10/25/2017
Internal FR

Upside risks to the inflation forecast

Regis Barnichon and Òscar Jordà

There is a risk that the labor market could eventually become over-heated, potentially creating an inflationary problem down the road that might be difficult to overcome [...]

Chair Yellen. September 26, 2017

Summary

- We highlight two upside risks to the inflation forecast:
  1. In response to aggregate demand shocks, inflation historically reacts with about a one-year lag relative to the labor market.
  2. Nonlinearities triggered when the unemployment rate exceeds its natural rate can accentuate the inflation response once it begins.

- Thus, while inflation may currently appear to drift slowly toward target, it could start increasing much more rapidly than current forecasts anticipate.

- The inflation/unemployment tradeoff has been traditionally described with the Phillips curve. However, estimating this relationship empirically is complicated by several issues: (1) lack of consensus on its specification; (2) counfounding from supply factors; (3) and issues related to the measurement of inflation expectations.

- Our estimate of the inflation/unemployment tradeoff circumvents these issues:
  1. We use aggregate demand shocks to identify the inflation-unemployment tradeoff.
  2. We account for the complex dynamics in inflation and unemployment by using a semi-parametric approach based on impulse response analysis.

- Our results are surprisingly robust to using a variety of demand shocks, such as Romer and Romer (2004) and recursively identified monetary shocks; as well as fiscal shocks based on Blanchard and Perotti's (2002) identification assumptions as well as Ramey and Zubairy's (2016) news shocks to defense spending.

- Subsample analysis focusing on the 1984-2007 period indicates an attenuation of the inflation/unemployment tradeoff 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.

1 Overview: Objectives and Approach

- The goal is to estimate the monetary-policy relevant inflation-unemployment (π–u) tradeoff.

- To this end, we take a new approach based on Functional Approximation of Impulse Responses (FAIR), by Barnichon and Matthes (2016). We directly estimate the responses of inflation and the unemployment rate

- Our series of aggregate demand shocks is derived as in Romer and Romer (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.

- As a robustness check, we conducted several other experiments with alternative aggregate demand shocks: (1) identification of monetary policy based on exclusion restrictions; (2) Blanchard and Perotti (2002) government spending shocks based on exclusion restrictions; and (3) news shocks to defense as recorded in Ramey and Zubairy (2016). The main results hold across all these alternatives. These results are reported in the appendix.

- 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 tradeoff with the ratio of monetary "multipliers" (where the monetary response is common to numerator and denominator and hence cancels out):

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

κ_h captures, in response to a common 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.

- We argue that κ_h is precisely the object of interest for the Fed: If the Fed changes monetary policy, how much will it respectively affect its two mandates, price stability and full employment? Similarly, how does an aggregate demand shock shift inflation and the unemployment rate away from equilibrium?

- Our multiplier 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 semiparametric framework, we think that we have a more accurate measurement of the π–u tradeoffs 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. Ball (1994), for example, defines it as the cumulative output/unemployment deviation from trend of a disinflationary episode defined as a change in trend inflation over a given period of time. To avoid confusion, we use the term tradeoff, which is less fraught.

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

- Hence we find that the unemployment rate responds faster than inflation, and a tighter labor market translates into higher inflation with about a 1-year delay. As a result, the tradeoff κ_h is initially small, but then increases markedly (in absolute value) as h grows over time.

- All else equal, we find that the tradeoff κ_h increases further (in absolute value) in tighter labor markets. Given that the current (CBO-based) gap $u - u^*$ is at -0.5, the tradeoff could worsen very quickly: a further 0.4 percentage point drop in the unemployment rate would almost double κ_h.

2 Functional Approximation of Impulse Responses FAIR

- Based on the intuition behind the Wold decomposition, any covariance-stationary vector time series process $\mathbf{y}_t$ can be characterized using the following structural vector moving average (SVMA) process (the constant is omitted for convenience):

$$\mathbf{y}_t = \sum_{h=0}^{H} \mathbf{\Psi}_h \varepsilon_{t-h} \tag{2}$$

  where ε_t is the vector of i.i.d. structural innovations with $E(\varepsilon_t) = \mathbf{0}$ and $E(\varepsilon_t \varepsilon_t') = \mathbf{I}$, H is the number of lags, which can be finite or infinite. $\mathbf{\Psi}_h$ is the matrix of lag coefficients, i.e., the matrix of impulse response coefficients at horizon h. We denote $\psi_j(h)$ the row-vector of $\mathbf{\Psi}_h$ corresponding to the j^{th} variable.

- Next, we represent each impulse response function as an expansion using basis functions, and we approximate $\psi_j(h)$ with one Gaussian basis function (more terms could be used but we find that one basis function approximates well the shapes typically encountered in practice), so that

$$\psi_j(h) \simeq a_j \exp\left\{-\frac{(h-b_j)^2}{c_j^2}\right\}, \quad \forall h > 0;\ j = \text{inflation, unemployment.} \tag{3}$$

- The SVMA can be estimated directly by maximum likelihood or using Bayesian methods (the approach used here), as long as one truncates the order of the MA structure. Theoretically, covariance-stationarity justifies consistency even after truncating if more terms are included as the sample size expands. The likelihood function can then be specified recursively for every observation in the sample.

- The FAIR approach thus straddles between the parametric parsimony of VARs and the flexibility of local projections. Moreover, nonlinearities are easily handled as we will show momentarily.

- The a, b and c coefficients in expression (3) have a direct economic interpretation as shown in Figure 1: the parameter a captures the height of the impulse-response, which corresponds to the maximum effect of a unit shock, the parameter b captures the timing of this maximum effect, and the parameter c captures the persistence of the effect of the shock—the amount of time τ required for the effect of a shock to be 50% of its maximum value (a half-life measure) is given by $\tau = c\sqrt{\ln 2}$.

Figure 1: Interpretation of the parameters in a Gaussian basis function (FAIR approach)

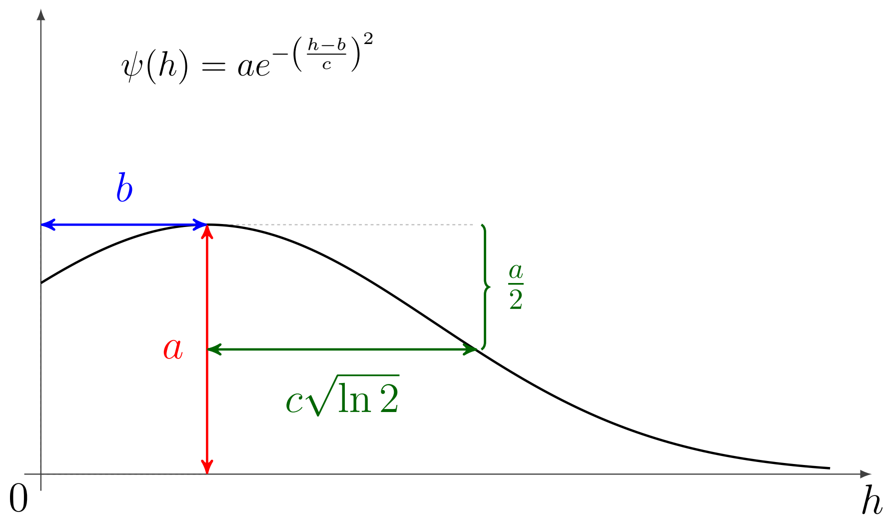

Notes: The parameter a measures the amplitude or maximum response, b measures the time to the maximum, and $c\sqrt{\ln 2}$ measures the half-life from the maximum.

3 Baseline results: Romer and Romer 2004 shocks

- Consider a system that includes the unemployment rate, u, measured by the usual U3 measure from the BLS; the annualized quarterly rate of PCE inflation, π; and the federal funds rate, i.
- We order the narratively identified Romer and Romer (RR) shocks series first and assume that they do not react contemporaneously to any other shocks. In other words, we are only assuming that the RR shocks are contemporaneously correlated with the true monetary shocks and are uncorrelated with other structural shocks.
- Our procedure allows the narrative shocks to contain measurement error, as long as the measurement error is independent of the structural shocks.
- 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 tradeoff). 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.

- 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.

- Next, Figure 3 plots the tradeoff 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

Figure 2: Unemployment and inflation responses to a monetary shock

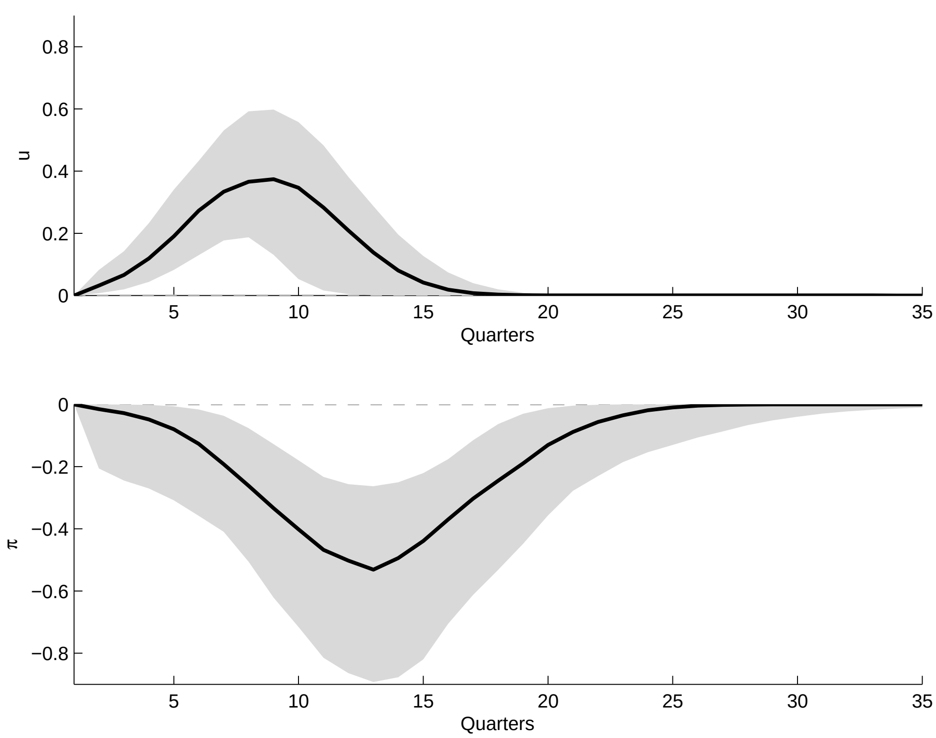

Notes: the sample is 1969q1–2007q4. 90% highest posterior intervals displayed. Inflation measured as annualized quarterly PCE inflation. The unemployment rate is the U3 measure from BLS. The system also includes the federal funds rate. Responses reported in percentage points to a 1% in the funds rate.

into only 0.5 ppt. higher inflation ($\kappa_h \approx 0.5$). However, as time goes by and changes accumulate, the tradeoff 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).

4 The inflation–unemployment tradeoff is nonlinear

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

$$\psi(x_{t-h}) = (\alpha + \beta x_{t-h})\, e^{-\left(\frac{h-b}{c}\right)^2}, \ \ \forall h > 0 \tag{4}$$

Figure 3: The dynamic inflation-unemployment tradeoff κ_h

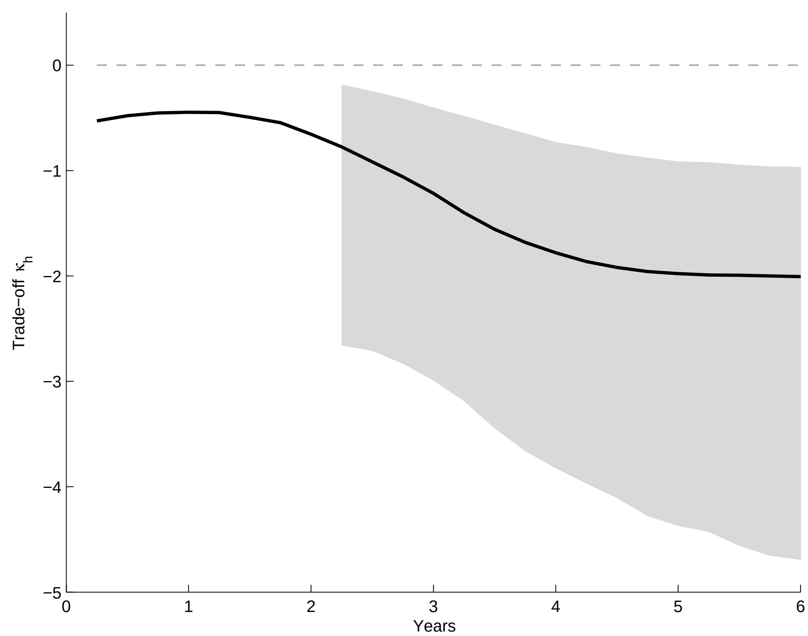

Notes: Romer and Romer narrative monetary shock. Sample: 1969q1–2007q4. 90% highest posterior intervals displayed. Inflation measured as annualized quarterly PCE inflation. The unemployment rate is the U3 measure from BLS. The system also includes the federal funds rate. Bands for small values of h are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

- 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.

Results

- 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 convexities in the inflation-unemployment tradeoff.
- Our main results are based on the RR shocks, available over 1969-2007, thus missing

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

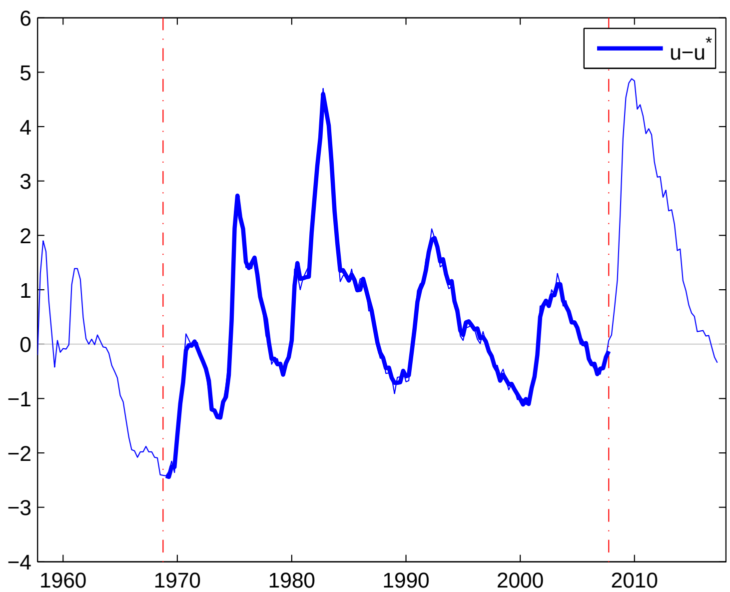

Notes: The dashed red lines denote the estimation sample (1969q1-2007q4) for the FAIR model based on the Romer-Romer narrative monetary shock.

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

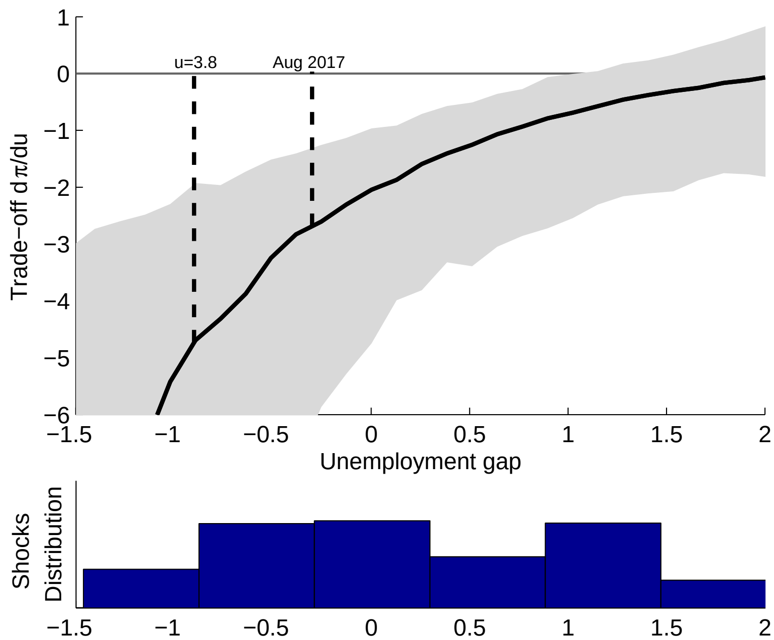

Notes: Top panel: Romer and Romer narrative monetary shock. Sample: 1969q1–2007q4. 90% highest posterior intervals displayed. System includes inflation, the unemployment rate and the federal funds rate. Bottom panel: histogram of the distribution of Romer and Romer shocks across different values of the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

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 tradeoff 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.

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

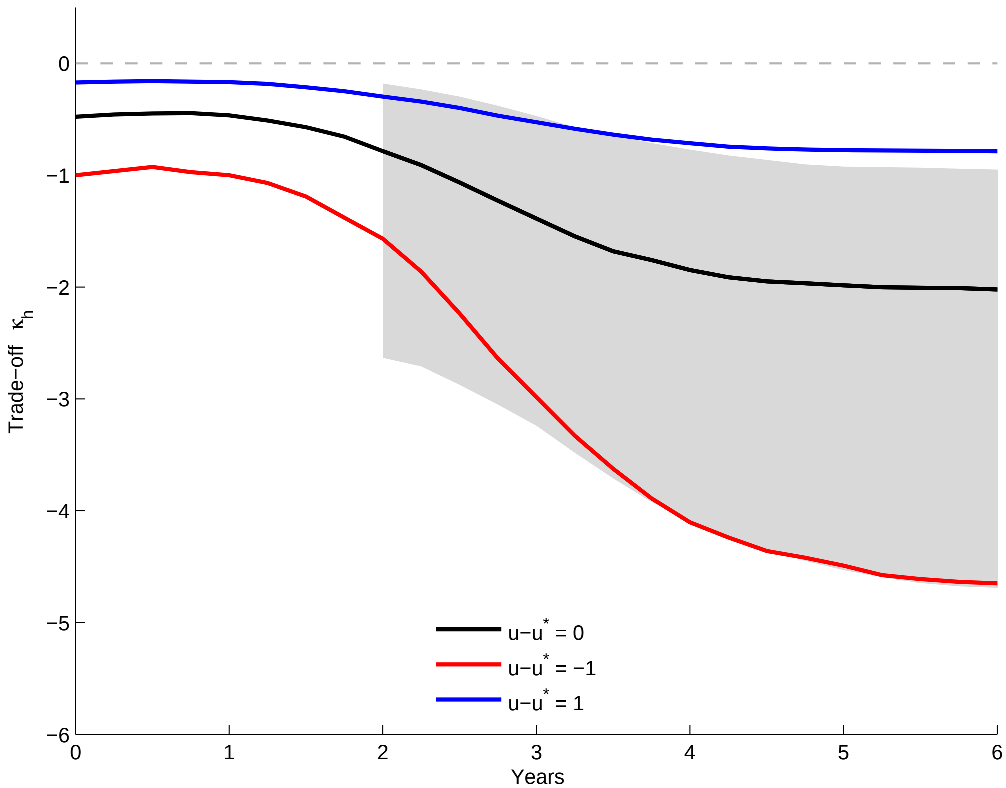

Notes: Romer and Romer narrative monetary shock. Sample: 1969q1–2007q4. 90% highest posterior intervals displayed. Inflation measured as annualized quarterly PCE inflation. The unemployment rate is the U3 measure from BLS. The system also includes the federal funds rate. Bands for small values of h are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

- To put our current labor market situation into perspective, Figure 5 highlights the current level of slack and its implied tradeoff. The current gap $u - u^*$ is -0.5. A further 0.4 ppt drop in the 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 tradeoff κ_h varies. The key takeaway is to notice that for low values of h, the tradeoff is very similar across labor market states. The difference becomes stark only as h increases.

5 Robustness checks

- We investigated whether the properties of the inflation-unemployment tradeoff 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.

5.1 Robustness to alternative aggregate demand shocks

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

  1. 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.
  2. 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 tradeoff are displayed in Figure C.1 and C.2).
  3. 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 tradeoff are displayed in Figure D.1 and D.2).

5.2 The flattening of the tradeoff

- In addition to these experiments, 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.

6 Implications for policy

- While currently in a low inflation/low unemployment environment, we highlighted two upside risks to the inflation forecast
  1. Inflation reacts with a one-year lag to labor market slack
  2. The inflation/unemployment tradeoff can increase rapidly (in absolute value) as the labor market tightens further

- To put some numbers of these two upside risks, consider the following two experiments based on our baseline estimation results:

  1. **The delayed response of inflation**
     Experiment #1: *An aggregate demand shock hit the US economy* ***a year ago*** *to lower unemployment from 4.7% in 2016Q3 to 3.8% in 2018Q4. PCE inflation is initially at 1.6%.*
     Given the gap then (0%) and abstracting from other factors, our results imply that unemployment would be at 4.2% today but inflation would have barely moved up (1.7%), –roughly consistent with what we see today–. However, with the dynamics of inflation in motion, we would soon see inflation starting to rise quickly and reach 2.1% by 2018Q4.
  2. **The non-linear effect of slack**
     Experiment #2: *An aggregate demand shock hits the US economy* ***today*** *and lowers unemployment from 4.2 in 2017Q4 to 3.8 % in 2019Q4. PCE inflation is initially at 1.5%.*
     Given the current gap (-0.5 %) and abtracting from other factors, our results imply that inflation would not react initially, but then it would rise faster and reach 2.4% in 2020Q4. Note that the non-linear effect of slack is crucial here: if we had done the same experiment with a gap of +0.5%, inflation would be only expected to increase to 1.8% in 2020Q4.

- Thus in the current environment of subdued inflation, there is a risk of inflation flaring up over time, all else equal.

- Nevertheless, the convexity of the tradeoff to a tight labor market appears to be somewhat more subdued in recent times. This statement comes with the warning that the recent post-1984 sample does not have data where the unemployment rate fell below the natural rate for a prolonged period of time as it did in the mid-1960s.

- Moreover, our sample (for the main Romer and Romer based results) ends in 2007 and therefore does not include data post-financial crisis. It is possible that the dynamics of inflation have changed in ways that our analysis has failed to capture. That said, the relative stability of the tradeoff across different samples and different demand shocks provides some indication that the fundamental relationship between inflation and unemployment may not have changed very much.

- All of these statements are *ceteris paribus*, meaning, that we have not factored in supply-side factors that may shift (usually temporarily) the entire tradeoff curve, much as they shift the Phillips curve.

10/25/2017
Internal FR

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.

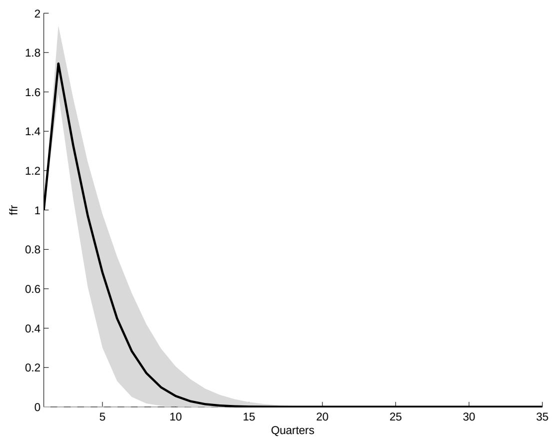

Figure A.2: Impulse Response Function of the fed funds rate. Romer and Romer identification, 1984q1-2007q4. 90% confidence interval.

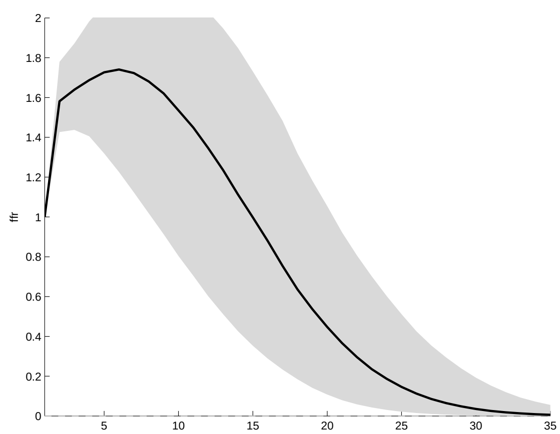

B Monetary shocks using recursive identification

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

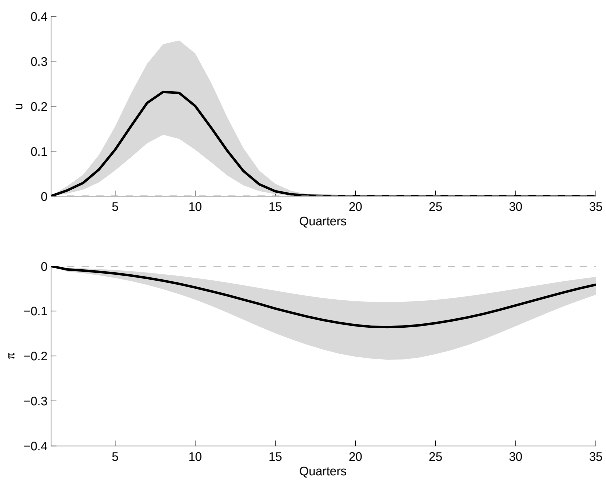

Figure B.2: The dynamic inflation-unemployment trade-off κ_h. Recursive identification of monetary shocks, 1959q1-2007q4. 90% confidence interval.

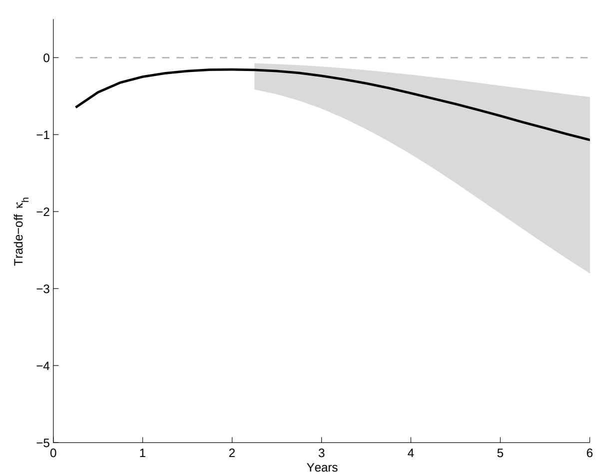

10/25/2017
Internal FR

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

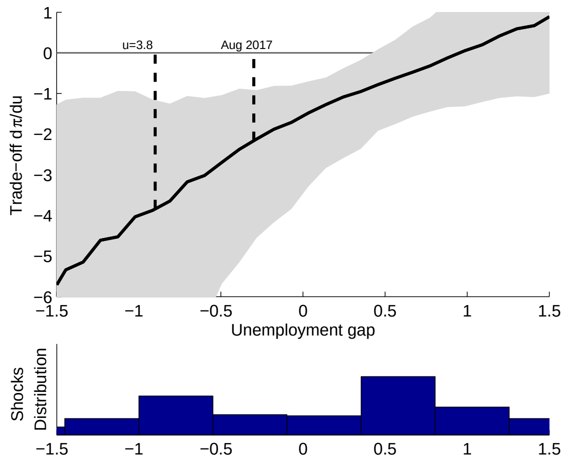

Figure B.4: The dynamic Inflation-Unemployment trade-off κ_h as a function of slack $u - u^*$: Recursive identification, 1959q1-2007q4. 90% confidence interval.

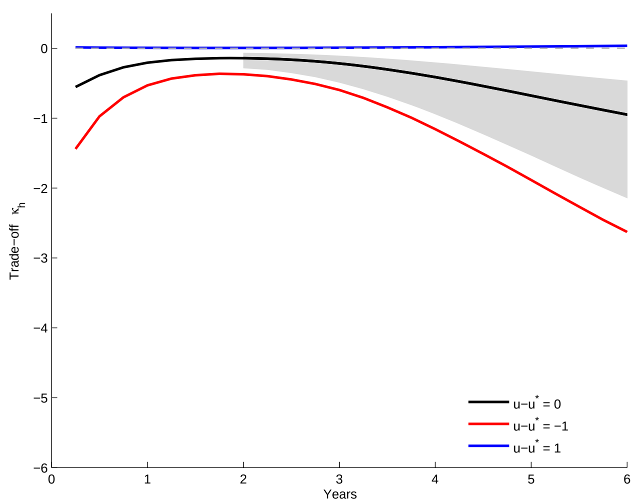

C Blanchard Perotti (2002) Government Spending Shocks

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

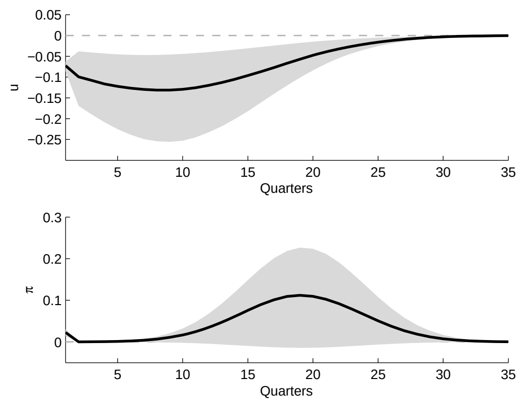

Figure C.2: The inflation-unemployment tradeoff κ_H as a function of $u - u^*$, , H=30 quarters.

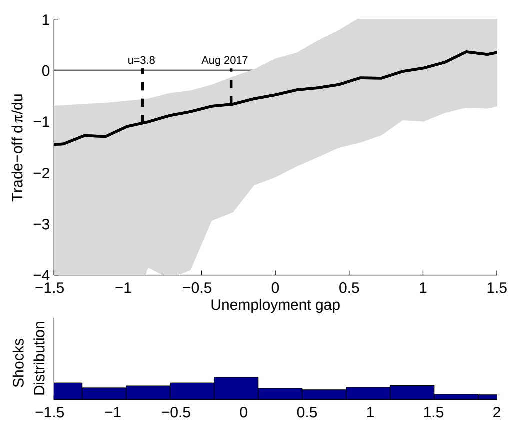

Notes: Top panel: Blanchard and Perotti (2002) recursively identified shocks to government spending. Sample: 1949q1–2015q4. 90% highest posterior intervals displayed. System includes government spending as a ratio to potential GDP, the unemployment rate, inflation, and real GDP as a ratio to potential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

10/25/2017
Internal FR

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

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

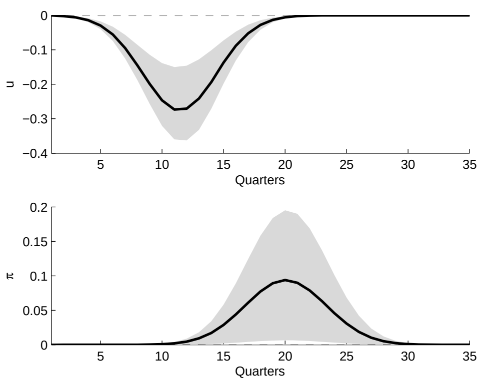

Figure D.2: The inflation-unemployment tradeoff κ_H as a function of $u - u^*$, , H=30 quarters

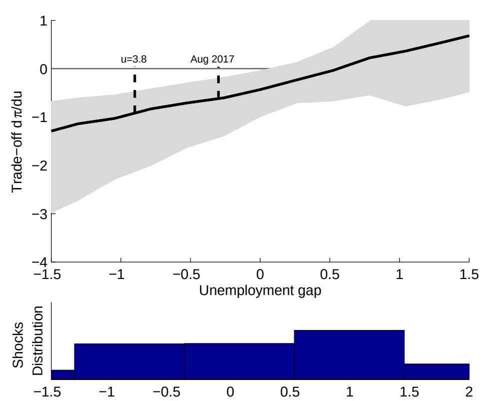

Notes: Top panel: Ramey and Zubairy (2016) news shocks to government spending. Sample: 1890q1–2015q4. 90% highest posterior intervals displayed. System includes inflation, the unemployment rate and the federal funds rate. Bottom panel: histogram of the distribution of the news shocks across different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

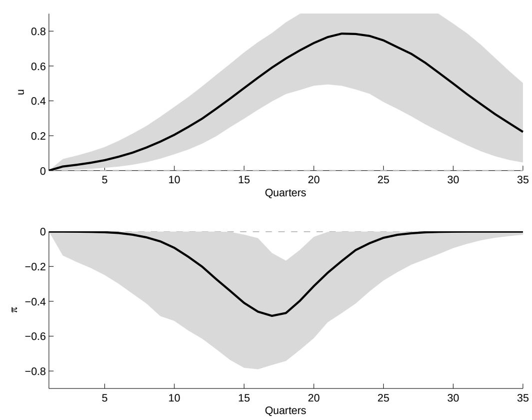

Figure E.2: Unemployment and inflation responses to a 1ppt funds rate shock: 1984q1–2015q4. 90% highest posterior intervals displayed. Response scale in ppt.

Recursive

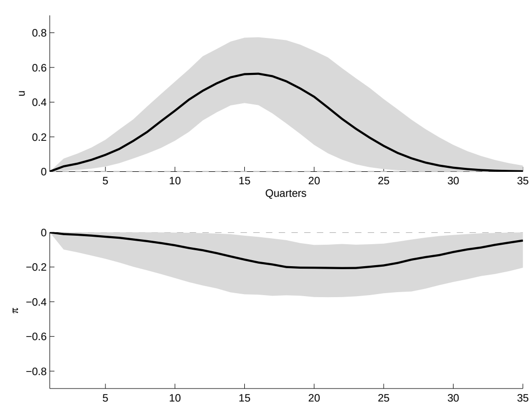

10/25/2017
Internal FR

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

Romer and Romer

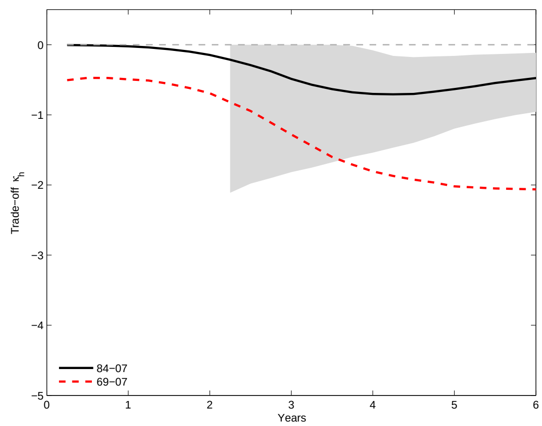

Figure E.4: The dynamic inflation-unemployment tradeoff κ_h: 1984q1–2015q4. 90% highest posterior intervals displayed. Full sample (1969-2007) tradeoff displayed as a red-dashed line.

Recursive

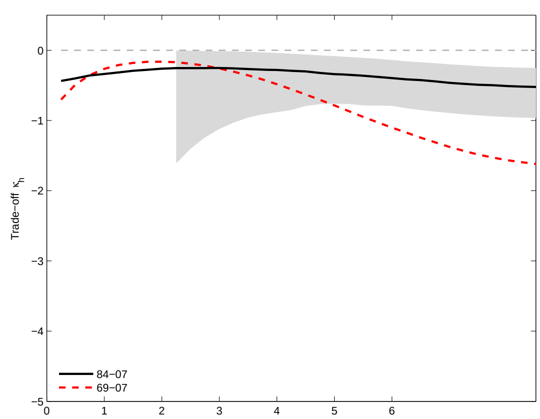

Figure E.5: The inflation-unemployment tradeoff κ_H as a function of $u - u^*$, H = 30 quarters. Sample: 1984q1–2015q4. 90% highest posterior intervals displayed

Romer and Romer

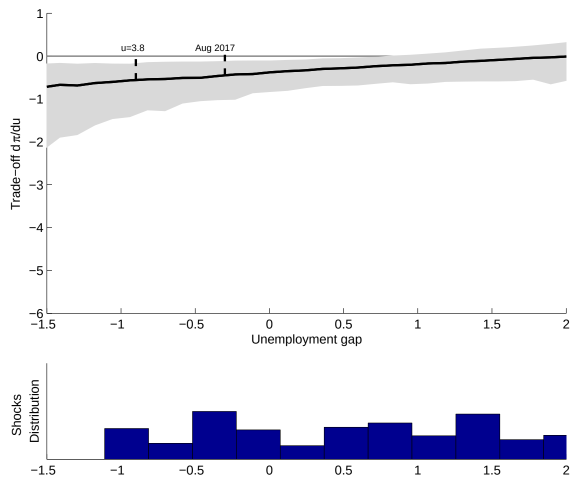

Figure E.6: The inflation-unemployment tradeoff κ_H as a function of $u - u^*$, H = 30 quarters. Sample: 1984q1–2015q4. 90% highest posterior intervals displayed

Recursive

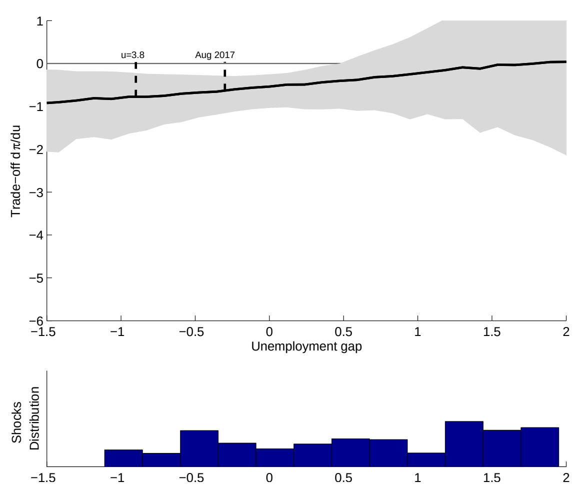

10/25/2017
Internal FR

F Assessing the significance of the effect of slack on κ_H

For each identification scheme considered in this briefing, Figure F.1 plots the joint marginal posterior density of $\kappa_H(x=-1)$ –a tight labor market (y-axis)– and $\kappa_H(x=1)$ –a slack labor market (x-axis)–. The dashed red line denotes identical κ_H across labor market states. In all cases, more than 95 percent of the mass lies below the 45 degree line, indicating that the inflation/unemployment tradeoff is singificantly larger in tight labor markets than in slack labor markets.

Figure F.1: Testing the significance of the non-linearity

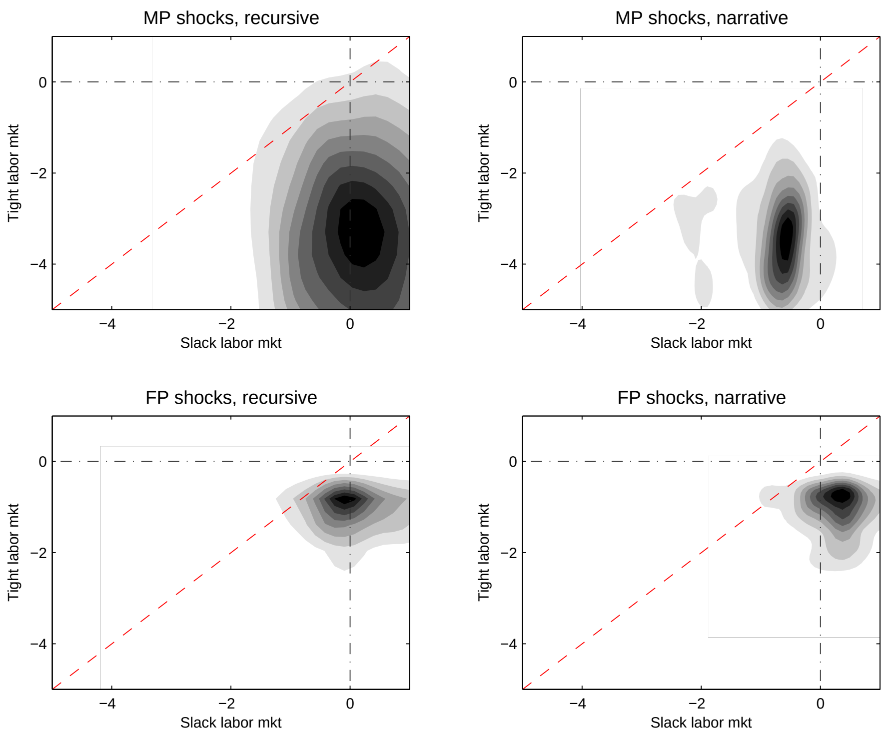

Notes: Joint marginal posterior distribution of $\kappa_H(x=-1)$ –tight labor market– and $\kappa_H(x=1)$ –slack labor market–. The dashed red line denotes no effect of slack on κ_H.

Bibliography

[1] Blanchard, O., and R. Perotti. "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output." Quarterly Journal of Economics 117 (4): 1329–68, 2002

[2] Coibion, Olivier. "Are the Effects of Monetary Policy Shocks Big or Small?. " American Economic Journal: Macroeconomics 4(2): 1-32, 2012

[3] Ramey V. and S. Zubairy. "Government Spending Multipliers in Good Times and in Bad: Evidence from U.S. Historical Data ," Journal of Political Economy, Forthcoming, 2016

[4] Romer, C., and D. Romer. "A New Measure of Monetary Shocks: Derivation and Implications," American Economic Review 94 (4): 1055–84, 2004

[5] Tenreyro, S. and G. Thwaites."Pushing on a String: US Monetary Policy Is Less Powerful in Recessions." American Economic Journal: Macroeconomics, 8(4): 43-74, 2016