Measuring the inflation-unemployment trade-off*

Regis Barnichon†　　Òscar Jordà‡

November 15, 2017

Abstract

The magnitude and properties of the inflation-unemployment trade-off –the Phillips curve– remain poorly understood, in part for two reasons: (i) Phillips curve coefficient estimates are biased by unobserved supply factors and measurement error, and (ii) there is little agreement on the correct Phillips curve specification. We propose a semi-parametric estimator of the inflation-unemployment trade-off that is independent of any specific Phillips curve specification, and we use aggregate demand shocks –monetary shocks and government spending shocks– as instruments to properly identify the trade-off. We obtain two main results. First, because of the inertial behavior of inflation, the inflation-unemployment trade-off can be temporarily muted. Second, the inflation-unemployment trade-off increases markedly (in absolute value) in very tight labor markets.

JEL classification codes: C14, C32, E32, E52

Keywords: Gaussian basis functions, local projections, instrumental variables, monetary policy, state-dependence, Phillips curve.

---

*Preliminary–Do not circulate. We thank Glenn Rudebusch and Dan Wilson for helpful comments. The views expressed in this paper are the sole responsibility of the authors and to not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

†Federal Reserve Bank of San Francisco and CEPR (regis.barnichon@sf.frb.org).

‡Federal Reserve Bank of San Francisco; and Department of Economics, University of California, Davis (oscar.jorda@sf.frb.org; ojorda@ucdavis.edu).

1 Introduction

The existence of an inflation-unemployment trade-off is a central concept in macroeconomics, both as a cornerstone of many macro models –the Phillips curve–, and as a guiding principle for central banks' policy actions.

The magnitude and properties of this trade-off remain poorly understood however, in part for two reasons. First, estimating the Phillips curve is notoriously difficult: (1) there is disagreement on its correct specification; (2) coefficient estimates are biased by confounding from supply factors and by measurement error in inflation expectations and in the natural rate of unemployment. Second, while there is a clear mapping between Phillips curve coefficients and the inflation-unemployment trade-off in the static formulation of the Phillips curve,[1] the mapping becomes much less trivial for more plausible Phillips curve specifications with lags of inflation and unemployment.[2] In fact, in those specifications, there is no agreed-upon definition of what the inflation-unemployment trade-off actually is.

In this paper, we propose a characterization of the inflation-unemployment trade-off that is independent of any specific Phillips curve specification, and we propose an estimation method that is largely immune to the empirical issues affecting the Phillips curve: it does not require a measure of inflation expectations, it is immune to the presence of aggregate supply shocks, and more generally it does not require the specification and estimation of a Phillips curve.

We define the inflation-unemployment trade-off as the change in inflation caused by an aggregate demand shock that lowers unemployment by one percentage point. Since the value of this trade-off depends on the time horizon after the shock, we define κ_h –the *dynamic* inflation-unemployment trade-off at horizon h– as the average change in inflation over the next h periods relative to an average change in unemployment of one percentage point over the same next h periods. By tracing the magnitude of κ_h from the short-run (small h) to the long-run (large h), we will be able to characterize the inflation-unemployment trade-off while taking into account the (possibly different) dynamics of inflation and unemployment.[3]

[1]In Phillips (1958) original formulation or in undergraduate textbooks, the Phillips curve is a static relation between inflation (or its change) and unemployment, and in that case the trade-off is simply the coefficient of unemployment.

[2]See King and Watson (1994), Staiger, Stock and Watson (1997*a*), Staiger, Stock and Watson (1997*b*).

[3]In effect, our notion of a dynamic inflation-unemployment trade-off generalizes the concepts of short-run

We measure this dynamic trade-off from the relative impulse response functions of inflation and unemployment to aggregate demand shocks. By doing so, we address the main challenges that have plagued Phillips curve estimation. First, using impulse response analysis allows us to capture the complex dynamics in inflation and unemployment in a flexible semi-parametric fashion, i.e., without imposing arbitrary a priori restrictions of relative dynamics of inflation and unemployment (as would be the case in a Phillips curve setting). Second, studying the relative behavior of inflation and unemployment *conditional* on aggregate demand shocks allows to address endogeneity issues, i.e., to avoid confounding from supply factors or measurement error.[4]

Our framework then allows us to explore additional properties of the inflation-unemployment trade-off, notably whether κ_h varies with the level of slack in the economy –going back to Phillips (1958) and Phelps (1968)'s original idea that the relation between unemployment and inflation may be highly non-linear–, or whether κ_h can change secularly over time –relating to a lively debate on the recent flattening of the Phillips curve (e.g., Blanchard (2016))– .

We consider two types on aggregate demand shocks: monetary policy shocks and government spending shocks. As baseline results, we use narratively identified shocks as instruments to identify κ_h, specifically the Romer and Romer (2004) monetary shocks and the Ramey and Zubairy (2016) news shocks to government (defense) spending. As robustness checks, we also consider internal instruments (in the language of Stock, 2008) and we use shocks identified from exclusion restrictions: recursively identified monetary shocks (Christiano et al., 2005) and recursively identified government spending shocks (Blanchard and Perotti (2002), Auerbach and Gorodnichenko (2012)).

We obtain two main results. First, because the unemployment rate responds faster than inflation following a demand shock, a tighter labor market translates into higher inflation with about a 1-year delay. As a result, the dynamic inflation-unemployment trade-off is small at short horizon –appearing temporarily muted–, but then increases markedly (in absolute value) over time. Quantitatively, for the first year after a shock ($h < 1$ year) $\kappa_h \approx .5$, so that a 1 ppt lower unemployment translate into only 0.5 ppt higher inflation. However, the trade-off increases (in absolute value) over time and at a 4 year horizon, $\kappa_h \approx 2$, so that a 1 ppt lower unemployment translate into 2 ppt higher inflation.

Second, the trade-off changes markedly with the level of slack: while the trade-off is small (little

(κ_0) and long-run (κ_∞) trade-offs proposed by King and Watson (1994).

[4]In other words, we use aggregate demand shocks as instrument variables to identify the inflation-unemployment trade-off. Intuitively, in the parlance of the static IS/LM-AD/AS framework, we measure the slope of the AS curve –the inflation-unemployment trade-off– using exogenous movements in the AD curve.

different from zero) in slack labor markets, the trade-off increases markedly (in absolute value) in tight labor markets. For instance, at the level of tightness found in the late 60s, the trade-off reaches -5, that is a 1 ppt lower unemployment translates into about 5 ppt higher inflation. In other words, our findings echo Phillips (1958) and Phelps (1969)'s original idea that the relation between unemployment and inflation may be highly non-linear.

These results have two implications for the current low inflation-low unemployment situation that have prompted many to question whether central banks now face a different inflation-unemployment trade-off than in the past. First, the currently muted inflation response in the face of a tightening labor market can be explained by the inertial behavior of inflation. Second, since the inflation-unemployment trade-off increases (in absolute value) with labor market tightness, the risks of a steep increase in inflation increase substantially as the unemployment rate continues to fall below its natural rate.

Subsample analysis focusing on the 1984-2007 period indicates an attenuation of the inflation-unemployment trade-off relative to the full sample results that is consistent with the well-documented flattening of the Phillips curve. This attenuation is also visible in the nonlinear response to a tight labor market, but with an important caveat. The estimation sample does not include observations from very tight labor markets such as the mid-1960s. This absence could unduly attenuate estimates of the nonlinearity over this sample.

To estimate the impulse response functions necessary to construct κ_h we propose a new method –Functional Local Projection (FLP)– that aims to straddle between the parametric parsimony of VARs and the flexibility of Local Projections (Jorda (2005)). FLP consists in running local projections with the additional constraint that the estimated impulse response function can be approximated by a family of basis functions. Intuitively, the idea behind FLP is to (i) directly estimate the impulse response functions to shocks (unlike the VAR approach, which imposes a finite VAR stucture to the data), and (ii) approximate the (high-dimensional) impulse response functions with a (small) number of basis functions, which then allows to explore non-linearities or time variation in the coefficient while preserving degrees of freedom (in contrast to non-parametric local projections). Importantly, we show that FLP estimation is straightforward, because it can be cast as a standard GMM problem.

While different basis functions are possible, we will use Gaussian basis functions to approximate impulse response functions. Intuitively, impulse responses are often found (or theoretically predicted) to be monotonic or hump-shaped,[5] and in such cases, one Gaussian basis function can provide a parsimo-

[5]See e.g., Christiano, Eichenbaum and Evans (1999) for the effects of monetary shocks. In New-Keynesian

nious and intuitive way to summarize an impulse response function with 3 parameters each capturing a separate and easily interpretable feature of the IRF: (a) the magnitude of the peak effect of a shock, (b) the time to that peak effect, and (c) the persistence of that peak effect. The parsimony and interpretability of FLP with Gaussian functions allow us to handle non-linearities or time-variation in κ_h in a controlled fashion.

While our paper relates to a very large literature on Phillips curve estimation (Mavroeidis, Plagborg-Mller and Stock (2014)), we intentionally shy away from direct estimation of a Phillips curve to avoid dealing with inflation expectations, insufficiently rich dynamics, and omitted supply confounders like oil shocks. Instead, we see our efforts to define and measure the inflation-unemployment trade-off as paralleling the fiscal literature (e.g., Ramey and Zubairy, 2016). Our inflation-unemployment trade-off is akin to a spending multiplier;[6] it captures how much inflation the central bank can "buy" with a one percentage higher unemployment rate. In this sense, our inflation/unemployment trade-off has an interpretation as the inverse of the sacrifice ratio although definitions in the literature can vary.

Section 2 discusses the pitfalls associated with Phillips curve estimates and proposes a new method to directly measure the inflation-unemployment trade-off. Section 3 presents FLP –our novel method to estimate impulse response functions– and Section 4 shows how FLP can easily accommodate non-linearities, Section 5 presents our results and our estimates of the dynamic inflation-unemployment trade-off.

2 Measuring the inflation-unemployment trade-off

2.1 The Phillips Curve

The Phillips curve is a cornerstone of macroeconomics as it captures the Aggregate Supply relation of an economy, that is, the relation linking the level of production and firms' pricing decisions. In its typical specification, the Phillips curve describes how inflation fluctuates in response to movements in labor market slack.

Since its first inception (Phillips (1958), Samuelson and Solow (1960)), the Phillips curve has been the subject of intense study. In its modern version, the Phillips curve postulates that inflationary pres-

models, the impulse response functions are generally monotonic or hump-shaped (see e.g., Walsh (2010)).

[6]The government spending multiplier is the ratio of change in national income to a change in government spending.

sures depend on the level of economic activity and inflation expectations. Specifically, a Friedman(1968) expectations-augmented Phillips curve takes the form

$$\pi_t - E_t\pi_{t+1} = c + \kappa x_t + \nu_t \tag{1}$$

where x_t is a measure of economic slack –typically the unemployment gap–, ν_t corresponds to aggregate supply (or cost-push) shocks, c is a constant and $E_t\pi_{t+1}$ is expected future inflation.

In (1), the slope of the Phillips curve –κ– is precisely the inflation-unemployment trade-off; it captures the trade-off between holding inflation stable and keeping real activity elevated and is thus an object of central importance for central banks and policy makers.

Estimating κ is notoriously difficult however: (1) there is disagreement on *the* correct Phillips curve specification; (2) confounding from supply factors bias coefficient estimates; (3) inflation expectations are difficult to measure and therefore prone to measurement error.

Disagreement on the correct Phillips curve specification

While the original Phillips curve (Phillips, 1958, Friedman 1968) or its modern New-Keynesian formulation (Clarida, Gali and Gertler, 1999, Gal (2011)) postulates a well defined relation between inflation, inflation expectations and the unemployment (or output) gap, these stylized models do not capture the rich dynamics that inflation and unemployment can display. As a result, most econometric study of the Phillips curve do allow for distributed lags of inflation and unemployment in (1) to capture the rich dynamics of inflation and unemployment (King and Watson (1994), Staiger, Stock and Watson (1997*a*), Staiger, Stock and Watson (1997*b*), 2001, 2014).[7]

The importance of dynamics create two problems for the estimation and interpretation of the Phillips curve. First, there is no consensus on the appropriate number of lags, and empirical studies typically provide little justification for their chosen number of lags.[8] Second, and most importantly, the presence of lags of inflation and unemployment in (1) destroys the simple mapping between the Phillips curve

[7]The need for lags might arise from lagged effects of the output gap on inflation, from the inertial dynamics of inflation or unemployment (see Barnichon and Nekarda, 2012) or from serially correlated measurement error in inflation, among other possibilities.

[8]For instance, in an effort to bring the New-Keynesian wage Phillips curve to the data, Gal (2011) assumes that unemployment follows an AR(2) at a quarterly frequency and thus allows for two lags of unemployment. Staiger, Stock and Watson (1997*a*) also allow for two (quarterly) lags in unemployment. In contrast, Staiger, Stock and Watson (1997*b*) or King and Watson (1994) use 12 (monthly) lags.

coefficients and the object of interest: the inflation-unemployment trade-off. In other words, with lags of inflation and unemployment in (1), the trade-off is no longer measured by the coefficient κ.

Traditionally, the literature has avoided the issue by measuring the trade-off from Phillips curve estimation that do not include lags of unemployment or lags of inflation apart from efforts to capture inflation expectations (e.g., Blanchard (2016)). A few studies however did identify the problem In particular, King and Watson (1994) do estimate a dynamic Phillips curve, i.e., with lag terms, and then define two different "slopes" of the Phillips curve: a "short-run" slope defined as the contemporaneous coefficient on the output gap and a "long-run" slope defined as the ratio of the sum of the lag coefficients on unemployment to one minus sum of the lag coefficients on inflation. Intuitively, this long-run measure captures by how much inflation increases in the long-run when the output gap is decreased permanently by one unit. As we will see, our approach builds on and extend King and Watson' intuition.

Inconsistency of OLS estimates

OLS estimates of Phillips curve coefficients are inconsistent for two reasons: (i) measurement error, and (ii) confounding from (typically unobserved) supply factors such as oil price shocks.[9] More generally, the two problems can be seen as an omitted variable bias.

Because of (i), there are back-and-forth in the literature on the most appropriate (but necessarily imperfect) measure of inflation expectation (e.g., Coibion and Gorodnichenko (2015)). To address (ii), researchers routinely augment their Phillips curve regressions with (necessarily imperfect) proxies for cost-push shocks, such as oil price shocks dummies, price control dummies (e.g., Staiger, Stock and Watson (1997*b*)).

2.2 Going beyond the Phillips curve

In this section, we propose a new approach to characterize and estimate the inflation/unemployment trade-off. Specifically, we propose a characterization of the inflation-unemployment trade-off that is independent of any specific Phillips curve specification, and we propose an estimation method that is largely immune to the issues affecting the Phillips curve: it does not require a measure of inflation expectations, it is immune to the presence of supply shocks, and more generally it does not require the specification and estimation of a Phillips curve.

[9] To give a concrete example of (ii), a positive cost-push shock will lead both slack and inflation to increase, leading to a bias in the OLS estimate of κ.

Our approach consists in (i) estimating the impulse response functions of inflation and unemployment to some aggregate demand shocks (for instance, monetary shocks or government spending shocks), and (ii) measuring the inflation-unemployment trade-off from these impulse response functions. Intuitively, the use of aggregate demand shocks allows us to address endogeneity issues, and the use of impulse response functions allows us to account for the complex dynamics in inflation and unemployment in a flexible semi-parametric fashion.

To make things concrete, imagine that we have a set of estimated impulse responses in hand (for instance, from a VAR, Local Projections or from the method we propose in the next section). Let $\psi_j(h)$ for $j = \pi, u$ denote the impulse response coefficient of the j^{th} variable, h periods after impact. We characterize the π–u trade-off with

$$\kappa_h = \frac{\frac{1}{h}\sum_{i=0}^{h} \psi_\pi(i)}{\frac{1}{h}\sum_{i=0}^{h} \psi_u(i)}; \text{ for } h = 1, .., H. \tag{2}$$

κ_h captures, in response to an aggregate demand shock, how much higher inflation will be over the next h periods (on average) relative to a decline in the unemployment rate of one percentage point (on average) over the same h periods. In the case of a monetary shock, κ_h can be seen as the ratio of monetary "multipliers" (where the monetary response is common to numerator and denominator of κ_h and hence cancels out). A similar interpretation works with fiscal shocks.

In essence, κ_h is a summary statistics that captures the trade-off existing between inflation and unemployment, while at the same time taking into account the dynamics present in the economy.[10] In fact, our dynamic inflation-unemployment trade-off can be seen as a generalization of the concepts of short-run and long-run trade-off proposed by King and Watson (1994), which in our case correspond to κ_0 (a short-run trade-off) and κ_∞ (a long-run trade-off). However, and different from previous work, we use aggregate demand shocks to properly identify the dynamic inflation-unemployment trade-off, that is we calculate the trade-off *conditional* on an aggregate demand shock.

κ_h is precisely the object of interest for Fed as it captures the rules of the game for the central bank; by how much the Fed affects its two mandates (low π and low u) when it changes monetary policy. More generally, κ_h captures how the two Fed's targets are affected when aggregate demand conditions

[10] Paralleling the recent literature on the government spending multiplier, κ_h is a akin to a spending multiplier; how much inflation can the Fed "buy" with a one percentage higher unemployment rate.

change.

κ_h can be seen as a semi-nonparametric estimate of the inflation-unemployment trade-off; it does not require the specification/estimation of any Phillips curve equation.

Finally, we note that our measure of the inflation/unemployment trade-off has a loose connection with two traditional economic concepts: the slope of the Phillips curve and the sacrifice ratio. We shy away from direct estimation of a Phillips curve to avoid dealing with inflation expectations, insufficiently rich dynamics, and omitted supply confounders (say an oil shock). By using only demand shifts to achieve identification, and by leaving dynamics unrestricted in our semi-parametric framework, we think that we have a more accurate measurement of the π–u trade-offs implicit in the slope of the Phillips curve. Clearly κ_h has an interpretation as the inverse of the sacrifice ratio although definitions in the literature can vary.

3 Functional Local Projections (FLP)

Our characterization of the inflation-unemployment trade-off relies on estimates of the impulse response functions of inflation and unemployment to shocks. In this section, we describe our proposed impulse response estimator –Functional Local Projections–, a procedure aimed at improving the efficiency of Local Projections while preserving their flexible nature.

3.1 Introduction

We are interested in estimating the average response of the outcome variable y to treatment in the variable x, given a vector of controls w that include exogenous and predetermined variables, such as lags of the outcome and the treated variables. In addition we have a vector of instrumental variables z that are correlated with exogenous movements in the treated variable x but are otherwise uncorrelated with outcome variable y. We will always include w into z when using instrumental variable (IV) methods.

Given this set-up, we want to estimate the following average response function:

$$\mathcal{R}(h, \delta) = E(y_{t+h}|x_t = x_0 + \delta; w_t) - E(y_{t+h}|x_t = x_0; w_t) \quad \text{for } h = 0, 1, ... \tag{3}$$

where x_0 is a reference value (usually in linear models the choice of this value is inconsequential but not when there is nonlinearity) and δ denotes the size of the experiment (say a $1ppt.$ increase in the funds rate; a one-standard deviation shock, etc.). Approximating these conditional expectations with their linear projections, an estimate of the response function can be easily obtained from two-stage least squares instrumental variables regression, specifically:

$$
\begin{aligned}
y_{t+h} &= \hat{x}_t \beta_h + w_t \gamma_h + u_{t+h} \qquad & h &= 0, 1, ... \\
x_t &= z_t \theta + \epsilon_t \qquad & \hat{x}_t &= z_t \hat{\theta};
\end{aligned}
\tag{4}
$$

where w_t is of dimension $1 \times k$ and includes the constant term; z_t is of dimension $1 \times (l + k)$ for $l \geq 1$ so that θ is of dimension $(l + k) \times 1$.

Expression (4) is nothing more than Jordà's (2005) local projections IV estimator. Jordà, Schularick and Taylor (2017) show how such an approach delivers estimates of the structural impulse response function with valid instruments. Moreover, under linearity and when instruments are not available, there is a one-to-one mapping with impulse responses estimated using an auxiliary vector autoregression. Identification achieved with zero exclusion restrictions is naturally accommodated in this set-up (such as the common Cholesky causal ordering). Finally, identification using external instruments as in Stock and Watson (2012) is more naturally accommodated with expression (4) as Ramey (2016) and Ramey and Zubairy (2016) show.

In order to accommodate a wide variety of identification schemes and to derive distributional results using general results from the theory of generalized method of moments (GMM) estimators, the next section shows how local projection estimators can be cast as a general GMM problem.

3.2 Local projections as a GMM problem

We find useful to define the vector $y_t(H) = (y_t, y_{t+1}, ..., y_{t+H})'$, which collects the values of the outcome variable for the horizons over which we want to calculate the average response $\mathcal{R}(h, \delta)$. Under conditions that will be more explicit below, the direct generalization of expression (4) relies on the following sample moment conditions:

$$g_T(H) = \frac{1}{T}\sum_{t=1}^{T} Z_t'(y_t(H) - X_t\beta - W_t\gamma) = 0 \tag{5}$$

where:

$$\underset{(H+1)\times(H+1)(l+k)}{Z_t} = \begin{pmatrix} z_t & 0 & \dots & 0 \\ 0 & z_t & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & z_t \end{pmatrix}; \quad \underset{(H+1)\times(H+1)}{X_t} = \begin{pmatrix} x_t & 0 & \dots & 0 \\ 0 & x_t & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & x_t \end{pmatrix};$$

$$\underset{(H+1)\times(H+1)k}{W_t} = \begin{pmatrix} w_t & 0 & \dots & 0 \\ 0 & w_t & \dots & 0 \\ \vdots & \vdots & \dots & \vdots \\ 0 & 0 & \dots & w_t \end{pmatrix}.$$

and

$$\underset{(H+1)\times 1}{\beta} = \begin{pmatrix} \beta_0 \\ \beta_1 \\ \vdots \\ \beta_H \end{pmatrix}; \quad \underset{(H+1)k\times 1}{\gamma} = \begin{pmatrix} \gamma_0 \\ \gamma_1 \\ \vdots \\ \gamma_H \end{pmatrix};$$

Accordingly, notice that $\beta' = (\beta_0, \dots, \beta_H)'$ is the vector of $(H+1)\times 1$ response coefficients; and $\gamma' = (\gamma_0', \dots, \gamma_H)'$ is a $(H+1)k \times 1$ vector of coefficients. The usual GMM objective function based on expression (5) has the usual form given by:

$$\min_{\beta,\gamma} Q_T(Y_t, Z_t, W_t; \beta, \gamma) = \min_{\beta,\gamma} g_T(H)'\, S_T\, g_T(H) \tag{6}$$

where S_T refers to the weighting matrix. Under standard regularity conditions and based on an optimal weighting matrix that adjusts for the serial correlation in $g_T(H)$, one would obtain consistent and asymptotically normal estimates of the responses coefficients, β. For now, we postpone a detailed discussion of the formal results.

Notice that the parameters β are estimated without any cross horizon restrictions. This feature has an advantage and a disadvantage. The advantage, as explained in Jordà (2005), is that local projections generate consistent estimates of the true impulse response coefficients under mild assumptions. The disadvantage is that the increased flexibility makes these estimates less efficient. The next section seeks a middle road.

3.3 Functional Approximation of Impulse Responses

A functional approximation of an impulse response consists in representing the impulse response function as an expansion using basis functions. Different families of basis functions are possible, and in this paper we use Gaussian basis functions and posit

$$\mathcal{R}(h) = \sum_{n=1}^{N} a_n e^{-(\frac{k-b_n}{c_n})^2}, \quad \forall h \geq 0 \tag{7}$$

with a_n, b_n, and c_n parameters to be estimated. S

A first advantage of using Gaussian basis functions is that a small number of Gaussians can already approximate a large class of impulse response functions, in fact most impulse responses encountered in macro applications. Intuitively, impulse response functions are often found (or theoretically predicted) to be monotonic or hump-shaped (e.g., Christiano, Eichenbaum, and Evans, 1999).[11] In such cases, one or two Gaussian basis functions can already provide a very good approximate description of the impulse response, so that Gaussian basis functions offer an "efficient" dimension reduction tool; reducing the number of parameters substantially while still allowing for a large class of impulse responses.

A second advantage of using Gaussian basis functions is that the estimated coefficients can have a direct economic interpretation. This can make the estimation results easier to interpret and use to inform theory and model building.

The ease of interpretation is most salient in a one-Gaussian basis function approximation model with

$$\mathcal{R}(h) \simeq a e^{-\frac{(h-b)^2}{c^2}}, \quad \forall h \geq 0,$$

where the a, b and c coefficients have a direct economic interpretation, and in fact capture three separate

[11] In New-Keynesian models, the impulse response functions are generally monotonic or hump-shaped (see e.g., Walsh, 2010).

characteristics of a hump-shaped impulse response: the peak effect of a shock, the time to peak effect, and the persistence of that peak effect. To see that graphically, the top panel of Figure 3.3 shows a hump-shape impulse response parametrized with a one Gaussian basis function: parameter a is the height of the impulse-response, which corresponds to the maximum effect of a unit shock, parameter b is the timing of this maximum effect, and parameter c captures the persistence of the effect of the shock, as the amount of time τ required for the effect of a shock to be 50% of its maximum value is given by $\tau = c\sqrt{\ln 2}$.

Figure 1: Interpretation of the parameters of a Gaussian basis function

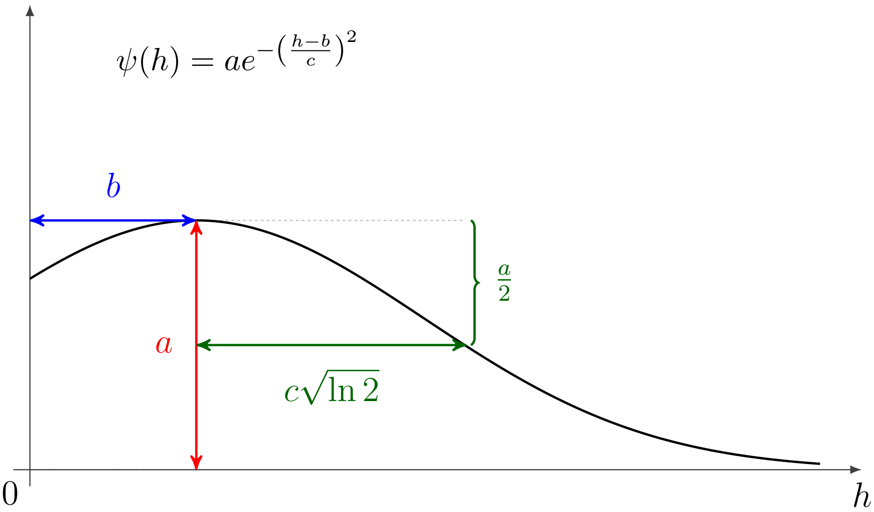

Notes: The parameter a measures the amplitude or maximum response, b measures the time to the maximum, and $c\sqrt{\ln 2}$ measures the half-life from the maximum.

3.4 Combining the parts: Functional Local Projections (FLP)

With a family of basis function in hand, one can easily incorporate the functional approximation into a local projection framework to obtain a flexible yet efficient estimator of responses as defined in expression (3). We will denote this estimator a Functional Local Projection (FLP) and to keep in mind that Gaussian basis function in only one possibly family of basis functions, we will denote FLP with Gaussian basis function as $\text{FLP}_{\mathcal{G}}$.

We now show how the estimation of FLP can be cast as a simple GMM problem.

For clarity, consider an example based on a single Gaussian basis function, that is a $\text{FLP}_{\mathcal{G}}(1)$

$$\mathcal{R}(h, a, b, c) = a \ \exp\left\{\frac{(h-b)}{c}\right\}$$

The simplest approach would be to use a standard two-step minimum estimator based on the GMM estimates derived from expression (6). That is, given $\hat{\beta}$ and its covariance matrix, say, $\hat{\Sigma}_\beta$, then the Gaussian basis function parameters $a,\ ,b$ and c can be estimated by minimizing:

$$\min_{a,b,c} (\hat{\beta} - \mathcal{R}(a, b, c))'\hat{\Sigma}_\beta^{-1}(\hat{\beta} - \mathcal{R}(a, b, c))' \tag{8}$$

under standard regularity conditions and standard minimum distance results, consistent and asymptotically normal estimates for $a,\ ,b$ and c could be easily obtained.

However, a more natural and direct approach is to simply cast the original GMM problem in expression (6) by specifying the sample moment conditions:

$$g_T^*(H) = \frac{1}{T}\sum_{t=1}^{T} Z_t'(y_t - X_t\mathcal{R}(h, a, b, c) - W_t\gamma)$$

as a replacement for condition (5).

[OJ: comment on nuisance parameters issues here]

3.5 Benchmark asymptotic results

Derive standard results, including standard errors corrected for serial correlation.

3.6 Nonlinearities and state-dependent impulse responses

A common assumption in Phillips curve estimates is that of linearity. However, in its original formulation (Phillips, 1958), the Phillips curve was estimated to be highly non-linear, with the slope of the Phillips curve increasing markedly in tight labor markets. While non-linearities in the Phillips curve have been somewhat overlooked since then (partly because the New-Keynesian foundations of the Phillips curve is often presented in its linearized version), a changing trade-off is of central importance of the Fed and its dual mandate. More generally, deliberation among FOMC members often highlight an implicit belief or worry in the existence of non-linearities, as a too tight labor maker is often perceived as running the danger of unleashing inflation (see e.g., Yellen, September 2017).

We now show how FLP allows us to easily generalize our approach to allow for non-linear effects of shocks, notably to allow for the possibility that labor market slack affect the impulse response functions of inflation and unemployment.

Specifically, we let the impulse responses of unemployment and inflation depend on the unemployment gap, $x_t = u_t - u_t^*$, as follows

$$mathcalR(h, x_{t-h}) = (\alpha + \beta x_{t-h})\, e^{-\left(\frac{h-b}{c}\right)^2}, \ \ \forall h > 0 \tag{9}$$

In this specification the amplitude of the impulse response depends linearly on the gap at the time of the shock. That is, the state of the cycle is allowed to stretch/contract the impulse response, but the shape of the impulse response is fixed (because b and c are not functions of x). This is a convenient and intuitive way to explore nonlinearities although others are clearly possible.

4 Measuring the US inflation-unemployment trade-off

4.1 Identification

For our baseline results, our series of aggregate demand shocks is derived as in Romer and Romer (RR, 2004), that is, changes in the fed funds rate orthogonal to the Greenbook/Tealbook staff forecasts. The series was extended to 2007 by Tenreyro and Thwaites (2015) and is available quarterly, 1969–2007.

To estimate impulse responses in response to monetary shocks while allowing for measurement error in the RR shocks series, we run FLP$_{\mathcal{G}}$ with one Gaussian basis function for inflation and unemployment using the RR shocks as instruments. In effect, our approach extends the LP-IV estimation strategy (Jorda and Taylor, 2015, Barnichon and Brownlees 2016, Stock and Watson 2017) to FLP.

4.2 Main results

We study two aspects of κ_h: (1) the dynamics of the π–u trade-off; and (2) whether the trade-off varies as a function of slack, that is whether $\kappa_h = \kappa_h(x)$ —a nonlinear Phillips curve. Think of x here as $u - u^*$, where u^* is an estimate of the natural rate of unemployment.

The dynamics of the π–u trade-off

Figure 2 plots the impulse responses of unemployment and inflation to 1 percentage point shock to i. The impulse response for the funds rate is reported in the appendix in Figure A.1 for both the full and post-1984 samples (we later use the post-1984 sample to evaluate the attenuation of the trade-off). Unemployment rate peaks at nearly 0.4 percentage points 2 years after impact, whereas inflation bottoms out by almost 0.5 percentage points about 3 years after impact. These impulse responses match quantitatively those reported in Coibion (2012), which also uses the Romer and Romer (2004) shocks.

Figure 2: Unemployment and inflation responses to a monetary shock

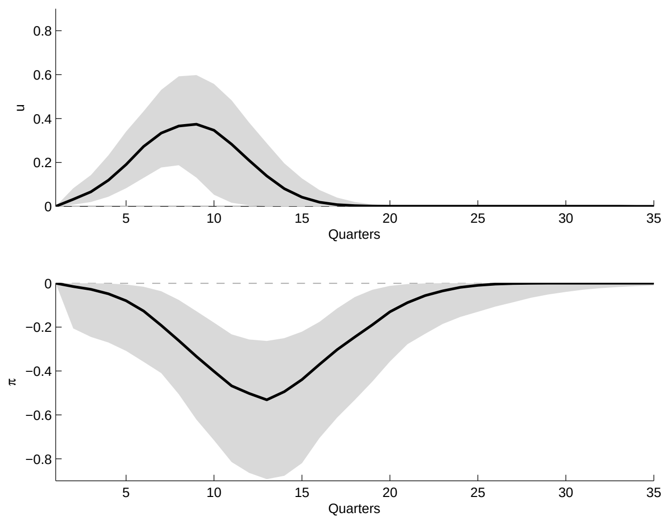

Notes: the sample is 1969q1–2007q4. 90% highest posterior intervals displayed. Inflation measured as annualized quarterly PCE inflation. The unemployment rate is the U3 measure from BLS. The system also includes the federal funds rate. Responses reported in percentage points to a 1% in the funds rate.

Interestingly note how the response of inflation lags the response of unemployment by about one year. Comparing the peak effects of the monetary shocks on inflation and unemployment, we estimate $b_\pi - b_u = 4.2$ quarters (1 year) with a 90% highest posterior interval $[1.7, 7.6]$ quarters, roughly between half a year and two years.

With the impulse responses in hand, we can calculate the inflation-unemployment trade-off κ_h. Figure 3 plots the trade-off κ_h as we increase the horizon h from 0 to 35 quarters. Recall this is the ratio of the cumulative change in inflation over the cumulative change in the unemployment rate. From

Figure 2, note that the impulse response of inflation is initially muted. Hence, over short horizons, 1 ppt lower unemployment translates into only 0.5 ppt. higher inflation ($\kappa_h \approx 0.5$). However, as time goes by and changes accumulate, the trade-off κ_h nearly quadruples from 0.5 ppt to about 2ppt.

An implication of Figure 3 is that, although inflation is quiescent now given how tight the labor market is, it could increase rapidly in the coming years (all else equal).

Figure 3: The dynamic inflation-unemployment trade-off κ_h

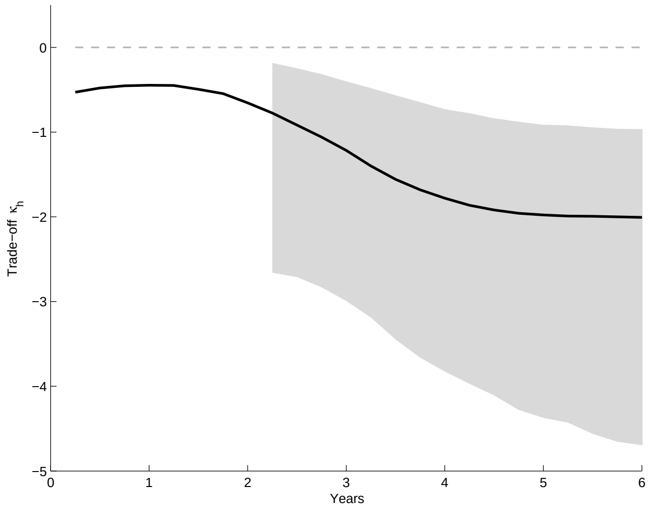

Notes: Romer and Romer narrative monetary shock. Sample: 1969q1–2007q4. 90% highest posterior intervals displayed. Inflation measured as annualized quarterly PCE inflation. The unemployment rate is the U3 measure from BLS. The system also includes the federal funds rate. Bands for small values of h are wide since responses are approximately zero and hence the ratio is indeterminate. They are omitted for clarity.

Non-linearities in the π–u trade-off

Next, we estimate whether the π–u trade-off varies as a function of slack, that is whether $\kappa_h = \kappa_h(x)$ —a nonlinear Phillips curve. Think of x here as $u - u^*$, where u^* is an estimate of the natural rate of unemployment.

To allow for state dependence in the impulse response, we proceed as in 9, and we use the Congressional Budget Office (CBO)'s estimate of the natural rate covering 1949-2017.

Figure 4 plots the unemployment gap. Notice that the US labor market has not experienced very tight conditions since the mid-1960s. Thus, any exercise that ignores this period will likely underestimate

Figure 4: The CBO's unemployment gap $u - u^{*}$, 1949-2017

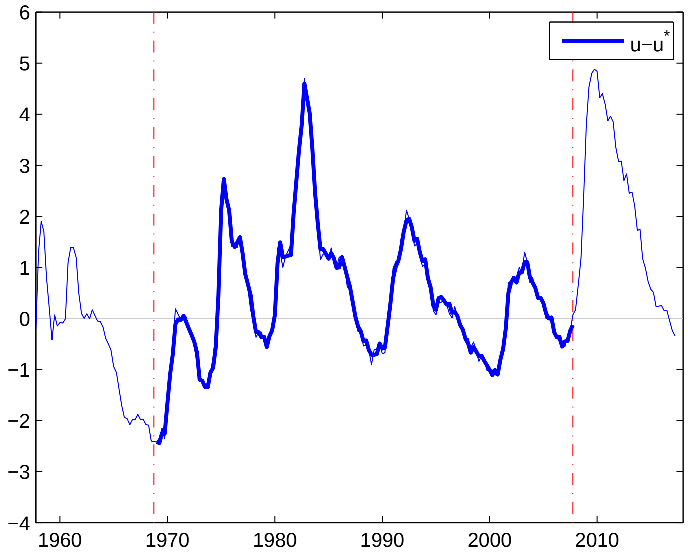

Notes: The dashed red lines denote the estimation sample (1969q1-2007q4) for the FAIR model based on the Romer-Romer narrative monetary shock.

convexities in the inflation-unemployment trade-off.

Our main results are based on the RR shocks, available over 1969-2007, thus missing some of the 1960s. We consider other (AD) shocks that include the 1960s. We show three such results in the Appendix: (i) using monetary shocks identified with a recursive ordering over 1959-2007, and (ii) using government spending shocks identified with a recursive ordering and covering 1949-2015 based on Blanchard and Perotti (2002), and (iii) using news shocks to defense spending and covering 1890-2015 based on Ramey and Zubairy (2016). We find similar nonlinearities with these alternative specifications.

Figure 5 shows $\kappa_H(x)$ as the unemployment gap x increases from -1.5 ppt to $+2$ ppt. The π–u trade-off displays substantial state dependence. When slack is loose, the π–u is close to zero ($\kappa_H \simeq 0$). But κ_H starts deteriorating rapidly as slack tightens. κ_H goes from -2.5 to -4 as the unemployment gap goes from zero to -0.5 percentage point. The deterioration of κ_H in tighter market comes (in roughly equal proportions) from a stronger response of inflation and a weaker response of unemployment.

To put our current labor market situation into perspective, Figure 5 highlights the current level of slack and its implied trade-off. The current gap $u - u^{*}$ is -0.5. A further 0.4 ppt drop in the

Figure 5: The inflation-unemployment trade-off κ_H as a function of $u - u^*$, H=30 quarters.

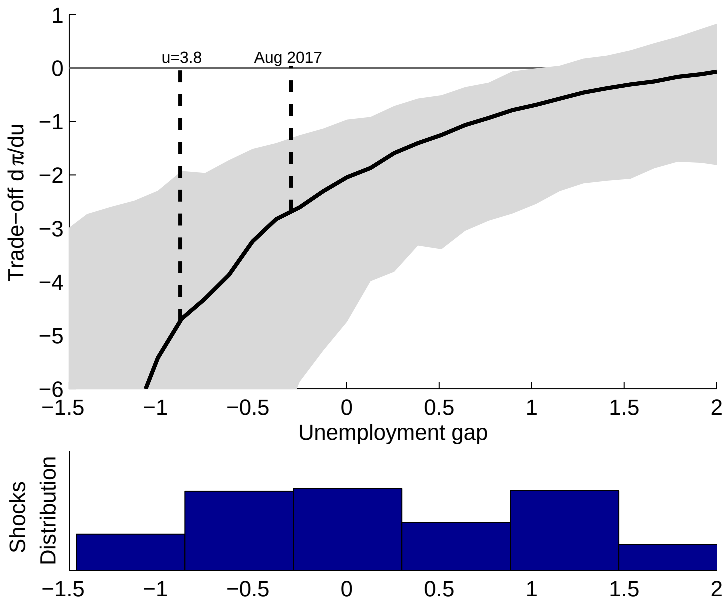

Notes: Top panel: Romer and Romer narrative monetary shock. Sample: 1969q1–2007q4. 90% highest posterior intervals displayed. System includes inflation, the unemployment rate and the federal funds rate. Bottom panel: histogram of the distribution of Romer and Romer shocks across different values of the unemployment gap.

unemployment rate (to the early-2000 low of 3.8 percent) would almost double κ_H.

Figure 6 illustrates the convexity of the π–u in a different manner. For given values of the unemployment gap, the figure shows how the trade-off κ_h varies. The key takeaway is to notice that for low values of h, the trade-off is very similar across labor market states. The difference becomes stark only as h increases.

4.3 Robustness

We investigated whether the properties of the inflation-unemployment trade-off estimated with the Romer and Romer shocks also hold with other aggregate demand shocks. These alternative specifications allow us to expand our sample size and include earlier episodes with tight labor markets.

In addition, we also investigate the sensitivity of our findings to the sample. An extensive literature has found that the Phillips curve has flattened in the past few decades. Hence we consider a shorter sample, 1984–2007.

Figure 6: The inflation-unemployment trade-off κ_h as a function of slack $u - u^*$

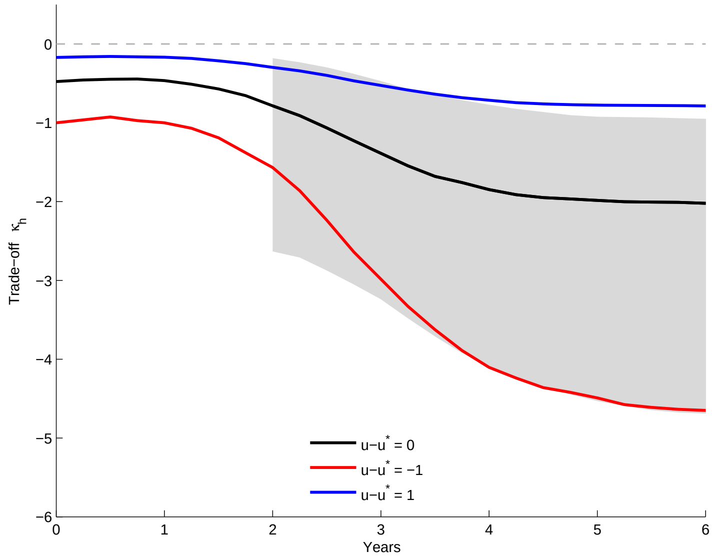

Notes: Romer and Romer narrative monetary shock. Sample: 1969q1–2007q4. 90% highest posterior intervals displayed. Inflation measured as annualized quarterly PCE inflation. The unemployment rate is the U3 measure from BLS. The system also includes the federal funds rate. Bands for small values of h are wide since responses are approximately zero and hence the ratio is indeterminate.

Robustness to alternative aggregate demand shocks

In results reported in the appendix, we replicated our baseline analysis using 3 alternative identification assumptions. In particular:

Monetary shocks estimated using exclusion restrictions (the typical Cholesky identification where the funds rate is ordered last). The sample is extended to 1959-2007 so as to include more observations where the unemployment gap is substantially negative. As with our baseline results, we assess the unemployment gap using the CBO estimate of the natural unemployment rate. The system includes u, π and i. The results are displayed in Figures B.1 to B.4.

Government spending shocks using exclusion restrictions as in Blanchard and Perotti (2002). The sample is 1949q1-2015q4 and the system includes g (government spending), u, π and output y, and we identify government spending shocks by assuming that g reacts with a lag to economic developments. As with our baseline results, we assess the unemployment gap using the CBO estimate of the natural unemployment rate. The impulse responses of inflation and unemployment and the effect of labor market slack on the inflation-unemployment trade-off are displayed in Figure C.1 and C.2).

News shocks to defense spending identified with a narrative approach by Ramey and Zubairy (2016) over 1890-2015. Since the CBO estimate of the natural unemployment rate is not available prior to 1949, we estimate the natural rate from low frequency movements in unemployment using an HP-filter with $\lambda = 10^5$. The impulse responses of inflation and unemployment and the effect of labor market slack on the inflation-unemployment trade-off are displayed in Figure D.1 and D.2).

The post-1984 flattening of the trade-off

We also investigated the sensitivity of our analysis to the sample chosen. Aware of the literature documenting the flattening of the Phillips curve, we repeated the analysis starting the sample in 1984 using RR shocks and recursive identification. The results are reported in Figures E.1 to E.6.

5 Conclusion

TBW

References

Auerbach, Alan J., and Yuriy Gorodnichenko. 2012. "Measuring the Output Responses to Fiscal Policy." *American Economic Journal: Economic Policy*, 4(2): 1–27.

Blanchard, Olivier. 2016. "The Phillips Curve: Back to the '60s?" *American Economic Review*, 106(5): 31–34.

Blanchard, Olivier, and Roberto Perotti. 2002. "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output." *The Quarterly Journal of Economics*, 117(4): 1329–1368.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 1999. "Monetary policy shocks: What have we learned and to what end?" In *Handbook of Macroeconomics.* Vol. 1 of *Handbook of Macroeconomics*, , ed. J. B. Taylor and M. Woodford, Chapter 2, 65–148. Elsevier.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics*, 7(1): 197–232.

Gal, Jordi. 2011. "The Return Of The Wage Phillips Curve." *Journal of the European Economic Association*, 9(3): 436–461.

Jorda, Oscar. 2005. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, 95(1): 161–182.

King, Robert G., and Mark W. Watson. 1994. "The post-war U.S. phillips curve: a revisionist econometric history." *Carnegie-Rochester Conference Series on Public Policy*, 41(1): 157–219.

Mavroeidis, Sophocles, Mikkel Plagborg-Mller, and James H. Stock. 2014. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature*, 52(1): 124–188.

Phelps, Edmund S. 1968. "Money-Wage Dynamics and Labor-Market Equilibrium." *Journal of Political Economy*, 76: 678–678.

Phillips, A. W. 1958. "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica*, 25(100): 283–299.

Ramey, Valerie A., and Sarah Zubairy. 2016. "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data." *Journal of Political Economy*, forthcoming.

Romer, Christina D., and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94(4): 1055–1084.

Samuelson, Paul A, and Robert M Solow. 1960. "Analytical aspects of anti-inflation policy." *The American Economic Review*, 50(2): 177–194.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997*a*. "The NAIRU, Unemployment and Monetary Policy." *Journal of Economic Perspectives*, 11(1): 33–49.

Staiger, Douglas O., James H. Stock, and Mark W. Watson. 1997*b*. "How Precise Are Estimates of the Natural Rate of Unemployment?" In *Reducing Inflation: Motivation and Strategy. NBER Chapters*, 195–246. National Bureau of Economic Research, Inc.

Stock, James H, and Mark W Watson. 2012. "Disentangling the Channels of the 200709 Recession." *Brookings Papers on Economic Activity: Spring 2012*, 81.

Walsh, Carl E. 2010. *Monetary Theory and Policy, Third Edition.* Vol. 1 of *MIT Press Books*, The MIT Press.

A Funds rate responses: 1969–2007 vs. 1984–2007

Figure A.1: Impulse Response Function of the fed funds rate. Romer and Romer identification, 1969q1-2007q4. 90% confidence interval.

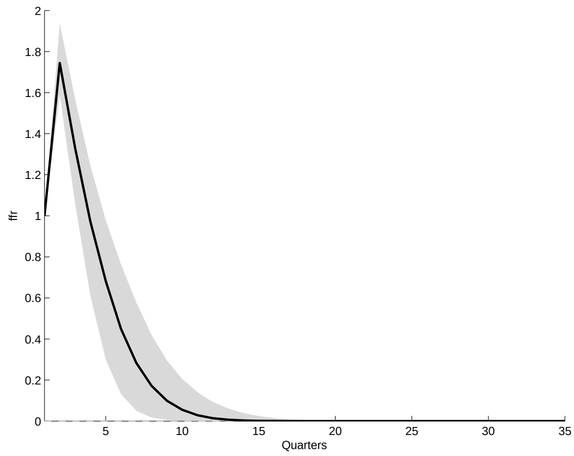

Figure A.2: Impulse Response Function of the fed funds rate. Romer and Romer identification, 1984q1-2007q4. 90% confidence interval.

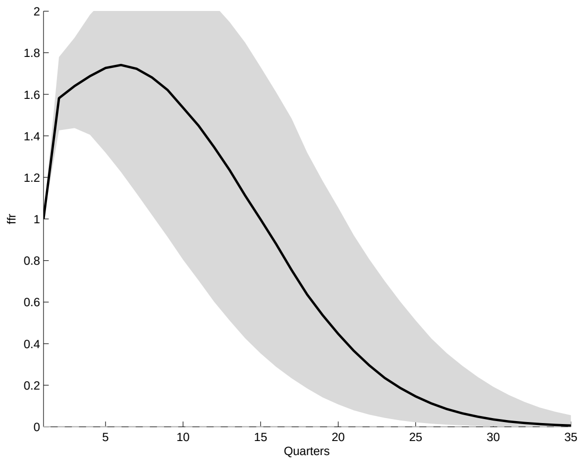

B Monetary shocks using recursive identification

Figure B.1: Impulse Response Functions of unemployment and inflation to a monetary shock. Recursive identification, 1959q1-2007q4. 90% confidence interval.

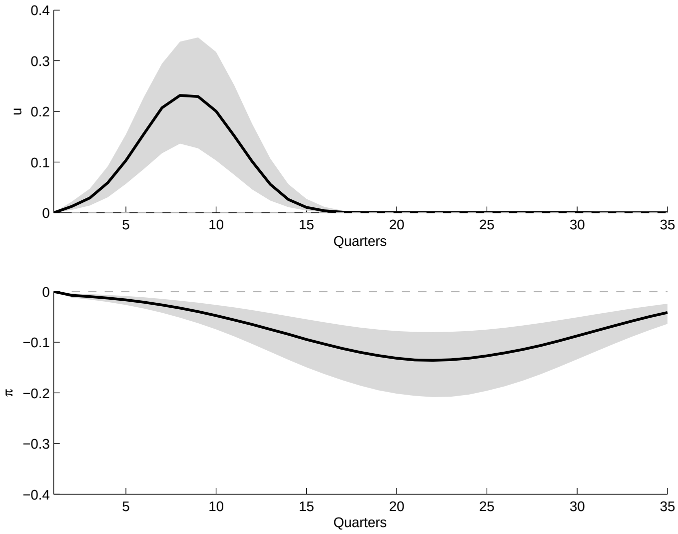

Figure B.2: The dynamic inflation-unemployment trade-off κ_h. Recursive identification of monetary shocks, 1959q1-2007q4. 90% confidence interval.

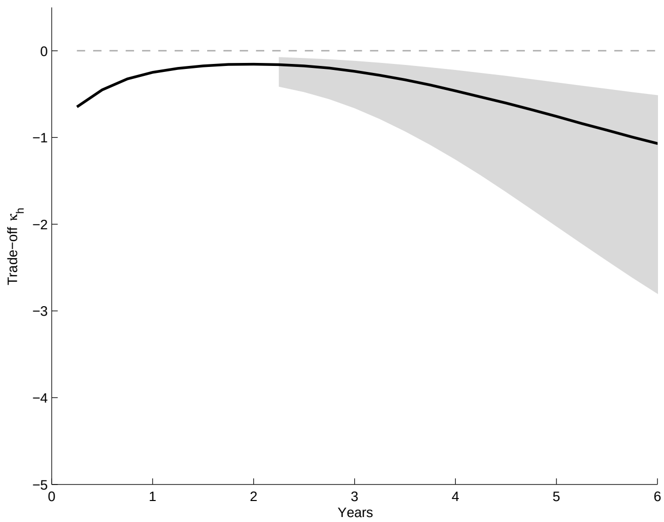

Figure B.3: The inflation-unemployment trade-off κ_H as a function of $u - u^*$, H=30 quarters, 1959q1-2007q4. 90% confidence interval.

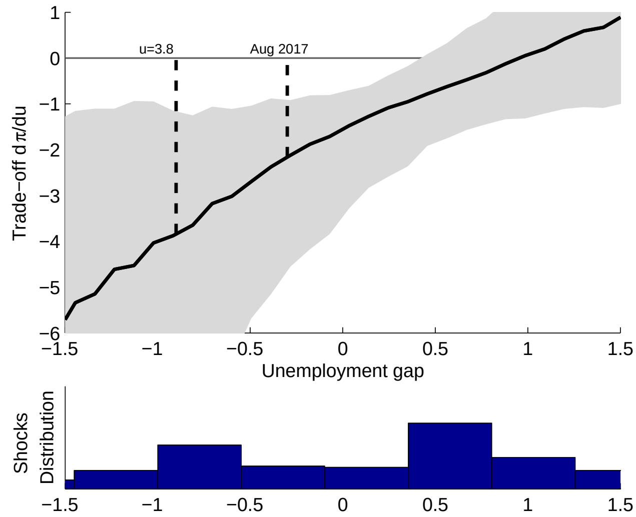

Figure B.4: The dynamic Inflation-Unemployment trade-off κ_h as a function of slack $u - u^*$: Recursive identification, 1959q1-2007q4. 90% confidence interval.

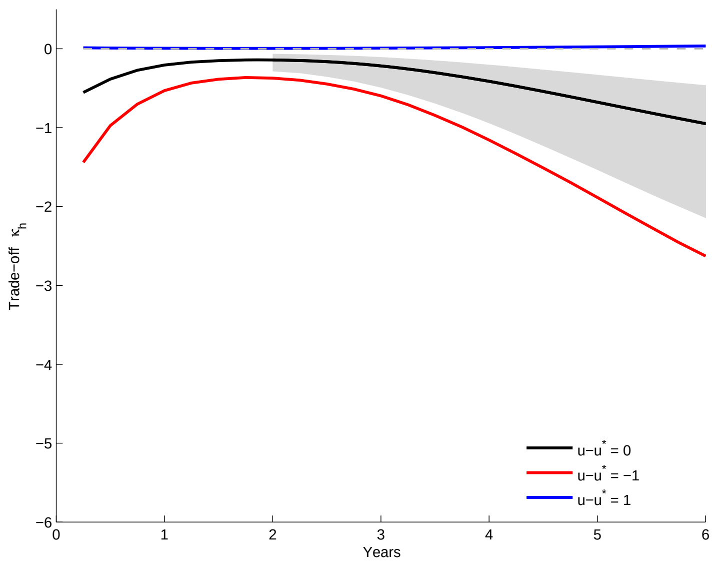

C Blanchard Perotti (2002) Government Spending Shocks

Figure C.1: Impulse Response Functions of unemployment and inflation to a government spending shock. Recursive identification, 1947q1-2015q4. 90% confidence interval.

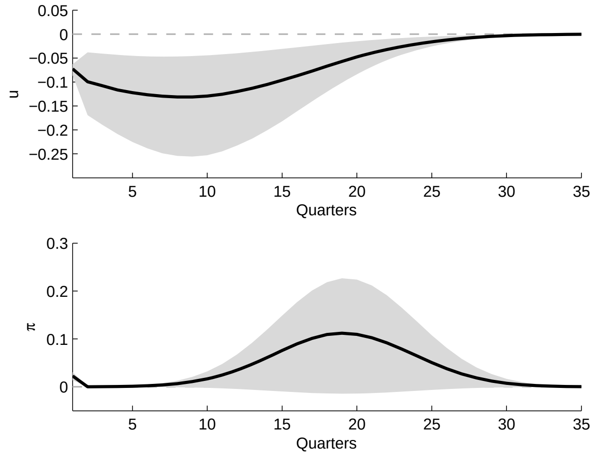

Figure C.2: The inflation-unemployment trade-off κ_H as a function of $u - u^*$, , H=30 quarters.

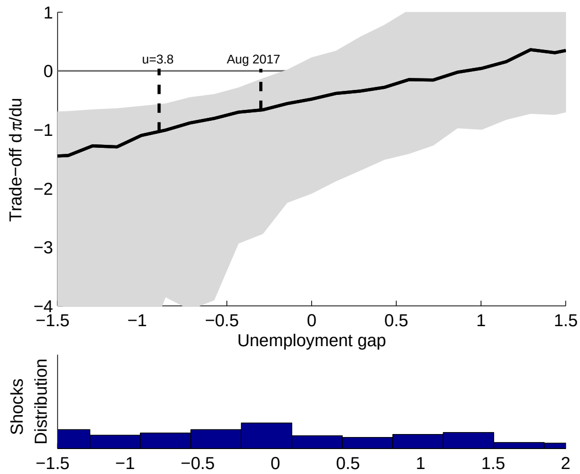

Notes: Top panel: Blanchard and Perotti (2002) recursively identified shocks to government spending. Sample: 1949q1–2015q4. 90% highest posterior intervals displayed. System includes government spending as a ratio to potential GDP, the unemployment rate, inflation, and real GDP as a ratio to potential GDP. Bottom panel: histogram of the distribution of the shocks across different values of the unemployment gap.

D Ramey and Zubairy (2016) news shocks to defence spending

Figure D.1: Impulse Response Functions of unemployment and inflation to a Ramey news shock to government spending. 1890q1-2015q4. 90% confidence interval.

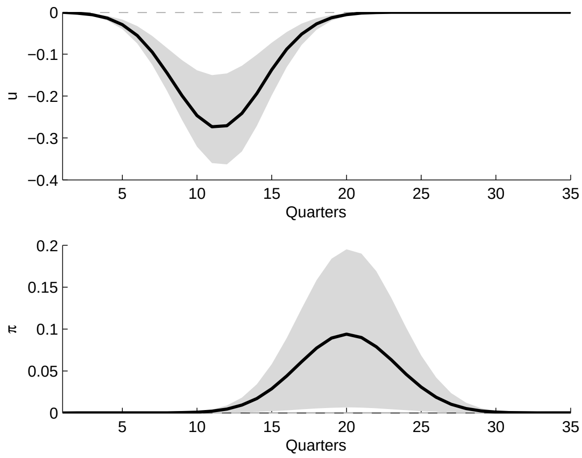

Figure D.2: The inflation-unemployment trade-off κ_H as a function of $u - u^*$, , H=30 quarters

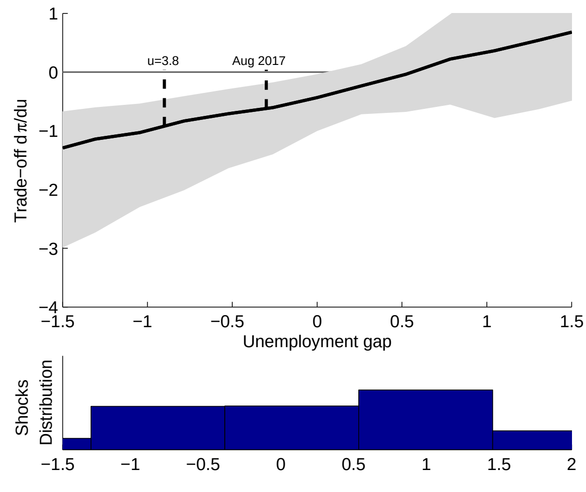

Notes: Top panel: Ramey and Zubairy (2016) news shocks to government spending. Sample: 1890q1–2015q4. 90% highest posterior intervals displayed. System includes inflation, the unemployment rate and the federal funds rate. Bottom panel: histogram of the distribution of the news shocks across different values of the unemployment gap.

E Subsample analysis: Romer and Romer versus recursively identified shocks: 1984–2007

Figure E.1: Unemployment and inflation responses to a 1ppt fed funds rate shock: 1984q1–2015q4. 90% highest posterior intervals displayed. Response scale in ppt.

Romer and Romer

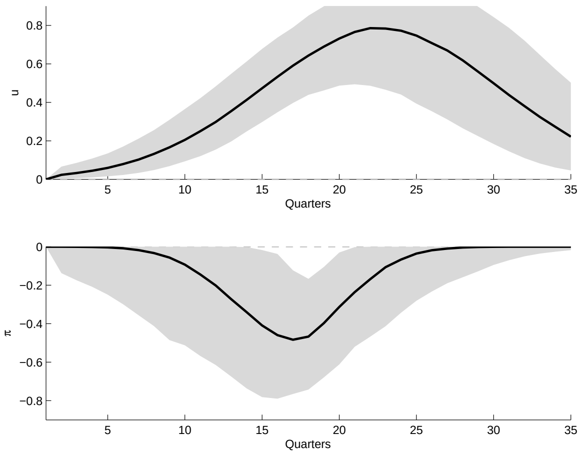

Figure E.2: Unemployment and inflation responses to a 1ppt funds rate shock: 1984q1–2015q4. 90% highest posterior intervals displayed. Response scale in ppt.

Recursive

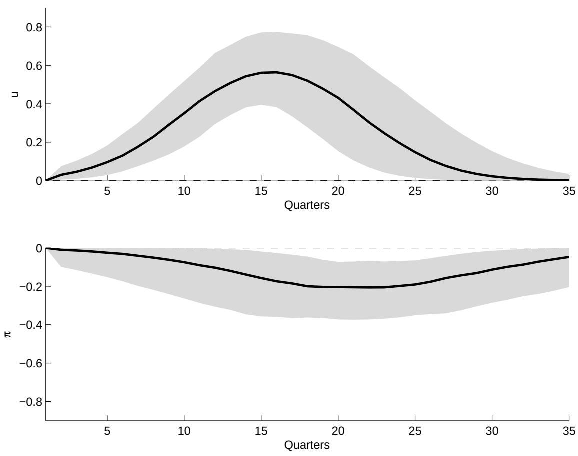

Figure E.3: The dynamic inflation-unemployment trade-off κ_h: 1984q1–2015q4. 90% highest posterior intervals displayed. Full sample (1969-2007) trade-off displayed as a red-dashed line.

Romer and Romer

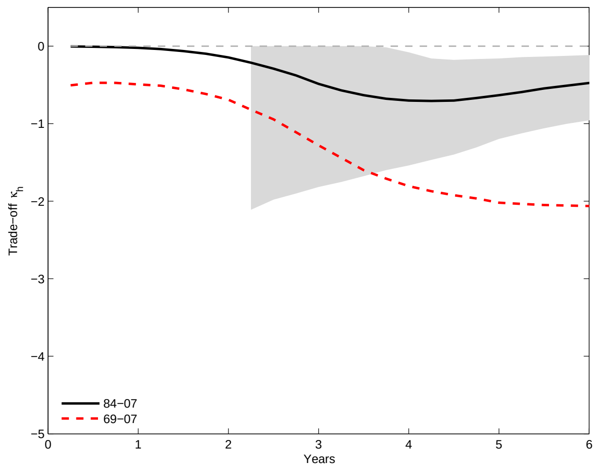

Figure E.4: The dynamic inflation-unemployment trade-off κ_h: 1984q1–2015q4. 90% highest posterior intervals displayed. Full sample (1969-2007) trade-off displayed as a red-dashed line.

Recursive

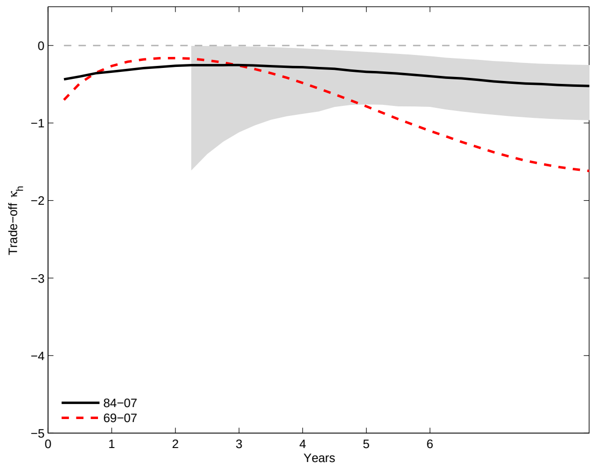

Figure E.5: The inflation-unemployment trade-off κ_H as a function of $u - u^*$, H = 30 quarters. Sample: 1984q1–2015q4. 90% highest posterior intervals displayed

Romer and Romer

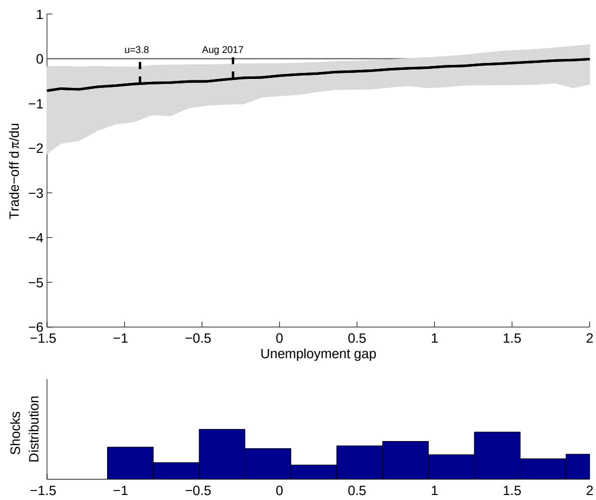

Figure E.6: The inflation-unemployment trade-off κ_H as a function of $u - u^*$, H = 30 quarters. Sample: 1984q1–2015q4. 90% highest posterior intervals displayed

Recursive