

CAMA

Centre for Applied Macroeconomic Analysis

The Evolution of the U.S. Output-Inflation Tradeoff

CAMA Working Paper 70/2013
October 2013

Benjamin Wong

Centre for Applied Macroeconomic Analysis (CAMA), ANU

Abstract

This paper proposes quantifying the evolution of the U.S. output-inflation tradeoff using a Time-Varying Parameter Structural VAR. This methodology circumvents issues with existing methods which tend to be either reduced form in nature or rely on more ad hoc assumptions regarding sample split dates and both trend output and trend inflation. Working through U.S. data since the 1970s reveals only a slight change in the tradeoff from around 1.70 to 1.75 percentage points of real output growth per percentage point increase in trend inflation. This contrasts with claims that the U.S. Phillips Curve has flattened dramatically.

Keywords

Sacrifice Ratio, Time-Varying Parameters

JEL Classification

C32, C54, E31, E66

Address for correspondence:

(E) cama.admin@anu.edu.au

[The Centre for Applied Macroeconomic Analysis](#) in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

The Evolution of the U.S. Output-Inflation Tradeoff

Benjamin Wong^{*†}

Abstract

This paper proposes quantifying the evolution of the U.S. output-inflation tradeoff using a Time-Varying Parameter Structural VAR. This methodology circumvents issues with existing methods which tend to be either reduced form in nature or rely on more *ad hoc* assumptions regarding sample split dates and both trend output and trend inflation. Working through U.S. data since the 1970s reveals only a slight change in the tradeoff from around 1.70 to 1.75 percentage points of real output growth per percentage point increase in trend inflation. This contrasts with claims that the U.S. Phillips Curve has flattened dramatically.

Keywords: Sacrifice Ratio, Time-Varying Parameters

JEL Classification: C32, C54, E31, E66

^{*}Centre for Applied Macroeconomic Analysis, The Australian National University. Canberra, ACT 0200, Australia. Email: ben.wong@anu.edu.au. Tel: (+61) 2 612 55183.

[†]I thank Renée Fry-McKibbin, Warwick McKibbin and Shaun Vahey for their guidance, Punnoose Jacob, Parvinder Kler, James Morley, Dan Thornton, Varang Wiriyawit, participants of the Econometric Society Australasia Meeting 2013 and CAMA Brownbag seminar at The Australian National University for extensive suggestions, helpful comments and discussions and the kind generosity of Gary Koop, Dimitri Korobolis and Andreas Tudyka, who provided code which I have adapted in parts. I acknowledge financial support from the Australian Government's International Postgraduate Research Scholarship and ARC Discovery Project DP120103443. Errors and omissions remain my own.

1 Introduction

The tradeoff between output and inflation is of key policy consideration. Given demand management policies induce output and inflation to move in the same direction, the relative benefits of demand policies depend crucially on the extent of this tradeoff. This tradeoff is essentially the slope of the short-run Phillips Curve.¹ There is enough evidence that policymakers regard this tradeoff has altered considerably over time and this change has ignited policy interest (see, for example, Yellen, 2006; Mishkin, 2007). This evolving tradeoff has non trivial policy implications. Some implications of a flatter Phillips Curve include the greater possibility of committing policy mistakes, with these errors also being more costly to repair (Yellen, 2006). Therefore, it is of interest to evaluate whether the Phillips Curve has indeed flattened and, if so, quantify to what extent.

The contribution of this paper is to propose a method of evaluating the evolution of the U.S. output-inflation tradeoff. While there is much work in the area of evaluating output-inflation tradeoffs, much of the evidence from current work rests upon either split sample analysis or does not account for the identification of demand shocks. Both are potentially problematic. The former can be reliant on nominating an explicit break date in the sample where there are reasons to believe results can be sensitive. The latter is problematic because of the classic issue of endogeneity. Given output and inflation are jointly determined, without an explicit identification of demand shocks, correlations between the real activity variable of interest and inflation can be misleading. A natural extension to the current literature is to therefore propose refinements which account for these two features explicitly. More specifically, the paper proposes the use of Time-Varying Parameter Vector Autoregression (TVP-VAR) models in the spirit of Cogley and Sargent (2005) and Primiceri (2005) and identifies demand shocks within the framework of Blanchard and Quah (1989) and Cecchetti (1994) to account for the two issues respectively.

Figure 1 presents scatterplots of U.S. inflation and the output gap, measured as the difference between real GDP and the Congressional Budget Office's estimate of real potential output, on the vertical and horizontal axes respectively. A regression line is fitted for subsamples. Focusing on the top panel (a), inflation has become less sensitive to fluctuations in real activity since the Volcker disinflation of the mid 1980s. Much empirical work statistically confirms this observation, concluding the Phillips Curve has become flatter.

Yet, there is reason to be skeptical. The bottom panel (b) changes the breakdate for the respective subsamples. It is now difficult to make a case that the Phillips Curve has flattened considerably. In fact, it appears the Phillips Curve is always flat! Admittedly, one can always make an argument where splitting the sample in 1980 affects the results given high inflation

¹While the Phillips Curve in its original formulation is strictly between unemployment and inflation, the tradeoff between real activity such as the output gap and inflation will be referred to as the slope of the Phillips Curve in this paper. These are equivalent as long as Okun's Law is stable.

persisted into 1981, and to some extent 1982, following the oil and food price shocks in 1979/80.² Evidence of the flattening Phillips Curve may thus rest on arbitrary choices of breakdates in the sample. If grouping the entire 1980-82 sample within the 1970s produces a steep Phillips Curve, then are these steep Phillips Curves a product of outliers? A useful metric in assessing the extent of the output-inflation tradeoff is the sacrifice ratio. The sacrifice ratio is the ratio of the percentage points of output growth foregone per percentage point decrease in inflation. Mishkin's (2007) reading, for example, suggests estimates within the literature imply a sacrifice ratio 40% larger today compared to the Great Inflation. The top panel of Figure 1 suggests such numbers may be a reasonable approximation. Yet, if all these numbers are mere artifacts of subsample breakpoints, the evidence encompassed within Figure 1 makes for a less compelling argument.

If this is a mere issue of sample breakpoints, the issue at hand would be a less complicated one. The evidence within Figure 1 is essentially reduced form correlations between inflation and the output gap. Recall that only demand policy induces a Phillips Curve tradeoff between output and inflation. Supply shocks do not. In essence, quantifying the slope of the Phillips Curve is also an identification problem. Neither the plotting of slopes as in Figure 1 or univariate regressions of inflation and the output gap can inform of the extent of the slope of the Phillips Curve. Identification, and more specifically identifying demand shocks, is mandatory for evaluating this tradeoff.

Conceptually, the paper extends the work of Cecchetti (1994). Cecchetti's approach identifies a generic demand shock within a VAR framework. The extension is to model time variation. The use of a TVP-VAR obviates the issue of finding breakdates to split the sample.

While the paper makes a methodological point, the main contribution of this paper is empirical. From this perspective, the paper contributes to a large body of applied work exploring changes in the Phillips Curve tradeoff. This body of work can be loosely demarcated onto two sides. On one hand, much evidence documenting changes in the output-inflation tradeoff is either reduced form or single equation in nature (see *inter alia*, Roberts, 2006; Benati, 2007; Kuttner and Robinson, 2010; Ball and Mazumder, 2011). On the other hand, important contributions using estimated structural macro models, often in the New Keynesian tradition but relying on split samples, also document flattening of the Phillips Curve (e.g. Galí and Gertler, 1999; Boivin and Giannoni, 2006; Smets and Wouters, 2007). Most of such fully structural models, though, are modelled in a way where the source of flattening Phillips Curves is through less flexible price adjustments resulting in inflation becoming less sensitive to shocks.³ In other words, there is no other dominant mechanism which flattens the Phillips Curve in these models. The approach taken explicitly accounts for identification issues relative to the reduced form or

²Stock and Watson (2008) report Phillips Curve type inflation forecast achieve their greatest success in the early 1980s disinflation. While they do not use the output gap as their real activity variable, the stylised example is consistent with their findings.

³To be more precise, it is an increase in the Calvo parameter, which measures price stickiness in New Keynesian models, that flattens the Phillips Curve in these environments.

single equation work. Conceptually, the approach here is faithful to the tradition of the more structural literature where identifying demand shocks in the model is mandatory for exercises in quantifying Phillips Curve slopes. On the other hand, a body of empirical work is interested in whether issues like openness (e.g. Temple, 2002; Daniels et al., 2005) or central bank independence (e.g Jordan, 1997; Gonçalves and Carvalho, 2009) may be sources of flattening the Phillips Curve. Therefore, it may be useful to have a model which allows for other possible sources. Relative to the more fully specified macro models, the contribution of this paper can be seen in a more semi-structural light. There is no attempt in distinguishing between the sources of a flattening Phillips Curve. The focus of this paper is much narrower. It merely evaluates whether the Phillips Curve has flattened, and if so, by how much. This cost, however, needs to be contrasted against the potential benefit of specifying a more flexible specification to model time variation. It must be stressed that almost all the estimated structural macro models need to rely on split samples and there is at least some reason to get away from such an approach, given the example illustrated earlier.

The results of this paper can be summarised as follows. Using U.S. data from 1970Q1 to 2012Q4, there has been a slight increase in the output-inflation tradeoff, partially supporting claims the Phillips Curve has indeed flattened. Nevertheless, it appears this flattening does not appear to be as large as previously thought. There is some international evidence to suggest there may be a global dimension with regards to the flattening of the Phillips Curve, with usual explanations such as openness and inflation targeting being possibilities.

Before moving on, a comment on terminology is needed. Terms like slope of the short-run Phillips Curve, sacrifice ratio and output-inflation tradeoff have been variously used in the literature, and also in the preceding passages. Broadly, they all appeal to the same concept and are the same within the framework of this paper. The metric of changes in the output-inflation tradeoff will be the sacrifice ratio. Note an increase in the sacrifice ratio simultaneously represents both a flattening of the Phillips Curve and an increase in the output-inflation tradeoff.

The rest of the paper is organised as follows. Section 2 discusses the conceptual framework and empirical methodology. Section 3 presents discusses the results. Section 4 offers some concluding remarks.

2 Methodology

To reiterate, this paper proposes using a TVP-VAR with stochastic volatility in the tradition of Cogley and Sargent (2005) and Primiceri (2005) in order to estimate the output-inflation tradeoff. Identification is similar to Cecchetti (1994). The estimation sample will be from 1970Q1 to 2012Q4.

2.1 Conceptual Framework

It is first useful to lay out the conceptual framework. It is now well accepted that while a short-run tradeoff between output and inflation exists, this relationship disappears in the long-run. Figure 2 provides a simple illustration of this idea. Suppose output is growing at trend level at inflation level π_0 . At t_0 , a disinflation occurs, perhaps as a result of deliberate demand policy. Output grows below its trend level until t_1 . At t_1 , inflation is down to π_1 and output is back at trend. The output-inflation tradeoff is the *cumulative output loss relative to trend*, namely the shaded area, divided by the change in inflation, $\pi_0 - \pi_1$. The resulting ratio k is the sacrifice ratio. It can be interpreted as for 1% decrease in inflation, the cost is $k\%$ of output. k would also be the slope of a short to medium run Phillips Curve, the type which Figure 1 presents.

Figure 2 illustrates a conceptual challenges with estimating the output-inflation tradeoff. Firstly, as previously noted, only demand shocks can offer the type of adjustment embodied within Figure 2. Supply shocks induce higher (lower) inflation with an output loss (gain), inconsistent with the ideas presented in Figure 2. Secondly, there is difficulty in quantifying both trend inflation and trend output. Both cannot be unobserved, but key concepts regarding any adjustments relate directly towards trend inflation and trend output. Mismeasurement of the trend potentially contaminates calculation of this tradeoff. Carlstrom and Fuerst (2008), for example, demonstrate, through a sequence of rolling univariate regressions, evidence of dramatic flattening Phillips Curves disappear if trend inflation is accounted for.

The empirical framework will therefore have to, and does, account for identification, trend inflation and trend output in order to quantify the output-inflation tradeoff. We turn our attention to the model.

2.2 A Time-Varying Parameter VAR

Let X_t be a vector of two variables which contains output growth and the change in the CPI inflation respectively. Denote this vector of variables $X_t = [\Delta y_t, \Delta \pi_t]$ where y_t is natural logarithm of real output and π_t is CPI inflation. Consider the following TVP-VAR model,

$$X_t = c_t + B_{1,t}X_{t-1} + \dots + B_{p,t}X_{t-p} + \Gamma_t \epsilon_t \quad (1)$$

where c is a 2×1 vector of time-varying constants. The B matrices are each a 2×2 matrix of time-varying coefficients with p the maximum lag order of the estimated VAR. Γ_t is of dimension 2×2 , sometimes referred to as the impact matrix for identification in Structural VAR models. ϵ_t are serially uncorrelated structural innovations with distribution $N(0, I_2)$. In state space form, equation (1) is the observation equation

Let

$$A_t = \begin{pmatrix} 1 & 0 \\ \alpha_t & 1 \end{pmatrix}, \Sigma_t = \begin{pmatrix} \sigma_{1,t} & 0 \\ 0 & \sigma_{2,t} \end{pmatrix}$$

A reduced form time-varying covariance matrix, Ω_t can be constructed using

$$\Omega_t = A_t^{-1} \Sigma_t \Sigma_t' A_t^{-1'}$$

Note that as the reduced form covariance matrix Ω has a t subscript, indicating time variation. The impact matrix for identification Γ thus also has a t subscript as it is also time-varying by construction. With information on the identification scheme, it is possible to recover Γ_t from Ω_t . Γ_t will be a valid decomposition of Ω_t of the form $\Omega_t = \Gamma_t \Gamma_t'$.

Stacking the matrices of time-varying parameters in a vectorised fashion where $B_t = \text{vec}(B_{i,t}), i \in \{1, 2, \dots, p\}$, and $\sigma_t = \text{diag}(\Sigma_t)$, the state equations can be written as

$$\begin{aligned} B_t &= B_{t-1} + v_t \\ \alpha_t &= \alpha_{t-1} + \zeta_t \\ \log \sigma_t &= \log \sigma_{t-1} + \eta_t \end{aligned}$$

All the time-varying parameters have a random walks structure, similar to Cogley and Sargent (2005) and Primiceri (2005). v, ζ and η are jointly normally distributed. A block diagonal matrix expresses the covariance matrix for the innovations

$$V \begin{pmatrix} \epsilon_t \\ v_t \\ \zeta_t \\ \eta_t \end{pmatrix} = \begin{pmatrix} I_2 & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{pmatrix}$$

where Q, S and W are all positive definite matrices. The model can be estimated using Markov Chain Monte Carlo methods. The estimation techniques and priors will be conventional relative to the existing TVP-VAR literature, with Primiceri (2005) a particularly useful reference. Details on the estimation are in the appendix.

2.3 Identification

The paper considers a time-varying version of Cecchetti's (1994) bivariate model which contains output growth and the change in the CPI inflation respectively. Using Galí and Gambetti's (2009) approach for imposing long-run restrictions in the Blanchard and Quah (1989) tradition within a TVP-VAR environment, we can rewrite (1) using lag polynomial notation

$$B_t(L) \begin{pmatrix} \Delta y_t \\ \Delta \pi_t \end{pmatrix} = \Gamma_t \begin{pmatrix} \epsilon_t^S \\ \epsilon_t^D \end{pmatrix} \quad (2)$$

where $B_t(L) = I - B_{1,t}L - B_{2,t}L^2 - \dots - B_{p,t}L^p$, with the t subscript once again denoting the parameters are time-varying. Use $B_t(L)$ to approximate for the infinite order Vector Moving

Average (VMA) process at every t . The innovations in this model are a supply and demand shock, denoted ϵ_t^S and ϵ_t^D respectively. Rewriting (2) in VMA form

$$X_t = B_t^{-1}(L)\Gamma_t\epsilon_t \quad (3)$$

$$= C_t(L)\epsilon_t \quad (4)$$

where $C_t(L) = B_t^{-1}(L)\Gamma_t$ and $\epsilon_t = [\epsilon_t^S, \epsilon_t^D]$. Expanding (4) and substituting in the elements of the vector X_t , we obtain

$$\begin{aligned} \begin{pmatrix} \Delta y_t \\ \Delta \pi_t \end{pmatrix} &= C_t(L) \begin{pmatrix} \epsilon_t^S \\ \epsilon_t^D \end{pmatrix} \\ &= \begin{pmatrix} C_{11,t}(L) & C_{12,t}(L) \\ C_{21,t}(L) & C_{22,t}(L) \end{pmatrix} \begin{pmatrix} \epsilon_t^S \\ \epsilon_t^D \end{pmatrix} \\ &= \begin{pmatrix} \sum_{i=0}^{\infty} c_{11,t}^i \epsilon_{t-i}^S + \sum_{i=0}^{\infty} c_{12,t}^i \epsilon_{t-i}^D \\ \sum_{i=0}^{\infty} c_{21,t}^i \epsilon_{t-i}^S + \sum_{i=0}^{\infty} c_{22,t}^i \epsilon_{t-i}^D \end{pmatrix} \end{aligned}$$

In the Blanchard and Quah (1989) tradition, demand shocks do not have a long-run impact on real output. This implies $C_{12,t}(1) = \sum_{i=0}^{\infty} c_{12,t}^i = 0$. The model is identified by imposing this one restriction. It is then straightforward to recover Γ_t .

Identifying the aggregate demand shock is key to constructing the output-inflation tradeoff. Let $\Theta_t(\tau)$ be the tradeoff at time horizon τ at time t , then

$$\begin{aligned} \Theta_t(\tau) &= \sum_{j=0}^{\tau} \frac{\partial y_{t+j}}{\partial \epsilon_t^D} \bigg/ \frac{\partial \pi_{t+\tau}}{\partial \epsilon_t^D} \\ &= \left[\left(\sum_{i=0}^0 c_{12,t}^i \right) + \left(\sum_{i=0}^1 c_{12,t}^i \right) + \dots + \left(\sum_{i=0}^{\tau} c_{12,t}^i \right) \right] \bigg/ \left(\sum_{i=0}^{\tau} c_{22,t}^i \right) \\ &= \frac{\sum_{i=0}^{\tau} \sum_{j=0}^i c_{12,t}^i}{\sum_{i=0}^{\tau} c_{22,t}^i} \end{aligned} \quad (5)$$

Calculating the output-inflation tradeoff from the expression in (5) is straightforward from the impulse response functions. Given inflation and output are at their respective trend levels once the system has converged, $\Theta_t(\infty)$ is the tradeoff between output and inflation. Of course, one cannot calculate this directly, but this can be approximated as long as τ is large enough. Note, $\Theta_t(\infty)$ is also the sacrifice ratio given the equivalence of the concepts being considered.

As motivated, the concept of trend inflation and trend output are important in computing

the output-inflation tradeoff. Given the model specifications, computing these ratios from impulse response functions eliminates the need to make assumptions on trend inflation and trend output levels because both inflation and output revert to their respectively (stochastic) trend level at long horizons. This is arguably an improvement on existing methods in calculating this tradeoff using sacrifice ratio methods without the need to make assumptions about when output is at trend (see, for example, the widely applied method by Ball, 1994). Inflation is an assumed non-stationary process, and modelled as such. Such a choice is deliberate. Suppose inflation is a mean-reverting stationary process, the entire concept of the tradeoff between output and trend inflation becomes meaningless because no policy can change the level of inflation in the long-run. To see this, inspect (5). If inflation is mean reverting, the denominator will converge to zero in the limit and $\lim \Theta_t(\tau)_{\tau \rightarrow \infty} = \infty$.

While the issue of whether inflation is non-stationary is indeed a controversial one, assuming trend inflation is non-stationary is commonly adopted in empirical work using trend cycle decompositions and unobservable components approaches for modelling trend inflation (e.g. Stock and Watson, 2007; Chan et al., 2013).⁴

A vertical long-run Phillips Curve implies aggregate demand shocks have no impact in the long-run. This is less controversial theoretically, but the time-varying nature of the problem does impose a vertical Phillips Curve at every point of time. Empirical verification of strong evidence of a vertical Phillips Curve in a time-varying framework by Koop et al. (2010) makes invoking such an assumption a less uncomfortable proposition.

The data series are taken from the FRED database.⁵ While standard TVP-VAR setups often only consider a lag order of two, the parsimonious setup will require a large lag order to approximate the dynamics, since any potential omitted variable may induce Moving Average terms. In this regard, the estimation uses eight lags, which is equivalent to two years worth of past dynamics.

3 Results

3.1 Evolution of the Output-Inflation Tradeoff

Figure 3 plots impulse response functions to a one standard deviation shock to aggregate demand for selected quarters.⁶ Two observations are immediately apparent. First, consistent with theory, the output and inflation responses are in the same direction. This assures that the

⁴Such models usually model trend inflation as an I(1) process in the state equation. This is usually augmented with a cycle component which is I(0) with inflation being I(1) since it is a linear combination of the trend and cycle component.

⁵The series are the consumer price index urban consumers, seasonally adjusted (CPIAUCSL) and GDP in Chained 2005 Dollars (GDPC1), seasonally adjusted.

⁶As the covariance matrix models time variation, the size of one standard deviation is also time-varying by construction. While construction of the tradeoff is invariant to scaling, the impulse response functions across quarters are only comparable qualitatively.

demand shock is satisfactorily identified. Second, the system settles down close to the steady state within six to eight quarters. Therefore, output-inflation tradeoffs will be computed by approximating the tradeoff at twenty quarters (i.e. $\Theta_t(20)$). This will be reasonably close to the slope of the short to medium run Phillips Curve.

Figure 4 presents estimates of the the output-inflation tradeoffs. The left subplot presents the posterior median estimate with the associated 16th and 84th percentile of the posterior distribution. The range of the tradeoff is between 0.8 to 3.1, with a median of around 1.7 percentage points of real output growth per percentage point change in trend inflation. These numbers are very much within the ballpark of Ball's (1994) estimates during 1974-1976 and Volcker disinflation of 1980-1983. This is assuring, as Ball's methods of constructing these ratios are often dependent variables for sacrifice ratios in empirical analysis. A second observation is that the tradeoff has been increasing over time. This can be most clearly observed with the creeping up of the 84th percentile of the posterior distribution. As the increase in the median is not as noticeable in this plot, the right subplot presents just the posterior median estimate of the sacrifice ratio. This ratio is slightly under 1.7 for most of the 1970s before steadily rising in a fairly monotonic fashion to its current estimate of 1.75.

That the increased tradeoff occurs only after 1985 suggests that this observation may not be merely coincidental. Afterall, this is a break often associated with the advent of the Great Moderation, coinciding with an era of low inflation. Empirical evidence by Ball and Mazumder (2011) and Benati (2007) document respectively U.S. and international evidence of a flattening of the Phillips Curve relation since the Great Moderation. Ball et al. (1988) make an early theoretical prediction where if the level of trend inflation is higher, price resetting occurs more frequently and thus nominal rigidities matter less. In this case, higher trend inflation will steepen the Phillips Curve. An increase in the sacrifice ratio is certainly consistent with this view. Another conjecture is with inflation expectations now well-anchored at a lower level, bringing down trend inflation further is likely to be costly. By the same token, accelerating inflation with well anchored expectations may also allow slightly larger policy gains given the adjustments of inflation expectations to the new level of trend inflation will take more time due to the credibility of the current policy regime.

Nevertheless, it must be stressed that while this increase in the output-inflation tradeoff is clear and coincides with the Great Moderation, the economic significance of this difference is probably marginal. While the evidence supports flattening of the U.S. Phillips Curve, it is also fair to say it has not flattened sufficiently to significantly alter the policy tradeoff.

3.2 Did the Monetary Policy Tradeoff Change?

A possible caveat of the presented flattening of the U.S. Phillips Curve is that the previously presented tradeoff does not pertain directly to monetary shocks. By identifying a generic demand shock, this assumes the impacts of all demand shocks are the same. Of interest to

many in the sacrifice ratio literature is the desire to compute the monetary policy tradeoff explicitly. Cecchetti and Rich (2001), for example, further studied models which decompose the Cecchetti (1994) aggregate demand shock.

In this spirit, the TVP-VAR is augmented to now allow a further decomposition of the aggregate demand shock into IS and LM shocks, where the LM component will be interpreted as the monetary shock.

To study this in detail, the model is augmented by the *ex post* real interest rate

$$B_t(L) \begin{pmatrix} \Delta y_t \\ \Delta \pi_t \\ i_t - \pi_t \end{pmatrix} = \Gamma_t \begin{pmatrix} \epsilon_t^S \\ \epsilon_t^{IS} \\ \epsilon_t^{LM} \end{pmatrix}$$

where $i_t - \pi_t$ is the *ex post* real interest rate.⁷ The demand shock is now separately identified as being IS and LM shock. Cecchetti and Rich (2001) also study such a specification, albeit in a time-invariant environment. With the interest rate, the specification looks very much like a standard three variable New Keynesian specification. This has been used in various contexts, often augmented with other variables by, for example, Shapiro and Watson (1988) and Galí (1992), among others. Expanding the size of the system now requires three, instead of one, identifying restrictions. The identifying assumption where demand shocks have no impact on real output in the long-run is retained. This implies both the IS and LM shocks have no long-run impact on real output. The model requires one more identifying restriction. This additional restriction will be made on a contemporaneous relationship, specifically IS shocks will have no contemporaneous impact on output. Contemporaneous identifying assumptions will always be contentious and this is no different. This is a direct cost of expanding the size of the system. It is therefore imperative to at least be convinced that the shocks are identified in a satisfactory manner.

Figure 5 presents the impulse responses of output and the nominal interest rate to one standard deviation LM shock in selected quarters, normalised by inflation rising upon impact. While not extremely precise, in general output rises and the nominal interest rate falls. Coupled with inflation normalised to rise, these are consistent with what one expects from a monetary expansion. This at least suggests that the shock is interpretable as a monetary shock. The output-inflation tradeoff from a monetary shock can thus be calculated directly from the impulse response functions of the output and inflation to a monetary shock, similar to the earlier two variable VAR exercise.

Figure 6 shows the output-inflation tradeoff associated with a monetary policy shock, calculated at the twenty quarter horizon. The left subplot displays these tradeoffs, once again with associated 16th and 84th percentiles of the posterior distribution with the right subplot presenting just the posterior median. As Cecchetti and Rich (2001) suggest, augmenting the model

⁷The nominal rate is the series 3-Month Treasury Bill: Secondary Market Rate (TB3MS) sourced from the FRED database. The model is estimated on four lags.

produces extremely imprecisely estimated tradeoffs. These results are no different. Nonetheless, one can make a few observations. These ratios are more precisely measured post Great Moderation, often with the 16th percentile clearing zero. This provides some suggestive evidence that the output-inflation tradeoff associated with monetary policy is statistically different from zero post Great Moderation, compared to the sample before. This suggests there was no evidence of a monetary policy tradeoff between inflation and output before the Great Moderation and this tradeoff was only restored post Great Moderation. The right subplot, presenting only the posterior median, suggests this tradeoff has been rising over time, though imprecision with the estimation is still an issue.

A large degree of this imprecision probably stems from the identification of the LM shock. However, the results are still fairly informative. In particular, an increase in this tradeoff occurred post Great Moderation, consistent with the pattern presented earlier. The timing and the patterns are not likely coincidental and strongly suggest that the output-inflation tradeoff has indeed increased since the Great Moderation. A large degree of that increase is likely to be due to the increased tradeoff faced by monetary policy.

Nonetheless, there are strong caveats to the estimated monetary tradeoff and therefore, focus should be on the more parsimonious benchmark two variable VAR case. First, larger VAR models need more identifying restrictions to identify shocks, with a greater number of restrictions making some potentially indefensible. The three variable VAR exercise is a clear example of this, with an additional restriction on IS shocks having no contemporaneous impact not particularly defensible, but producing a system which appears theoretically consistent with the identified monetary shock.

Second, it is not clear that the identifying restrictions on monetary shocks should hold throughout the entire sample. In particular, there is sufficient reason to suspect changes in the interest rate may have had nothing to do with monetary policy (see, for example, Rudebusch, 1998). In particular, there was a change to targeting the nonborrowed reserves after 1979 until a change back to interest rate targeting at an indeterminate date during the 1980s. Therefore, it is generally unclear whether it was nonborrowed reserves or the interest rate was linked to monetary policy.⁸ Moreover, it is particularly true the Federal Reserve monitored an eclectic mix of variables, with the interest rate only being one of them. If this is the case, identifying monetary shocks with the interest rate is flawed. This is also the case in the post Global Financial Crisis environment where interest rates have hit a zero lower bound and it is not clear whether using the interest rate is the correct way of identifying monetary shocks. This is not just an issue specific to this paper, but to the wider VAR literature identifying monetary shocks. More generally, the issues here are not new and are entirely associated with the typical issue associated with monetary identification in the VAR literature.

Third, it is not clear theoretically why the tradeoff between different demand shocks, including monetary shocks, should differ. For example, the slope of the Phillips Curve in the

⁸See Thornton (2006) for a narrative account of the policy targets in the 1980s.

workhorse New Keynesian models only depends on deep parameters (e.g. the frequency of price adjustments, the intertemporal elasticity of substitution etc.). Such models do not suggest monetary shocks induce a different tradeoff from other demand shocks like preference or fiscal shocks.⁹

For such reasons, while the paper can loosely find evidence of an increasing tradeoff associated with monetary shocks, in support with the earlier evidence of the flattening of the Phillips Curve, focusing on the results in the benchmark two variable case is preferable.

3.3 International Evidence

Some commentators (e.g. Borio and Filardo, 2007; Benati, 2007) argue the flattening Phillips Curve as a more global phenomenon. A natural extension is to extend the analysis by estimating the benchmark TVP-VAR specification in order to uncover whether this observation pertains specifically only to the U.S.

The model was estimated for Australia, United Kingdom and Canada. The choice of these three countries is entirely driven by the availability of a long stretch of consistent inflation and output data in order to match up with the sample period for the model with U.S. data. There are two broad incidental features regarding these three countries. First, they are often regarded as small open economies in the open economy macroeconomic literature. Second, they all are inflation targeting regimes, with Australia, Canada and the United Kingdom adopting inflation targeting in 1994, 1991 and 1997 respectively.

Figure 7 presents the posterior medians for all three countries.¹⁰ While Canada and the United Kingdom display output-inflation tradeoffs evolving similarly to those estimated for the U.S., the evidence appears more mixed for Australia. One can also observe that any form of flattening for the Phillips Curves in these countries appears to be of a magnitude much larger to those estimated for the U.S. For example, these range from the tradeoff increasing by approximately 0.2 and 1.5 percentage points of growth per percentage point decrease in trend inflation for the United Kingdom and Canada respectively.

An interesting feature emerges where the introduction of inflation targeting appears to be associated with a fall in this tradeoff prior to a rise thereafter. This appears to coincide quite precisely with the evidence during the Volcker disinflation period in the model estimated with U.S. data, though the magnitudes associated with the Volcker disinflation are economically small. Interestingly, a literature has evolved attempting to link inflation targeting with less

⁹Consider the workhorse New Keynesian Phillips Curve which the graduate textbook by Galí (2008), among others, formulate. $\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa \tilde{y}_t$ where π is inflation, \tilde{y} is the output gap, β is the discount factor and κ is a reduced form parameter which is a function of deep parameters like the Calvo parameter and discount factor. The slope of the Phillips Curve is κ . Since this equation has to *always* be satisfied in a New Keynesian model, the relation between the output gap and inflation is stable. Therefore, all demand shocks induce the same output-inflation tradeoff, κ . From this perspective, there is no compelling reason why different demand shocks have to be considered to identify the slope of the Phillips Curve.

¹⁰The Canada sample only starts in 1971Q3 while the other two start in 1970Q1.

costly disinflation (e.g. Gonçalves and Carvalho, 2009), and the sudden fall in this tradeoff appears to consistent with this hypothesis.

Globalisation has also been variously hypothesised to coincide with flatter Phillips Curves. No direct claims with regards to this hypothesis can be made, apart from observing the larger output-inflation tradeoffs for all these three small open economies, especially since the mid 1990s, and the rising tradeoff for the U.S. These patterns may be worthwhile exploring for future research.

3.4 Policy Discussion

The evidence presented in this paper does support the claim that while the output-inflation tradeoff has changed in recent decades, these tradeoffs have not changed as sharply as suggested by previous studies. This has some contemporary policy implications.

With sluggish growth a key concern of policymakers in the aftermath of the Global Financial Crisis, some have contended whether to raise the inflation target to four percent from the currently implicitly assumed non-pernicious inflation rate of around two percent (i.e. Blanchard et al., 2010; Krugman, 2012). While their motivations of a higher target are much deeper than exploiting the output-inflation tradeoff, raising the inflation target inevitably generates discussion regarding some of the short-term benefits of allowing higher inflation. This point is elucidated clearly by statements from the Federal Reserve. For example, Federal Reserve Chairman Ben Bernanke, in response to reporters after the Federal Open Market Committee meeting on 25 April 2012, stated, “The question is, does it make sense to actively seek a higher inflation rate in order to achieve a slightly faster reduction in the unemployment rate.” The Richmond Federal Reserve’s President Jeffrey Lacker, said on May 2 2012, “It would be very dangerous for the Federal Reserve to adopt a higher inflation target to spur growth in the hopes that it could bring inflation down later when it needed to.”

Recall that a flattening of the Phillips Curve implies these growth benefits are also going to be larger. However, if one took the estimates from this paper at face value, there are suggestions these benefits are not going to be so substantially different in terms of magnitude, compared to running similar demand policies in the 1970s. The Phillips Curve has simply not flattened substantially enough to gain an additional so-called “credibility bonus” to exploit. This may not be true though from the international evidence. For example, the analysis here suggests countries like Canada and Australia may potentially have a much larger exploitable tradeoff.

At a deeper level though, it is useful to ask whether one can literally take the estimates here for policy purposes. This will depend primarily on what has driven this flattening. To see this more carefully, consider the following model, which is an adaptation of Barro and Gordon (1983). Suppose an expectations augmented Phillips Curve

$$\pi - \pi^e = \gamma(y - \tilde{y}) \tag{6}$$

where \tilde{y} is the natural rate of output and γ the speed of adjustment of inflation towards it's expected level π^e . Assuming a quadratic loss function for the policymaker

$$L(\pi, y) = \frac{1}{2} [(\pi - \pi^*)^2 + \lambda(y - y^*)^2] \quad (7)$$

where the asterisk indicates some desired level of inflation and output for the policymaker.

Substituting the expectation augmented Phillips Curve into the loss function and taking a first order condition with respect to π yields the short-run Phillips Curve

$$\pi = \pi^* + \frac{\lambda}{\gamma}(y^* - \tilde{y}) \quad (8)$$

The issue is whether, if the degree of tradeoff between inflation and output, namely $\frac{\lambda}{\gamma}$ in this illustrative model, is invariant to the Lucas Critique. Now, if the flattening is caused by structural parameters which influence price resetting, perhaps through more openness or less frequent resetting of prices, these presumably impact γ in the simple illustration. Then it is probably true that as long as these structural parameters are stable, the tradeoff between output and inflation will merely be moving along the short-run Phillips Curve in the simple illustrative model. In this case, such estimates can be used for policy calculations and experiments. If, for example, the monetary policy regime changes, perhaps through the preference of the policymaker impacts directly on the tradeoff, which could be the relative weight of inflation and output in the policymaker's loss function, λ , then the short-run tradeoff is probably going to change according to the policy regime. Using the estimated tradeoff will then run afoul of the Lucas Critique.

This simple illustrative model demonstrates that while the empirical methodology allows the estimation methodology of identifying the output-inflation tradeoff through demand shocks, there are varying levels of "reduced form". The model is structural in the sense demand and supply shocks are identified, but is silent, and perhaps a "reduced form" in terms of what has been flattening this slope. In other words, the exercise here identifies what is loosely $\frac{\lambda}{\gamma}$ in this illustrative model, but does not identify these two coefficients separately. Being able to now ascertain the degree of evolution with the output-inflation tradeoff, future research needs to evaluate whether it is worthwhile explaining sources of this change.

4 Concluding Remarks

This paper proposes a methodology of evaluating the output-inflation tradeoff. An analysis of the U.S. output-inflation tradeoff reveals only a mild increase from around 1.70 to 1.75 percentage points decrease of real output growth associated per percentage point fall in trend inflation. Therefore, while it supports the intuition that the U.S. Phillips Curve may indeed have flattened, this flattening is of a much smaller magnitude than previously thought. This is mainly driven by accounting explicitly for the identification of demand shocks and making

less *ad hoc* assumptions regarding sample split date, trend inflation and trend output. Part of this increase in the output-inflation tradeoff is suggestively supported by a corresponding increase in the tradeoff associated with monetary policy. International evidence also indicates three small open economies exhibit an increase in the output-inflation tradeoff, after they all adopted inflation targeting. This ties in neatly with a literature which suggests trade openness increases the sacrifice ratio and thus flattens the Phillips Curve (e.g. Temple, 2002; Daniels et al., 2005).

A particularly interesting extension for future research is to consider adopting the proposed methodology here to produce estimates for sacrifice ratio regressions. At this current point, empirical work of sacrifice ratio regressions construct their dependent variable using Ball's (1994) approach (e.g. Jordan, 1997; Temple, 2002; Daniels et al., 2005; Gonçalves and Carvalho, 2009; Mazumder, 2012). A potential downside to the Ball's (1994) approach is the sacrifice ratio can only be identified for large disinflation episodes given the conceptual difficulties in isolating demand contractions. In the absence of large disinflation episodes, there is no way of quantifying the sacrifice ratio using Ball's (1994) approach. This has three distinct disadvantages. First, to the extent countries without large disinflation episodes share some common characteristics, systematic omission of such observations potentially renders sacrifice ratio regressions misleading. Second, inferences from such samples are an artefact of restricting to only disinflation phrases. Even if the sample is representative, there are several papers in the sacrifice ratio literature where non-rejection of the null may be due to low power in tests with small samples (e.g. Temple, 2002). The methodology outlined within this paper offers a potential solution to identifying the sacrifice ratio in the absence of large disinflation episodes. Third, given disinflations are only temporary episodes, this obscures the ability of the Ball (1994) approach of pinning down a sacrifice ratio in typical circumstances. In other words, after disinflation, we have no idea what happens to the slope of the Phillips Curve. The outlined methodology once again offers future research the opportunity to study evolutions of these tradeoffs pre and post disinflation. These are especially helpful in pinning down the precise tradeoff in the presence of policy regime changes.

APPENDIX

A Priors

The priors of Q , W and blocks of S have an independent inverse Wishart distribution. This places a normal distribution on the initial conditions of the time-varying coefficients B_0 , A_0 and $\log \sigma_0$. A typical step is to use a training sample to calibrate the initial conditions on the prior, a practice this paper adopts. The range of the training sample is 1947Q2 to 1969Q4. Given 8 lags of the VAR are estimated, the training sample thus has 82 observations. Esti-

mating a constant coefficient VAR using least squares on the training sample provides the priors.

The priors are

$$\begin{aligned}
B_0 &\sim N(\hat{B}_{OLS}, V(\hat{B}_{OLS})) \\
A_0 &\sim N(\hat{A}_{OLS}, 4.V(\hat{A}_{OLS})) \\
\log\sigma_0 &\sim N(\log(\hat{\sigma}_{OLS}), I_N) \\
Q &\sim IW(k_Q.82.V(\hat{B}_{OLS}), 82) \\
W &\sim IW(k_W.3.I_N, 3) \\
S &\sim IW(k_S.2.V(\hat{A}_{OLS}), 2)
\end{aligned}$$

with subscript *OLS* referring to the least squares estimates from a constant coefficient VAR of the training sample. The paper will adopt relatively uncontroversial priors introduced by Primiceri. k_Q , k_S and k_W are set to 0.0001, 0.01 and 0.0001 respectively. The degrees of freedom on the covariance matrix of the VAR coefficient is set to the size of the training sample, 82.

B Description of Posterior Simulator

A Metropolis within Gibbs Sampler is used to simulate the posterior distribution. The outline of the steps are

1. Draw from $p(B_t^{(i)}|Y, A_t^{(i-1)}, \Sigma_t^{(i-1)}, V^{(i-1)})$
2. Draw from $p(A_t^{(i)}|Y, B_t^{(i)}, \Sigma_t^{(i-1)}, V^{(i-1)})$
3. Draw from $p(\Sigma_t^{(i)}|Y, B_t^{(i)}, A_t^{(i)}, V^{(i-1)})$
4. Draw from $p(V^i|Y, B_t^{(i)}, A_t^{(i)}, \Sigma_t^{(i)})$
5. Repeat steps 1 to 4

where the superscript denotes the i^{th} draw of the simulation. Steps 1 to 4 consist of time-varying parameters and are estimated through the standard Kalman filtering and smoothing techniques for time-varying parameter models as per Carter and Kohn (1994), except for imposing a restriction requiring the VAR roots to lie within the unit circle in step 1. The posterior simulator achieves this by utilising an algorithm developed by Koop and Potter (2011) to simulate the VAR coefficients in step 1.

The index for the order of the simulation draw will now be dropped and replaced with a superscript T to denote the parameters are time-varying.

Drawing elements of B_t . The density of the time-varying parameters can be factored

$$p(B^T|Y^T, A^T, \Sigma^T, V) = \prod_{t=1}^T p(B_t|B_{t-1}, Y^T, A^T, \Sigma^T, V)\mathbb{1}(B_t \in \Phi)$$

where $\mathbb{1}(\cdot)$ is an indicator function and Φ denotes the entire set of VAR coefficients which satisfy the stability condition by having roots lying within the unit circle. Without imposing stable roots, the time-varying VAR coefficients can normally be drawn using a Kalman filtering and smoothing technique like Carter and Kohn (1994) where

$$B_t|(B_{t+1}, Y^T, A^T, \Sigma^T, V) \sim N(B_{t|t+1}, P_{t|t+1})$$

$P_{t|t+1}$ is the precision matrix of the Kalman filter. Filtering forward of the Kalman filter and smoothing backwards allows computation of $B_{t|t+1}$ and $P_{t|t+1}$.

As Koop and Potter (2011) demonstrate, an accept/reject algorithm like Cogley and Sargent (2005) is not exactly correct due to ignoring the integrating constant. To account for this, the usual step of taking a draw by using the Kalman filter is first done. This will be a proposal draw for the metropolis step. If any of the time-varying VAR coefficients imply an unstable system, the algorithm immediately rejects this draw and stays with the previous draw. Otherwise, if the condition of stable VAR roots is achieved, an acceptance probability as per Koop and Potter (2011) is computed. The algorithm is allowed to switch between the multi and single move algorithm. This is particularly helpful if the multi move algorithm is rejecting a long sequence of draws which contain explosive roots. Depending on whether the algorithm is a multi or single move, the acceptance probability is computed using equations (9) and (10) or (16) respectively in their paper.

Drawing elements of A_t . Conditional on Y^T , B^T and Σ^T . Rewriting $A_t \hat{Y}_t \equiv A_t(Y_t - X_t' B_t) = \Sigma_t \epsilon_t = u_t$ where

$$\begin{aligned} \hat{Y}_{1,t} &= u_{1,t} \\ \hat{Y}_{2,t} &= -\alpha_t \hat{Y}_{1,t} + u_{2,t} \end{aligned} \tag{9}$$

and $\hat{Y}_t \equiv [\hat{Y}_{1,t}, \hat{Y}_{2,t}]$. The model has now been transformed into a series of time-varying regression models.

The α 's can now be drawn using the Kalman filter algorithm similar to the one described in drawing the time-varying VAR coefficients one equation at a time.

Drawing elements of Σ_t . Consider $A_t(y_t - X_t' B_t) = y^* = \Sigma_t \epsilon_t$. Taking the square and log of every element leads to

$$y_t^{**} = 2h_{i,t} + e_{i,t}$$

where $y_t^{**} = \log(y^* + c)^2$ where c is a small constant set to 0.0001 to prevent taking the log of zero. $e_{i,t} = \log(\epsilon_i^2)$ and $h_{i,t} = \log \sigma_{i,t}$. A mixture of normals as described by Kim et al. (1998) is used to approximate the transformed innovations, which are distributed $\log \chi^2(1)$. First, the expression is demeaned by utilising the relationship $\mathbb{E}(e_{i,t}) = -1.2704$. To define a mixture of seven normal probabilities with component probabilities q_j mean $m_j - 1.2704$ and variances ν_j^2

, $j \in (1, 2, \dots, 7)$, let $s^T = [s_1 \dots s_T]'$, a matrix of indicator variables selecting the mixtures of normal approximation at every point of time. We can sample the h 's as per the Kalman filter described previously and

$$Pr(s_{i,t} = j | y_t^{**}, h_{i,t}) \propto q_j f_N(y_t^{**} | 2h_{i,t} + m_j - 1.2704, \nu_j^2), j \in (1 \dots 7), i \in (1 \dots n)$$

Kim et al. (1998) compute the weights of the seven component mixtures. These same weights are also used in the estimation.

Drawing elements of V. R , W and the blocks of S have an inverse Wishert distribution. They can be drawn as the innovations are observable.

The posterior simulated takes sufficient draws for the different country samples and the three variable VAR specification with the objective of achieving a valid posterior distribution for inference. For the benchmark two variable TVP-VAR for the U.S. data in the paper, the posterior simulator takes 70,000 draws, with the first 30,000 discarded as a burn in sample. Thereafter, 1 in 40 draws is taken to thin the Monte Carlo sample and break the autocorrelation between draws. Therefore, this sample of 1,000 is taken as the posterior distribution for inference. The Monte Carlo samples are all checked for inefficiency factors and use the Geweke (1992) Z test to check for convergence to ensure valid inferences.

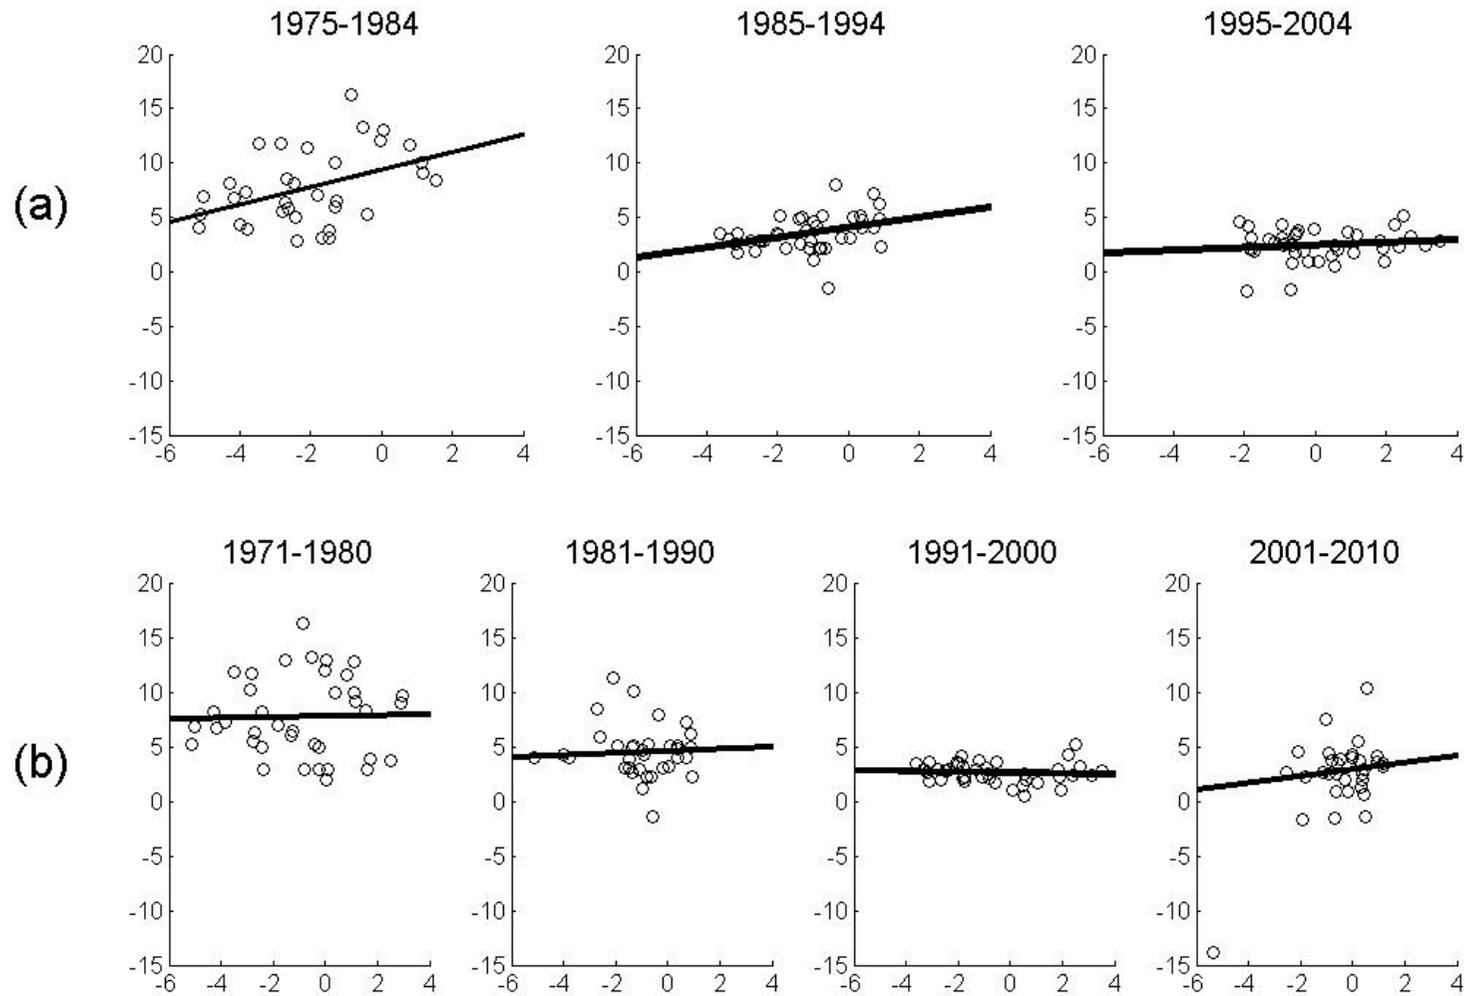
References

- Ball L. 1994. What determines the sacrifice ratio? In Mankiw NG (ed.) *Monetary Policy*, NBER Chapters. National Bureau of Economic Research, Inc, 155–182.
- Ball L, Mankiw NG, Romer D. 1988. The new Keynesian economics and the output-inflation trade-off. *Brookings Papers on Economic Activity* **19**: 1–82, doi: 10.2307/2534424.
- Ball L, Mazumder S. 2011. Inflation dynamics and the great recession. *Brookings Papers on Economic Activity* **42**: 337–405, doi: 10.1353/eca.2011.0005.
- Barro RJ, Gordon DB. 1983. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy* **91**: 589–610, doi: 10.1086/261167.
- Benati L. 2007. The time-varying Phillips correlation. *Journal of Money, Credit and Banking* **39**: 1275–1283, doi: 10.1111/j.1538-4616.2007.00066.x.
- Blanchard OJ, Dell’Ariccia G, Mauro P. 2010. Rethinking macroeconomic policy. *Journal of Money, Credit and Banking* **42**: 199–215, doi: 10.1111/j.1538-4616.2010.00334.x.
- Blanchard OJ, Quah D. 1989. The dynamic effects of aggregate demand and supply disturbances. *American Economic Review* **79**: 655–73.
- Boivin J, Giannoni MP. 2006. Has monetary policy become more effective? *The Review of Economics and Statistics* **88**: 445–462, doi: 10.1162/rest.88.3.445.

- Borio CEV, Filardo A. 2007. Globalisation and inflation: New cross-country evidence on the global determinants of domestic inflation. BIS Working Papers 227, Bank for International Settlements.
- Carlstrom CT, Fuerst TS. 2008. Explaining apparent changes in the Phillips curve: Trend inflation isn't constant. *Economic Commentary* .
- Carter C, Kohn R. 1994. On Gibbs sampling for state space models. *Biometrika* **81**: 541–553, doi: 10.1093/biomet/81.3.541.
- Cecchetti SG. 1994. Comment. In Mankiw NG (ed.) *Monetary Policy*, NBER Chapters. National Bureau of Economic Research, Inc, 188–193.
- Cecchetti SG, Rich RW. 2001. Structural estimates of the U.S. sacrifice ratio. *Journal of Business and Economic Statistics* **19**: 416–427, doi: 10.1198/07350010152596664.
- Chan JC, Koop G, Potter SM. 2013. A new model of trend inflation. *Journal of Business and Economic Statistics* **31**: 94–106, doi: 10.1080/07350015.2012.741549.
- Cogley T, Sargent TJ. 2005. Drift and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics* **8**: 262–302, doi: 10.1016/j.red.2004.10.009.
- Daniels JP, Nourzad F, Vanhoose DD. 2005. Openness, central bank independence, and the sacrifice ratio. *Journal of Money, Credit and Banking* **37**: 371–379, doi: 10.1353/mcb.2005.0020.
- Galí J. 1992. How well does the IS-LM model fit postwar U.S. data. *The Quarterly Journal of Economics* **107**: 709–38, doi: 10.2307/2118487.
- Galí J. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.
- Galí J, Gambetti L. 2009. On the sources of the great moderation. *American Economic Journal: Macroeconomics* **1**: 26–57, doi: 10.1257/mac.1.1.26.
- Galí J, Gertler M. 1999. Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* **44**: 195–222, doi: 10.1016/S0304-3932(99)00023-9.
- Geweke J. 1992. Evaluating the accuracy and sampling-based approaches to the calculation of posterior moments. In Bernardo JM, Berger JO, Dawid P, Smith AF (eds.) *Bayesian Statistics*. Oxford University Press, 169–193.
- Gonçalves CES, Carvalho A. 2009. Inflation targeting matters: Evidence from OECD economies' sacrifice ratios. *Journal of Money, Credit and Banking* **41**: 233–243, doi: 10.1111/j.1538-4616.2008.00195.x.
- Jordan TJ. 1997. Disinflation costs, accelerating inflation gains, and central bank independence. *Weltwirtschaftliches Archiv* **133**: 1–21, doi: 10.1007/BF02707673.
- Kim S, Shephard N, Chib S. 1998. Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies* **3**: 361–93, doi: 10.1111/1467-937X.00050.
- Koop G, Leon-Gonzalez R, Strachan RW. 2010. Dynamic probabilities of restrictions in state space models: An application to the Phillips curve. *Journal of Business & Economic Statistics* **28**: 370–379, doi: 10.1198/jbes.2009.07335.

- Koop G, Potter S. 2011. Time varying VARs with inequality restrictions. *Journal of Economic Dynamics and Control* **35**: 1126–1138, doi: 10.1016/j.jedc.2011.02.001.
- Krugman PR. 2012. *End This Depression Now!* W. W. Norton, New York, 1 edition. ISBN 9780393088779.
- Kuttner K, Robinson T. 2010. Understanding the flattening Phillips curve. *The North American Journal of Economics and Finance* **21**: 110–125, doi: 10.1016/j.najef.2008.10.003.
- Mazumder S. 2012. Determinants of the sacrifice ratio-evidence from OECD and non-OECD countries. *Wakefield Forest College Working Paper no. 106* .
- Mishkin FS. 2007. Inflation dynamics. *International Finance* **10**: 317–334, doi: 10.1111/j.1468-2362.2007.00205.x.
- Primiceri G. 2005. Time varying structural vector autoregressive and monetary policy. *Review of Economic Studies* **72**: 821–852, doi: 10.1111/j.1467-937X.2005.00353.x.
- Roberts J. 2006. Monetary policy and inflation dynamics. *International Journal of Central Banking* **2**: 193–230.
- Rudebusch GD. 1998. Do measures of monetary policy in a VAR make sense? *International Economic Review* **39**: 907–931, doi: 10.2307/2527346.
- Shapiro M, Watson M. 1988. Sources of business cycles fluctuations. In *NBER Macroeconomics Annual 1988, Volume 3*, NBER Chapters. National Bureau of Economic Research, Inc.
- Smets F, Wouters R. 2007. Shocks and frictions in us business cycles: A Bayesian DSGE approach. *American Economic Review* **97**: 586–606, doi: 10.1257/aer.97.3.586.
- Stock JH, Watson MW. 2007. Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking* **39**: 3–33, doi: 10.1111/j.1538-4616.2007.00090.x.
- Stock JH, Watson MW. 2008. Phillips curve inflation forecasts. NBER Working Papers 14322, National Bureau of Economic Research, Inc.
- Temple J. 2002. Openness, inflation, and the Phillips curve: A puzzle. *Journal of Money, Credit and Banking* **34**: 450–68, doi: 10.1353/mcb.2002.0049.
- Thornton DL. 2006. When did the FOMC begin targeting the federal funds rate? What the verbatim transcripts tell us. *Journal of Money, Credit and Banking* **38**: 2039–2071, doi: 10.1353/mcb.2007.0010.
- Yellen JL. 2006. Monetary policy in a global environment. *Speech at The Euro and the Dollar in a Globalized Economy Conference, U.C. Santa Cruz , Santa Cruz, CA, May 27, 2006* .
 URL <http://www.frbsf.org/our-district/press/presidents-speeches/yellen-speeches/2006/may/monetary-policy-in-a-global-environment/060527.pdf>

Figure 1: Scatterplot of U.S. inflation and output gap



Annualised quarter on quarter inflation is on the vertical axis.

Natural log of the output gap on the x axis, calculated as the difference between real GDP and the CBO's measure of potential real GDP.

Figure 2: Conceptual illustration of the tradeoff between output and inflation

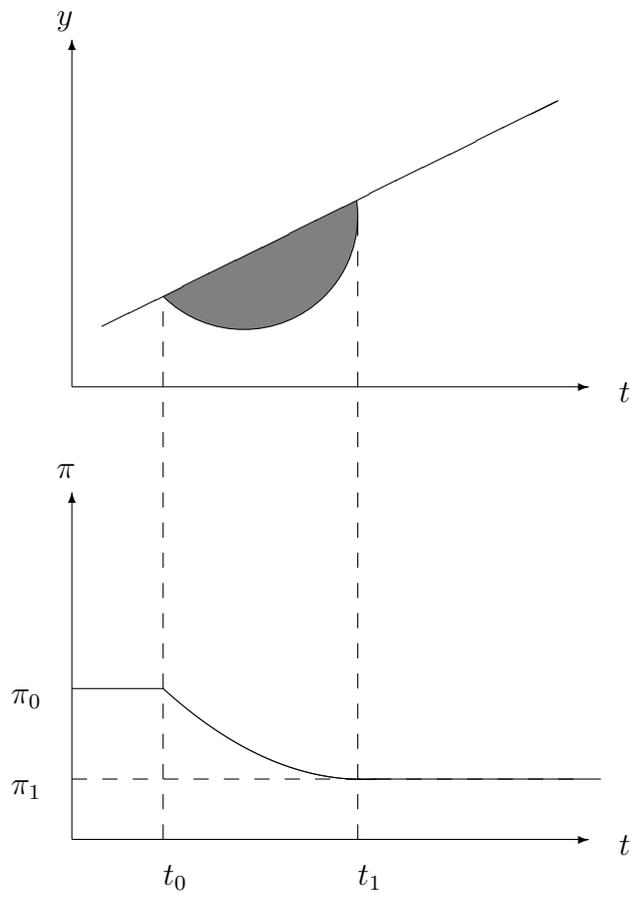
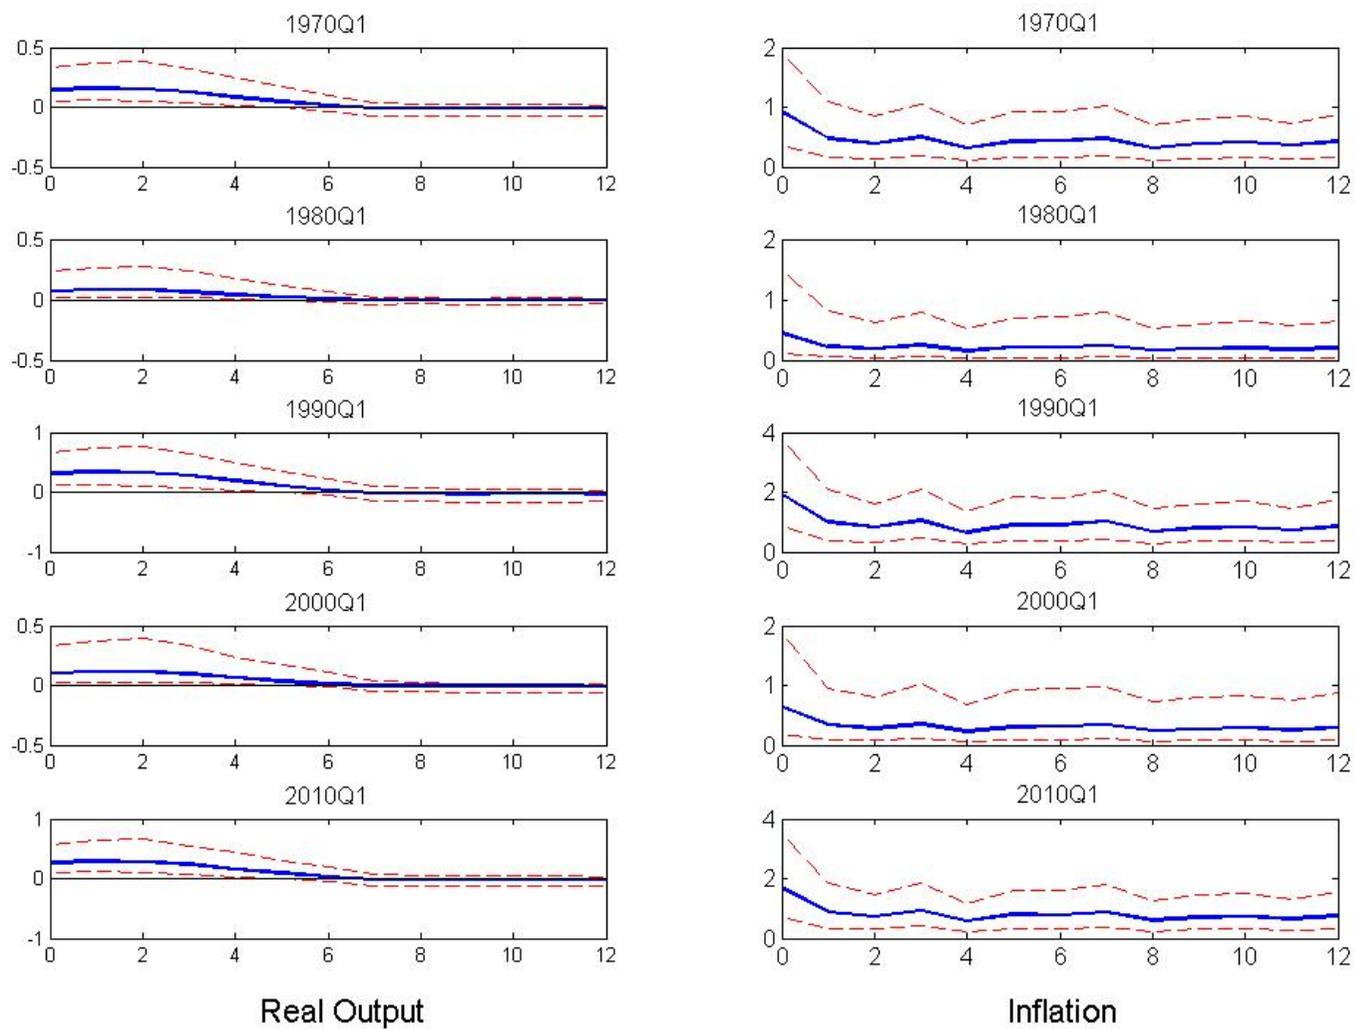


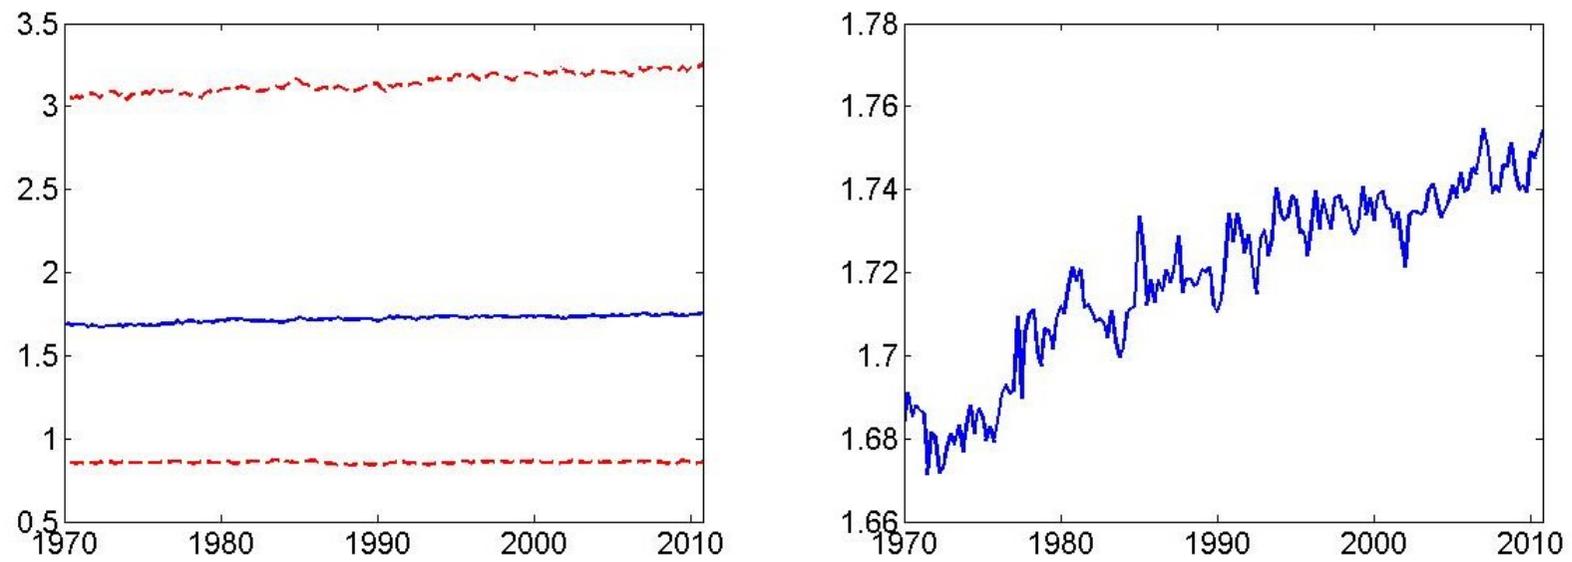
Figure 3: Impulse response function to one standard deviation demand shock



23

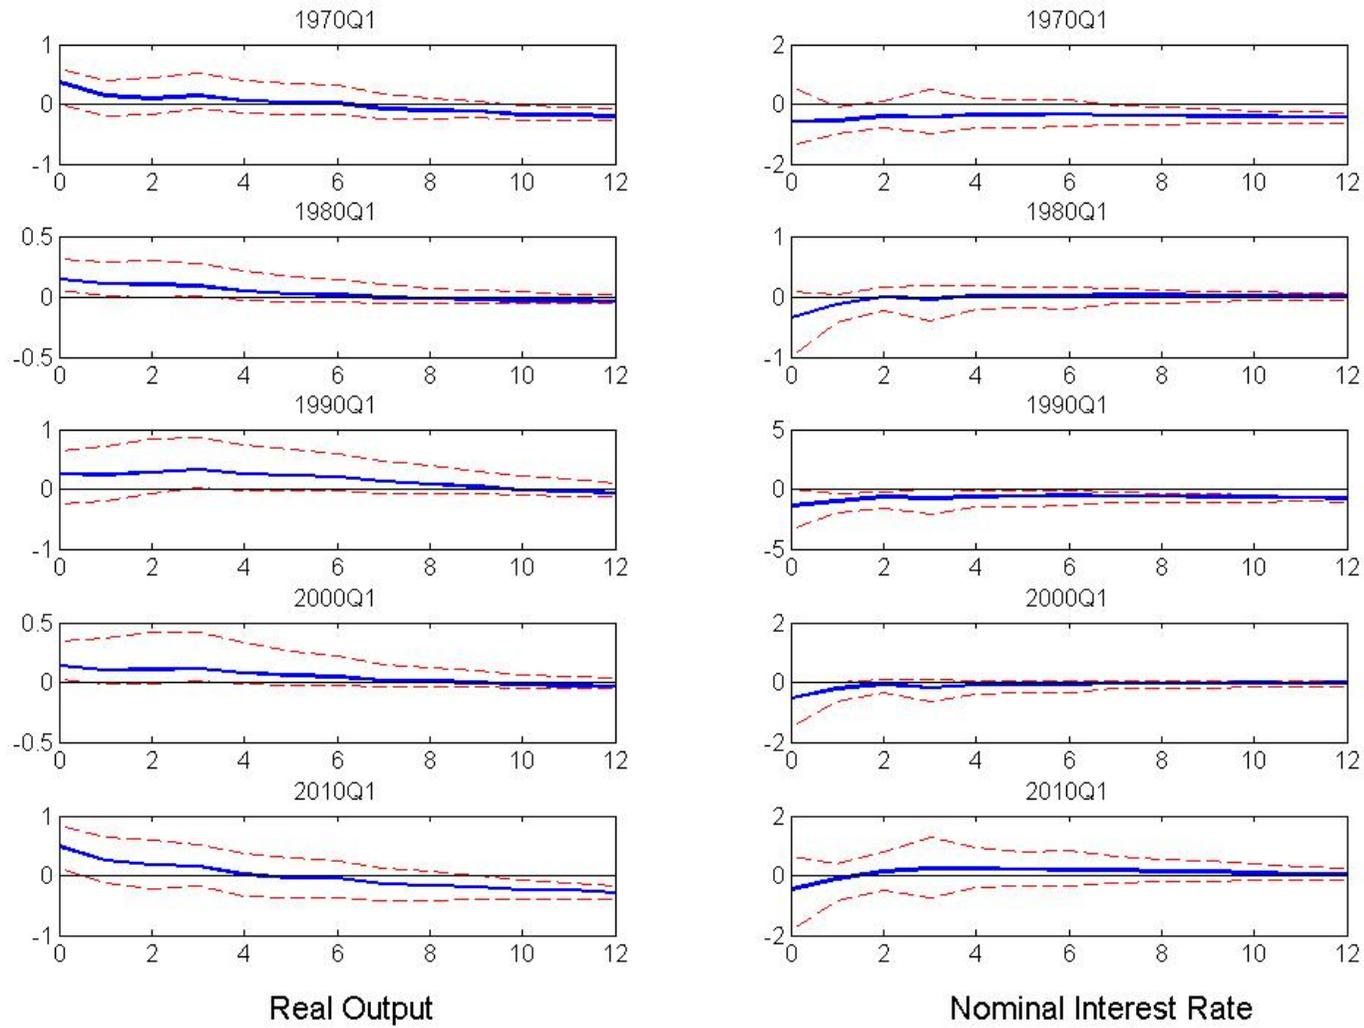
Posterior median with 68% credible sets.

Figure 4: Output-inflation tradeoff



*The left subplot presents the posterior median with the 68% credible sets.
The right subplot displays the posterior median.*

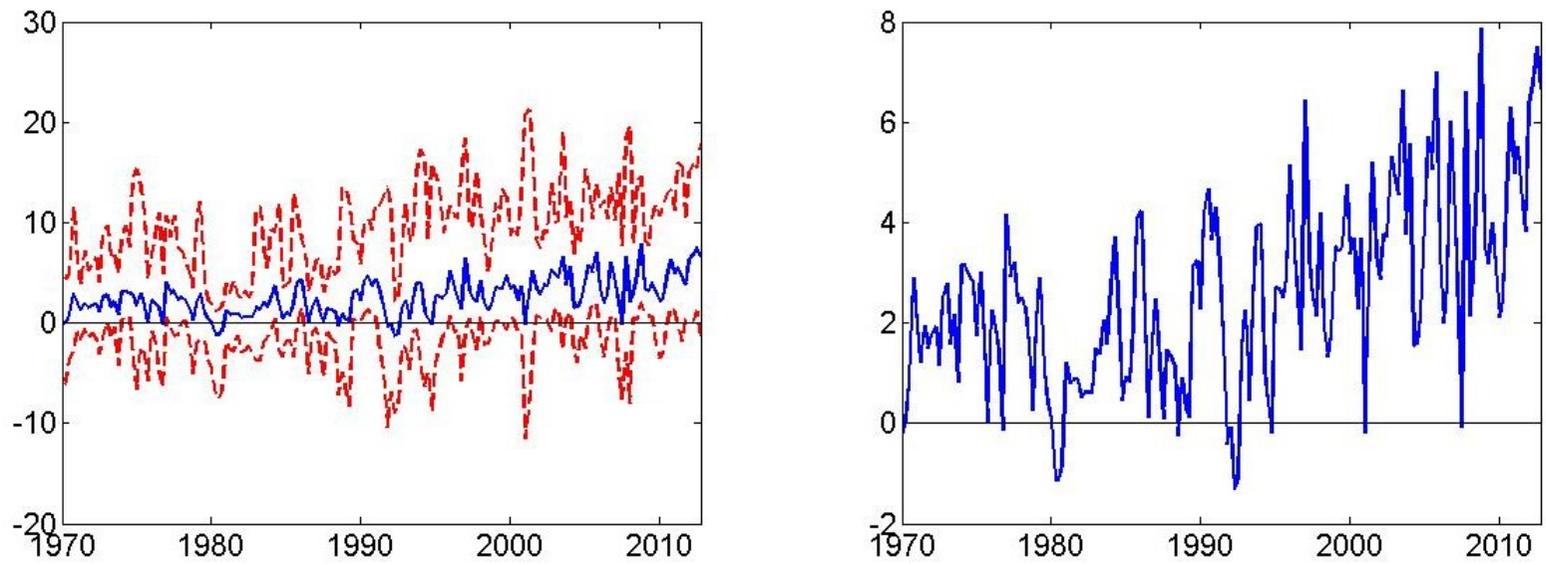
Figure 5: Impulse response function to one standard deviation LM shock



25

Posterior median with 68% credible sets.

Figure 6: Output-inflation tradeoff to monetary shock



*The left subplot presents the posterior median with the 68% credible sets.
The right subplot displays the posterior median.*

Figure 7: International comparisons of output-inflation tradeoff

