

The Case of the Reappearing Phillips Curve: A Discussion of Recent Findings

Asha Bharadwaj and [Maximiliano Dvorkin](#)

The Phillips curve seems to have flattened over time. In this article, we use a simple New Keynesian model to analyze potential pitfalls in the estimation of the slope of the structural Phillips curve. Changes in the conduct of monetary policy or in the relative importance of supply and demand shocks may bias simple estimations of the slope of the Phillips curve. Recent proposals have favored estimations using regional or city data in an effort to overcome these issues. We use a simple model of a monetary union with a continuum of economies and find that some of the drawbacks of the aggregate model are still present in a cross-section of many regions in a monetary union. The relative importance of the demand and supply shocks largely determines the empirical relation between unemployment and inflation in both the aggregate and the cross-section of regions. Our analysis shows potential pitfalls in estimating the slope of the Phillips curve, even if using regional data. (JEL E12, E31, E58, R13)

Federal Reserve Bank of St. Louis *Review*, Third Quarter 2020, 102(3), pp. 313-37.
<https://doi.org/10.20955/r.102.313-37>

1 INTRODUCTION

Central banks around the world intervene in financial markets by setting the short-term nominal interest rates to stimulate economic activity and control inflation. There are many theories that suggest that changes in interest rates due to monetary policy have effects on real activity: An increase in interest rates is associated with a tightening of the economy and a decrease in the real output, while a decrease in interest rates is associated with an increase in real output. Interest rates can affect output through several channels. For instance, an increase in interest rates implies an increase in the return on savings. In this case, individuals have an incentive to save more and forego current consumption, which translates to a decrease in current real output, all else equal. Clearly, the opposite is true in the case of a decrease in interest rates. Other possible mechanisms through which movements in interest rates affect real activity include borrowing costs and investment. As interest rates decline it becomes cheaper to borrow and invest, so business investment goes up, increasing total output.

Asha Bharadwaj is a research associate and Maximiliano Dvorkin is a senior economist at the Federal Reserve Bank of St. Louis.

© 2020, Federal Reserve Bank of St. Louis. The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Articles may be reprinted, reproduced, published, distributed, displayed, and transmitted in their entirety if copyright notice, author name(s), and full citation are included. Abstracts, synopses, and other derivative works may be made only with prior written permission of the Federal Reserve Bank of St. Louis.

While these theories shed light on the effects of monetary policy on the real economy, there is no consensus on the link between interest rates and inflation. A widespread view among policymakers is that there is a trade-off between real activity and inflation and that monetary policy decisions on interest rates, by affecting the real economy, ultimately affect inflation. This trade-off is known as the Phillips curve.

The Phillips curve was popularized by A.W. Phillips in 1958, when he showed a statistically significant negative relation between the unemployment rate and the growth rate of nominal wages—that is, wage inflation. Based on this empirical relationship, Samuelson and Solow (1960) argued that a looser monetary policy could reduce the unemployment rate by allowing inflation to rise. This then implied that monetary authorities could exploit this trade-off. Since it was first discovered empirically, the Phillips curve has guided discussions of monetary policy and has shaped our understanding of the transmission of monetary policy to prices. More recently, several theories on price setting by firms can rationalize the existence of a Phillips curve in an economic model.

Over the years, the Phillips curve has received several criticisms. Recent articles have argued that inflation can be approximated by statistical processes unrelated to the amount of slack in the economy (Atkeson and Ohanian, 2001; Cecchetti et al., 2017; and Stock and Watson, 2007). Moreover, the lack of a stable relationship between inflation and various measures of slack has led several articles to conclude that the Phillips curve has weakened over the years (Blanchard, Cerutti, and Summers, 2015; and Coibion and Gorodnichenko, 2015).

Several articles have pushed back on this criticism and have attempted to “recover” the Phillips curve. Fitzgerald and Nicolini (2014) argue that aggregate data are uninformative about the true structural relationship between unemployment and inflation, and that in fact, under a specific definition of inflation targeting, the evolution of equilibrium inflation is a random walk. They then show that regional data can be used to identify the structural relationship between unemployment and inflation. The main intuition is that monetary policy typically reacts to the aggregate state of the economy, but not to regional conditions. Thus, it is possible to use the deviations of regional economic activity relative to the aggregate and the deviation of inflation relative to the aggregate to recover the relationship between unemployment and prices.

A recent article by McLeay and Tenreyro (2019) also supports this view. The authors argue that it is difficult to identify the slope of the Phillips curve empirically, even if a negative relationship does hold true in the underlying model. This is because monetary policy will react to economic shocks in order to stimulate output when it is below potential and reduce inflation when it is above target. The actions of the monetary authority will typically affect the empirical slope of the Phillips curve. McLeay and Tenreyro (2019) then use a simple New Keynesian model to highlight this estimation bias due to the endogeneity of monetary policy. They propose several solutions including using regional Phillips curves to circumvent this identification problem.

In our article, we follow the approach suggested by McLeay and Tenreyro (2019) closely and use a simple New Keynesian model to highlight the issues that arise with the identification of the empirical slope of the aggregate Phillips curve. We then use a simple New Keynesian

model of a monetary union, as given by Gali and Monacelli (2008), to attempt to recover the Phillips curve at the regional level. However, we find that this approach is not sufficient to overcome the identification issues highlighted by McLeay and Tenreryo (2019), since even at the regional level our model fails to recover the slope of the Phillips curve. We argue that several factors, including the relative importance of the demand and cost-push shocks, affect the estimation of the slope of the Phillips curve, at both the aggregate and regional levels.

The structure of this article is as follows. Section 2 briefly reviews existing literature on the Phillips curve. Section 3 presents empirical evidence on the aggregate relationship between unemployment and inflation. Section 4 introduces a simple New Keynesian model of optimal policy with the Phillips curve and describes the empirical relationships we obtain when the model is used as a data-generating process. Section 5 discusses empirical challenges with regional Phillips curves. Section 6 presents a simple model of a monetary union with a continuum of regions and discusses the results we obtain when we use the model to simulate data. Section 7 concludes.

2 LITERATURE REVIEW AND BACKGROUND ON THE PHILLIPS CURVE

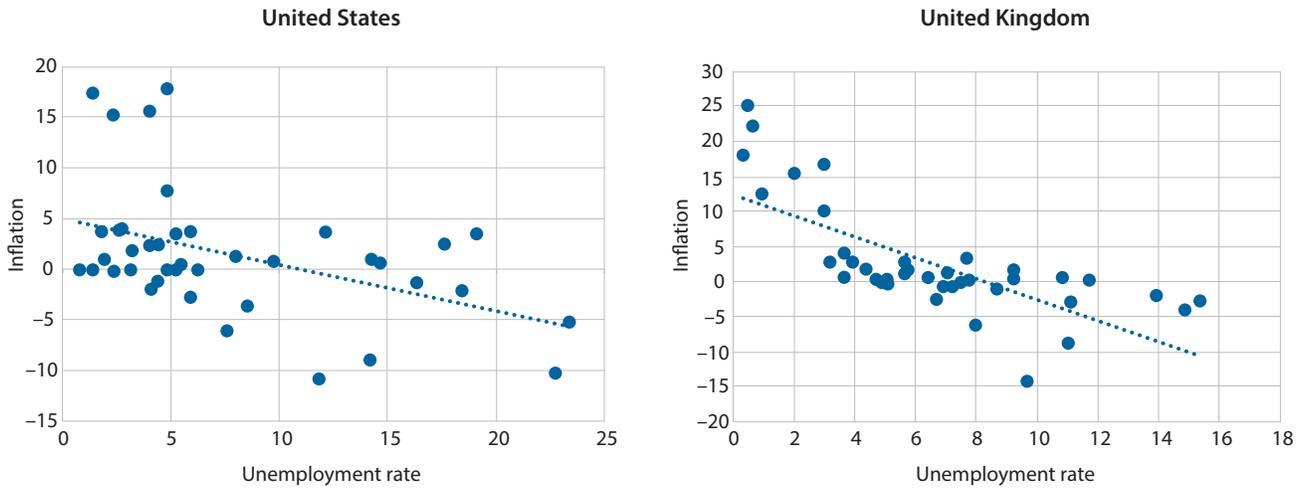
Phillips (1958) showed, for the United Kingdom and for the years 1861-1913, a statistically significant negative relationship between the unemployment rate and the growth rate of nominal wages held in the data. This result led to an outpouring of work on this topic, including the work by Samuelson and Solow (1960), who argued that policy could exploit the trade-off between inflation and the unemployment rate: Looser monetary policy would lead to an increase in inflation and a decrease in the unemployment rate. Figure 1 portrays the relationship between inflation and unemployment for the United States and United Kingdom, between 1900 and 1940, and we observe a clear negative trend between the level of unemployment and the growth rate of prices.

However, in the 1960s and 1970s, the relationship between unemployment and inflation started to change and was no longer a robust negative relationship. Figure 2 represents the relationship between unemployment and inflation in the 1960s for three countries—the United Kingdom, the United States, and France—and we observe that the negative relationship we would expect is no longer consistent across countries. It is negative for the United States, while it is positive for France and the United Kingdom. Lucas (1972) argued that the statistical relationship between unemployment and inflation depends on the parameters governing monetary policy and, therefore, is not “structural.” Lucas provided a theoretical model consistent with the existence of the observed statistical relationship between inflation and unemployment, but in which systematic attempts by policy to exploit the trade-off are not successful and change the statistical relation of these two variables.

Gali and Gertler (1999) developed a structural model of the Phillips curve with micro-foundations, based on initial work by Calvo (1983). They argued that real marginal costs and inflation expectations are significant determinants of inflation, as opposed to the backward-looking behavior that had been considered quantitatively important in the existing literature. Several articles then used this microfounded Phillips curve, also called the New Keynesian

Figure 1

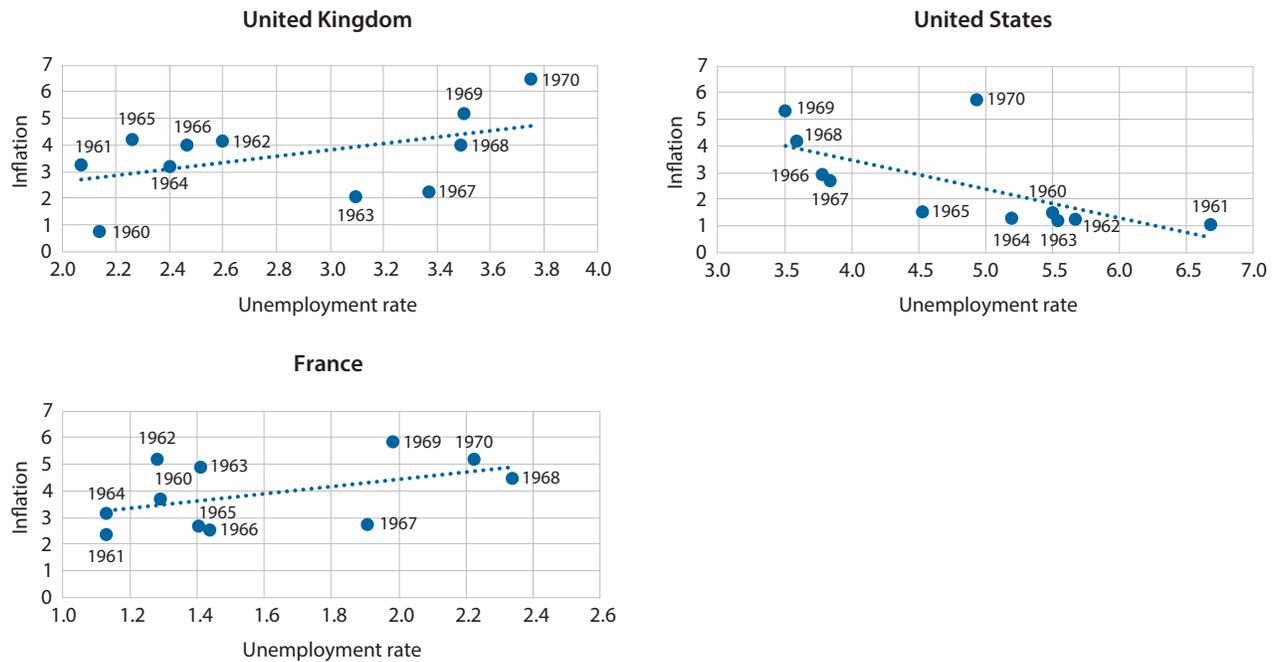
The Phillips Curve for the United States and United Kingdom, 1900-40



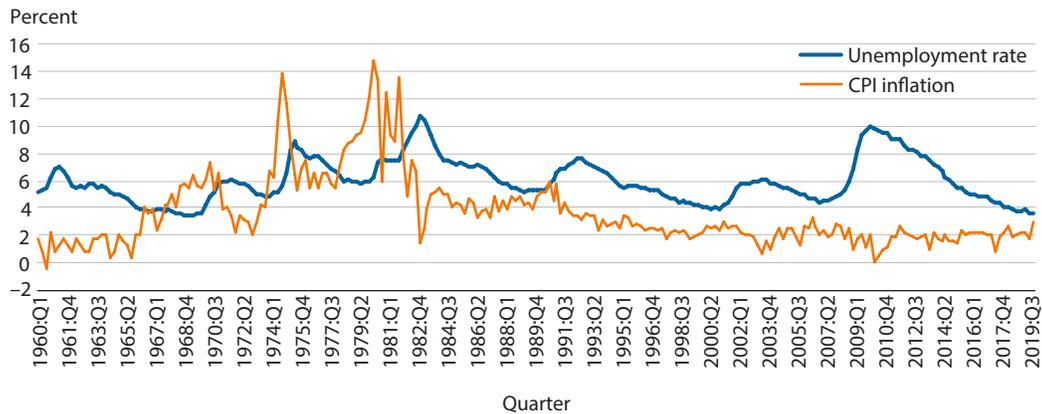
SOURCE: FRED®, Federal Reserve Bank of St. Louis; Annual Estimates of Unemployment in the United States, National Bureau of Economic Research; and Monthly Labor Review, Census Bureau.

Figure 2

The Phillips Curve in the 1960s for Select Countries



SOURCE: FRED®, Federal Reserve Bank of St. Louis.

Figure 3**Evolution of Inflation and Unemployment Rate in the United States, 1955-2019**

SOURCE: FRED®, Federal Reserve Bank of St. Louis and authors' calculations.

Phillips curve, to argue that successful monetary policy is responsible for flattening the slope of the Phillips curve by anchoring inflation expectations (Williams, 2006; Bernanke, 2007; and Mishkin, 2007). Several articles also support the argument that even in a purely static setting without expectations, the structural relationship between unemployment and inflation can be masked by the conduct of monetary policy (Bullard, 2018; Roberts, 2006; and Krogh, 2015).

3 EMPIRICAL PHILLIPS CURVE FOR THE UNITED STATES

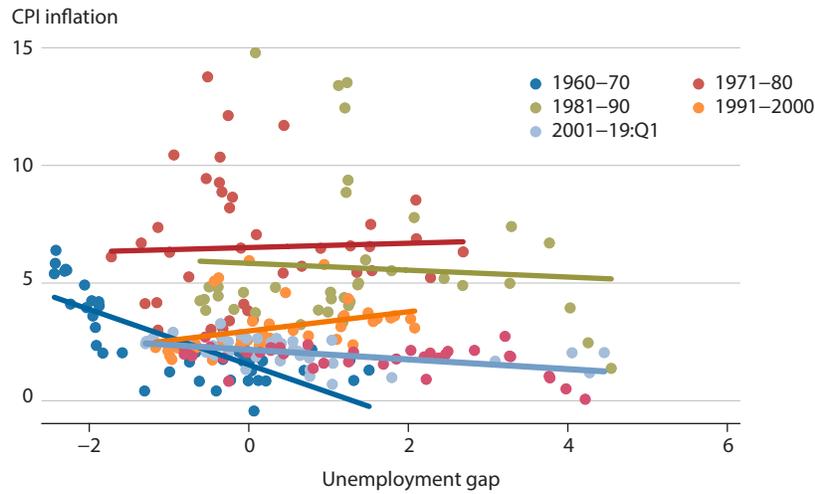
We begin by looking at the evolution of U.S. inflation and the unemployment rate over the past 60 years. As we can see from Figure 3, the times when inflation falls below trend coincide with the times with higher unemployment. This is suggestive of a negative correlation between the series. In the 1960s, for instance, there seems to be a clear negative correlation, as well as in the 1980s. We next use some scatterplots to understand the correlation between these variables better.

Figure 4 shows the behavior of the Phillips curve in the United States. Each dot in the graph represents a quarter. From the graph, it is clear that the slope of the Phillips curve has not been stable over the past 60 years. In the 1960s, there was a sizeable negative correlation between inflation and the unemployment gap, which has since flattened. We observe that over the past two decades, the negative relationship has returned somewhat, but the slope is not as steep as it was in the 1960s.

We formally estimate the empirical Phillips curve using aggregate data. The equation that we estimate is

$$\pi_t = \alpha + \beta_1(U_t - U_t^*) + \beta_2 E_t \pi_{t+1} + \sum_{k=1}^3 \beta_{2+k} \pi_{t-k} + \varepsilon_t,$$

Figure 4
The U.S. Phillips Curve



NOTE: The unemployment gap is defined as the difference between the unemployment rate and the CBO estimate of natural rate of unemployment.
 SOURCE: BLS, CBO, and authors' calculations.

Table 1
Phillips Curve Using Aggregate U.S. Data (1990:Q1-2018:Q1)

Inflation	Bivariate	With lags	New Keynesian	Hybrid
Unemployment gap	-0.222*** (-3.96)	-0.0171 (-0.41)	-0.159*** (-3.95)	-0.112* (-2.46)
Inflation expectations			1.006*** (12.37)	0.634*** (3.91)
<i>Inflation lags</i>				
First lag		0.424*** (4.32)		0.186 (1.65)
Second lag		0.461*** (4.28)		0.137 (1.36)
Third lag		0.0884 (0.86)		-0.000435 (-0.00)
Constant	2.587*** (26.32)		-0.339 (-1.51)	-0.111 (-0.47)
Observations	118	118	107	107
Adjusted R ²	0.085	0.957	0.517	0.534

NOTE: The t-statistics are in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.
 SOURCE: BLS, CBO, the Survey of Professional Forecasters, and authors' calculations.

where π_t is the inflation rate and $(U_t - U_t^*)$ is the unemployment gap. We compute inflation as the annualized change in quarterly core consumer price index (CPI) inflation (obtained from the Bureau of Labor Statistics [BLS]); the unemployment gap is defined as the difference between the unemployment rate (from the BLS) and the natural rate of unemployment (from the Congressional Budget Office [CBO]). Inflation expectations are obtained from the Survey of Professional Forecasters.¹ We estimated four specifications of this Phillips curve: a bivariate model, in which we simply regress inflation on the unemployment gap ($\beta_2 = \beta_3 = \beta_4 = \beta_5 = 0$); a model with lags, in which we regress inflation on the unemployment gap and three lags of inflation ($\beta_2 = 0$); a model with inflation expectations, in which we regress inflation on the unemployment gap and expectations ($\beta_3 = \beta_4 = \beta_5 = 0$); and a hybrid model where both past inflation and inflation expectations can influence the inflation rate.

Across all specifications, the regression coefficient for the unemployment gap is negative, which agrees with our previous discussion on the Phillips curve. The estimates in Table 1 show that the steepest slope is with the bivariate specification (-0.2), while the one with lags yields the flattest slope (-0.02). Also worth noting is the fact that inflation expectations are significant and positive in the third regression, indicating that if inflation is expected to increase in the future, current inflation responds by increasing as well.

4 THE BASIC NEW KEYNESIAN FRAMEWORK

To better understand the problems and limitations of a simple empirical estimation of the slope of the Phillips curve, we now study the issue using a macroeconomic model: the New Keynesian model. We first describe the canonical New Keynesian model and then discuss its equilibrium conditions. We follow closely Clarida, Gali, and Gertler (1999), as well as Chapter 3 of Gali (2008).

Within the New Keynesian macro-model framework, monetary policy affects the real economy in the short run. The basic model consists of identical households that make decisions about labor supply, consumption, and savings. Households are risk averse and dislike fluctuations in consumption. Households can borrow or save in one-period bonds that pay an interest rate. These bonds are in zero net supply. Firms produce a range of differentiated goods using only labor with a constant returns-to-scale production function. Firms set prices for their goods, taking into account that they face competition from close substitutes of their goods and that they are unable to change prices every period and able to do so only after some random period. In the basic model, we abstract from investment and capital and from government consumption and interactions with the rest of the world. The total production cannot be stored, and total consumption has to equal total production. We also use a version of the model that abstracts from money holdings, sometimes referred to as the cashless limit economy. Since all households are identical, in equilibrium the interest rate must be such that households are indifferent between saving and consuming their income.²

The dynamic conditions that characterize equilibrium in this model can be approximated by two equations: a dynamic IS curve that relates the output—in deviations from trend—to the real interest rate, and a Phillips curve that relates inflation positively to the output. It is

important to note this dynamic aggregate behavior evolves from optimization by firms and households. Moreover, these equations approximate the dynamics around a stationary equilibrium.

For any variable z , let \hat{z} denote the deviation of the variable from its steady state or long-run value. The New Keynesian Phillips curve is given by

$$(1) \quad \hat{\pi}_t = \beta E_t [\hat{\pi}_{t+1}] + \kappa \hat{y}_t + \varepsilon_t.$$

Here, π_t refers to the aggregate inflation at time t , \hat{y}_t refers to the output in deviations from trend or potential at time t , and ε_t is a cost-push shock that follows an AR(1) process ($\varepsilon_t = \rho \varepsilon_{t-1} + e_t$).

This equation reflects the aggregate relationship between the inflation rate—in deviations from a long-run value—and output, and it tells us that an increase in output is associated positively with a positive change in inflation. This equation evolves from the individual firm's problem: Firms are monopolistically competitive, and each firm chooses its price to maximize profits subject to constraints on the frequency of future price adjustments. The aggregate Phillips curve is simply capturing the aggregation of individual pricing decisions by a log-linear approximation. Another way of interpreting this equation is by thinking in terms of marginal costs. Excess demand (output above the potential) is associated with marginal costs above average, and the reverse is true in the case of excess supply. Thus, firms set nominal prices based on their expectations of future marginal costs, as well as current marginal costs. The cost-push shock is meant to capture other forces that may affect inflation not already captured in the model.

The next ingredient of the model is the dynamic IS curve, which inversely relates the aggregate output (or consumption since they are the same in this model) to the real interest rate, and tells us that aggregate output evolves as a function of the expected future output, the interest rate, the expected inflation level, and demand shocks:

$$\hat{y}_t = E_t [\hat{y}_{t+1}] - (i_t - E_t [\hat{\pi}_{t+1}] - z_t).$$

Here, z_t is a demand shock that follows an exogenous AR(1) process with persistence ρ_z (i.e., $z_t = \rho_z z_{t-1} + \gamma_t$). This equation is a result of log linearizing the consumption Euler equation that arises from optimizing the household's utility. The term $i_t - E_t [\hat{\pi}_{t+1}]$ is the (expected) real interest rate since i_t is the nominal rate. The mechanism by which the real interest rate affects output is somewhat intuitive: Higher interest rates imply higher borrowing costs or a higher return to savings, which lead to lower consumption today and thus lower current output. This is also known as the intertemporal substitution of consumption. The reason that expected future output affects current output is that individuals prefer a smooth path of consumption. If individuals expect higher income and consumption in the next period, they prefer to consume more today, thus resulting in an increase in output today. Overall, output demand depends negatively on the real interest rate and positively on expected future output. One possible interpretation for the demand shock z_t is a change in households' preferences for consuming today versus tomorrow, like a sudden increase in impatience.

This simple model does not have a direct link to unemployment. We add an exogenous relationship between the output and the unemployment rate, or, more precisely, the unemployment rate gap, which is best known as Okun’s law:

$$\hat{u}_t = \alpha \hat{y}_t.$$

Finally, as is usual in most New Keynesian models, we add a monetary authority with the objective to control inflation and promote high levels of employment. The monetary authority sets the nominal interest rate to achieve these objectives. In particular, we assume that the short-term nominal interest rate is set according to a Taylor rule of the form

$$i_t = r^* + \pi^* + \phi_\pi \hat{\pi}_t + \phi_y \hat{y}_t,$$

where r^* is the long-run, or natural, real interest rate, and π^* is a long-run value of inflation, both of which we assume are exogenous. We assume that the coefficients ϕ_π and ϕ_y are positive. This policy function implies that the monetary authority will increase the nominal interest rate if inflation is above its long-run value—that is, when $\hat{\pi}_t$ is positive—or when output is above its long-run value—that is, when \hat{y}_t is positive. A solution to this model includes decision rules for inflation, output, interest rate, and unemployment, as a function of the two exogenous shocks, demand shocks, and cost-push shocks. Assuming that the Taylor principle holds and that the parameters of the model are within normal ranges, we can solve for the equilibrium inflation and output by plugging this optimal policy rule into the Phillips curve. We then find that both the equilibrium inflation and output gap are functions of the demand shock and cost-push shock:

$$\hat{y}_t = c_{ye} \mathcal{E}_t + c_{yz} Z_t,$$

$$\hat{i}_t = c_{ie} \mathcal{E}_t + c_{iz} Z_t,$$

$$\hat{\pi}_t = c_{\pi e} \mathcal{E}_t + c_{\pi z} Z_t, \text{ and}$$

$$\hat{u}_t = c_{ue} \mathcal{E}_t + c_{uz} Z_t,$$

where coefficients in the equations are constants that depend on model parameters. In the analysis that follows, we focus on the last two equations that say that the equilibrium values of inflation and unemployment at time t depend on the realization of shocks. In particular, it can be shown that both of the coefficients in the inflation decision rule are positive; thus, inflation increases with positive demand shocks or positive cost-push shocks. On the other hand, unemployment decreases with positive demand shocks but increases with cost-push shocks, as these last shocks will induce a contraction in economic activity.

We now use this model as a data-generating process to simulate shocks to the economy. For this, we need to specify values for the parameters of the model. We use the same calibration as Gali (2008).³ We can use the New Keynesian Phillips curve in equation (1), together

with the equation for Okun's law, to derive a relationship between inflation and unemployment, which we can estimate. In particular, we have

$$(2) \quad \hat{\pi}_t = \beta E_t[\hat{\pi}_{t+1}] + \delta \hat{u}_t + \varepsilon_t,$$

where the coefficient δ is a combination of parameters κ and α . A direct estimation of equation (2) is problematic. As the decision rules clearly show, a simple regression of inflation on the unemployment gap would give us biased estimates because the unemployment gap is correlated with the cost-push shock ε_t . In a similar way, the expected value of inflation in the next period would also be correlated with today's cost-push shock if the shocks were persistent.

McLeay and Tenreyro (2019) discuss several assumptions and estimation methods under which estimates would be correct. First, in an economy without cost-push shocks and in which expected inflation and unemployment are observed, a regression would recover the slope δ . Alternatively, if there are good instruments for both the unemployment rate and the expected inflation that are uncorrelated to the cost-push shock, then an instrumental variables regression would be valid. Finally, they propose a different way to tackle the problem, which we discuss later.

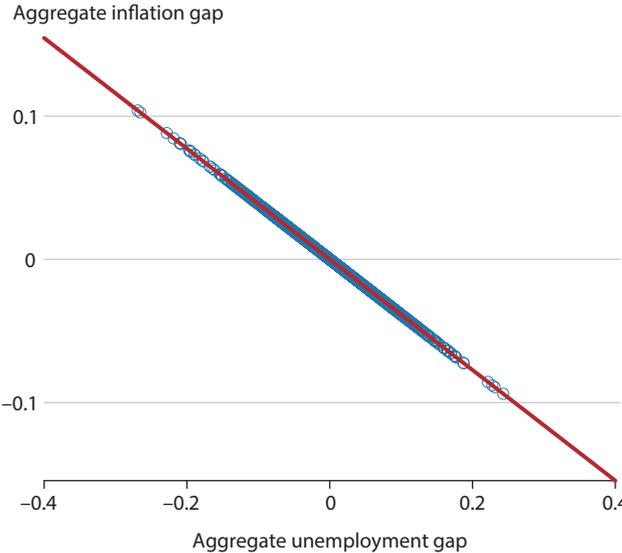
We now use the model to understand possible shortcomings in the empirical estimation of Phillips curves. The exercises that we describe next use the decision rules for inflation and unemployment together with simulated values for the demand and cost-push shocks. These give a simulated time series for the two variables in the Phillips curve. Figure 5 shows the results of simulating various combinations of shocks. To simplify the analysis and make our point, we assume that the shocks are not persistent. In other words, we assume that $\rho = \rho_z = 0$, which would then imply that $E_t[\pi_{t+1}]$ is not affected by the current realization of the shocks.⁴

4.1 Case I: Only Demand Shocks

We now start our analysis by focusing on the effects of different forces on the slope of the Phillips curve, one by one. In this case, we allow for only demand shocks. As Figure 5 shows, all the different equilibrium values of inflation and unemployment reflect a typical negative relationship between inflation and unemployment.

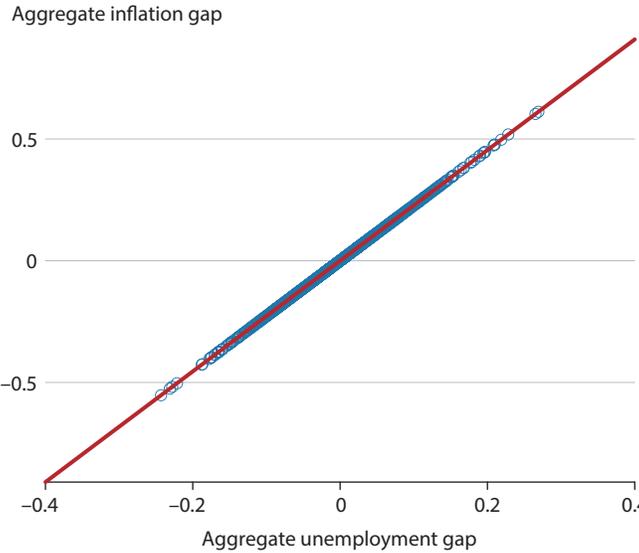
This is what we expect to see given the decision rules obtained before and that only demand shocks are allowed. When a demand shock hits, output increases, lowering unemployment but also increasing inflation. The monetary authority counteracts positive demand shocks by increasing the interest rate, which pushes both output and inflation to an otherwise lower level. In this case, the monetary authority does not face a trade-off between output and inflation. Note, however, that the slope of the empirical Phillips curve in Figure 5, in general, will not be equal to the slope of the structural Phillips curve—that is, the coefficient δ in equation (2). The reason is that the actions of the monetary authority affect the movements in inflation and the unemployment rate, which in turn affect the empirical slope, as we highlight in Case III.

Figure 5
Simulated Inflation and Unemployment with Only Demand Shocks

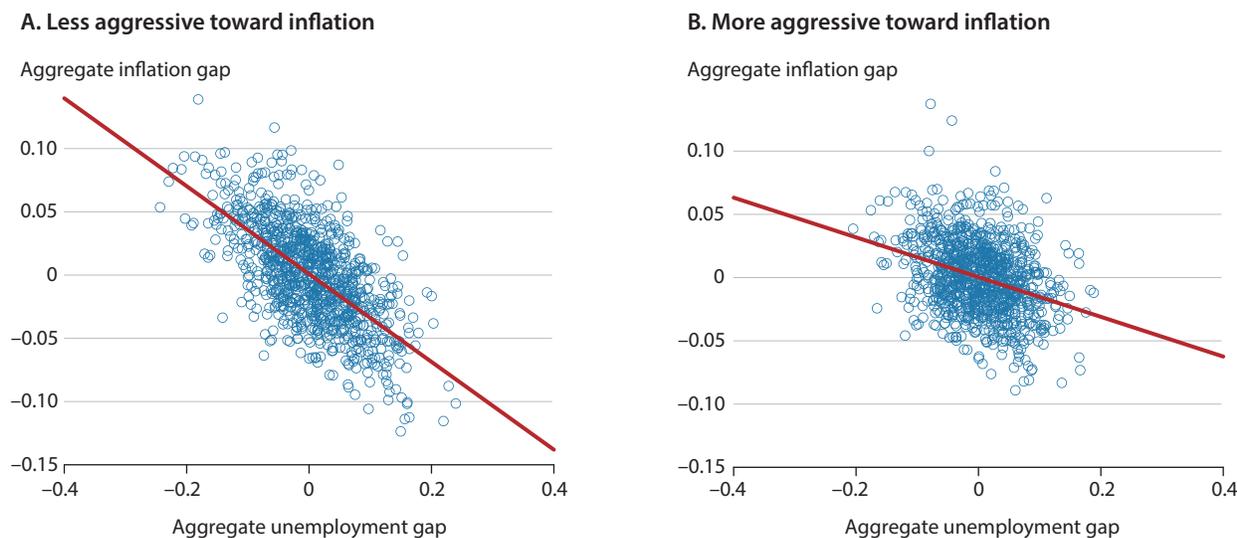


SOURCE: Authors' calculations.

Figure 6
Simulated Inflation and Unemployment with Only Cost-Push Shocks



SOURCE: Authors' calculations.

Figure 7**Simulated Inflation and Unemployment with Different Policy Parameters**

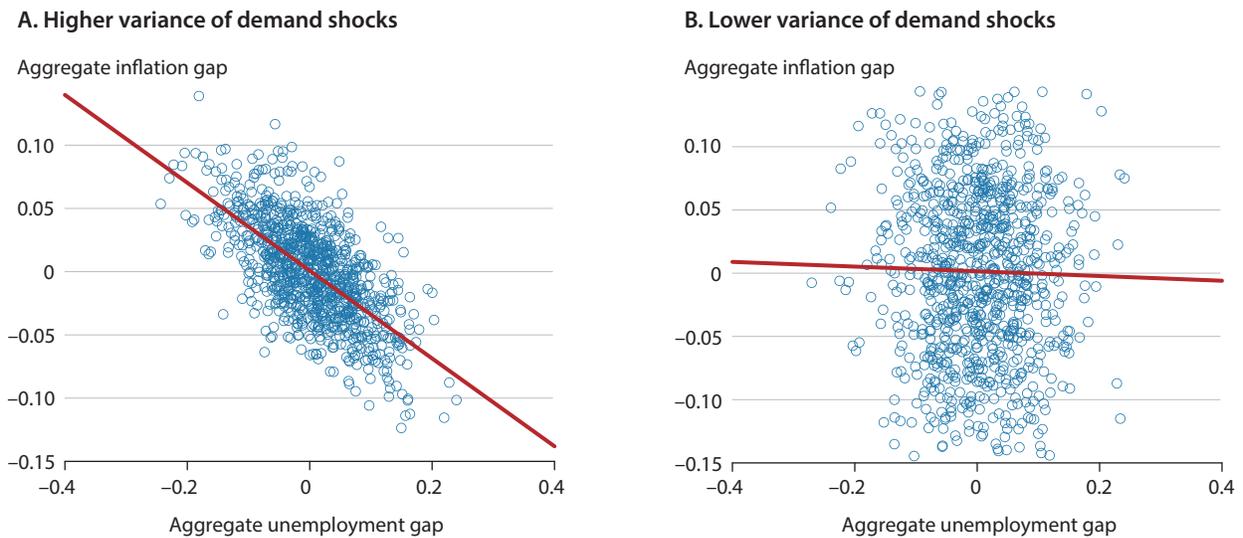
SOURCE: Authors' calculations.

4.2 Case II: Only Cost-Push Shocks

We next set all demand shocks equal to zero in our simulations and allow for only cost-push shocks. Figure 6 shows the results of our simulations for this case. In the decision rules, the coefficient that multiplies the cost-push shock is positive for both variables—unemployment and inflation. Thus, with a positive cost-push shock, both inflation and unemployment increase and we observe a positive relationship between inflation and the unemployment gap in the graph. In the background, it is the optimal decisions of agents and the monetary authority that drive this behavior. In particular, the monetary authority is aggressive at fighting inflation resulting from the cost-push shock, increasing the interest rate, thus lowering total consumption and production and increasing unemployment.

4.3 Case III: Changes in the Policy Rule Over Time

In a recent discussion, Bullard (2018) studies how the empirical Phillips curve would look in a simple New Keynesian model if the monetary authority reacts more aggressively toward increases in inflation. Figure 7 shows this comparison. Panel A presents simulations using a value of $\phi_\pi = 1.5$, and Panel B uses a value of $\phi_\pi = 5.5$, which implies a more aggressive stance toward increases in inflation. As the figure shows, with a more aggressive policy, the empirical slope of the Phillips curve is flatter. The main takeaway of this exercise is that if the monetary authority becomes more aggressive toward inflation over time, then, all else equal, the empirical Phillips curve will be flatter. However, this flattening of the empirical Phillips curve occurs even if the structural Phillips curve of equation (1) did not change, as we are keeping δ constant.

Figure 8**Simulated Inflation and Unemployment with Different Variance of Cost-Push Shocks**

SOURCE: Authors' calculations.

This exercise illustrates an important pitfall in empirical analysis using reduced-form estimations of the Phillips curve: The actions of the monetary authority affect estimates of the slope of the Phillips curve. Even if the slope of the structural Phillips curve remains constant (parameter δ), changes in monetary policy over time affect the empirical slope of the curve, and a more aggressive stance will make the empirical curve flatter.

4.4 Case IV: Changes in the Variance of Shocks

As Cases I and II show, the slope of the empirical Phillips curve will be affected by the type of shock that impacts the economy. We next simulate both shocks and compute the corresponding time series behavior for inflation and unemployment according to the decision rules. Figure 8 presents the results for two different sets of parameters. In particular, Panel A shows the simulations for an economy with a relatively higher variance of the cost-push shock, while Panel B shows the same simulations but for a lower variance of these shocks. As is clear from the figure, an economy with larger cost-push shocks will display a flatter (or even an upward-sloping) empirical Phillips curve. Once again, it is worth stressing that the flattening of the empirical Phillips curve would occur even if the structural Phillips curve of equation (2) does not change and δ remains constant, which also shows a potential pitfall in empirical analysis using reduced-form estimations of the Phillips curve. To the extent that empirical estimates cannot control properly or instrument correctly for cost-push shocks, differences in the variances of the shocks will be reflected in the estimates of the slope.

Case III highlights that the empirical Phillips curve may flatten over time if the monetary authority adopts a more aggressive stance toward inflation. Case IV, on the other hand, highlights that the flattening could be due to a decrease in the importance (size) of demand shocks relative to cost-push shocks.

5 REGIONAL PHILLIPS CURVES

As we mentioned before, McLeay and Tenreyro (2019) discuss several avenues to properly recover the slope of the Phillips curve. The first method they suggest is to control for cost-push shocks, to the extent that these are observable. In essence, this would allow us to replicate Case II from the previous section by removing the effect of cost-push shocks from both inflation and the unemployment rate, such that any remaining variation in the unemployment gap and inflation gap must be due to movements in aggregate demand. There are articles, such as Roberts (1995), which use oil prices to control for cost-push shocks. This approach, while reasonable, is difficult to implement in practice since there is a large number of potential cost-push shocks that would need to be controlled for, which in many cases may not be observable.

Another method is to use an instrumental variable estimation for the unemployment gap. The instrument should be correlated with unemployment but uncorrelated with cost-push shocks. This is also hard to do in practice because it is difficult to find good instruments—that is, a macroeconomic variable that is truly exogenous to cost-push shocks but correlated with the unemployment rate and inflation expectations. Moreover, strong temporal dependence or persistence in the shocks may also affect the validity of the instruments.⁵

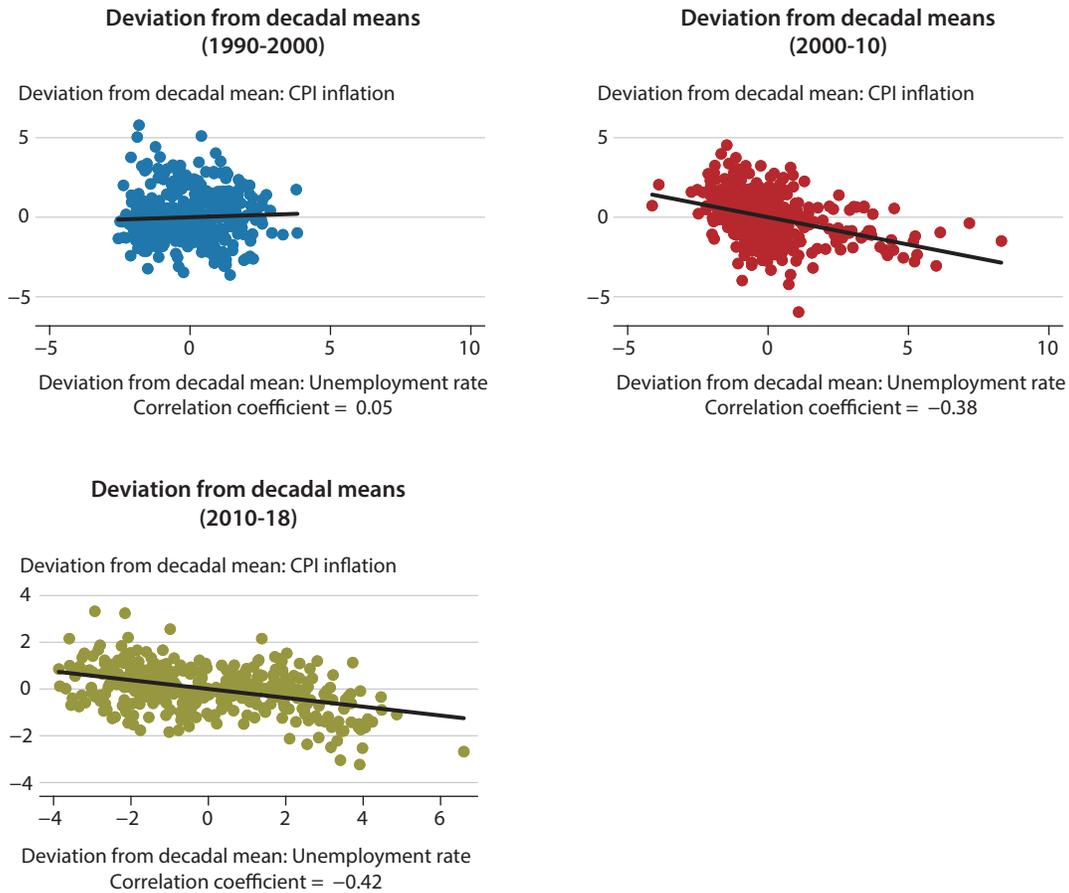
Finally, the third method they mention, and that we analyze in this article, is to exploit the cross-sectional heterogeneity in regional data. Fitzgerald and Nicolini (2014) also implement this approach in their article. As highlighted in Case III, the actions of the monetary authority affect the movements in unemployment and inflation and thus affect the estimates of the slope of the Phillips curve, even if properly instrumenting for cost-push shocks. However, in a large country like the United States, comprising several states and cities, the monetary authority only reacts to fluctuations in aggregate inflation and unemployment, and monetary policy is independent of local conditions. Thus, the main idea behind a regional approach to estimating the Phillips curve is to use fluctuations in regional economic conditions, in deviations from the aggregate, which the monetary authority will not influence by its actions. In this way, by properly instrumenting or controlling for cost-push shocks, we should be able to recover the structural relationship between inflation and the unemployment rate as we are also controlling for monetary policy. We first discuss empirical evidence and then use a model to analyze this approach.

5.1 Empirical Analysis

To explore the empirical implications of this regional approach to the Phillips curve estimation, we use semi-annual CPI inflation and unemployment data from the BLS for 20 metropolitan statistical areas (MSAs). We divide our sample into three periods: 1990-2000, 2000-10,

Figure 9

Empirical Relation between the Inflation Gap and Unemployment Gap for Different MSAs



SOURCE: BLS and authors' calculations.

and 2010-19. For each decade, we compute the mean inflation and unemployment for each MSA and use this as a proxy for the steady-state level of unemployment and inflation for each MSA. In Figure 9 we use simple scatterplots of the deviation of inflation from the decadal mean versus the deviation of unemployment from the decadal mean for each MSA. We interpret these as the unemployment and inflation gaps for each MSA.

As can be seen from Figure 9, the slope of the scatterplots varies over time. The curve was positively sloped in the 1990s, which contradicts our understanding of the theoretical relationship between these variables.

Formally, we estimate

$$\pi_{it} = \alpha_i + \beta_1 E_t \pi_{it+1} + \beta_2 U_{it} + \gamma_t + \varepsilon_{it},$$

Table 2**Estimates of the Empirical Phillips Curve Using MSA-Level Data, 1990-2018**

Inflation	Pooled	Metro area FE only	Year FE only	Metro area and year FE
Unemployment rate	-0.141*** (-7.20)	-0.168*** (-6.57)	-0.215*** (-5.81)	-0.406*** (-5.91)
<i>Inflation lag</i>				
First lag	0.390*** (8.00)	0.367*** (6.97)	0.127* (2.58)	0.0570 (1.18)
Within R ²	0.252	0.254	0.435	0.446
Observations	1,191	1,191	1,191	1,191
MSA FE	No	Yes	No	Yes
Year FE	No	No	Yes	Yes
Seasonal dummies	Yes	Yes	Yes	Yes

NOTE: FE, fixed effects. The *t*-statistics are in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

SOURCE: BLS.

where subindex i denotes an MSA and subindex t denotes the time period. As before, π is the inflation rate and U is the unemployment rate. The regression includes both time fixed effects and MSA fixed effects. Including MSA fixed effects allows us to remove the effect of the state-specific mean unemployment rate, which can be seen as a proxy for the equilibrium rate of unemployment for each MSA. The time fixed effects help control for time-varying changes in the aggregate equilibrium unemployment rate and help us overcome the bias caused by the correlation between the regional unemployment rate and aggregate unemployment rate.⁶ In Table 2, we present results from estimating the above regression, with a combination of fixed effects.

Column (2) shows the results of regressing inflation on unemployment without controlling for time fixed effects or seasonality. However, once we begin to control for these factors, we observe that the coefficient increases in magnitude; from Column (5), we observe that the coefficient is highly negative and significant.

5.2 A Simple New Keynesian Model of a Monetary Union

To understand the factors that influence the empirical estimates of the slope of the Phillips curve using regional data, we turn to a simple model of a monetary union with a common monetary authority, specified by Gali and Monacelli (2008). With this model, we proceed in a similar fashion as before, finding decision rules and then simulating shocks and time series for inflation and unemployment, in this case at both the aggregate and regional levels.

Gali and Monacelli (2008) model the currency union as a continuum of small (atomistic) open economies that are subject to imperfectly correlated shocks. Each of these small economies, which we call regions, share identical preferences, technology, and market structures.

Let P_t^j represent region j 's price index for goods produced locally, and let P_t^f be the price index for goods purchased (imported) from other regions. In this way, the CPI for a region is the geometric weighted average of domestic and imported price indexes, with weights given by consumption shares $P_{c,t}^j = (P_t^j)^{1-\alpha}(P_t^*)^\alpha$. For the monetary union as a whole, the price index is the aggregator across all regions.

The bilateral terms of trade between regions j and k can be defined as $S_{f,t}^j = \frac{P_t^f}{P_t^j}$, which is the price of region k 's goods in terms of region j 's goods. Then, the CPI inflation rate of the region ($\pi_{c,t}^j$) can be approximated as follows:

$$\pi_{c,t}^j = \pi_t^j + \alpha \Delta s_t^j.$$

The model assumes complete markets and perfect risk-sharing across all regions, together with a zero net savings for all regions. This means that regional consumption and consumption for the whole country can be approximated as

$$c_t^j = c_t^* + (1-\alpha) s_t^j,$$

where c_t^* is the (log) aggregate consumption for the whole country.

A log-linear approximation of the model around a symmetric steady state for all regions leads to similar dynamic equilibrium conditions as before. For region j , the Phillips curve is

$$\pi_t^j = \beta E_t[\pi_{t+1}^j] + \kappa \hat{y}_t^j + \varepsilon_t^j.$$

In addition, in each region an IS equation holds:

$$c_t^j = E[c_{t+1}^j] - (i_t^* - \pi_{c,t}^j - z_t^j),$$

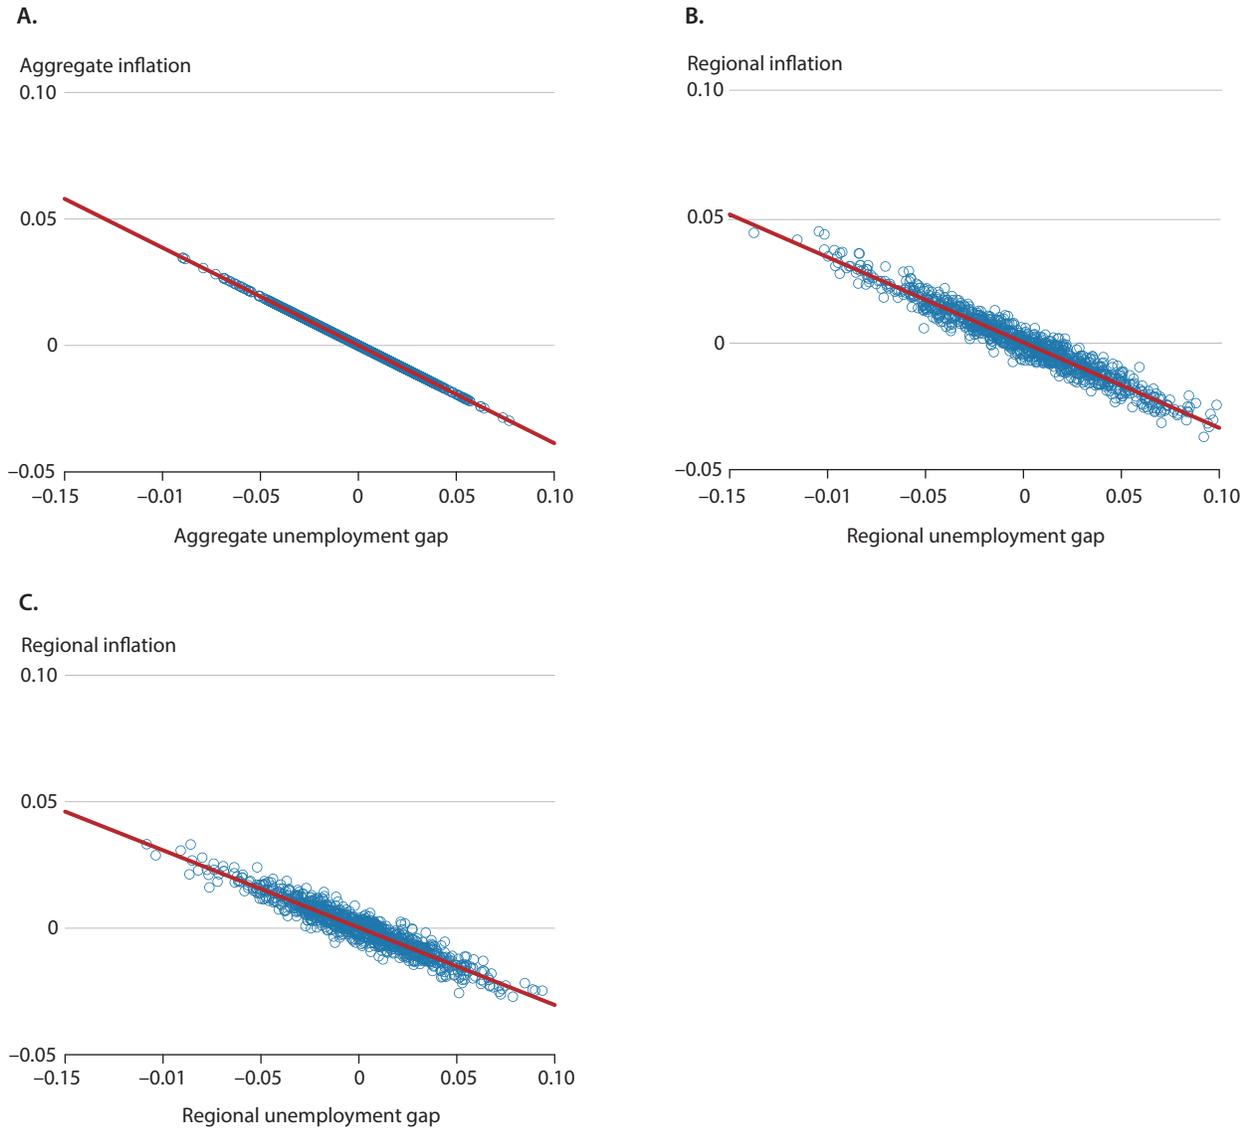
where i_t^* is the interest rate for the whole country. Aggregating these expressions across regions can form a Phillips curve and an IS curve for the whole country, which leads to the same expressions as in Section 4. We further assume that an expression like Okun's law holds in every region and that the monetary authority follows the same Taylor rule as in Section 4, reacting only to fluctuations in aggregate inflation and the aggregate output gap. As before, we use this model to simulate shocks to the economy.⁷

5.3 Case V: Only Demand Shocks—Regional Analysis

As in Case I, here we allow for only demand shocks. We assume a few possible scenarios. First, we allow for only demand shocks, both at the aggregate level and for individual regions. Second, we allow for only demand shocks at the regional level, but no aggregate shocks of any kind. In this last case, aggregate inflation and unemployment are constant, and the monetary authority will not change the interest rate. Figure 10 presents the results. Panel A shows the effect of aggregate demand shocks at the aggregate level—that is, for aggregate inflation and the aggregate unemployment rate. Panel B shows the effect of aggregate and regional demand shocks at the regional level. That is, each dot in the figure shows the inflation rate and the

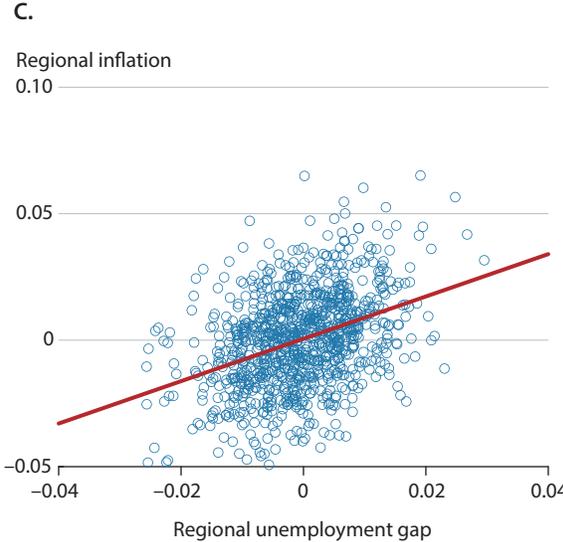
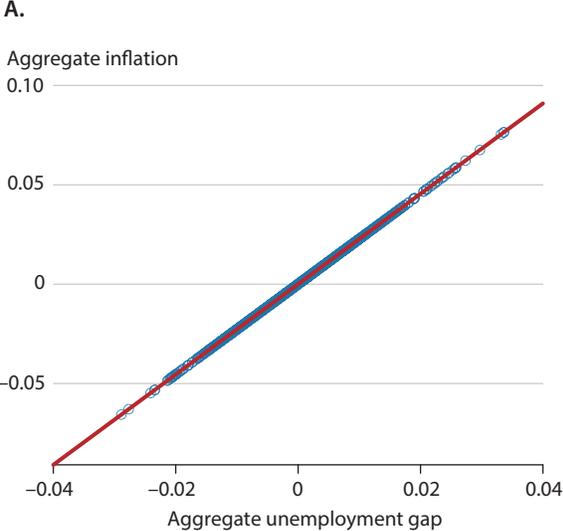
Figure 10

Effects of Regional and Aggregate Demand Shocks



SOURCE: Authors' calculations.

Figure 11
Effects of Regional and Aggregate Cost-Push Shocks



SOURCE: Authors' calculations.

unemployment rate in one individual region. Finally, Panel C shows the effect of only regional demand shocks at the regional level. In this case, the aggregate inflation and unemployment rate are constant at their steady-state value.

Clearly, for aggregate inflation and unemployment displayed in Panel A, allowing for only demand shocks shows a similar picture as Case I (shown in Figure 5). Adding regional shocks on top of this (Panel B) introduces some noise around the empirical relationship, but the main pattern still holds. More interesting, Panel C shows the relationship in the absence of aggregate shocks. In this case, we can see that regional demand shocks can help identify a negative empirical slope between inflation and unemployment at the regional level.

5.4 Case VI: Only Cost-Push Shocks—Regional Analysis

Similar to the previous case, we now analyze the effects of allowing for only cost-push shocks. Figure 11 shows the results. In Panel A, and similar to Figure 6, cost-push shocks generate a positive-sloping empirical relationship between inflation and unemployment for the aggregate economy. While the monetary authority intervenes to stabilize the economy, it faces the trade-off between lower inflation at the expense of a larger contraction and higher unemployment. At the regional level, the effects of both aggregate and regional cost-push shocks, shown in Panel B, generate a more dispersed pattern. Yet, the empirical slope between unemployment and inflation is still positive when using only regional data. Finally, Panel C shows the results for an economy with only regional cost-push shocks. In this case, aggregate variables do not move and there is no intervention by the monetary authority; the positive slope is more pronounced.

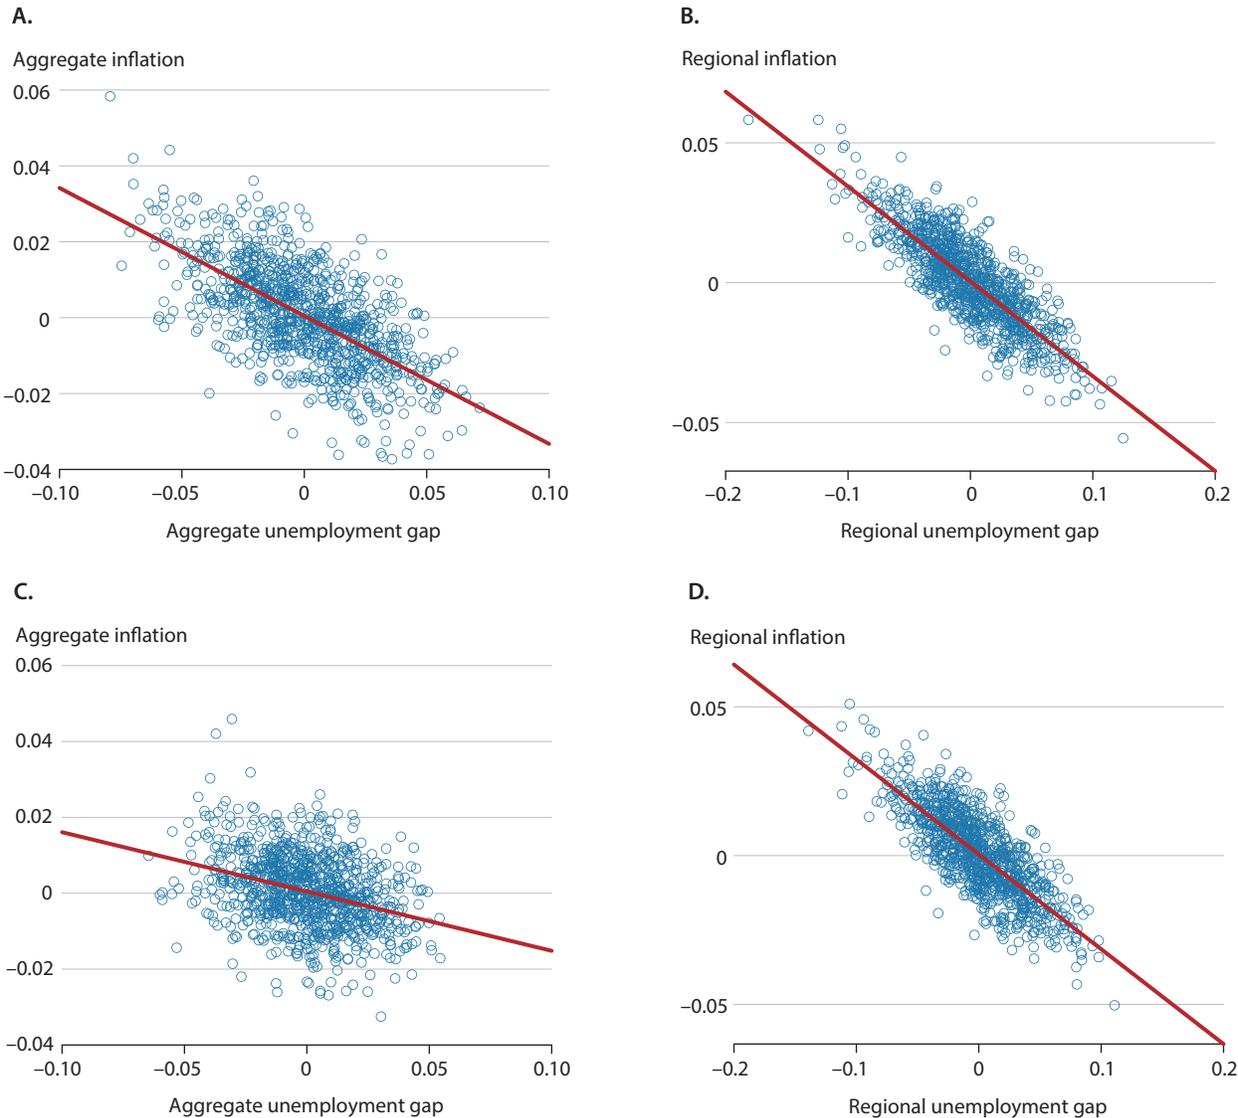
This exercise leads to an important conclusion. The empirical analysis using regional data must also control or instrument for regional cost-push shocks. Otherwise, estimates of the Phillips curve will be affected.

5.5 Case VII: Changes in the Policy Rule Over Time—Regional Analysis

Similar to our methods in the simple model, here we ask if changes in the monetary authority's preferences for how aggressively to fight inflation affect the empirical slope of the Phillips curve estimated using regional data. For this we use an economy where all shocks are active (aggregate and regional; demand and cost-push), and we compare whether and how the empirical slope of the Phillips curve estimated using regional data changes. We show this comparison in Figure 12, where we use simulated data from the model.

Panels A and C show the relation between unemployment and inflation at the aggregate level for two economies with different degrees of monetary policy aggressiveness, with Panel C showing the results for an economy with a more aggressive stance. These graphs are similar to those in Case III analyzed above, where the empirical slope of the Phillips curve was affected by changes in the monetary authority's preferences. Panels B and D show the relationship between unemployment and inflation using regional data, where Panel D shows the case of a monetary authority more aggressive toward inflation. In this case, the changes in the policy rule do not translate into noticeable changes in the slope of the empirical Phillips curve at the regional level. In other words, while the slope estimated with regional data does change with

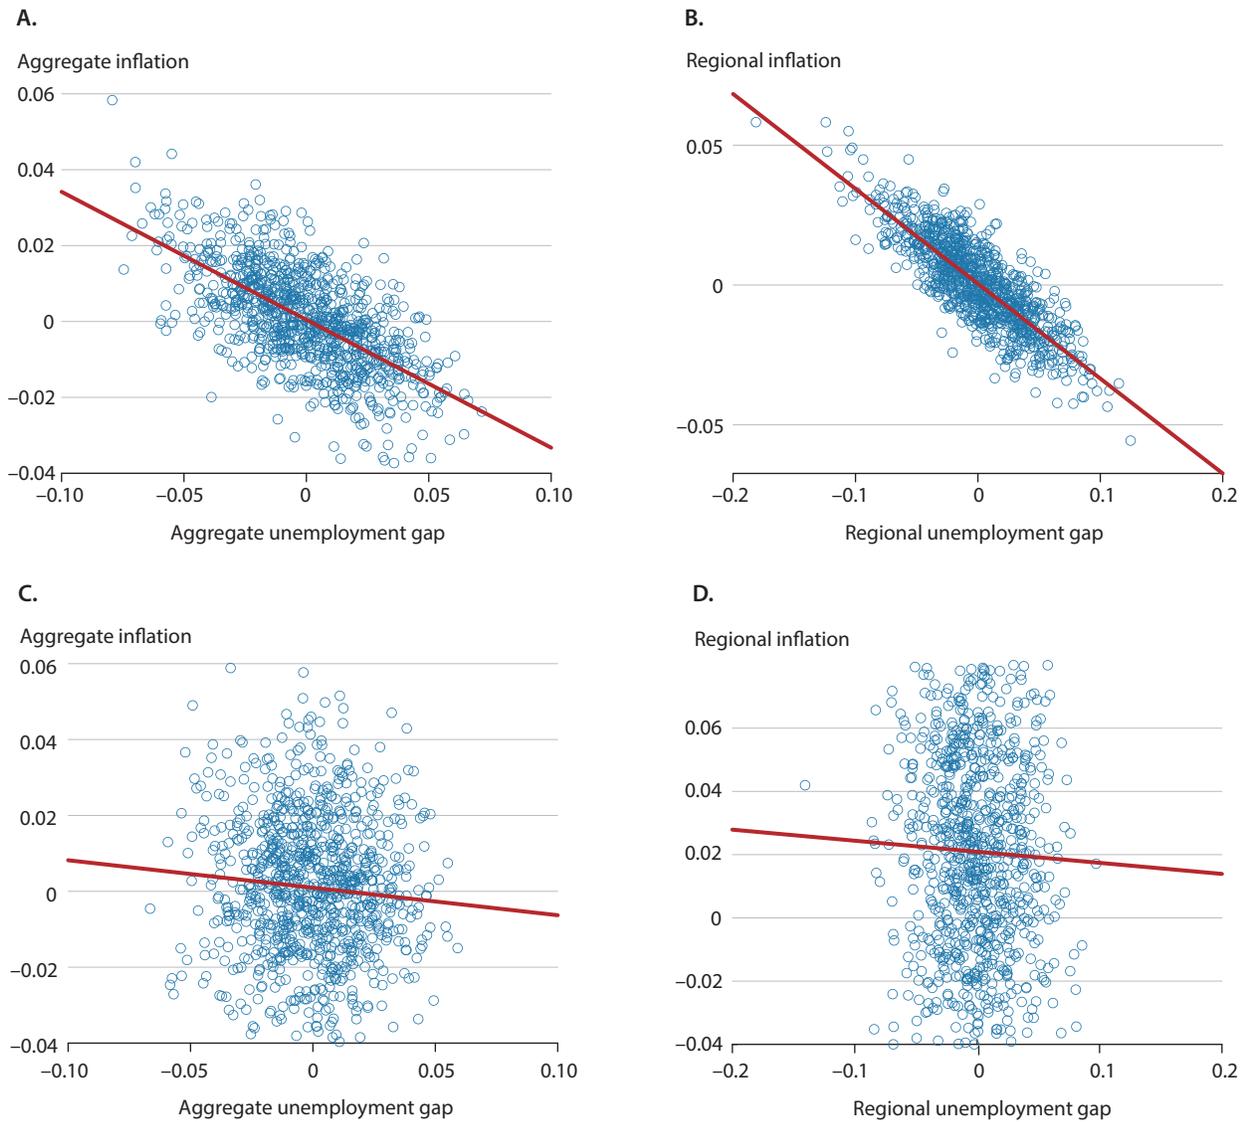
Figure 12
Effects of Changes in the Policy Rule at Aggregate and Regional Levels



SOURCE: Authors' calculations.

Figure 13

Effects of Changes in the Variance of Shocks at Aggregate and Regional Levels



SOURCE: Authors' calculations.

changes in monetary policy, the changes are very small. This result lends some support to the use of regional data for estimating the slope of the curve, as it is almost invariant to changes in policy.

5.6 Case VIII: Changes in the Variance of Shocks—Regional Analysis

Finally, we analyze the effect of changes in the relative variance of demand and supply shocks to answer the following question: Can regional data help estimate the slope of the empirical Phillips curve if the relative variance of shocks changes? We proceed in a similar way as in Case IV presented earlier and increase the variance of cost-push shocks relative to the variance of demand shocks, at both the aggregate and regional levels. Figure 13 shows the results of this exercise. Panels A and B display the baseline case for aggregate and regional variables, respectively; and Panels C and D have the same information but for an economy with a larger variance of cost-push shocks. Comparing the top and bottom panels makes it clear that an increase in the relative variance of cost-push shocks will affect the estimated slope of the empirical Phillips curve, whether estimated using aggregate data or regional data.

This highlights an important limitation with the use of regional data to estimate Phillips curves. The empirical slope of the curve does not need to capture the structural slope in equation (2). Moreover, the empirical slope may flatten over time even if the structural slope did not change.

6 CONCLUSION

We use a simple New Keynesian framework to illustrate the main problems in estimating the slope of the structural Phillips curve. Some of these problems arise due to the actions of monetary policy, as it affects economic activity to fight inflation, which leads to biases in simple estimations. Recent proposals have favored estimations using regional or city data in an effort to overcome these issues, as monetary policy will react to only aggregate economic conditions and not regional. We use a simple model of a monetary union with a continuum of economies and a common monetary policy authority that reacts to aggregate conditions. When we use this model as a data-generating process, we find that the main drawbacks of the aggregate model are still present in a cross-section of many regions in a monetary union. The relative importance of the demand and supply shocks will largely determine the empirical relation between unemployment and prices in both the aggregate and the cross-section of regions. Our analysis shows potential pitfalls in estimating the slope of the Phillips curve, even if using regional data. ■

NOTES

- ¹ We follow McLeay and Tenreyro (2019) and use 10-year-ahead mean forecasts for CPI inflation as our measure of inflation expectations. In their words, “we use five to ten year ahead inflation expectations, as suggested by Bernanke (2007) and Yellen (2015) as having a stronger empirical fit with the data” (p. 32).
- ² For further details of the model, see Gali (2008).
- ³ The specific values we use in our quantitative analysis are not important. Our goal is to make a qualitative point, not a quantitative one.
- ⁴ Note, however, that this assumption is not without loss of generality, since, as we mentioned before, with persistent shocks future expected inflation is correlated with the current realization of the shocks.
- ⁵ And, as said before, one must also instrument for expected inflation.
- ⁶ Note that in the regional Phillips curve case, we abstract from inflation expectations. The reason is that we do not have data on inflation expectations at the regional level, since the Survey of Professional Forecasters contains information on inflation for the whole U.S. economy.
- ⁷ Note that we assume regions within the country are small and do not affect aggregate variables. This assumption is not without loss of generality. In fact, this assumption gives the best chance of success to the hypothesis that using regional data would recover the slope of the Phillips curve. Our point is that, even in this case, the assumption may fail in some circumstances.

REFERENCES

- Atkeson, Andrew and Ohanian, Lee E. “Are Phillips Curves Useful for Forecasting Inflation?” Federal Reserve Bank of Minneapolis *Quarterly Review*, 2001, 25(1), pp. 2-11; <https://doi.org/10.21034/qr.2511>.
- Bernanke, Ben. “Inflation Expectations and Inflation Forecasting.” Speech delivered at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, MA, July 19, 2007; <https://www.federalreserve.gov/newsevents/speech/bernanke20070710a.htm>.
- Blanchard, Olivier; Cerutti, Eugenio and Summers, Lawrence. “Inflation and Activity—Two Explorations and their Monetary Policy Implications.” Working Paper No. 21726, National Bureau of Economic Research, 2015; <https://doi.org/10.3386/w21726>.
- Bullard, James. “The Case of the Disappearing Phillips Curve.” Speech delivered at the 2018 ECB Forum on Central Banking, Sintra, Portugal, June 19, 2018; (accompanying slides) https://www.stlouisfed.org/~media/files/pdfs/bullard/remarks/2018/bullard_ecb_sintra_june_19_2018.pdf.
- Calvo, Guillermo A. “Staggered Prices in a Utility-Maximizing Framework.” *Journal of Monetary Economics*, 1983, 12(3), pp. 383-98; [https://doi.org/10.1016/0304-3932\(83\)90060-0](https://doi.org/10.1016/0304-3932(83)90060-0).
- Cecchetti, Stephen G.; Feroli, Michael; Hooper, Peter; Kashyap, Anil K. and Schoenholtz, Kermit. “Deflating Inflation Expectations: The Implications of Inflation’s Simple Dynamics.” CEPR Discussion Paper No. DP11925, Centre for Economic Policy Research, March 2017.
- Clarida, Richard; Gali, Jordi and Gertler, Mark. “The Science of Monetary Policy: A New Keynesian Perspective.” *Journal of Economic Literature*, 1999, 37(4), pp. 1661-707; <https://doi.org/10.1257/jel.37.4.1661>.
- Coibion, Olivier and Gorodnichenko, Yuriy. “Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation.” *American Economic Journal: Macroeconomics*, 2015, 7(1), pp. 197-232; <https://doi.org/10.1257/mac.20130306>.
- Fitzgerald, Terry J. and Nicolini, Juan Pablo. “Is There a Stable Relationship Between Unemployment and Future Inflation? Evidence from U.S. Cities.” Working Paper 713, Federal Reserve Bank of Minneapolis, Research Department, May 2014; <https://doi.org/10.21034/wp.713>.
- Gali, Jordi. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, 2008.

- Gali, Jordi and Monacelli, Tommaso. "Optimal Monetary and Fiscal Policy in a Currency Union." *Journal of International Economics*, 2008, 76(1), pp. 116-32; <https://doi.org/10.1016/j.jinteco.2008.02.007>.
- Lucas Jr., Robert E. "Expectations and the Neutrality of Money." *Journal of Economic Theory*, 1972, 4(2), pp. 103-24; [https://doi.org/10.1016/0022-0531\(72\)90142-1](https://doi.org/10.1016/0022-0531(72)90142-1).
- McLeay, Michael and Tenreyro, Silvana. "Optimal Inflation and the Identification of the Phillips Curve." NBER Working Paper 25892, National Bureau of Economic Research, 2019; <https://doi.org/10.3386/w25892>.
- Mishkin, Frederic S. "Inflation Dynamics." *International Finance*, 2007, 10(3), pp. 317-34; <https://doi.org/10.1111/j.1468-2362.2007.00205.x>.
- Phillips, Alban W. "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica*, 1958, 25(100), pp. 283-99; <https://doi.org/10.2307/2550759>.
- Roberts, John M. "New Keynesian Economics and the Phillips Curve." *Journal of Money, Credit, and Banking*, 1995, 27(4), pp. 975-84; <https://doi.org/10.2307/2077783>.
- Roberts, John M. "Monetary Policy and Inflation Dynamics." *International Journal of Central Banking*, September 2006, 2(3).
- Samuelson, Paul A. and Solow, Robert M. "Analytical Aspects of Anti-Inflation Policy." *American Economic Review*, 1960, 50(2), pp. 177-94.
- Stock, James H. and Watson, Mark W. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 2007, 39, pp. 3-33; <https://doi.org/10.1111/j.1538-4616.2007.00014.x>.
- Williams, John C. "Inflation Persistence in an Era of Well-Anchored Inflation Expectations." Federal Reserve Bank of San Francisco *Economic Letter*, Number 2006-27, October 2006; <https://www.frbsf.org/economic-research/publications/economic-letter/2006/october/inflation-persistence-in-an-era-of-well-anchored-inflation-expectations/>.

