$$

Indeed, if ξ_t is randomly assigned (conditional on w_t) then it is uncorrelated with the error terms $e_{h,t+h}^\pi$ and $e_{h,t+h}^u$ and the coefficients β_h^π and β_h^u capture the causal effect of the aggregate demand shock on inflation and unemployment. In particular, we have that $\mathcal{R}_\pi(h) = \beta_h^\pi$ and $\mathcal{R}_u(h) = \beta_h^u$.

Similar to Stock and Watson (2018) we find it convenient to first project out the control variables. To this extent, let π_t^\perp , u_t^\perp and ξ_t^\perp denote the residuals from the population projection onto $\{w_t\}$. Then the population coefficients of (9) are given $\beta_h^\pi = \text{E}((\xi_t^\perp)^2)^{-1} \text{E}(\xi_t^\perp \pi_{t+h}^\perp)$ and $\beta_h^u = \text{E}((\xi_t^\perp)^2)^{-1} \text{E}(\xi_t^\perp u_{t+h}^\perp)$. Now the Phillips multiplier can be written as

$$\mathcal{P}_h = \frac{\frac{1}{h} \sum_{j=0}^h \beta_j^\pi}{\frac{1}{h} \sum_{j=0}^h \beta_j^u} = \frac{\sum_{j=1}^h \text{E}(\xi_t^\perp \pi_{t+j}^\perp)}{\sum_{j=1}^h \text{E}(\xi_t^\perp u_{t+j}^\perp)} = \frac{\text{E}\left(\xi_t^\perp \sum_{j=1}^h \pi_{t+j}^\perp\right)}{\text{E}\left(\xi_t^\perp \sum_{j=1}^h u_{t+j}^\perp\right)}.$$

The last expression on the right hand side is exactly the population coefficient implied by regression (4) of the main text after projecting out the control variables.

Appendix B: Anderson-Rubin type confidence sets

In this section we detail the construction of the weak instrument robust confidence bounds based on inverting the Anderson and Rubin (1949) statistic. We following the serial correlation robust implementation discussed in Andrews et al. (2019) Section 5.1. Consider equation (4) here restated in more convenient notation

$$y_{t+h} = x_{t+h} \mathcal{P}_h + w_t' \gamma_h + e_{t+h},$$

where $y_{t+h} = \sum_{j=0}^h \pi_{t+j}$ and $x_{t+h} = \sum_{j=0}^h u_{t+j}$. Since, the AR-statistic is constructed the same way for all h we detail the procedure while omitting the subscript h . Let y_t^\perp , x_t^\perp and e_t^\perp denote the residuals from the projection on $\{w_t\}$. We have

$$y_t^\perp = x_t^\perp \mathcal{P} + e_t^\perp$$

and to construct the confidence bands for \mathcal{P} we invert test statistics for the null hypothesis $H_0 : \mathcal{P} = \mathcal{P}^0$ for different values of \mathcal{P}^0 . The AR test statistic is based on the moment condition

$$\text{E}(\xi_t e_t^\perp(\mathcal{P}^0)) = \text{E}(\xi_t (y_t^\perp - x_t^\perp \mathcal{P}^0)) = 0,$$

where $e_t^\perp(\mathcal{P}^0) = y_t^\perp - x_t^\perp \mathcal{P}^0$ and ξ_t is the monetary policy shock. We define its empirical counterpart as follows $\bar{g}(\mathcal{P}^0) = \frac{1}{n} \sum_{t=1}^n \xi_t (y_t^\perp - x_t^\perp \mathcal{P}^0)$ and the AR test statistic becomes

$$AR(\mathcal{P}^0) = \left(\frac{\bar{g}(\mathcal{P}^0)}{\hat{\sigma}(\mathcal{P}^0)} \right)^2$$

where $\hat{\sigma}^2(\mathcal{P}^0)$ is the serial correlation and heteroskedasticity robust variance estimate for $\text{Var} \left(\frac{1}{\sqrt{n}} \sum_{t=1}^n \xi_t (y_t^\perp - x_t^\perp \mathcal{P}^0) \right)$. In practice we use the Newey and West (1994) estimator to obtain the variance estimate $\hat{\sigma}^2(\mathcal{P}^0)$. Note that unusual form of the $AR(\mathcal{P}^0)$ statistic is due to the fact that we use only one instrument, this leads to the scalar form of the AR statistic. Under mild assumptions, that do not involve the strength of the instruments, we have that $AR(\mathcal{P}_0) \xrightarrow{d} \chi^2(1)$, see Andrews et al. (2019). Confidence bounds for \mathcal{P} are obtained by inverting the AR statistic for different values of \mathcal{P}^0 . The full procedure is summarized in the following steps.

- For $h = 0, 1, \dots, H$:
 - (i) Consider a set of plausible values for \mathcal{P}_h , denoted by \mathcal{D} ;
 - (ii) For each $\mathcal{P}_h^0 \in \mathcal{D}$ compute the $AR(\mathcal{P}_h^0)$ statistic;
 - (iii) Evaluate,
 - if $AR(\mathcal{P}_h^0) < c_{1-\alpha, \chi^2(1)}$ include \mathcal{P}_h^0 in the horizon h confidence set;
 - if $AR(\mathcal{P}_h^0) > c_{1-\alpha, \chi^2(1)}$ exclude \mathcal{P}_h^0 from the horizon h confidence set;
 where $c_{1-\alpha, \chi^2(1)}$ is the level $1 - \alpha$ critical value of of the $\chi^2(1)$ distribution.

In practice we take $\mathcal{D} = [-10, 10]$ and $\alpha = 0.05$.

Appendix C: Model-based estimation of the trade-off

In this section, we discuss the traditional structural model-based approach to infer the inflation-unemployment trade-off faced by policy makers, and we highlight its two main challenges: (i) mis-specification issues, and (ii) endogeneity issues.

A model-based approach needs to take a stand on how monetary policy affects inflation and unemployment. This is traditionally done by specifying a Phillips curve, linking inflation to real activity, and (often only implicitly) an (IS) curve, linking real activity to the short-term policy rate.

To help with the discussion, consider the basic three-equation New-Keynesian model

$$\begin{cases} \pi_t - \pi_{t+1|t}^e = -\kappa \tilde{u}_t + \nu_t \\ \tilde{u}_t = \tilde{u}_{t+1|t}^e - \sigma^{-1}(i_t - \pi_{t+1|t}^e - r_t^n) \\ i_t = \phi_\pi \pi_t + \phi_u \tilde{u}_t + \varepsilon_t^i \end{cases} \quad (10)$$

The first equation is a Phillips curve linking inflation to expected future inflation $\pi_{t+1|t}^e$ and the unemployment gap $\tilde{u}_t = u_t - u_t^*$ (also referred to as the “forcing variable”) with u_t^* the natural rate of unemployment, $\tilde{u}_{t+1|t}^e$ the expected future output gap and ν_t a disturbance

term.²⁸ The second equation is the dynamic (IS) curve linking the unemployment gap to the real interest rate gap $r_t - r_t^n$ with $r_t = i_t - \pi_{t+1|t}$ the real interest rate and r_t^n is the natural real rate of interest.²⁹ The third equation is a Taylor rule linking the nominal interest rate i_t to inflation and the unemployment gap with ε_t^i a monetary shock.

Mis-specification issues

We discuss three related mis-specification issues that can affect the ability of model-based approaches to learn about the inflation-unemployment trade-off faced by the central bank.

- The trade off and the slope of the Phillips curve:
In the basic New-Keynesian model sketched above, the inflation-unemployment trade off, labeled T_h^{nk} , reduces to the slope of the Phillips curve, i.e.,³⁰

$$T_h^{nk} = -\kappa, \forall h \geq 0. \quad (11)$$

In line with (11), the literature has traditionally treated the slope of the Phillips curve as the relevant measure of the inflation-unemployment trade-off faced by the central bank. However, this result is not a general one, and small plausible changes to the specification of the (IS) curve can break it. In general, the inflation-unemployment trade-off depends on the full dynamic specification of both the Phillips and the (IS) curves, and the slope of the Phillips curve alone is not enough to characterize the trade-off.

To give a concrete example, we take the conclusions of Fuhrer and Rudebusch (2004) at face value and consider an (IS) curve of the form

$$\tilde{u}_t = \rho_u \tilde{u}_{t-1} - \sigma^{-1} (i_t - E_{t-1} \pi_{t+1} - r_t^n). \quad (12)$$

²⁸The forcing variable is sometimes the unemployment gap (Gordon (2011), Gali (2011), Coibion and Gorodnichenko (2015)), sometimes the output gap (Mavroeidis et al. (2014)), or some more direct measure of real marginal cost (Gali and Gertler (1999)). Given our focus on the trade-off between unemployment and inflation, we treat the forcing variable as the unemployment gap.

²⁹The (IS) curve is typically expressed in terms of the output gap rather than the unemployment gap. An implicit assumption in (10) is thus the existence of a fourth equation—an Okun-type law—linking the output gap to the unemployment gap: $x_t = -c(u_t - u_t^*)$. Dynamic mis-specification issues may also affect the Okun’s law and thus compound the specification issues discussed in this section.

³⁰To see that, we can use the method of undetermined coefficients as in Galí (2015) and set non-monetary shocks to zero, which gives

$$\begin{cases} \tilde{u}_t = \Lambda \varepsilon_t^i \\ \pi_t = -\kappa \Lambda \varepsilon_t^i \end{cases}, \quad \forall t$$

with $\Lambda = \frac{1}{\sigma + \phi_y + \kappa \phi_\pi}$.

This implies

$$\begin{cases} \frac{\partial \tilde{u}_t}{\partial \varepsilon_t^i} = \Lambda \\ \frac{\partial \pi_t}{\partial \varepsilon_t^i} = -\kappa \Lambda \end{cases}, \quad \forall t$$

and $\frac{\partial \tilde{u}_{t+h}}{\partial \varepsilon_t^i} = 0, \frac{\partial \pi_{t+h}}{\partial \varepsilon_t^i} = 0$ for $h > 0$, which then gives

$$MRT_h = \frac{\sum_{j=0}^h \left. \frac{\partial \pi_{t+t+j}}{\partial i_t} \right|_{\varepsilon_t^i=1}}{\sum_{j=0}^h \left. \frac{\partial \tilde{u}_{t+t+j}}{\partial i_t} \right|_{\varepsilon_t^i=1}} = -\frac{\sum_{j=0}^h \kappa \Lambda}{\sum_{j=0}^h \Lambda} = -\kappa.$$

With such backward-looking (IS) curve, T_h becomes³¹

$$T_h = \frac{-\kappa}{1 - \rho_u} \neq -\kappa$$

so that the inflation-unemployment trade-off depends on *both* the slope of the Phillips curve *and* on the (IS) curve coefficients.³²

- **Dynamic mis-specification:**
While plausible empirical specifications of (10) likely involve numerous lags of the right-hand side variables, there is considerable disagreement on the correct dynamic specification (e.g., King and Watson (1994), Fuhrer and Rudebusch (2004), Gordon (2011)).
- **Omitted variable bias:**
Another specification challenge involves the set of relevant explanatory variables to include in the Phillips and (IS) curves. For instance, for the Phillips curve a number of additional determinants of inflation have been proposed. For instance, the cost channel of monetary policy (e.g., Barth and Ramey (2001)) implies that inflation is also directly influenced by the level of the interest rate, so that the Phillips curve should include the interest rate as a right-hand side variable. Similarly, commodity prices can also directly influence inflation (e.g., Gordon (2011)). Subsuming these additional determinants into the residual can lead to an omitted variable bias.

Endogeneity issues

In addition, the Phillips and (IS) curves suffer from a number of endogeneity biases, as argued for instance by Mavroeidis et al. (2014) and Fuhrer and Rudebusch (2004).³³ Taking the Phillips curve as an example, endogeneity issues arise because of (i) unobserved inflation expectations, (ii) unobserved unemployment gap \tilde{u}_t and (iii) confounding from supply shocks.³⁴ Similar issues affect the (IS) curve.

³¹ To see that, combine 12 with the Phillips curve $\pi_t = E_t \pi_{t+1} - \kappa \tilde{u}_t$ to get

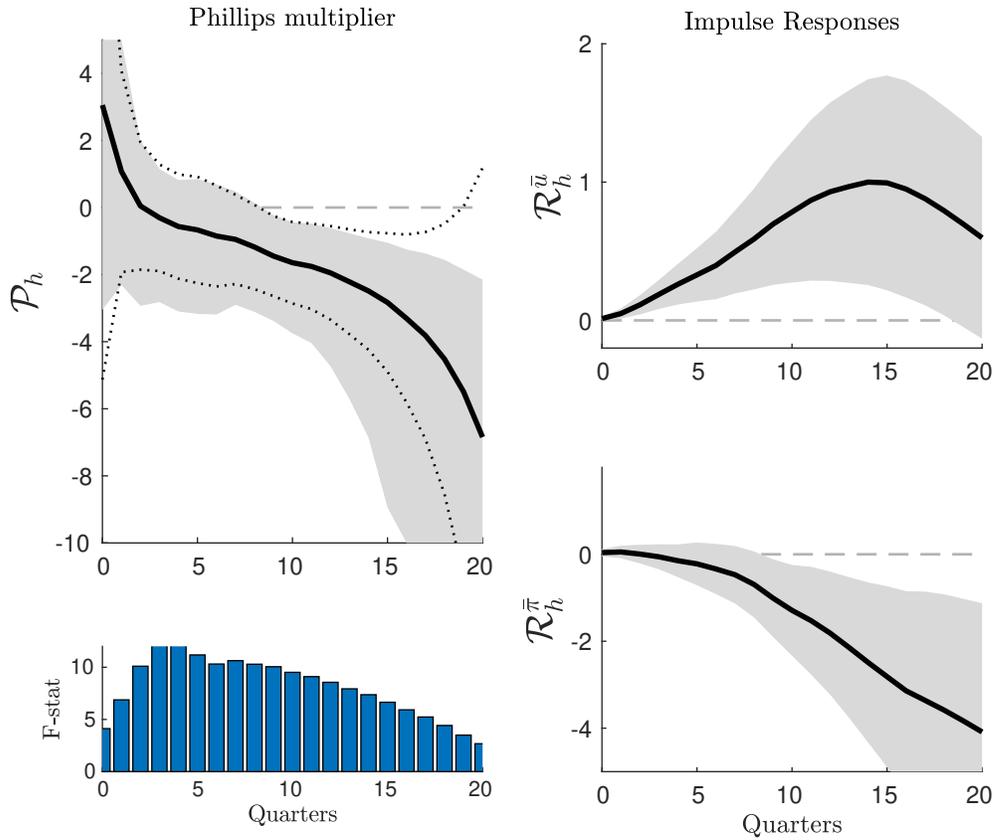
$$\begin{aligned} \frac{\partial \pi_t}{\partial \varepsilon_t^i} &= -\kappa E_t \sum_{l=0}^{\infty} \frac{\partial \tilde{u}_{t+l}}{\partial \varepsilon_t^i} = -\kappa \sum_{l=0}^{\infty} \rho_u^l \\ &= -\frac{\kappa}{1 - \rho_u}. \end{aligned}$$

³²To get some intuition, note that the dynamic trade-off between inflation and unemployment depends on the relative speed with which inflation and unemployment react to a change in monetary policy. Varying ρ_u affects T_h , because it changes the speed at which unemployment reacts to a change in policy, but it has little effect on the reaction of inflation (since inflation is forward-looking in this model).

³³By modeling the behavior of all relevant variables, a full information approach based on a DSGE model has the potential to address these endogeneity issues. However, the full-information approach introduces an even greater risk of misspecification, magnifying the specification issues discussed above, inducing bias or inconsistency of the parameters of interest (Mavroeidis et al. (2014)).

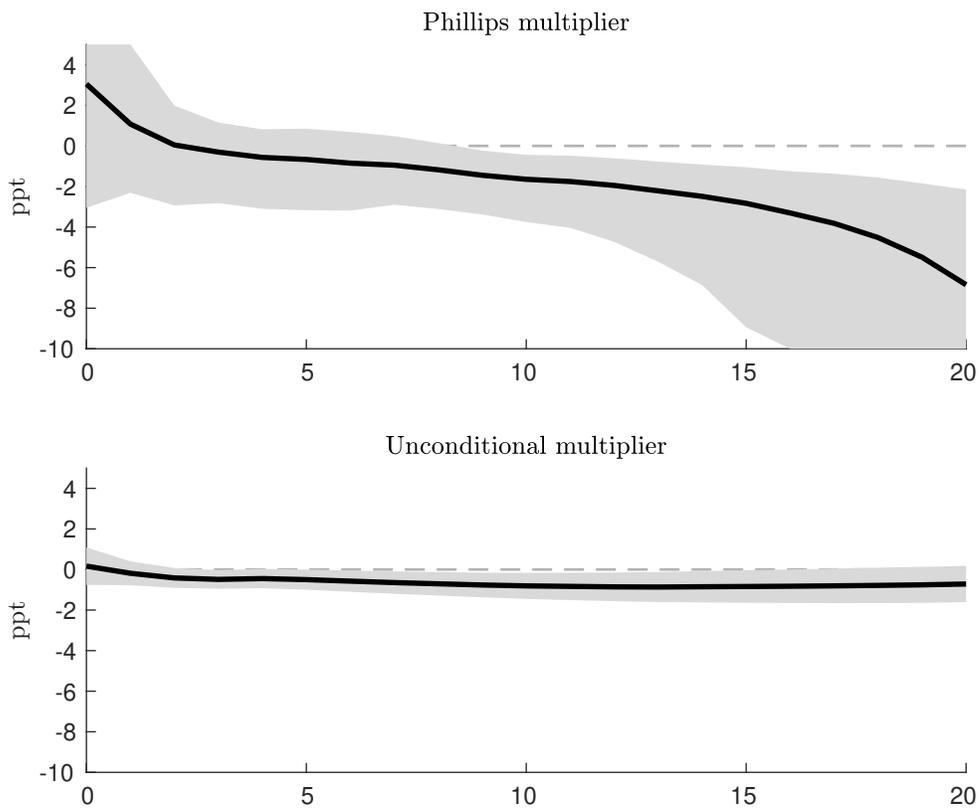
³⁴More precisely, (i) and (ii) are measurement error biases. Expected future inflation can only be imperfectly proxied by future inflation or some survey measure, while the natural rate of unemployment u^* is also unobserved and imperfectly measured by some natural rate proxy like that of the CBO, e.g. Coibion and Gorodnichenko (2015). Note that the measurement error in the gap is likely not even iid.

Figure 1: The US Phillips multiplier \mathcal{P}_h — 1969-2007, RR



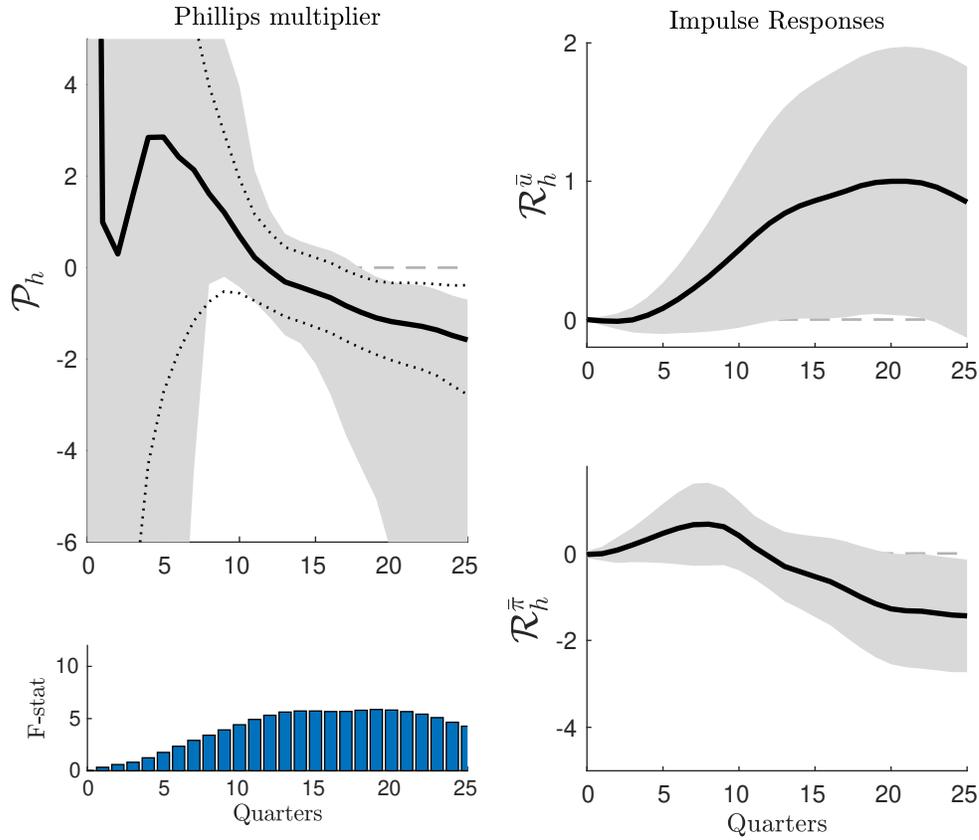
Notes: US Phillips multiplier estimated using Romer and Romer (RR) narrative monetary shocks as instruments. The sample is 1969q1-2007q4. For the multiplier (upper-left) the shaded area corresponds to the weak instrument robust 95% confidence interval that is obtained by point-wise inversion of the AR-statistic, see Appendix B. The dotted lines correspond to the confidence bounds implied by the normal limiting distribution of the 2SLS estimator. The *average* impulse responses (right panels) for inflation and unemployment are obtained from the OLS regressions (5) where the confidence bounds are computed using Newey and West (1994). The *F*-statistics (bottom left) are computed as discussed in Olea and Pflueger (2013).

Figure 2: The Phillips multiplier \mathcal{P}_h vs the unconditional multiplier



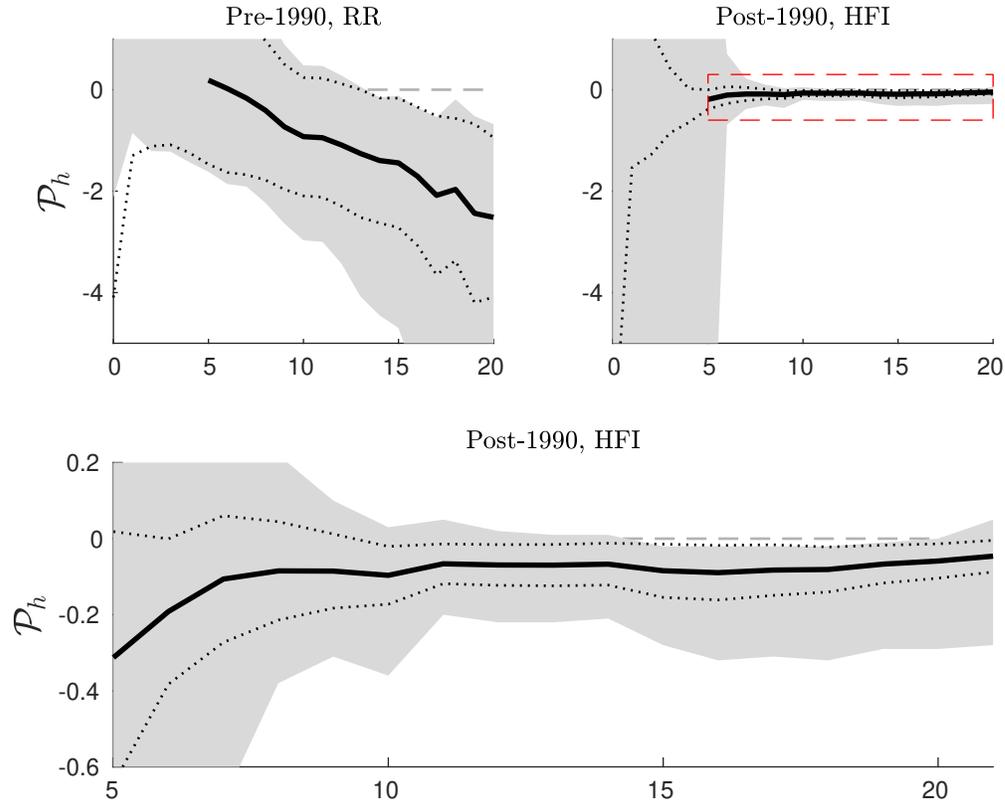
Notes: Top panel: US Phillips multiplier estimated using Romer and Romer (RR) narrative monetary shocks as instruments. The sample is 1969q1-2007q4. The 95% weak instrument robust confidence bands (shaded area) are obtained by point-wise inversion of the AR-statistic, see Appendix B. Bottom panel: The unconditional multiplier estimated using OLS where the confidence bounds are computed using Newey and West (1994).

Figure 3: The UK Phillips multiplier \mathcal{P}_h — 1975-2007, CH



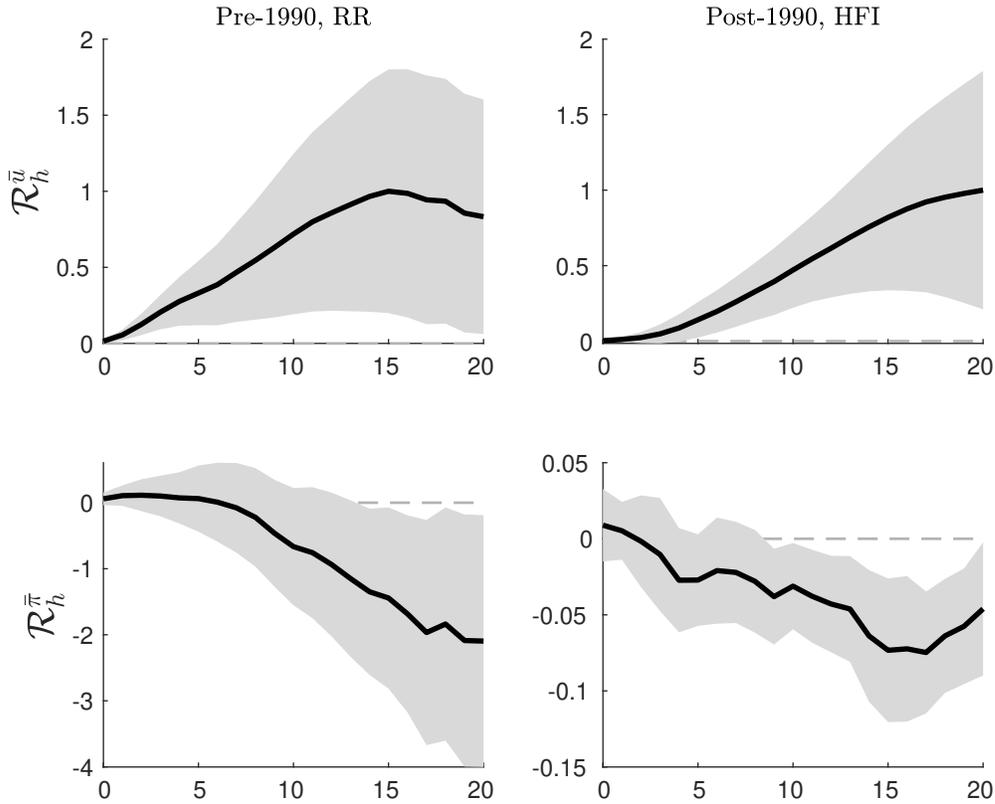
Notes: UK Phillips multiplier estimated using Cloyne and Hurtgen (2016) (CH) narrative monetary shocks as instruments. The sample is 1975q1-2007q4. For the multiplier (upper-left) the shaded area corresponds to the weak instrument robust 95% confidence interval that is obtained by point-wise inversion of the AR-statistic, see Appendix B. The dotted lines correspond to the confidence bounds implied by the normal limiting distribution of the 2SLS estimator. The *average* impulse responses (right panels) for inflation and unemployment are obtained from the OLS regressions (5) where the confidence bounds are computed using Newey and West (1994). The *F*-statistics (bottom left) are computed as discussed in Olea and Pflueger (2013).

Figure 4: The Phillips multiplier \mathcal{P}_h over time — RR/HFI



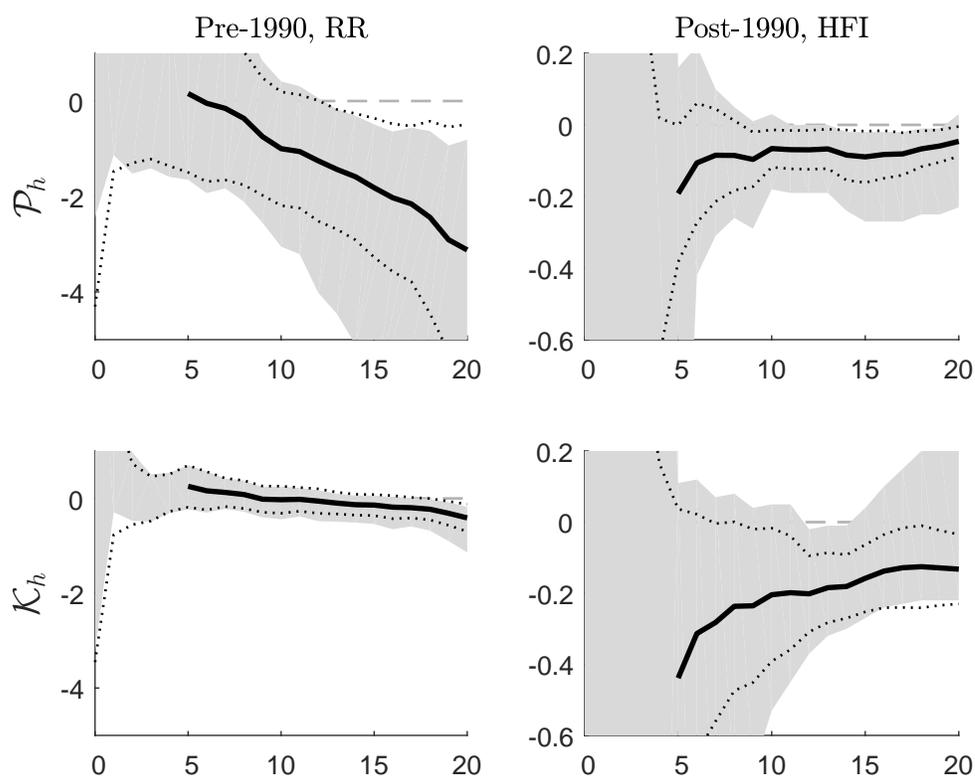
Notes: Upper left: Phillips multiplier estimated using Romer-Romer (RR) monetary shocks over 1969q1–1989q4. Upper right: Phillips multiplier estimated using High-Frequency Identified (HFI) monetary surprises (“FF4”, the three month ahead monthly Fed Funds futures) over 1990q1–2008q4. Lower panel: Phillips multiplier estimated using High-Frequency Identified (HFI) monetary surprises (part corresponding to the red dashed rectangle in the upper right panel). The shaded areas indicate the 95% weak instrument robust confidence regions computed by point-wise inverting the AR-statistic, see Appendix B. The dotted lines correspond to the confidence bounds implied by assuming the normal limiting distribution of the 2SLS estimator.

Figure 5: IRs of *average* inflation and unemployment — RR/HFI



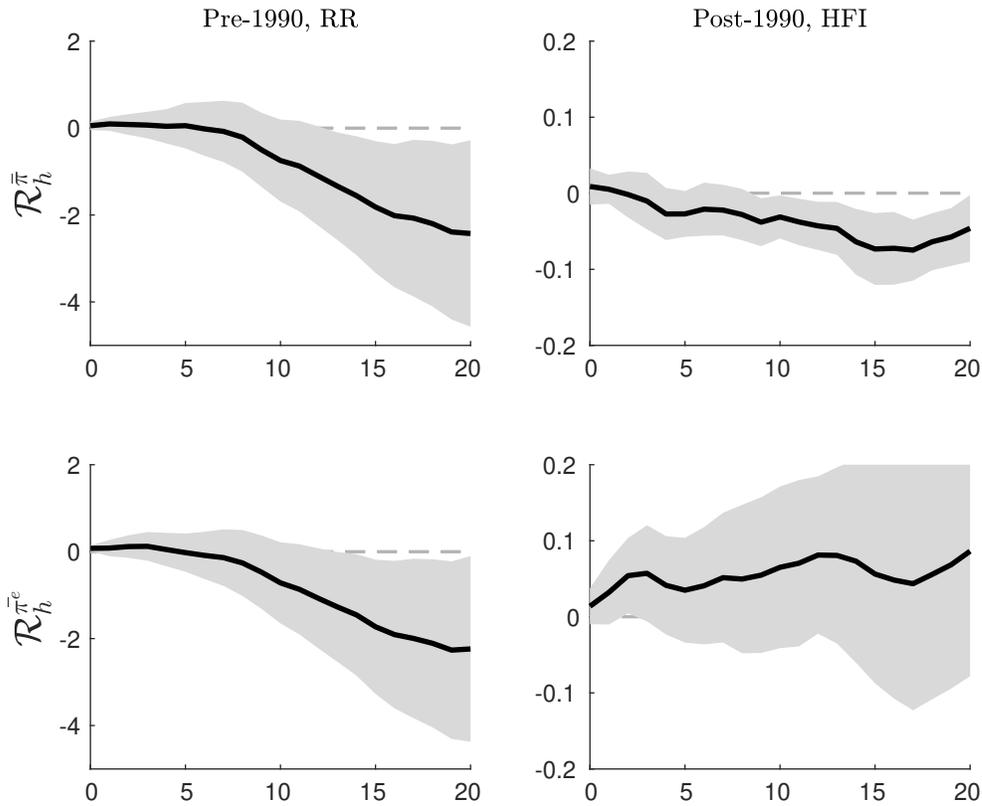
Notes: Left column: Impulse responses for $\bar{u}_{t:t+h}$ and $\bar{\pi}_{t:t+h}$ computed using the Romer and Romer narrative monetary shocks over 1969q1–1989q4. Right column: Impulse responses for $\bar{u}_{t:t+h}$ and $\bar{\pi}_{t:t+h}$ computed using High-Frequency Identified (HFI) monetary surprises (“FF4”, the three month ahead monthly Fed Funds futures) over 1990q1–2008q4. 95% Newey and West (1994) based confidence intervals for the OLS estimates are displayed. For ease of comparison, the size of the monetary shock is set such that the impulse response of average unemployment peaks at 1ppt in both cases.

Figure 6: \mathcal{P}_h and \mathcal{K}_h over time



Notes: Top panels: The Phillips multiplier pre- and post-1990, see Figure 4 for details. Bottom left: \mathcal{K}_h , the expectation-augmented Phillips multiplier, estimated using Romer-Romer (RR) monetary shocks over 1969q1–1989q4. Bottom right: \mathcal{K}_h estimated using High-Frequency Identified (HFI) monetary surprises (“FF4”, the three month ahead monthly Fed Funds futures) over 1990q1–2008q4. The shaded areas indicate the 95% weak instrument robust confidence regions computed by point-wise inverting the AR-statistic, see Appendix B. The dotted lines correspond to the confidence bounds implied by assuming the normal limiting distribution of the 2SLS estimator.

Figure 7: IRs of *average* inflation and inflation expectation over time — RR/HFI



Notes: Left column: Impulse responses for $\bar{\pi}_{t:t+h}$ and $\bar{\pi}_{t:t+h}^e$ computed using the Romer and Romer narrative monetary shocks over 1969q1–1989q4. Right column: Impulse responses for $\bar{u}_{t:t+h}$ and $\bar{\pi}_{t:t+h}$ computed using High-Frequency Identified (HFI) monetary surprises (“FF4”, the three month ahead monthly Fed Funds futures) over 1990q1–2008q4. 95% confidence intervals for the OLS estimates are displayed. For ease of comparison, the size of the monetary shock is set such that the impulse response of average unemployment peaks at 1ppt.