

Optimal Inflation and the Identification of the Phillips Curve*

Michael McLeay[†] Silvana Tenreyro[‡]

First draft: 26 April 2018

Abstract

This note explains why inflation follows a seemingly exogenous statistical process, unrelated to the output gap. In other words, it explains why it is difficult to empirically identify a Phillips curve. We show why this result need not imply that the Phillips curve does not hold – on the contrary, our conceptual framework is built under the assumption that the Phillips curve always holds. The reason is simple: if monetary policy is set with the goal of minimising welfare losses (measured as the sum of deviations of inflation from its target and output from its potential), subject to a Phillips curve, a central bank will seek to increase inflation when output is below potential. This targeting rule will impart a negative correlation between inflation and the output gap, blurring the identification of the (positively sloped) Phillips curve.

*The views expressed in this paper are those of the authors and not necessarily those of the Bank of England. This paper was motivated by a conversation with Ben Broadbent and Jan Vlieghe. We would like to thank Francesco Caselli, Andy Haldane, Richard Harrison, Per Krussell, Clare Macallan, Chris Redl, Martin Seneca, Jan Vlieghe and Matt Waldron for helpful comments and Oliver Ashtari Tafti for superb research assistance.

[†]Bank of England (michael.mcleay@bankofengland.co.uk).

[‡]Bank of England, LSE, CFM and CEPR (s.tenreyro@lse.ac.uk).

1. INTRODUCTION

A number of recent papers have pointed out that inflation can be approximated (and forecast) by statistical processes unrelated to the amount of slack in the economy ([Atkeson and Ohanian, 2001](#); [Stock and Watson, 2007, 2009](#); [Dotsey, Fujita and Stark, 2017](#); [Cecchetti et al., 2017](#); [Forbes, Kirkham and Theodoridis, 2017](#)). The empirical disconnect between inflation and various measures of slack has been interpreted by some commentators as evidence that the Phillips curve (a positive relation between inflation and the output gap) has weakened or even disappeared ([Ball and Mazumder, 2011](#); [IMF, 2013](#); [Blanchard, Cerutti and Summers, 2015](#)).^{1,2}

Considering that most macroeconomic models used by central banks embed the Phillips curve as one of their main building blocks, the empirical elusiveness of the Phillips curve challenges the models' wisdom. Or does it?

In this paper we use a standard conceptual framework to show why:

1. the empirical disconnect between inflation and slack is a result to be expected when monetary policy is set optimally; and
2. it is also perfectly consistent with an underlying stable and positively sloped Phillips curve.

More specifically, our framework is built under the assumption that the Phillips curve always holds; that is, inflation depends positively on the degree of slack in the economy. We also allow for cost-push shocks that can lead to deviations from the curve, but without altering its slope. Monetary policy is set with the goal of minimising welfare losses (measured as the sum of the quadratic deviations of inflation from its target and of output from its potential), subject to the Phillips curve or aggregate supply relationship. In that setting a central bank will seek to increase inflation when output is below its potential. This targeting rule imparts a negative correlation between inflation and the output gap, blurring the identification of the (positively sloped) Phillips curve.³

The paper is extended along four dimensions. First, it studies differences in the solutions between discretion – our baseline case in which the monetary authority cannot commit to a future

¹For a selection of the vast media comment on the issue, see articles in the *Financial Times*, *Wall Street Journal* and *The Economist* or opinion pieces by [Alan Blinder](#) and [Larry Summers](#).

²The output gap is defined as the deviation of output from its potential; in the original paper of [Phillips \(1958\)](#), the focus was the negative relationship between wage inflation and unemployment

³This result follows straightforwardly from the basic New Keynesian model as derived in [Clarida, Gali and Gertler \(1999\)](#), while similar results would obtain in the classic setting of [Barro and Gordon \(1983\)](#).

path of inflation and the output gap – and the case of commitment, in which the authority credibly commits to a future plan. We find that the main intuition goes through in both cases. The difference lies in the implied properties of the statistical process for inflation generated by the optimal policy in each case. In the simple framework studied here, the greater degree of inertia under optimal commitment also offers one potential solution to the identification problem.

A second extension introduces shocks to the targeting rule. These shocks can be interpreted as lags in monetary transmission; as shocks to the monetary policy instrument rule; or, in a multi-region setting, as idiosyncratic demand shocks affecting different regions or countries within a monetary union. We show that the relative variance of these shocks vis-a-vis the cost-push shocks is key for the empirical identification of the Phillips curve using standard regression analysis. This result also rationalises the findings of the vast empirical literature that uses identified monetary policy shocks to estimate the transmission of monetary policy. Effectively, well-identified monetary policy shocks should help in retrieving the Phillips curve.

A third extension studies a multi-region (multi-country or multi-sector) setting with a common central bank and discusses conditions under which regional (or sectoral) data can help mitigate the bias from the endogeneity of monetary policy. The discussion, however, also underscores some of the limitations faced by the regional analysis.

The final extension turns to the estimation of a wage-Phillips curve and compares the identification challenges with those faced in the price-Phillips curve.

That the empirical Phillips curve may vary with monetary policy was one of the examples given by [Lucas \(1976\)](#), and similar points have been echoed in different forms by other authors since, often focusing on the effect of monetary policy on inflation expectations.⁴ Similarly, [Mankiw, Ball and Romer \(1988\)](#) showed how increases in average inflation rates, by changing the frequency with which firms reset prices, could change the slope of the Phillips curve. Others have modelled a situation when policymakers themselves set policy based on a misspecified or unidentified Phillips Curve ([Haldane and Quah, 1999](#); [Primiceri, 2006](#); [Sargent, Williams and Zha, 2006](#)). In these papers, mistakes or imperfect information on the part of policymakers can lead to changes in inflation expectations that cause the reduced-form Phillips curve to disappear.

⁴For a recent example, see the explanation in [Del Negro, Giannoni and Schorfheide \(2015\)](#) of the ‘missing disinflation’ following the financial crisis.

In contrast, [Haldane and Quah \(1999\)](#), [Roberts \(2006\)](#), [Mishkin \(2007\)](#) and [Carlstrom, Fuerst and Paustian \(2009\)](#) show how good monetary policy can also endogenously offset some of the variation in real activity, which leads to the identification issue we focus on here: the Phillips curve disappears when monetary policy is successful. Despite these contributions, a surprisingly bulky literature has not resisted the temptation to keep searching, against all odds, for the Phillips curve in the data. This paper sets out, in the simplest framework possible, the identification challenge that must be addressed, while also rationalising findings in various strands of the empirical literature and discuss some possible solutions.

The paper is organised as follows. Section 2 introduces a simple model of optimal policy embedding the Phillips curve and illustrates the ‘exogeneity result’ or disconnect between equilibrium inflation and output gap under the assumption that the monetary authority cannot commit to a future path of inflation (discretion). Section 3 illustrates the empirical identification problem. Section 4 presents and discusses extensions of the model and explains how some of those extensions map into practical solutions to the identification problem. Section 5 contains concluding remarks.

2. OPTIMAL INFLATION IN THE BASIC NEW KEYNESIAN MODEL

This section uses an optimal monetary policy framework to illustrate why, in equilibrium, one should expect inflation to follow a seemingly exogenous process, unrelated (or negatively related) to measures of slack.

To explain the intuition as starkly as possible, we use the canonical New Keynesian model, as derived in [Clarida, Gali and Gertler \(1999\)](#), [Woodford \(2003\)](#) and elsewhere. Here we closely follow the textbook exposition from [Galí \(2008\)](#). For now, we dispense with the usual IS equation determining aggregate demand. This equation is necessary only to determine how policy is implemented. In the basic model it does not constrain equilibrium outcomes, so we can equivalently consider the policymaker as directly choosing the output gap as their policy instrument. Our model therefore consists of just two equations: a Phillips curve and a description of optimal monetary policy.

The (log-linearised) New Keynesian Phillips curve is given by

$$\pi_t = \beta E_t \pi_{t+1} + \kappa x_t + u_t \quad (1)$$

where π_t is the deviation of inflation from its target; x_t is the output gap, measured as the difference between output and its potential level⁵ and u_t is a cost-push shock that follows an exogenous AR(1) process with persistence ρ ($u_t = \rho u_{t-1} + \epsilon_t$, where ϵ_t are i.i.d. and mean zero). We *assume* that the Phillips curve has a strictly positive slope, denoted by $\kappa > 0$.

The Phillips curve is evidently alive and well in the model: it is the *only* equation making up its non-policy block. By construction, we have a positively sloped Phillips curve. Increases in the output gap clearly increase inflation and falls in the output gap reduce it. Nonetheless, once we augment the model with a description of optimal monetary policy, this relationship will not be apparent in the data. Inflation will instead inherit the properties of the exogenous shock process u_t .

To show this, we assume that the policymaker sets monetary policy optimally under discretion. Period by period, she minimises the following quadratic loss function

$$L_t = \pi_t^2 + \lambda x_t^2$$

subject to the constraint (1) and taking expectations of future inflation as given.⁶ The solution to the minimisation problem is the policymaker's optimal targeting rule

$$\pi_t = -\frac{\lambda}{\kappa} x_t \quad (2)$$

When faced with a positive cost-push shock that creates a trade-off between the inflation and output stabilisation objectives, the policymaker balances them, creating a negative output gap to reduce the degree of above-target inflation. The relative weight placed on each objective depends on the policymaker's preference parameter λ .

The Phillips curve (1) and optimal targeting rule (2) together completely determine the path of

⁵In the full model derived in Galí (2008), this is the welfare-relevant gap between output and its efficient level.

⁶Clarida, Gali and Gertler (1999) show how minimising such a loss function is equivalent to maximising the welfare of the representative agent in the model. But it can alternatively be motivated as a simple way to capture the preferences enshrined in the mandates of modern (flexible) inflation targeting central banks: see Carney (2017), for example.

inflation in the model. We can solve for equilibrium inflation by using (2) to substitute out for x_t in (1), and by iterating forward to obtain

$$\pi_t = \frac{\lambda}{\kappa^2 + \lambda(1 - \beta\rho)} u_t \quad (3)$$

In equilibrium, inflation deviations are at all times perfectly proportional to the exogenous cost-push shock. In other words, with a constant target, equilibrium inflation itself behaves as an exogenous process. In the limit, when the monetary authority does not put any weight on the output gap ($\lambda = 0$), inflation equals the target rate, a point previously made by [Haldane and Quah \(1999\)](#).

This behaviour is entirely consistent with recent empirical work by [Cecchetti et al. \(2017\)](#) and [Forbes, Kirkham and Theodoridis \(2017\)](#) suggesting that inflation data in the US and the UK can be modelled as an exogenous statistical process, unrelated to measures of slack.⁷ But crucially, the basic theory is also built under the assumption that monetary policy is at all times constrained by a working Phillips curve. There is no discrepancy between the two results. The Phillips curve may be the correct structural model of the inflation process, but that does not mean that one should observe it in the empirical relationship between (equilibrium levels of) inflation and the output gap.

The reason is simple. The policymaker in the model is able to set policy to achieve any desired level of the output gap. Successful monetary policy should lean against any undesirable deviations in output from potential, which would otherwise cause inflationary or deflationary pressures. Precisely because monetary policy can be used to offset the effect of such output gaps on inflation, their effect should not be visible in the data.

Optimal monetary policy does not seek to eliminate all output volatility: from (2) we can see that in response to cost-push shocks, the policymaker will prefer to tolerate output deviations from potential. But such shocks impart a negative correlation between inflation and output, rather than a positive one. Again, the more successful monetary policy is in managing any trade-offs between inflation and output, the more it will blur the underlying positive Phillips curve correlation.

To summarise, we have shown that with an optimizing monetary policy, equilibrium levels of inflation are unrelated to measures of slack. Instead, inflation inherits the statistical properties of


⁷It is also consistent with the observation that in larger DSGE models such as [Smets and Wouters \(2007\)](#), inflation is largely explained by exogenous markup shocks ([King and Watson, 2012](#)).

exogenous cost-push shocks. This does not necessarily tell us that the Phillips curve is not present. In the model, the Phillips curve exists and policymakers are completely aware of its existence. But because they know exactly how the curve operates, they are able to perfectly offset its effects on equilibrium inflation.⁸

3. PHILLIPS CURVE IDENTIFICATION

As may already be apparent from the discussion in Section 2, regression analysis will have difficulty in recovering the Phillips curve. Figure 1 shows data simulated from the model described by (1) and (2), with parameters calibrated as in Galí (2008). Specifically, the slope of the Phillips curve is set at $\kappa = 0.1275$, the policymaker’s weight on output deviations relative to quarterly inflation is set as $\lambda = 0.0213$, or around one-third relative to annualised inflation. The discount factor is set to $\beta = 0.99$ and the persistence of the cost-push shock to $\rho = 0.5$.

Figure 1: Inflation/output gap correlation in model-simulated data



Notes. 1000 periods of data are simulated from the model described by (1) and (2). We draw each ϵ_t from a standard normal distribution.

⁸Stock and Watson (2009) raise the possibility that, despite its failure to forecast or explain the data, the Phillips curve is still useful for conditional forecasting. They pose the question ‘...suppose you are told that next quarter the economy would plunge into recession, with the unemployment rate jumping by 2 percentage points. Would you change your inflation forecast?’

Of course, there is no Phillips curve visible in the simulated data. As can be seen from the line of best fit, a naive OLS regression of inflation on the output gap,

$$\pi_t = \gamma_1 x_t + \varepsilon_t \quad (4)$$

will produce a negative parameter estimate, $\hat{\gamma}_1 = -\frac{1}{6}$, reflecting the targeting rule (2), rather than a consistent estimate of the positive slope of the Phillips curve. Many papers have focused on the difficulty of controlling for inflation expectations in Phillips curve estimation, but the problem here is a more straightforward one.⁹

The issue here is a simple case of simultaneity bias. The regressor x_t is correlated with the error term ε_t . The naive econometrician does not observe the Phillips curve in the data. Rather, he or she observes equilibrium inflation and output gap outturns: which are the intersection of the Phillips curve (1) and the targeting rule (2). In fact, the case here is an extreme one: the regressor and the error are perfectly negatively correlated.¹⁰ The issue is completely analogous to the classic case of simultaneity bias: jointly determined supply and demand equations.

To show the identification challenge, we first plot the two model equations in Figure 2.¹¹ The Phillips curve (1) is in blue, the optimal targeting rule (2) in red, while the black circles index the policymaker's loss function at different levels of loss. The observed inflation-output gap pairs are the equilibrium where the two lines intersect. With no cost-push shocks to the Phillips curve, the first-best outcome of at target inflation and no output gap is feasible, so the lines intersect at the origin.

When the upward sloping Phillips curve is subject to cost-push shocks, the equilibrium shifts to different points along the optimal targeting path, shown in Figure 3. But with monetary policy set optimally, there are no shifts along the Phillips curve: at all times the equilibrium remains on the negatively sloped optimal targeting rule line. As a result, the simulated data trace out the optimal targeting rule, not the Phillips curve. The estimated coefficient is $\hat{\gamma}_1 = -\frac{\lambda}{\kappa} = -\frac{1}{6}$.

The issue is that the Phillips curve is not identified. Our simple set-up has no exogenous

⁹See [Nason and Smith \(2008\)](#), [Mavroeidis, Plagborg-Møller and Stock \(2014\)](#) and [Krogh \(2015\)](#) for discussions.

¹⁰Using (3) to substitute out for π_t in (2) gives the equilibrium evolution of the output gap $x_t = -\frac{\kappa}{\kappa^2 + \lambda(1 - \beta\rho)} u_t$. While the regression error term is equal to $\varepsilon_t = u_t + \beta E_t \pi_{t+1} = (1 + \frac{\rho\lambda}{\kappa^2 + \lambda(1 - \beta\rho)}) u_t$

¹¹This graphical illustration of optimal discretionary policy is from [Seneca \(2018\)](#): we are grateful to him for making it available to us. A similar graphical exposition also appears in [Carlin and Soskice \(2005\)](#).

variables shifting monetary policy. Worse, the only shocks are to the equation of interest, so the estimated parameter is almost entirely unrelated to the slope of the Phillips curve.¹² The problem is the same one that arises when trying to identify a supply curve while only observing equilibrium quantities and prices. Without any exogenous demand shifter, there is no way of doing so.

4. EXTENSIONS TO THE BASIC MODEL AND SOLUTIONS TO THE ESTIMATION CHALLENGE

In this section we study a number of extensions to the basic model. For each extension, we discuss whether and how it can help solving the Phillips Curve’s empirical identification problem. In subsection 4.1, we discuss the case in which the monetary authority can commit to a path of inflation and output gap. In subsection 4.2, we allow for shocks to the targeting rule and we discuss how they link to the identified monetary policy shocks in the monetary policy transmission literature. In subsection 4.3, we study a multi-region setting. In 4.4 we discuss the mapping into a wage Phillips curve and comment on the Phillips curve in bigger models.

¹²Other than the fact that the slope of the Phillips curve happens to appear in the optimal targeting rule.

Figure 2: Graphical illustration of optimal monetary policy under discretion





Figure 3: Graphical illustration of optimal discretionary policy in response to cost-push shocks



4.1. Commitment

First, we show that our main results are unchanged when the monetary policymaker is able to commit to a future plan for inflation and the output gap. In Sections 2 and 3 we assumed that the policymaker was unable to commit. There are a range of practical issues that may make commitment difficult: monetary policy committees often have changes in membership and future policymakers may not feel bound by prior commitments and perhaps relatedly, successful commitment requires that promises are credible, even when they are time inconsistent. Nonetheless, the optimal commitment policy is able to achieve better outcomes in the face of cost-push shocks than optimal policy under discretion, so it is important to know how this affects our results.

It turns out that the same intuition holds, although the precise details slightly differ. Again following Galí (2008), when the policymaker instead minimises the loss function

$$L = E_0 \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda x_t^2) \quad (5)$$

subject to the sequence of Phillips curves given by (1) for each period. This gives a pair of optimality conditions

$$\pi_0 = -\frac{\lambda}{\kappa}x_0 \quad (6)$$

$$\pi_t = -\frac{\lambda}{\kappa}(x_t - x_{t-1}) \quad (7)$$

These can be combined to give the targeting rule under commitment

$$p_t = -\frac{\lambda}{\kappa}x_t \quad (8)$$

where p_t is the log deviation of the price level from its level in period -1 . Substituting $p_t - p_{t-1}$ for π_t in (1) and substituting out x_t using (8) gives a difference equation in p_t . Galí (2008) shows the solution for this in terms of the previous period's price level and the current period cost-push shock. Iterating backwards and then taking the first difference gives equilibrium inflation

$$\pi_t = \frac{\delta}{1 - \delta\beta\rho} \left(u_t - (1 - \delta) \sum_{i=0}^{t-1} \delta^{t-1-i} u_i \right) \quad (9)$$

where $\delta \equiv \frac{((\lambda(1+\beta)+\kappa^2)-((\lambda(1+\beta)+\kappa^2)^2-4\beta\lambda^2)^{0.5})}{2\lambda\beta}$.


Equilibrium inflation under optimal commitment policy depends solely on the cost-push shock process. The equilibrium path is quite different to that under discretion, however. At any point in time inflation displays history dependence, depending on the entire history of cost-push shocks rather than just the one in the current period.

Simple regressions will again fail to uncover the Phillips curve. The only difference is that under commitment, the optimal targeting rule imposes a negative correlation between the output gap and the *price level*. The relationship between inflation and the output gap in the simulated data shown in Figure 4 is noisier, but shows no sign of the Phillips curve embedded in the model. The OLS estimate of γ in (4) gives the coefficient $\hat{\gamma}_1 = -0.085$.

At least in the simple framework here, the history-dependence of optimal commitment policy also suggests a straightforward solution to the identification problem. From (7), the equilibrium output gap will be correlated with its own lagged values, which can therefore be used as an

instrument. Intuitively, the policymaker chooses to create an output gap even after the cost-push shock has disappeared. They commit to do so in order to achieve better inflation outcomes when the shock originally occurs. The policymaker therefore optimally reintroduces the positive Phillips curve relation that is absent under optimal discretion. As a result, in the simple case here, a suitable choice of instrument will be able to recover the true Phillips curve slope.

Figure 4: *Inflation/output gap correlation in model-simulated data: optimal commitment*



Notes. 1000 periods of data are simulated from the model described by (1) and (7). We draw each ϵ_t from a standard normal distribution.

4.2. Shocks to the targeting rule

The previous sections have illustrated how successful monetary policy might mask the underlying structural Phillips curve in the data. We now show that the opposite is also true in our model: if monetary policy is set far from optimally, the Phillips curve is likely to reappear.

So far we have assumed policymakers can implement monetary policy by directly choosing their desired observable output gap each period. But alas in practice, policymaking is not quite so simple. In empirical studies we observe lags between changing policy and its impact on the output gap and inflation, which means that in practice central banks are inflation *forecast* targeters (Svensson, 1997; Haldane, 1998). Forecast errors will therefore inject noise into the targeting rule.

Potential output is unobservable, so the output gap must be estimated (with error). And the effect of the policy instruments actually available (typically the central bank policy rate and forward guidance on its future path; as well as quantitative easing) on the target variables is also unknown. Errors from any of these sources will insert noise into the desired balance between inflation and output gap deviations. These various shocks to the targeting rule correspond closely to the typical interpretations of identified monetary policy shocks in the vast empirical literature on this topic (Christiano, Eichenbaum and Evans, 1996, 1999; Romer and Romer, 2004; Olivei and Tenreyro, 2007; Cloyne and Hürtgen, 2016). That literature is able to identify a positively correlated response of inflation and the output gap to monetary policy shocks, in line with the results below.

Returning to optimal policy under discretion, we model implementation errors by simply adding an AR(1) shock process e_t to the targeting rule (2) to give

$$\pi_t = -\frac{\lambda}{\kappa}x_t + e_t \quad (10)$$

where $e_t = \rho_e e_{t-1} + \zeta_t$ and ζ_t is a zero-mean i.i.d. with variance σ_e^2 .¹³ We can show that equilibrium inflation and the output gap now both have an additional term proportional to e_t . Respectively, they are given by $\pi_t = s_1 \lambda u_t + s_2 \kappa e_t$ and $x_t = -s_1 \kappa u_t + s_2 (1 - \rho_e) e_t$, where $s_1 \equiv \frac{1}{\lambda(1-\rho_e) + \kappa^2}$ and $s_2 \equiv \frac{\kappa}{\lambda(1-\rho_e) + \kappa^2}$.

With shocks to the targeting rule, neither equation is identified. The equilibrium values of inflation and the output gap both depend on a combination of both shocks. Consequently, if either equation is estimated by OLS, its regressor will be correlated with the regression error term and the resulting parameter estimate inconsistent. In particular, it follows from substituting the equilibrium values of π_t and x_t into the definition of the OLS estimator in the regression (4) that


$$\text{plim}(\hat{\gamma}) = \frac{\text{plim}\left(\frac{1}{T} \sum_{t=1}^T x_t \pi_t\right)}{\text{plim}\left(\frac{1}{T} \sum_{t=1}^T x_t^2\right)} = \frac{\frac{-\lambda s_1}{\kappa} \frac{\sigma_u^2}{s_2 \sigma_u^2 + \sigma_e^2} + (1 - \beta \rho_e) \kappa \frac{\sigma_e^2}{\sigma_u^2 + \sigma_e^2}}{\frac{s_1}{s_2} \frac{\sigma_u^2}{\sigma_u^2 + \sigma_e^2} + (1 - \beta \rho_e)^2 \frac{\sigma_e^2}{\sigma_u^2 + \sigma_e^2}} \quad (11)$$

The size of the simultaneity bias to each equation depends on the relative variances of the shocks. Figure 5 plots simulated data for three cases. We set $\rho_e = 0.5$ and set the other parameters as before.

¹³Clarida, Gali and Gertler (1999) and Svensson and Woodford (2004) show in the basic New Keynesian model that when there are policy control lags that mean all variables are predetermined in advance, up to an unforecastable shock, the optimal targeting rule will take exactly this form, where e_t is the forecast error.

First, the red circles show the case where the cost-push shock has a variance 100 times larger than the targeting rule shock. These look almost identical to the case with only a cost-push shock: the circles trace out the targeting rule. Second, the green circles show the case when the shocks have equal variance. The slope is still negative, but flatter. The final case gives the cost-push shock a variance 100 times smaller than the targeting rule shock, and the data trace out a positively sloped line.

Figure 5: Inflation/output gap correlation in model-simulated data: optimal discretion with shocks to the targeting rule



Notes. 1000 periods of data are simulated from the model described by (1) and (10). The green circles show the case when each ϵ_t and ζ_t is drawn from a standard normal distribution. The blue circles show the case when each ϵ_t is drawn from an $N(0,10)$ distribution and the red circles each ζ_t is instead drawn from an $N(0,10)$ distribution.

Looking at the regression coefficients in Table 1, in the first two cases these are both strongly influenced by the endogenous policy response embodied in the optimal targeting rule. It also makes little difference whether or not the econometrician correctly controls for inflation expectations, which also enter the Phillips curve. In the third case however, the regression coefficient turns positive. The estimate is actually upward biased in specification 5, which omits inflation expectations. Once these are controlled for, the bias becomes very small. The regression correctly identifies the slope of the Phillips curve to four decimal places.

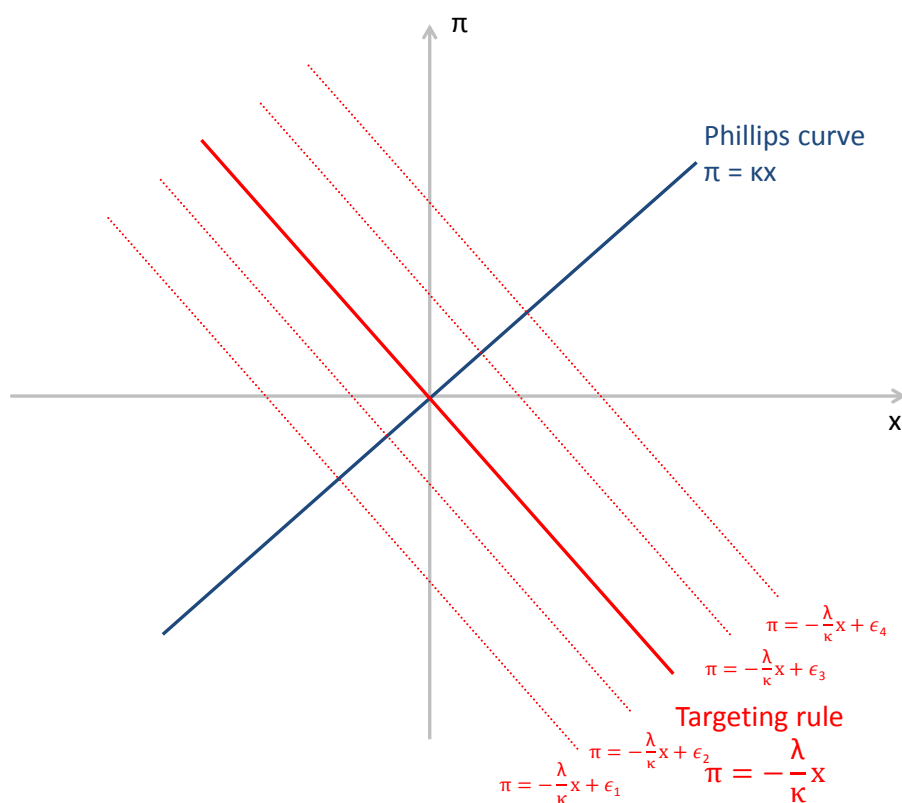
The reason the bias disappears is straightforward. When cost-push shocks have a relatively low

Table 1: OLS regressions of inflation on the output gap in the simulated data

	(1)	(2)	(3)	(4)	(5)	(6)
	$\frac{\sigma_y^2}{\sigma_\varepsilon^2} = 100$	$\frac{\sigma_y^2}{\sigma_\varepsilon^2} = 100$	$\frac{\sigma_y^2}{\sigma_\varepsilon^2} = 1$	$\frac{\sigma_y^2}{\sigma_\varepsilon^2} = 1$	$\frac{\sigma_y^2}{\sigma_\varepsilon^2} = 0.01$	$\frac{\sigma_y^2}{\sigma_\varepsilon^2} = 0.01$
LHS variable	π_t	$\pi_t - E_t\pi_{t+1}$	π_t	$\pi_t - E_t\pi_{t+1}$	π_t	$\pi_t - E_t\pi_{t+1}$
x_t	-0.1667	-0.1805	-0.0873	-0.0792	0.2523	0.1275

Notes. Table shows the OLS regression coefficients of OLS for the shock distributions described in the notes to Figure 5. Specifications (2), (4) and (6) (perfectly) control for inflation expectations by subtracting from the dependent variable the true value of $\beta E_t\pi_{t+1}$. The true slope of the Phillips curve is $\kappa = 0.1275$, while the true slope of the optimal targeting rule is $-\frac{\lambda}{\kappa} = -0.1667$.

Figure 6: Graphical illustration of optimal discretionary policy in response to targeting-rule shocks



variance, most of the variation in the simulated data arises from the shocks to the targeting rule. With the Phillips curve stable, these movements in the targeting rule now trace out the Phillips curve, as shown graphically in Figure 6.

4.3. Regional Phillips curves

Partly to avoid the difficulties associated with identifying the Phillips curve at the national level, a number of authors have estimated Phillips curves at a more disaggregated, regional or sectoral level

(Fitzgerald and Nicolini, 2014; Kiley, 2015; Tuckett, 2018; Vlieghe, 2018). In this subsection we show that in an extended version of the basic model, this may also help the econometrician to identify the aggregate Phillips curve.

The key to identification is that at the regional level, the endogenous response of monetary policy to demand shocks is switched off, ameliorating the simultaneity bias in estimating aggregate Phillips curves. This point was made by Fitzgerald and Nicolini (2014) as motivation for their estimation of Phillips curves at a regional level. The same logic can explain why the Phillips curve may be more evident in countries within a monetary union such as the euro area.

We assume that the aggregate Phillips curve (1) continues to hold, but that aggregate inflation and the aggregate output gap also depend on the weighted average of inflation and the output gap in each of n regions

$$\pi_t = \sum_{i=1}^n \alpha_i \pi_t^i \quad (12)$$

$$x_t = \sum_{i=1}^n \alpha_i x_t^i \quad (13)$$

where regional inflation is determined by a regional Phillips curve analogous to (1)

$$\pi_t^i = \beta E_t \pi_{t+1}^i + \kappa_i x_t^i + u_t^i \quad (14)$$

where the idiosyncratic cost-push shocks $u_t^i = \rho u_{t-1}^i + \epsilon_t^i$ and ϵ_t^i are zero-mean and i.i.d over time, but potentially correlated across regions. We must also specify how idiosyncratic demand shocks and aggregate monetary policy affect the regional output gap with an equation analogous to the IS curve in the basic New Keynesian model, given by

$$x_t^i = E_t x_{t+1}^i - \sigma^{-1} (i_t - E_t \pi_{t+1}^i - r_t^i) \quad (15)$$

where the idiosyncratic demand shocks are given by $r_t^i = \rho_r r_{t-1}^i + e_t^i$ and e_t^i are zero-mean and i.i.d over time, but potentially correlated across regions. The equations can be aggregated together to

give the usual aggregate IS relation

$$x_t = E_t x_{t+1} - \sigma^{-1}(i_t - E_t \pi_{t+1} - r_t) \quad (16)$$

We therefore allow inflation and the output gap are determined partly by idiosyncratic shocks to each region, but restrict the monetary policy rate i_t to be the same across all n regions.

We next denote for any regional variable its (log) deviation from the aggregate as $\hat{z}_t^i = z_t^i - \sum_{i=1}^n \alpha_i z_t^i$. We can then subtract (1) from (14) to give a Phillips curve in terms of log deviations from aggregate inflation.

$$\hat{\pi}_t^i = \beta E_t \hat{\pi}_{t+1}^i + \kappa_i \hat{x}_t^i + \hat{u}_t^i \quad (17)$$

Subtracting (16) from (15) gives an equivalent IS curve

$$\hat{x}_t^i = E_t \hat{x}_{t+1}^i + \sigma^{-1}(E_t \hat{\pi}_{t+1}^i + \hat{r}_t^i) \quad (18)$$

Monetary policy is set (under discretion) by minimising the same aggregate period loss function as in Section 2, subject to the aggregate Phillips curve (1).¹⁴ Policy therefore follows the same targeting rule (2) depending solely on aggregate variables.¹⁵

The crucial difference to the identification problem at the regional level is that while monetary policy perfectly offsets the aggregate demand shocks, $r_t = \sum_{i=1}^n r_t^i$, it does not respond at all to the idiosyncratic regional deviations from that average, \hat{r}_t^i . The regressor in the Phillips curve equation \hat{x}_t^i is now affected by exogenous demand shocks that do not enter the regional Phillips curve. As a result, the endogeneity problem is mitigated.

For each region, we can verify that one solution to the model described by (17) and (18) is

$$\hat{\pi}_t^i = c_1(1 - \rho)\hat{u}_t^i + c_2\kappa_i\hat{r}_t^i \quad (19)$$

¹⁴This differs from the monetary policy that would be welfare-optimal in the model, since welfare would also be lowered by dispersion in prices within a region, even if average inflation was zero. Clarida, Gali and Gertler (2001) show in the context of an open economy model that the welfare-optimal policy would minimise a loss function that included the sum across countries of the squared deviations of inflation, rather than the square of the sum of deviations.

¹⁵Although to ensure determinacy, the policymaker's instrument rule will need to respond to idiosyncratic variables.

and

$$\hat{x}_t^i = c_1 \rho \sigma^{-1} \hat{u}_t^i + c_2 (1 - \rho_r \beta) \hat{r}_t^i \quad (20)$$

where $c_1 \equiv \frac{1}{(1-\rho)(1-\rho\beta)-\rho\kappa_i\sigma^{-1}}$ and $c_2 \equiv \frac{\sigma^{-1}}{(1-\rho_r)(1-\rho_r\beta)-\rho_r\kappa_i\sigma^{-1}}$.¹⁶ Unlike aggregate inflation, which evolves in line with the exogenous shocks to the Phillips curve, regional inflation also depends on idiosyncratic demand shocks. In the simplest case when the shocks are entirely transitory ($\rho = \rho_r = 0$), the equilibrium output gap will be independent of the idiosyncratic cost-push shocks \hat{u}_t^i and a simple regression of $\hat{\pi}_t^i$ on \hat{x}_t^i will give a consistent estimate of κ_i .

Away from that special case, inflation expectations present a separate challenge to identifying regional Phillips curves. With $\rho > 0$ or $\rho_r > 0$, there will be omitted variable bias unless the econometrician can control for the effect of regional inflation expectations. While possible in principle, reliable data are likely to be less readily available than at the national level. There is perhaps likely to be more chance of success when estimating at the country level within a single multi-country monetary authority. A second difficulty at the regional level is that while the specification will help mitigate the bias from the endogeneity of national monetary policy, insufficient cross-sectional variation in the regional data will lead to imprecise estimates of κ_i .

4.4. The wage Phillips curve

While identification of the price Phillips curve is complicated by the endogenous response of optimal monetary policy, the focus of the original Phillips study was the correlation between wage inflation and unemployment in the UK. Especially in the UK, a number of researchers have had more success in finding the latter relationship. In this subsection we comment on how optimal monetary policy maps into the original wage Phillips curve relationship between wage inflation and unemployment. Intuitively, one might expect the wage Phillips curve to be less vulnerable to identification issues related to the endogeneity of monetary policy, since wage inflation is one step removed from the price-inflation targeting remit of most central banks.

As well as a different dependent variable (wage inflation rather than price inflation), the

¹⁶While this is one solution, depending on how policy is implemented, there may be a multiplicity of equilibria. It is beyond the scope of this paper to study those, so we assume that the policymaker's instrument rule is able to rule them out. In practice, this will involve responding to deviations of regional inflation or regional output gaps from their equilibrium values, even when those deviations have no impact on aggregate inflation or the aggregate output gap.

typical wage Phillips curve attempts to explain inflation using variation in unemployment or the unemployment gap, rather than the output gap. Using unemployment in the equation is unlikely to solve the identification issues arising from the behaviour of monetary policy for at least two reasons.

First, many central banks' remits explicitly specify unemployment or employment as one of their (secondary or dual) target variables. As such, they will optimally set policy to close any gap between unemployment and its natural rate, unless there is a trade-off between that goal and their inflation targets, in which case they will seek to balance the two goals, as was the case with the output gap in Section 2. Monetary policy will therefore blur the structural relationship between inflation and the unemployment gap in a similar way. Second, even for central banks without an explicit mandate to minimise fluctuations in employment, when there is comovement between the output gap and the unemployment gap, policy will often implicitly seek to stabilise employment.¹⁷

There are, however, reasons to think that using wage inflation as the dependent variable might lessen some of the identification problems. Nominal wage rigidities can be incorporated into the basic model in an analogous way to price rigidities, as introduced by [Erceg, Henderson and Levin \(2000\)](#). With both wage and price stickiness, some shocks, such as innovations to firms' desired price-markups, will lead to a wedge between the rate of price inflation and the output gap, but not between the rate of wage inflation and the output gap. Since inflation targeting central banks typically target price inflation, policymakers may respond by adjusting the output gap to achieve their desired trade-off with price inflation. But doing so would lead to variation in wage inflation operating via the wage Phillips curve. Put differently, the equilibrium output gap will be a function of the exogenous shocks hitting the economy. But if those shocks only directly affect the price Phillips curve and not the wage Phillips curve, then the output gap will be correlated with the error term in the former but not the latter, which will be consistently estimated.

The wage Phillips curve may not face quite as severe problems, but there remain limits to how easily it can be identified under optimal monetary policy. First, while there may be some shocks that only affect the price Phillips curve, there are likely to be several more that affect both curves (for a given output gap). Wage mark-up shocks will increase both price and wage inflation relative to the prevailing output gap. [Erceg, Henderson and Levin \(2000\)](#) show that shocks to household

¹⁷[Galí \(2011\)](#) shows how the basic framework can be easily extended to include unemployment in a way that closely resembles the output gap in the basic model.

consumption or leisure preferences, or to total factor productivity, will conversely move price and wage inflation in opposite directions for a given output gap. Since the inflationary impact of these shocks will lead policymakers to attempt to lean against them via the output gap, this will induce a correlation between the output gap and the shocks affecting the wage Phillips curve (for a given output gap). The direction of the bias will differ according to the shock, but the equation will in general not be identified.

Second, even if price inflation shocks are particularly prevalent, many typical examples of such shocks, such as changes in oil prices, have relatively transitory effects on price inflation. Since monetary policy is typically thought to have its peak effect on inflation with some lag, attempting to offset very transitory shocks may not be possible. As a result, policymakers are perhaps less likely to respond to the very shocks that would otherwise have helped econometricians identify the wage Phillips curve. Conversely, when transitory shocks are affecting price inflation, wage inflation can sometimes give a better signal of underlying price pressures, which may lead policymakers to behave at times as if they were targeting wage inflation.¹⁸

In addition to nominal wage rigidities, larger macroeconomic models of the type used for policy analysis in central banks usually have a range of other frictions, additional factors of production and a richer dynamic structure.¹⁹ The same logic we have outlined in this paper also complicates single-equation identification of the Phillips curve in data simulated from larger models. Given their richer structure and wider variety of shocks, if one were able to estimate the full structural model and there was enough variation in the data, then it would be possible to recover any Phillips curve relationship. But precisely because we do not know the true model of the economy, such an approach may be less robust to misspecification. Single equation OLS estimation, meanwhile, will be biased by the effect of monetary policy just as in the simple model. Monetary policy that is successful in targeting inflation and in closing the output gap will be definition be successful in offsetting the effect of any shocks that affect inflation via the Phillips curve.²⁰

¹⁸In addition, the welfare optimal policy in models with sticky wages typically involves placing a positive weight on avoiding wage inflation (Erceg, Henderson and Levin, 2000). But we are not aware of any central banks who officially target wage inflation in practice.

¹⁹See for example Edge, Kiley and Laforge (2010); Burgess et al. (2013); Brubakk and Sveen (2009); Adolfson et al. (2013) for descriptions of models used respectively at the Federal Reserve Board, the Bank of England, Norges Bank and the Riksbank.

²⁰Identifying a Phillips curve is further complicated in larger models by the fact that a stable, structural relationship between the output gap and firms' marginal costs (and therefore inflation) may not exist. The reduced-form Phillips curve correlation typically varies for different shocks.

5. CONCLUSION

We use standard analytical framework to explain why inflation follows a seemingly exogenous statistical process, or, in other words, why the Phillips curve cannot be easily identified with macroeconomic data. In the framework, a monetary authority minimizes welfare losses, measured as deviations of inflation and output from their targets, subject to a Phillips curve. This leads the authority to follow an optimal targeting rule in which it seeks to increase inflation when the output gap decreases. This imparts a negative relation between inflation and the output gap that blurs the identification of the positively sloped Phillips curve. In equilibrium, inflation inherits the statistical properties of any cost-push shocks affecting the Phillips curves (e.g., energy price shocks, exchange rate changes, and so on).

We show that shocks to the targeting rule are key for the identification of the Phillips curve. These targeting shocks can take the form of monetary policy shocks in a Taylor rule or, in a multi-region setting or a multi-country monetary union, idiosyncratic demand shocks affecting the various regions or countries in different ways. In a univariate regression analysis, if the relative variance of these shocks is sufficiently high, vis-a-vis the variance of cost-push shocks, the slope of the Phillips curve can be identified. Similarly, in a VAR specification that allows for the identification of monetary policy shocks, the positive relationship between inflation and output gap can be distilled.

We have also shown how the simple framework here can jointly rationalise several empirical findings on the Phillips curve. First, it should be weaker in periods when there are large cost shocks – such as the 1970s – *and* when monetary policy is relatively successful in achieving its targets – as in the inflation targeting era. Second, wage Phillips curves should be more evident in the data than price Phillips curves. And third, the Phillips curve relationship should appear stronger in disaggregated panel data than in aggregate data.

To summarise, the paper explains the identification problem posited by the estimation of Phillips curves; rationalises findings in the empirical literature and discusses possible solutions to the identification problem.

REFERENCES

- Adolfson, Malin, Stefan Laseen, Lawrence Christiano, Mathias Trabandt, and Karl Walentin.** 2013. "Ramses ii-model description." *Sveriges Riksbank Occasional Paper Series*, 12: 1009.
- Atkeson, Andrew, and Lee E Ohanian.** 2001. "Are Phillips curves useful for forecasting inflation?" *Federal Reserve Bank of Minneapolis. Quarterly Review-Federal Reserve Bank of Minneapolis*, 25(1): 2.
- Ball, Laurence, and Sandeep Mazumder.** 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity*, 42(1 (Spring)): 337–405.
- Barro, Robert J, and David B Gordon.** 1983. "A positive theory of monetary policy in a natural rate model." *Journal of political economy*, 91(4): 589–610.
- Blanchard, Olivier, Eugenio Cerutti, and Lawrence Summers.** 2015. "Inflation and activity—Two explorations and their monetary policy implications." National Bureau of Economic Research.
- Brubakk, Leif, and Tommy Sveen.** 2009. "NEMO—a new macro model for forecasting and monetary policy analysis." *Economic Bulletin (Norges Bank)*, 80(1).
- Burgess, Stephen, Emilio Fernandez-Corugedo, Charlotta Groth, Richard Harrison, Francesca Monti, Konstantinos Theodoridis, and Matt Waldron.** 2013. "The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models." Bank of England Bank of England working papers 471.
- Carlin, Wendy, and David Soskice.** 2005. "The 3-equation New Keynesian Model—a graphical exposition." *Contributions in Macroeconomics*, 5(1).
- Carlstrom, Charles T, Timothy S Fuerst, and Matthias Paustian.** 2009. "Inflation persistence, monetary policy, and the great moderation." *Journal of Money, Credit and Banking*, 41(4): 767–786.
- Carney, Mark.** 2017. "Lambda." *speech at the London School of Economics, available online (January 2017)*.

- Cecchetti, Stephen G, Michael Feroli, Peter Hooper, Anil K Kashyap, and Kermit L Schoenholtz.** 2017. "Deflating Inflation Expectations: the Implications of Inflation's Simple Dynamics."
- Christiano, Lawrence J, Martin Eichenbaum, and Charles Evans.** 1996. "THE EFFECTS OF MONETARY POLICY SHOCKS: EVIDENCE FROM THE FLOW OF FUNDS." *The Review of Economics and Statistics*, 78(1): 16–34.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 1999. "Monetary policy shocks: What have we learned and to what end?" In . Vol. 1 of *Handbook of Macroeconomics*, 65 – 148. Elsevier.
- Clarida, Richard, Jordi Gali, and Mark Gertler.** 1999. "The science of monetary policy: a new Keynesian perspective." *Journal of economic literature*, 37(4): 1661–1707.
- Clarida, Richard, Jordi Gali, and Mark Gertler.** 2001. "Optimal monetary policy in open versus closed economies: an integrated approach." *American Economic Review*, 91(2): 248–252.
- Cloyne, James, and Patrick Hürtgen.** 2016. "The Macroeconomic Effects of Monetary Policy: A New Measure for the United Kingdom." *American Economic Journal: Macroeconomics*, 8(4): 75–102.
- Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide.** 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics*, 7(1): 168–96.
- Dotsey, Michael, Shigeru Fujita, and Tom Stark.** 2017. "Do Phillips Curves Conditionally Help to Forecast Inflation?"
- Edge, Rochelle M., Michael T. Kiley, and Jean-Philippe Laforte.** 2010. "A comparison of forecast performance between Federal Reserve staff forecasts, simple reduced-form models, and a DSGE model." *Journal of Applied Econometrics*, 25(4): 720–754.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin.** 2000. "Optimal monetary policy with staggered wage and price contracts." *Journal of Monetary Economics*, 46(2): 281–313.
- Fitzgerald, Terry J, and Juan Pablo Nicolini.** 2014. "Is There a Stable Relationship between Unemployment and Future Inflation? Evidence from US Cities."

- Forbes, Kristin, Lewis Kirkham, and Konstantinos Theodoridis.** 2017. "A trendy approach to UK inflation dynamics."
- Gali, J.** 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.
- Gali, Jordi.** 2011. "The return of the wage Phillips curve." *Journal of the European Economic Association*, 9(3): 436–461.
- Haldane, Andrew, and Danny Quah.** 1999. "UK Phillips curves and monetary policy." *Journal of Monetary Economics*, 44(2): 259–278.
- Haldane, Andrew G.** 1998. "On inflation targeting in the United Kingdom." *Scottish Journal of Political Economy*, 45(1): 1–32.
- IMF.** 2013. "The Dog that Didnt Bark: Has Inflation Been Muzzled or Was it Just Sleeping? Chapter 3." *World Economic Outlook (IMF)*.
- Kiley, Michael T.** 2015. "An evaluation of the inflationary pressure associated with short-and long-term unemployment." *Economics Letters*, 137: 5–9.
- King, Robert G, and Mark W Watson.** 2012. "Inflation and Unit Labor Cost." *Journal of Money, credit and Banking*, 44(s2): 111–149.
- Krogh, Tord S.** 2015. "Macro frictions and theoretical identification of the New Keynesian Phillips curve." *Journal of Macroeconomics*, 43: 191–204.
- Lucas, Robert E.** 1976. "Econometric policy evaluation: A critique." *Carnegie-Rochester Conference Series on Public Policy*, 1: 19 – 46.
- Mankiw, N.G., Laurence Ball, and David Romer.** 1988. "The New Keynesian Economics and the Output-Inflation Trade-off." *Brookings Papers on Economic Activity*, 1: 1–65.
- Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H Stock.** 2014. "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature*, 52(1): 124–88.

- Mishkin, Frederic S.** 2007. "Inflation dynamics." *International Finance*, 10(3): 317–334.
- Nason, James M, and Gregor W Smith.** 2008. "Identifying the new Keynesian Phillips curve." *Journal of Applied Econometrics*, 23(5): 525–551.
- Olivei, Giovanni, and Silvana Tenreyro.** 2007. "The Timing of Monetary Policy Shocks." *American Economic Review*, 97(3): 636–663.
- Phillips, Alban W.** 1958. "The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957." *economica*, 25(100): 283–299.
- Primiceri, Giorgio E.** 2006. "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy*." *The Quarterly Journal of Economics*, 121(3): 867–901.
- Roberts, John M.** 2006. "Monetary Policy and Inflation Dynamics." *International Journal of Central Banking*.
- Romer, Christina D, and David H Romer.** 2004. "A new measure of monetary shocks: Derivation and implications." *American Economic Review*, 94(4): 1055–1084.
- Sargent, Thomas, Noah Williams, and Tao Zha.** 2006. "Shocks and Government Beliefs: The Rise and Fall of American Inflation." *American Economic Review*, 96(4): 1193–1224.
- Seneca, Martin.** 2018. "A graphical illustration of optimal monetary policy in the New Keynesian framework."
- Smets, Frank, and Rafael Wouters.** 2007. "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American economic review*, 97(3): 586–606.
- Stock, James, and Mark W. Watson.** 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 39(s1): 3–33.
- Stock, James, and Mark W. Watson.** 2009. "Phillips Curve Inflation Forecasts." *Understanding Inflation and the Implications for Monetary Policy*, , ed. Jeffrey Fuhrer, Yolanda Kodrzycki, Jane Little and Giovanni Olivei, 99–202. Cambridge:MIT Press.

- Svensson, Lars E.O.** 1997. "Inflation forecast targeting: Implementing and monitoring inflation targets." *European Economic Review*, 41(6): 1111 – 1146.
- Svensson, Lars EO, and Michael Woodford.** 2004. "Implementing optimal policy through inflation-forecast targeting." 19–92.
- Tuckett, Alex.** 2018. "Regional data and the Phillips Curve." Bank of England Working Paper.
- Vlieghe, Gertjan.** 2018. "From asymmetry to symmetry: changing risks to the economic outlook." *speech*.
- Woodford, Michael.** 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.