

Model uncertainty and the direction of fit of the postwar U.S.

Phillips curve(s)

Francesca Rondina[†]

University of Ottawa

Preliminary and incomplete

This version: June 2018

Abstract

This paper proposes a model uncertainty framework that accounts for the uncertainty about both the specification of the Phillips curve and the identification assumption to be used for parameter estimation. I first study the evolution of the model posterior probabilities, which can be interpreted as a measure of the econometrician's real-time beliefs over the prevailing model of the Phillips curve. I then characterize the optimal policy rule within each model, and I analyze alternative policy recommendations that incorporate model uncertainty. As expected, different directions of fit of the same model of the Phillips curve imply very different optimal policy choices, with the "Classical" specifications typically suggesting low and stable optimal inflation rates. I also find that allowing rational agents to incorporate model uncertainty in their expectations does not change the optimal or robust policies. On the other hand, I show that the models' fit to the data and the robust policy recommendations are affected by the specific price index that is used to measure inflation.

Keywords: Phillips curve, Model Uncertainty, Robust Policy, Bayesian Model Averaging, Expectations.

JEL codes: C52; E37; E52; E58

*Contact: Department of Economics, University of Ottawa, Social Sciences Building, 120 University, Ottawa (ON) Canada, K1N6N5; email: frondina@uottawa.ca

[†]This paper is part of a larger project on robust policy-making under different forms of model uncertainty in a Phillips curve framework. I would like to thank Steven Durlauf, Giacomo Rondina, and Emmanuele Bobbio for valuable discussions on this topic, and the participants to the lunch workshops at the Institute for Economic Analysis (IAE-CSIC) and University of Ottawa for their suggestions. I have also benefited from the comments received at the 2013 CEF conference, the 2015 NAWM of the Econometric Society, the 2015 CEA Annual Meeting, and the 2018 Carleton Macro-Finance Workshop (CMFW). The usual disclaimer applies.

1 Introduction

This paper is concerned with the issue of model uncertainty in a Phillips curve representation of the economy, and in particular with the uncertainty about the direction of fit of the empirical model to be estimated. More specifically, I propose a framework that incorporates uncertainty about both the specification of the Phillips curve and the identification assumption to be employed for parameter estimation. I use this framework to address three questions. First, I assess the extent to which the different models of the Phillips curve considered in the analysis fit the data over time. Second, I discuss the role that different forms of model uncertainty might play in policy decisions. Third, I study the impact of private agents' perception of model uncertainty on the models' fit to the data and on the robust policy recommendations.

The negative correlation between inflation and unemployment was first noted by Phillips (1958) in the U.K. data, and has subsequently been studied for many different countries and different time periods. For the U.S., Samuelson and Solow (1960) showed the existence of a negative trade-off similar to the one studied by Phillips (1958), and suggested an empirical model that was able to reproduce this relationship. However, a few years later, the Phillips curve model advanced by Samuelson and Solow (1960) seemed to break down when the rise in inflation that emerged in the late 1960s was not followed by a decrease in unemployment. The failure of the Samuelson-Solow model to rationalize the behavior of the data in these years encouraged other researches to develop alternative representations of the Phillips curve (in particular, this paper will consider the models introduced by Solow, 1968; Tobin, 1968; Lucas, 1972; Sargent, 1973, and a version of the New Keynesian Phillips curve introduced by Galí and Gertler, 1999). According to Cogley and Sargent (2001) and Sims (2001), it should have become clear shortly thereafter that the model of the Phillips curve offering the best fit to the data was one incorporating the natural rate hypothesis.¹

One of the fundamental characteristics of the Phillips curve is that it was discovered as a purely empirical relationship. As such, the original model does not imply any assumptions about the direction of causality between inflation and unemployment. However, in order to be able to estimate the parameters of the Phillips curve equation, the econometrician will need to make an assumption about the direction of fit of the empirical model. In other words, the econometrician will need to

¹Sims (2008) offers an interesting and thought provoking discussion of the Phillips curve and its evolutions over time.

decide whether to use inflation or unemployment as the left-hand side variable in the model to be estimated. In the absence of an underlying structural framework, this choice is, to a large extent, arbitrary. This issue is well known in the economic literature, and several works have discussed the impact that the direction of fit of the Phillips curve has on the estimation and interpretation of the trade-off between inflation and unemployment (see, for instance, King and Watson, 1994; Sargent, 2001; Cogley and Sargent, 2005).

This paper provides a contribution to this literature in a new direction. More specifically, this work proposes a framework in which the uncertainty over the direction of fit of the Phillips curve is included in the analysis in the same way as the uncertainty about the version of the model that best approximates the true data generating process. The models of Phillips curve that I consider are those studied by Cogley and Sargent (2005), with the addition of the New Keynesian Phillips curve. However, two of the models, namely the Samuelson-Solow and the Solow-Tobin Phillips curves, will be estimated using different directions of fit, each of which will be treated as a separate model. To my knowledge, this is the first paper that considers the direction of fit of the Phillips curve as a form of model uncertainty, and that incorporates it in a model uncertainty type of environment. In this sense, this work also adds to the literature on model uncertainty in economics, and more specifically to the branch following the Bayesian Model Averaging approach introduced by Brock, Durlauf and West (2003, 2007).²

The model uncertainty framework that I define in the first part of the paper is then be used for three purposes. First, I study the patterns of the posterior probabilities of each specification of the Phillips curve using U.S. data for the postwar period. Following the standard approach of the literature on uncertainty and learning in economic policy³, these probabilities are computed in each period based on the data available up to that point, and they are updated over time as additional observations become available. This implies that the models' posteriors can be interpreted as real-time estimates of the extent to which each model approximates the true data generating process. Second, I add a monetary policy equation to the baseline framework, and I use it to study the optimal policy rule within each model and alternative policy choices under model uncertainty. Third, I examine the role of private agents' expectations, and in particular I investigate whether allowing

²An alternative approach to model uncertainty is the Minimax Robust Control method developed by Hansen and Sargent (2008).

³See, among the others, Cogley and Sargent (2005), Primiceri (2006), Sargent, Williams and Zha (2006).

rational agents to account for model uncertainty affects the models' posteriors and the robust policy recommendations.

The results of the paper are as follows. In terms of model posterior probabilities, I find that after the mid 1960s, the data quickly favors models of the Phillips curve that imply (long-run) neutrality between inflation and unemployment. However, I show that the exact specification that predominates at each point in time is to some extent sensitive to the price measure used to compute inflation, and to whether private agents account for model uncertainty when forming expectations. With respect to the direction of fit, the results clearly show that the same specification of the Phillips curve can exhibit very different posterior probabilities depending on the direction in which the model is estimated. In the policy exercise, I verify the claim of Cogley and Sargent (2005) that adding uncertainty over the direction of fit does not alter the robust policy if the more unstable models of the economy (the "worst-case" scenarios) were already included in the original model space. I also confirm their conclusion that in a model uncertainty environment, a robust policy type of argument can rationalize the pattern of inflation in the postwar U.S. data, while an "average" policy approach cannot do so. I show that this result extends to the period 2002 – 2017, which was not covered in their original study. Finally, I show that allowing rational agents to incorporate model uncertainty in their expectations can change the patterns of the models' posteriors to some extent, but it does not change the robust policy rules.

Overall, the main take-away of this paper is that accounting for the uncertainty over the direction of fit of the Phillips curve is important for our understanding of the data in real-time and, as a consequence, for our ability to interpret the current behavior of the variables and to predict their future patterns. However, if our main interest is the analysis of robust policy-making under model uncertainty, then a larger model space might be unnecessary, as the recommended policy choices will be affected almost exclusively by the specifications of the Phillips curve that are more difficult to stabilize.

Before moving on with the discussion, I would like to remark that this paper is in the spirit of the literature on economic decisions in real-time, and the analysis is performed from the perspective of an econometrician (or policymaker) who has to form his opinions based on the currently available information. The main purpose of this work is to discuss some possible approaches to decision-making in a model uncertainty environment that econometricians can employ using the data that

they observe, and without having to make any assumptions about the true model of the economy. This implies, among other things, that the patterns of the models' posteriors that I report in the paper are the econometrician's estimates of the model that best approximates the data over time, and they do not (directly) measure changes in the underlying true model of the Phillips curve. Thus, even if these estimates can change as new data becomes available, no assumption is made on whether the underlying true model of the Phillips curve is actually time-varying.⁴

The remainder of the paper is organized as follows. Section 2 describes the models of the Phillips curve considered in the analysis and characterizes the model space. Section 3 studies the patterns of the models' posterior probabilities in the postwar U.S. data. Section 4 discusses the impact of the uncertainty over the direction of fit of the Phillips curve on policy decisions, and analyses the role of private agents' expectations. Section 5 concludes.

2 Models of the Phillips Curve

I assume that the true model of the economy is unknown, and I consider a small number of alternative empirical models, which represent possible approximations to the true data generating process. More specifically, the model space that I consider in this paper is composed of five alternative empirical models, which originate from three different specifications of the Phillips curve. The true data generating process does not need to be among the specifications included in the model space and, as mentioned above, the different models considered in the analysis will all be treated as just approximations to the true underlying representations of the world.

The four approximating models of the Phillips curve that I employ in this paper include those previously studied by Cogley and Sargent (2005) and the New Keynesian Phillips Curve. All the models are in the spirit of the empirical frameworks previously adopted in the macroeconomic literature addressing similar questions (King, Stock and Watson, 1995; Rudebusch and Svensson, 1999; Primiceri, 2006; Brock, Durlauf and West, 2007). Each model postulates its own version of the relationship between the inflation rate π_t , and the unemployment rate u_t . The first model is a version of the Phillips curve proposed by Samuelson and Solow (1960), and allows for a long-run

⁴Some works assuming a time-varying model of the Phillips curve are Conway and Gill (1991), Stock and Watson (2010), Matheson and Stavrev (2013), Blanchard, Cerutti and Summers (2015), and Blanchard (2016).

trade-off between unemployment and inflation:

$$\pi_t = \gamma_0^{SS} + \gamma_\pi^{SS}(L) \pi_{t-1} + \gamma_u^{SS}(L) u_t + \eta_t^{SS} \quad (1)$$

The shock η_t^{SS} is assumed to be i.i.d. $N(0, \sigma_{SS}^2)$.

The second model is inspired by Solow (1968) and Tobin (1968):

$$\Delta\pi_t = \gamma_\pi^{ST}(L) \Delta\pi_{t-1} + \gamma_u^{ST}(L) (u_t - u_t^*) + \eta_t^{ST} \quad (2)$$

Here, u_t^* is the natural rate of unemployment and η_t^{ST} is i.i.d. $N(0, \sigma_{ST}^2)$. This model is a special case of (1), and is characterized by a short-run trade-off between inflation and unemployment, and long-run neutrality.

The third model is in the spirit of Lucas (1972) and Sargent (1973):

$$u_t - u_t^* = \gamma_\pi^{LS} (\pi_t - E_{t-1}\pi_t) + \gamma_u^{LS}(L) (u_{t-1} - u_{t-1}^*) + \eta_t^{LS} \quad (3)$$

where $E_{t-1}\pi_t$ represents private agents' rational expectations of time t inflation, given the information available at time $t - 1$, and η_t^{LS} is i.i.d. $N(0, \sigma_{LS}^2)$. This model is characterized by long-run neutrality and by the assumption that only unexpected inflation is able to affect the unemployment rate. Notice that in order to be able to estimate the fit of the Lucas-Sargent model to the data, I will need to compute the value of $E_{t-1}\pi_t$, which in turns means that I will need to make further assumptions about how π_t is determined in this model. Following the previous literature using similar models (Sargent, 2001; Cogley and Sargent, 2005; Sargent, Williams and Zha, 2006), I will assume that the monetary authority is able to control the inflation rate to some extent, and that private agents will set their expectations based on this information. Further details will be provided in section 3.

The last model included in the analysis is the New Keynesian Phillips Curve. The specification that I adopt is the same previously used by Nason and Smith (2008a, 2008b) and has been referred

to as the *hybrid* New Keynesian Phillips Curve:⁵

$$\pi_t = \gamma_0^{NK} E_{t-1} \pi_t + (1 - \gamma_0^{NK}) \pi_{t-1} + \gamma_u^{NK} (u_t - u_t^*) + \eta_t^{NK} \quad (4)$$

with η_t^{NK} assumed to be i.i.d. $N(0, \sigma_{NK}^2)$. In this model, $(u_t - u_t^*)$ measures the “slack” in the economy and η_t^{NK} is usually interpreted as a cost-push shock. As for (3), the model includes private agents’ expectations of inflation.

2.1 The Direction of Fit of the Phillips Curve

As previously mentioned, the Phillips curve was first discovered as an empirical relationship between inflation and unemployment. While alternative models of the Phillips curve have been proposed over time, several of them remain empirical relationships, and do not aim at making explicit statements about the direction of causality between these two variables. This implies that the econometrician will need to decide which orthogonality assumption to impose in order to be able to estimate the parameters of the model. Using the notation introduced above, Phillips curves in the form of (1) and (2) have been estimated using π_t as the dependent variable and assuming that the shock η_t^i is orthogonal to u_t , or vice-versa using u_t as the dependent variable and assuming that the shock η_t^i is orthogonal to π_t . Sargent (2001) called these two alternative identification assumptions “Keynesian” and “Classical”. In the Keynesian identification assumption, the shock η_t^i is assumed to be orthogonal to current unemployment (in addition to all the other right hand side variables, such as past inflation and unemployment). On the contrary, in the Classical identification assumption the shock η_t^i is assumed to be orthogonal to current inflation (again, in addition to all the other right hand side variables).

The choice of the Keynesian or Classical identification assumption in a Phillips curve framework is not without consequences. As previously mentioned, King and Watson (1994) and Sargent (2001) show that the identification assumption used to estimate the parameters of the Phillips curve affects the nature of the trade-off between inflation and unemployment, and implies a different optimal policy response to changes in the variables of interest. In addition, Cogley and Sargent (2005) show that the assumption used to identify the coefficients of the Phillips curve affects the estimated sacrifice

⁵See Nason and Smith (2008a, 2008b) for a discussion of a number of underlying models of the economy that could originate Phillips curves in the form of (4).

ratios, which measure the cost of reducing inflation in terms of unemployment. In more detail, they show that the Samuelson-Solow and Solow-Tobin models described by (1) and (2) produce high sacrifice ratios if estimated using the Keynesian assumption, while they suggest a much lower cost of disinflation if estimated using the Classical assumption. This in turn implies that the optimal inflation rate recommended by these two models will typically be lower and more stable when they are estimated using the Classical direction of fit instead of the Keynesian.

In the analysis of the robust policy choice in the face of model uncertainty, Cogley and Sargent (2005) argue that it is unnecessary to account for the uncertainty over the direction of fit of the Phillips curve if the specifications that represent the worst-case scenario of the economy are already included in the model space.⁶ However, researchers might be interested in adopting a model uncertainty approach to address other type of questions. For instance, they might want to investigate which specification of the Phillips curve provides the best fit to the data given the information that they have available. As it will be discussed below, this exercise might offer interesting insights about the true underlying model of the economy. In addition, the knowledge of the model that best describes the data at each point in time might be useful to interpret the behavior of inflation and unemployment and to predict their future patterns. Finally, the monetary authorities might be interested in analyzing the optimal policies recommended by alternative specifications of the Phillips curve, and they might want to study how far apart these policies are.⁷

This paper suggests that, for the reasons just mentioned, it might be of interest to researchers and policymakers to consider a model space that incorporates uncertainty over the direction of fit of the Phillips curve in addition to specification uncertainty. In particular, this second form of model uncertainty will be accounted for in the models of Phillips curve for which this issue was originally pointed out in the literature, i.e. (1) and (2).⁸ Thus, the space of models of the Phillips curve under study in this paper will be composed of 6 model specifications: the Samuelson-Solow model, estimated using the Keynesian and the Classical identification assumptions (from now on denoted as $SS - K$ and $SS - C$), the Solow-Tobin model estimated using the Keynesian and the Classical

⁶Section 4 of this paper does indeed confirm this argument.

⁷Brock, Durlauf and West (2007) call this type of analysis “action dispersion”.

⁸In the LS model, only the unexpected part of inflation is able to affect the unemployment rate, so this model can only be estimated using $u_t - u_t^*$ as independent variable. The NK Phillips Curve, on the other hand, is typically only estimated using the inflation rate as independent variable, as suggested by the underlying structural models that are usually employed to obtain this equation.

identification assumptions (denoted as $ST - K$ and $ST - C$), the Lucas-Sargent model (denoted as LS), and finally the New Keynesian model (denoted as NK). In the $SS - K$ model, equation (1) will be estimated using inflation as the dependent variable, with the identification assumption $E(u_t \eta_t^{SS-K}) = 0$. On the other hand, in the $SS - C$ model the same equation will be estimated using unemployment as the dependent variable, with the identification assumption $E(\pi_t \eta_t^{SS-C}) = 0$. Similarly, in the $ST - K$ model equation (2) will be estimated using the identification assumption $E(u_t \eta_t^{ST-K}) = 0$, while in the $ST - C$ model the same equation will be estimated under the assumption $E(\pi_t \eta_t^{ST-C}) = 0$.

To provide some further intuition about the relationship between the Keynesian and Classical versions of the same model of the Phillips curve, I conclude this section with an example which uses the Samuelson-Solow model described by (1).⁹ Let assume that inflation admits the following Moving Average (MA) representation: $\pi_t = f(L) v_t$, where v_t could be a combination of structural shocks of the economy, with $E(v_t) = 0$.¹⁰ Under the Classical identification assumption, $E(\pi_t \eta_t^{SS-C}) = 0$ implies that $E(v_t \eta_t^{SS-C}) = 0$. On the other hand, under the Keynesian identification assumption $E(\pi_t \eta_t^{SS-K}) \neq 0$, which implies that $E(v_t \eta_t^{SS-K}) \neq 0$. Let

$$\eta_t^{SS-K} = \rho v_t + \tilde{\eta}_t^{SS-C}$$

where $\rho = cov(v_t, \eta_t^{SS-K}) / var(v_t)$. Notice that, by construction, $E(v_t \tilde{\eta}_t^{SS-C}) = 0$, which also implies that $E(\pi_t \tilde{\eta}_t^{SS-C}) = 0$. Thus, we can rewrite:

$$\pi_t = \gamma_0^{SS-K} + \gamma_\pi^{SS-K}(L) \pi_{t-1} + \gamma_u^{SS-K}(L) u_t + \rho v_t + \tilde{\eta}_t^{SS-C}$$

and we can further substitute $v_t = \tilde{f}(L) \pi_t$, where $\tilde{f}(L) = f^{-1}(L)$, to obtain:

$$\pi_t = \gamma_0^{SS-K} + \gamma_\pi^{SS-K}(L) \pi_{t-1} + \gamma_u^{SS-K}(L) u_t + \rho \tilde{f}(L) \pi_t + \tilde{\eta}_t^{SS-C} \quad (5)$$

⁹This example is very much in the spirit of the discussion presented in King and Watson (1994) and Sargent (1976).

¹⁰This MA representation for π_t could be interpreted, for instance, as originating from the policy rule that the monetary authority uses to control the inflation rate.

Finally, we can rewrite (5) in a Classical form as:

$$u_t = \tilde{\gamma}_0^{SS-C} + \tilde{\gamma}_\pi^{SS-C} (L) \pi_t + \tilde{\gamma}_u^{SS-C} (L) u_{t-1} + \gamma_\eta \tilde{\eta}_t^{SS-C} \quad (6)$$

where $E(\pi_t \tilde{\eta}_t^{SS-C}) = 0$, and all the parameters can be obtained as a function of ρ and the parameters in (5).

This example shows that if inflation has the MA representation assumed above, then from a model of the Phillips curve estimated using the Keynesian identification assumption, we can obtain an equivalent model that has the same characteristics of the Classical version of the same Phillips curve. This conclusion can be extended to more general representations for the processes of inflation and unemployment (for more details, see King and Watson, 1994). Yet, as I will show later on, the Classical and Keynesian versions of the same model are very different in terms of fit to the data and policy implications.

2.2 Empirical Approach

The empirical approach that I use to estimate the parameters of each model and to compute their posterior probabilities is the same as in Cogley and Sargent (2005). For each Phillips curve specification, I estimate the model parameters using Bayesian methods. Given the selected identification assumption, (1) - (4) are simple regression models. For each model, the prior distribution of the parameters is assumed to be of the Normal-Inverse Gamma family. The Phillips curve residuals η_t^i in (1) - (4) are assumed to be i.i.d. and conditionally normal given the regressors for each of the models. These assumptions ensure that the conditional likelihood function is Gaussian. Given a Normal-Inverse Gamma prior and a Gaussian conditional likelihood, the posterior joint distribution will also be of the Normal-Inverse Gamma family, with parameters that can be updated recursively. More details, together with the updating formulas for the parameters, are given in Appendix 2.

The fit of each specification of the Phillips curve to the data is evaluated based on the model's posterior probability. Let $i = \{SS - K, SS - C, ST - K, ST - C, LS, NK\}$. The posterior probability of model i given data up to time t can be defined according to Bayes's theorem as:

$$p(M_i | Y^t, X^t) \propto m_{it} \cdot p(M_i) \quad (7)$$

where $p(M_i)$ and $p(M_i | Y^t, X^t)$ are the prior and posterior probabilities of model i , while m_{it} is the marginalized likelihood function for model i at time t . Here, X^t represents the history up to time t of the right-hand variables of the model, and Y^t the history up to time t of the left-hand variable. Notice that the set of right-hand variables and left-hand variables will be different for each of the specifications of the Phillips curve under analysis.

The expression for the models' posteriors can be used to compute the normalized posterior probabilities, defined as:

$$\alpha_{it} = \frac{w_{it}}{\sum_i w_{it}} \quad (8)$$

where $w_{it} \equiv m_{it} \cdot p(M_i)$. Given the assumptions on the posterior distribution of the parameters of each model, the marginalized likelihood function m_{it} can be computed analytically, and the values of w_{it} can be updated recursively.¹¹ Again, Appendix 2 provides more details and reports the updating formulas.

3 The Postwar U.S. Phillips Curve(s)

I start by employing the model uncertainty framework presented in the previous section to study the patterns of the model posterior probabilities in the U.S. data from 1960 to 2017. The purposes of this exercise are two. First, I want to examine whether accounting for the uncertainty over the direction of fit of the Phillips curve affects the pattern of the model posteriors. Cogley and Sargent (2005) studies a model space that only includes the $SS - K$, $ST - K$ and LS Phillips curves, and report the normalized posterior probabilities for these models in the period 1960 – 2002. By comparing the results obtained in this paper with those arising from the model space of Cogley and Sargent (2005), we can verify whether the model of the Phillips curve that best fits the data over time changes when the larger model space is used for the analysis. Second, I am interested in investigating the patterns of the posterior probabilities during the last few years of the sample, and in particular during the 2008 – 2009 recession. A recent literature in economics suggests that there might have been a change in the true underlying relationship between inflation and unemployment during the global financial crisis (see, for instance, Coibon and Gorodnichenko, 2015; Friedrich, 2014). If this is the case, then it is possible that the patterns of the models' posteriors will show larger adjustments around 2008.

¹¹See Cogley and Sargent (2005) for a more extensive explanation of this result.

For the interpretation of the results, it is worth remarking that the normalized posterior probabilities reported in this paper are a measure of the models' fit to the data up to the point in time at which they are computed. In other words, for each time t , they measure the models' relative ability to describe the data in the sample period going from time 1 to time t . This implies that these posteriors will change over time because the additional data that becomes available will improve the estimates of each model. However, the posteriors could also change because of changes in the underlying true model of the economy. As mentioned above, one of the central features of the model uncertainty approach adopted in this paper is that no assumptions is made on the actual data generating process, and the specifications of the Phillips curve considered in the analysis could all be just approximations of the true model of the economy. As a consequence, different specifications could fit the data better in some periods rather than others, especially if the underlying data generating process is not stable over time. If this is the case, then the relative posterior probabilities will keep adjusting. Notice that, for this reason, it is not necessarily the case that there will be one dominant model for which the posterior probability will converge to one in the long run.

3.1 The Prevailing Model of the Phillips Curve

I estimated the model parameters and posterior probabilities using U.S. quarterly data for inflation and unemployment. Unemployment is the quarterly average of the monthly Civilian Unemployment rate. In the baseline exercise, inflation is computed from the Personal Consumption Expenditure (PCE) chain-type price index, which is the primary price measure used by the Federal Reserve for policy decisions.¹² However, I also report the results obtained using the GDP chain-type price index, the GDP implicit price deflator, and the CPI. The data employed in the estimation goes from 1960 : I to 2017 : IV , while observations from 1949 : I to 1959 : IV are used to set the parameters in the prior distributions and the initial values in the updating formulas. Appendix 1 gives further details about the assumptions and the initial settings used in the estimation procedure.

The Keynesian and Classical versions of the SS and ST models are the same in everything, except for the direction of fit of the equation that is estimated. The ST , LS , and NK models include the variable u_t^* , which denotes the natural rate of unemployment. As known, this variable

¹²In the February 2000 "Monetary Policy Report to the Congress", the Federal Open Market Committee (FOMC) stated that its main measure of inflation was changing from CPI inflation to inflation computed using the PCE chain-type price index.

is not observed by the econometrician, and needs to be either estimated or approximated. While I believe that it would be very interesting to treat u_t^* as unobservable, and estimate its history jointly with all the other parameters of the models, this approach would significantly complicate the computations, particularly for the model posteriors.¹³ For this reason, in this paper I decided to approximate u_t^* using the same formula as in Cogley and Sargent (2005)

$$u_t^* = u_{t-1}^* + g(u_t - u_{t-1}^*) \quad (9)$$

with gain parameter set as $g = 0.075$. The last part of this section will provide a discussion of the sensitivity of the results to alternative approximating formulas for u_t^* .

As discussed above, in order to estimate equations (3) and (4) I will need to specify how the variable $E_{t-1}(\pi_t)$ is determined in this framework. Following the previous literature using similar models of the Phillips curve, I assume that the monetary authority is able to control the inflation rate to some extent, and that private agents will set their expectations based on this information. Given this assumption, two possible approaches could be followed in a model uncertainty environment. The approach that I adopt in the baseline analysis is the same as in Cogley et al. (2011), and it is based on the idea that expectations are formed “internally” within the model. In other words, private agents in the *LS* model are assumed to believe that the true model of the economy is the *LS* model and that the value of inflation will be equal to the optimal rate based on this model. Similarly, in the *NK* model private agents believe that the true model of the economy is the *NK* model and form their expectations accordingly. In a model uncertainty environment, an alternative to this assumption is the approach employed by Cogley and Sargent (2005), which set $E_{t-1}(\pi_t) = x_{t|t-1}^E$. Here, $x_{t|t-1}^E$ is the “Encompassing” policy choice, which takes into account policymakers’ uncertainty over the true model of the economy. Under this second approach, private agents’ expectations will incorporate model uncertainty through its impact on the decisions of policymakers. My choice of setting expectations within the model for the *LS* and *NK* frameworks was motivated by a preference for leaving all the models self-contained in their implications and predictions, in order to let the data discriminate more clearly between them. However, in order to have a better understanding of the impact of private agents’ expectations on the conclusions of this paper, I will later also discuss the

¹³In this case, we would not be able to use the analytical formulas reported in Appendix 2 to compute the model weights.

results and the policy implications under the alternative approach $E_{t-1}(\pi_t) = x_{t|t-1}^E$ in both the *LS* and *NK* models.

For the choice of the models' priors, I followed the same reasoning as in Cogley and Sargent (2005). Since the *ST*, *LS*, and *NK* models were developed after 1960 : *I*, I set the prior to 0.97 for the *SS* model, and to 0.01 each for the *ST* *LS* and *NK* models. For the *SS* and *ST* models, I split the prior equally between the Classical and the Keynesian specifications. Thus, the *SS* – *K* and *SS* – *C* models received a prior of 0.485 each, the *ST* – *K* and *ST* – *C* models a prior of 0.005 each, and the *LS* and *NK* models a prior of 0.01 each. Again, I will discuss the sensitivity of the results to this choice.

Figure 1 reports the patterns of the model posteriors in the 1960 – 2017 period for the full model space considered in this paper. The figure shows that the *SS* model estimated under the Keynesian assumption remains the predominant model in the first few years of the sample, while the *LS* Phillips curve becomes the model with the highest posterior in the last half of the 1960s. From the early 1970s until the end of the sample, two models maintain roughly 1/2 of the posterior probability each, they are the *ST* – *C* and the *LS* Phillips curves. At different points, one of these models can predominate on the others for a few years, but their posteriors remain always substantially larger than zero.

With respect to the direction of fit of the Phillips curve models, there are some interesting insights that we can observe from Figure 1. The first one is the importance of the Classical versions of the *SS* and *ST* models. The posterior attached to the *SS* – *C* model is substantial in the first part of the sample, while the *ST* – *C* model exhibits a relatively high posterior from the mid 1970s until today. Overall, the Classical models of the Phillips curve, i.e. those estimated with unemployment as the left-hand side variable, seem to fit the data better than their Keynesian counterparts in almost the entire sample period under analysis. A second interesting results is that none of the specifications with the inflation as the dependent variable seems to fit the data particularly well in the sample period under analysis, with the exception of the *SS* – *C* model in the very first years of the sample. Finally, it is also worth noting that the results shown in Figure 1 are somehow different from those reported by Cogley and Sargent (2005). More specifically, the *ST* – *C* and *NK* Phillips curves, which were not included in the framework of Cogley and Sargent (2005), seem to capture part of the posterior probability that was estimated to be attached to the *LS* model in this contribution.

Thus, accounting for the uncertainty over the direction of fit of the *SS* and *ST* Phillips curve can change policymakers' views over the models of the economy that best fit the data over time.

A final interesting result that we can observe from Figure 1 is that the patterns of the model posteriors do change considerably in the years of the 2008-2009 crisis. More specifically, the posterior attached to the *ST* – *C* model increased rapidly right after the beginning of the crisis at the expenses of the *LS* model. Since 2010, the posteriors of the two models have been slowly moving back towards their previous values.

One possible interpretation for the behavior of the posterior probabilities for the *ST* – *C* and *LS* Phillips curves is in terms of private agents' expectations of inflation. Let consider the *LS* Phillips curve reported in (3). We can relax the assumption of rational expectations and assume instead that expectations are formed using some general rule in the form:

$$E_{t-1}(\pi_t) = \gamma\pi_t + (1 - \gamma)\gamma^e(L)\pi_{t-1} \quad (10)$$

with $\left[\gamma + (1 - \gamma)\sum_{j=0}^J\gamma_j^e\right] = 1$.¹⁴ Under this alternative assumption, it is straightforward to see that the *LS* model in (3) can be rewritten as:

$$u_t - u_t^* = \tilde{\gamma}_\pi^{LS}(\pi_t - \pi_{t-1}) + \tilde{\gamma}^e(L)(\pi_{t-1} - \pi_{t-2}) + \gamma_u^{LS}(L)(u_{t-1} - u_{t-1}^*) + \eta_t^{LS} \quad (11)$$

where $\tilde{\gamma}_\pi^{LS}$ and $\tilde{\gamma}^e(L)$ can be obtained from γ_π^{LS} , γ and $\gamma^e(L)$. Notice that the assumption on the parameters γ and $\gamma^e(L)$ in the expression for private agents' expectations implies that $\left[\tilde{\gamma}_\pi^{LS} + \sum_{j=0}^J\tilde{\gamma}_j^e\right] = 0$. Thus, if expectations are formed according to (10), equation (3) becomes exactly the same as the equation for the *ST* model estimated under the Classical identification assumption. The purpose of this discussion is to point out that, in fact, the *ST* – *C* and the *LS* specifications can be interpreted as two different versions of the same model. In one of these versions (*LS*) rational expectations are imposed, while in the other (*ST* – *C*) a more general expectation formation process is assumed instead.¹⁵

In the light of this argument, we can attempt some conjectures about the changes in the true

¹⁴This condition enforces expectations to be correct in the long-run.

¹⁵We could even extend this interpretation to the *SS* – *C* model, which can be viewed as a version of (3) where expectations are formed based on (10) but the condition $\left[\gamma + (1 - \gamma)\sum_{j=0}^J\gamma_j^e\right] = 1$ does not hold.

underlying data generating process based on the patterns of the model posteriors reported in Figure 1. The posteriors for the $ST - C$ and LS models suggest that, starting from the mid 1970s, the data has been favoring a model of the economy that resembles (3), but in which expectations were not always rationally formed. The assumption of rational expectations seems to fit the data slightly better for a long period of time, specifically from the mid 1970s until around 2007. However, during the recent financial crisis, a version of (3) in which private agents were allowed to form expectations using a more general rule was able to describe the data to a much better extent. The study the true model of the Phillips curve and its possible changes over time is not among the main purposes of this paper; nonetheless, the patterns of the model posteriors reported in Figure 1 provides some interesting insights that researchers could use to investigate this issue further.

3.2 Sensitivity Analysis

The patterns reported in Figure 1 are robust to changes in several of the assumptions used in the empirical implementation of the model uncertainty framework. Different values of the parameter g in (9) did not change the pattern of the posterior probabilities reported in Figure 1. Replacing (9) with the updating formulas employed by Primiceri (2006) did not substantially altered the results either.¹⁶ In general, alternative assumptions generating a smoother series for u_t^* resulted in the $SS - C$ model maintaining a high posterior for a longer period at the beginning of the sample, but the patterns of the posterior probabilities in the later portion of the sample were never affected in a significant way. Similarly, changes in the models' priors resulted in some changes in the first part of the sample (where the priors have a relatively larger impact, since the number of observations is still small), but the normalized posteriors were again very similar to those reported in Figure 1 in the later part of the sample. The results were also unaffected by changes in the length of the training sample, and in the approach used to approximate $E_{t-1}(\pi_t)$ in the LS and NK models within the training sample (see Appendix 1 for more details).

3.2.1 Measuring Inflation

There is no widespread consensus about which measure of inflation should be used to estimate the Phillips curve. The main measure used in this paper is the PCE chain-type index which, as

¹⁶Primiceri (2006) estimates the current value of the natural rate of unemployment using a constant gain learning approach. For the specific formulas, see Primiceri (2006).

mentioned, is the primary indicator used by the Federal Reserve for policy purposes. However, in order to assess the robustness of the results and to compare them with those of previous contributions in this area, I repeated the analysis using a few alternative measures of inflation. These measures are the GDP deflator chain-type index (as in Cogley and Sargent, 2005), the GDP implicit price deflator (as in Primiceri, 2006), and the CPI for all items (as in Blanchard, 2016). Figure 2 compares the model posteriors computed using these alternative measures with the baseline values obtained using the PCE index and reported in Figure 1.

The patterns of the models' posteriors are similar for all measures, but there are some differences, particularly on the relative fit to the data of the $ST - C$ and LS models. These two specifications of the Phillips curve are very similar in terms of policy implications but, as I will discuss in the next section, some of the differences in the patterns of the models' posteriors might actually originate different robust policy recommendations in a model uncertainty framework. A second consideration that we can make from Figure 2 is that the shift in posteriors that happened around 2009 is the most evident when the PCE index is used to compute inflation. The shift is still present, but not as pronounced, when the CPI index is used, while in the case of the GDP implicit price deflator the relative weight of the $ST - C$ model seems to have been slowly increasing since the mid 1970s.

Overall, even if there are some differences in the four panels shown in Figure 4, the main message that they all convey is still the same: starting from the early 1970s, only two of the six models of Phillips curve considered in the analysis exhibit posteriors that are (essentially) different from zero. Both of these models assume long-run neutrality between inflation and unemployment. In addition, they imply a very steep (or vertical) short-run Phillips curve, in which the incentives to exploit the unemployment-inflation trade-off are low, because a reduction in unemployment would require a very large cost in terms of inflation. In a model uncertainty environment, what alternative policies could a "robust" policymaker implement? The answer to this question is discussed in the next section.

4 Policy Implications

In a model uncertainty environment, there are several directions in which policy analysis can be discussed. In this section, I will first focus on the optimal policy choices within each of the models of the Phillips curve described in the first part of the paper, and then I will move into the analysis of

policies that account for model uncertainty. The main approach that I will use in this second group of policies is the robust regulator problem proposed by Cogley and Sargent (2005).

One thing that is worth remarking is that, in a model uncertainty environment, the monetary authorities will need to consider, at least to some extent, the effects of policy changes in all the models that they believe to be possible approximations of the underlying data generating process. In this sense, a model uncertainty framework will urge policymakers to address the issue of Lucas' critique. A model of the Phillips curve with time-varying parameters will perhaps be more flexible and it might be able to capture changes in the data to a better extent, but it will probably be unable to incorporate the impact of policy changes as precisely. In a model uncertainty framework, on the other hand, policymakers will be able to explicitly take into account that alternative model specifications imply different assumptions on the channels through which policy affects the economy and on the set of parameters that remains unaffected by policy changes.

In order to be able to discuss policy implications, the models of the Phillips curve described in section 2 need to be augmented of a monetary policy equation. The general assumption employed in the Phillips curve literature is that the central bank can control the inflation rate to some extent.¹⁷ In particular, I will follow the general approach employed by Cogley and Sargent (2005) and Sargent, Williams and Zha (2006), and assume that the policymaker can set the value of the policy instrument $x_{t|t-1}$, which is related to inflation according to the expression:

$$\pi_t = x_{t|t-1} + \xi_t \tag{12}$$

where ξ_t is an i.i.d. normal shock with mean zero and variance σ_ξ^2 . The subscript on the policy variable x implies that the policymaker will choose the value of the policy instrument for time t using the information available at time $t - 1$. Thus, the value of $x_{t|t-1}$ is assumed to be decided in real-time, based on the observable data and on the value of the parameters estimated using this data.

I assume that the policymaker chooses the value of the policy instrument optimally by solving a linear quadratic dynamic programming problem. More specifically, the policymaker aims at

¹⁷This could happen directly, as in the $SS - C$, $ST - C$, and LS models, or through their ability to affect the unemployment rate or the unemployment gap, as in the $SS - K$, $ST - K$, and NK models.

minimizing the loss function:

$$\mathcal{L} = E_t \sum_{j=0}^{\infty} \beta^j \left[(u_{t+j} - u^*) + \lambda (\pi_{t+j} - \pi^*)^2 \right] \quad (13)$$

where β is the discount factor, u^* and π^* are the target values of the variables, and λ represent the weight attached to inflation relative to unemployment. As common in the literature on learning in macroeconomics (see, for instance, Evans and Honkapohja, 2001), the policymaker is assumed to consider the estimated parameters of each model as if they were true values, which will not change over time. This assumption allows to disregard parameter uncertainty, and makes the decision rules depend only on point estimates rather than on the whole posterior distributions of the coefficients. In addition, it allows the policymaker to engage in a process of learning that is passive, in the sense that it takes into account the new information that becomes available but ignores the effects of current decisions on the future values of the variables of interest.¹⁸ This approach has been called “anticipated utility” decision-making by Kreps (1998).

The method used to compute the optimal and robust policy rules is described in detail in Appendix 3. All the exercises performed in this section use a target inflation rate $\pi^* = 2\%$, and a target unemployment rate equal to the value of u_t^* in the period in which the optimal policy is computed. However, none of these assumptions substantially matters for the results, as the main conclusions of the policy analysis remain unaltered if the values of the targets are changed in a reasonable manner, or if the targets are eliminated altogether.

Figure 3 reports the optimal policy recommendations for each specification of the Phillips curve included in the model space. The middle panel focuses on the Keynesian specifications (including the *NK* model), while the bottom panel shows the optimal policies for the Classical specifications (including the *LS* model). The patterns of the posterior probabilities already presented in Figure 1 are also reported in the top panel to facilitate the interpretation of the results. As already discussed before, in the *LS* model the optimal value of $x_{t|t-1}^{LS}$ is the target inflation rate. As expected, the other Classical models also recommend low inflation rates, which are very close to the 2% target. On the other hand, the optimal policy rates for the Keynesian models are much higher. These models

¹⁸In a learning environment, the connection between today’s decisions and tomorrow’s information might create the incentive for experimentation, which is not allowed in this framework. Cogley, Colacito and Sargent (2007) and Svensson and Williams (2007) discuss the benefits that policymakers could obtain from experimentation by exploiting this connection.

estimate small Phillips curve slopes, which suggest that unemployment can be maintained low at the expenses of the inflation rate. An alternative interpretation of the policy recommendations originating from the Keynesian models is the one discussed in Cogley and Sargent (2005) for the $SS - K$ and $ST - K$ specifications. These models turn out to be unstable in some parts of the sample if low inflation policies are implemented (this is true, in particular, in the first part of the sample until the mid 1980s). Thus, the patterns of $x_{t|t-1}^{SS-K}$ and $x_{t|t-1}^{ST-K}$ also reflect the fact that policymakers will need to accept higher inflation rates in order to stabilize these models.

As shown in Figure 3, the empirical models of the Phillips curve included in the model space imply very different optimal policy recommendations. So, what specific policy should be selected in a model uncertainty environment? Figure 4 provides three possible answers to this question. This figure reports three different policy measures. The top panel shows the weighted average optimal policy, computed as the sum of the optimal policies weighted by the models' posteriors. The middle panel reports the optimal policy for the model with the highest posterior in each quarter. Finally, the bottom panel reports the robust policy recommendation obtained using the approach proposed by Cogley and Sargent (2005).¹⁹ In all panels, the actual inflation rate is also depicted for comparison.

A large literature in economics has focused on policymaker's changing beliefs over the true model of economy as the main explanation for the rise and fall of the U.S. inflation in the postwar period (see, for instance, DeLong, 1997; Sargent, 2001; Cogley and Sargent, 2005; Sargent and Williams, 2005; Primiceri, 2006). The results reported in Figures 3 and 4 provide some interesting insights in this respect. It is clear from Figure 3 that the alternative models of the Phillips curve included in the analysis recommend very different optimal inflation rates. In addition, the results offer a clear picture of the extent to which the direction of fit used to estimate the SS and ST models matters for monetary policy. As previously discussed, the identification assumption used to compute the parameters of the Phillips curve has a large impact on the perceived cost of disinflation and, consequently, on the monetary authority's optimal policy decisions. This argument is evident in the middle and bottom panels of Figure 3, which show that for the same specification of the Phillips curve, a different direction of fit of the estimated empirical model implies completely different policy recommendations.²⁰

¹⁹Again, see Appendix 3 for more details about this approach.

²⁰Figure 3 is consistent with the statement of Cogley and Sargent (2005) that for the results of their paper, "the direction of fit matters more than the qualitative nature of the tradeoff" (Cogley and Sargent, 2005, p. 547).

The results reported in Figure 4 also seem to confirm the conclusion of the previous literature that, in a model uncertainty environment, the model (or models) that represents the worst-case scenario of the economy is going to affect the robust policy choice to a very large extent. As the Keynesian versions of the Samuelson-Solow and Solow-Tobin models are less stable and imply significantly higher sacrifice ratios²¹ than the corresponding Classical versions, Cogley and Sargent (2005) suggest that these specifications are the models against which the policymaker will want to protect. For this reason, they argue that including the Classical versions of the same models in the optimal control problem would not alter the robust policy recommendations. The results reported here do indeed support this argument, as the policies depicted in the bottom panel of Figure 4 are very similar to the robust policies obtained by Cogley and Sargent (2005) for a reduced model space that consists of the $SS - K$, $ST - K$, and LS models only. In addition, this argument seems to be confirmed by the three panels of Figure 4. Indeed, by comparing the first two panels (which show the weighted average optimal policy and the optimal policy from the highest posterior model) with the third one, it is clear that the robust inflation rate departs from the optimal policy in the most likely model or from an average of optimal policies. In this sense, we can conclude that what matters the most to the robust policymaker is the magnitude of the trade-off between inflation and unemployment in the model(s) in which the cost to reduce inflation in terms of unemployment is the highest.

4.1 Private Agents' Expectations

The policies shown in Figure 3 and Figure 4 were computed under the assumption that private agents use the optimal policies within the model to set their expectations, i.e. $E_{t-1}(\pi_t) = x_{t|t-1}^{LS}$ in the LS model and $E_{t-1}(\pi_t) = x_{t|t-1}^{NK}$ in the NK model. As I explained above, this assumption implies that private agents form expectations internally within the model, without taking model uncertainty into account. Cogley and Sargent (2005), on the other hand, assume that private agents form expectations by setting $E_{t-1}(\pi_t) = x_{t|t-1}^E$, where $x_{t|t-1}^E$ is the robust policy under model uncertainty, computed using the approach described in Appendix 3. This assumption implies that private agents account for model uncertainty in their expectations through its impact on policymakers' decisions.

In order to investigate the impact of private agents' expectations on the conclusions of the

²¹In terms of the increase in unemployment required to reduce inflation.

paper, I computed the patterns of the posterior probabilities again using $E_{t-1}(\pi_t) = x_{t|t-1}^E$ instead of $E_{t-1}(\pi_t) = x_{t|t-1}^{LS}$. I repeated this exercise for all the measures of inflation that were employed to produce Figure 2. The results are reported in Figure 5.

The different way in which agents are assumed to form expectations can potentially change the fit to the data of the *LS* and *NK* models and, as a consequence, the normalized posterior probabilities. In fact, for some of the measures of inflation employed in the analysis, the normalized posteriors reported in Figure 5 are slightly different from those shown in Figure 2. In particular, assuming that private agents account for model uncertainty in their expectations improves the fit of the *LS* model, especially when inflation is measured using the PCE index or the CPI. Despite the different patterns of the posterior probabilities reported in Figure 5, the story that all these figures suggests is similar to the one depicted in Figure 2, the only difference seems to be the relative weight of the two Classical models of the Phillips curve that maintain a non-zero posterior after 1970. In terms of robust policy, it is possible to show that for all the measures of inflation considered in Figures 2 and 5, the robust policy recommendations are almost unaffected if private agents' expectations are assumed to be set as $E_{t-1}(\pi_t) = x_{t|t-1}^E$. This result is consistent with the discussion of the previous section, as both the *ST - C* and the *LS* models imply very similar optimal policies.

5 Concluding Remarks

This paper studied a model uncertainty framework that explicitly accounts for the uncertainty about the specification of the Phillips curve and the identification assumption to be used for the estimation of its parameters. This framework was employed to analyze the changes in the fit of the different models over time, and to understand the implications of these changes for monetary policy.

The main messages of the paper are three. First, different directions of fit of the same specification of the Phillips curve fit the postwar U.S. data in a very different way, imply very different trade-offs between inflation and unemployment, and suggest very different optimal policies. In general, "Classical" Phillips curves fit the data better than their Keynesian counterparts; these models typically imply a very steep (or vertical) slope of the Phillips curve, and recommend optimal policies in which the inflation rate is low and stable. Second, in rational expectations models of the Phillips curve (the *LS* and *NK* models), assuming that private agents incorporate policymakers'

model uncertainty in their expectations can change the fit of the model to the data, but in a model uncertainty framework, it does not alter the robust policy recommendation. Third, the price index used to compute the inflation rate matters for the fit of the models to the data, and for the robust policy choices.

The analysis developed in this paper presents several suggestions for future research. I am particularly interested in extending the discussion in two directions. First, I believe that the issue of private agents' expectations in a model uncertainty environment should be investigated further. The literature in this area has focused on policymakers' uncertainty about the true model of the economy, while private agents are typically assumed to disregard model uncertainty (as in Cogley et al., 2011) or to incorporate it through their observation of policy decisions (as in Cogley and Sargent, 2005). A different approach is the one of Brock, Durlauf and West (2007), in which the form of model uncertainty that the monetary authorities face is precisely the uncertainty about the way in which private agents form expectations. I think that it would be very interesting to extend this analysis towards the study of environments in which private agents are also assumed to be uncertain about the true model of the economy. In this context, I believe that it would be of great relevance to assess the impact of monetary policy in different scenarios in which private agents might or might not share the same form of model uncertainty as policymakers.

The second direction in which I would like to extend the framework presented in this paper is in the analysis of the role of the natural rate of unemployment. As mentioned, the literature has suggested that the monetary authority's incorrect beliefs about the natural rate of unemployment is a possible explanation of the high inflation rates in the U.S. in the late 1970s and early 1980s. For this reason, I believe that it would be interesting to develop a framework in which the natural rate of unemployment is treated as an unobservable variable that needs to be estimated together with the parameters of the model. However, in a model uncertainty environment this assumption requires some additional considerations, as it implies that the estimated natural rate of unemployment could be different in each of the alternative models of the Phillips curve included in the model space.

Appendix 1

Data description and initial settings

Unemployment u_t is measured from the monthly Civilian Unemployment rate, averaged to obtain quarterly data. The baseline measure used to compute the inflation rate π_t is the quarterly Personal Consumption Expenditure (PCE) chain-type price index. The alternative measures employed in the paper are the GDP chain-type price index, the GDP implicit price deflator, and the quarterly CPI. The sample goes from 1948 : *I* to 2017 : *IV*; the data from 1960 : *I* to 2017 : *IV* is used in the empirical exercises, while the observations from 1948 : *I* to 1959 : *IV* are used as training sample to set the initial values in the updating recursions and the parameters in the prior distributions. All the data was obtained from the Federal Reserve Bank of St. Louis website (FRED).

The natural rate of unemployment is approximated as:

$$u_t^* = u_{t-1}^* + g(u_t - u_{t-1}^*)$$

where the gain parameter is set as: $g = 0.075$. The value of u_0^* was set as: $u_0^* = u_0$, where $t = 0$ is 1948 : *I*.

The *NK* model is estimated exactly as specified in (4). On the other hand, the number of lags used in the estimation of the *SS*, *ST*, and *LS* models is as in Cogley and Sargent (2005). In more detail, the *SS* model included 4 lags of inflation and 2 lags of unemployment in addition to the contemporaneous values of the variables. The *ST* model included 3 lags of inflation and 2 lags of the unemployment gap in addition to the contemporaneous values of the variables. The *LS* model included 2 lags of the unemployment gap in addition to the difference between inflation and its expected value.

The parameters for (3), (4), and (1) and (2) for both the Keynesian and the Classical identification assumptions, were estimated using Bayesian methods. For each model i , let σ_i^2 denote the variance of the Phillips curve residuals and θ_i the vector of coefficients of the model. The prior distribution

of the parameters is assumed to be a Normal-Inverse Gamma distribution:

$$\begin{aligned}
 p(\theta_i, \sigma_i^2) &= p(\theta_i | \sigma_i^2) p(\sigma_i^2) \\
 &= N\left(\theta_{i,0}, \sigma_i^2 P_{i,0}^{-1}\right) IG(s_{i,0}, v_{i,0})
 \end{aligned}
 \tag{14}$$

The parameters of this distribution were computed using data from the training sample. More specifically, for each model i , the vector $\theta_{i,0}$ was set as the point estimate of the coefficients obtained from an OLS regression. The initial value of the other variables in the prior distributions of the parameters was set as follows: $P_{i,0} = X'_{i,T_0} X_{i,T_0}$, where X_{i,T_0} is a matrix including the training sample observations of the right-hand variables for model i ; $s_{i,0}$ is the sum of squared residuals from the initial regression; $v_{i,0}$ is the difference between the number of observations and the number of estimated coefficients in the initial regression. After the initial beliefs on $\theta_{i,0}$, $P_{i,0}$, $s_{i,0}$ and $v_{i,0}$ are set, 1959 – IV becomes $t = 0$, and the model parameters and weights are updated starting from 1960 – I using the approach described below in Appendix 2.

For the LS and NK models, the initial settings require to approximate the evolution of $E_{t-1}(\pi_t)$ in the training sample. I followed Cogley and Sargent (2005) and I approximated $E_{t-1}(\pi_t)$ by the value x_t , which was obtained by exponentially smoothing the current inflation rate according to the formula:

$$x_t = x_{t-1} + 0.075(\pi_t - x_{t-1})$$

with $x_0 = \pi_0$.

Finally, in the policy exercises performed in section 4, the discount factor β was set so that the annual discount rate is 4%, and the weight on inflation was set to $\lambda = 16$, reflecting an equal weight with unemployment.

Appendix 2

Parameters updating

Let σ_i^2 denote the variance of the residuals and θ_i the vector of coefficients for each specification of the Phillips curve included in the model space, $i = \{SS - K, SS - C, ST - K, ST - C, LS, NK\}$.

Let Z^t summarize the joint history of both right-hand and left-hand variables in each of the models up to time t . At time t , the econometrician is assumed to have information on the variables up to time $t - 1$, while the data for time t is still unknown. The prior on the model's parameters at time t is $p(\theta_i, \sigma_i^2 | Z_i^{t-1})$, while the posterior after the data for time t is observed is $p(\theta_i, \sigma_i^2 | Z_i^t)$. As in (14), we have:

$$p(\theta_i, \sigma_i^2 | Z_i^{t-1}) = p(\theta_i | \sigma_i^2, Z_i^{t-1}) p(\sigma_i^2 | Z_i^{t-1})$$

$$p(\theta_i | \sigma_i^2, Z_i^{t-1}) = N(\theta_{i,t-1}, \sigma_i^2 P_{i,t-1}^{-1}) \quad (15)$$

$$p(\sigma_i^2 | Z_i^{t-1}) = IG(s_{i,t-1}, v_{i,t-1}) \quad (16)$$

The values $\theta_{i,t-1}$, $P_{i,t-1}$, $s_{i,t-1}$ and $v_{i,t-1}$ are estimated based on data through period $t - 1$. In each period, before the data is observed, the value of θ_i can be estimated by $\theta_{i,t-1}$, while the value of σ_i^2 can be estimated by $s_{i,t-1}/v_{i,t-1}$. When information about data at time t becomes available, the econometrician will update the parameters in (15) and (16) to obtain the posterior distribution $p(\theta_i, \sigma_i^2 | Z_i^t)$:

$$\begin{aligned} P_{i,t} &= P_{i,t-1} + X_{i,t} X_{i,t}' \\ \theta_{i,t} &= P_{i,t}^{-1} (P_{i,t-1} \theta_{i,t-1} + X_{i,t} Y_{i,t}) \\ s_{i,t} &= s_{i,t-1} + Y_{i,t}' Y_{i,t} + \theta_{i,t-1}' P_{i,t-1} \theta_{i,t-1} - \theta_{i,t}' P_{i,t} \theta_{i,t} \\ v_{i,t} &= v_{i,t-1} + 1 \end{aligned}$$

Here $X_{i,t}$ is the vector of right-hand variables and $Y_{i,t}$ is the left-hand variable for model i . The posterior for date t becomes the prior for date $t + 1$, and given the new data that becomes available at $t + 1$, the same procedure can be used to compute $P_{i,t+1}$, $\theta_{i,t+1}$, $s_{i,t+1}$ and $v_{i,t+1}$.

Model posteriors

The approach used to compute the normalized model posteriors is the same as in Cogley and Sargent (2005). Given the assumptions on the distribution of the parameters, the posterior (7) can be obtained analytically, and w_{it} can be written recursively as:

$$\log w_{i,t+1} = \log w_{i,t} + \log p(Y_{i,t+1} | X_{i,t+1}, \theta_i, \sigma_i^2) - \log \frac{p(\theta_i, \sigma_i^2 | Z_i^{t+1})}{p(\theta_i, \sigma_i^2 | Z_i^t)}$$

where $p(Y_{i,t+1} | X_{i,t+1}, \theta_i, \sigma_i^2)$ is the conditional likelihood for observation $t + 1$ and $p(\theta_i, \sigma_i^2 | Z_i^t)$

is the posterior density computed using observations up to time t , $Z_i^t = (Y_i^t, X_i^t)$. Analytical expressions for all these terms are available from Cogley and Sargent (2005). The recursion for w_{it} was started by setting $w_{i,0} = \alpha_{i,0}$, with the values of $\alpha_{i,0}$ reported in the main text: $\alpha_{SS-K,0} = \alpha_{SS-C,0} = 0.97/2$; $\alpha_{ST-K,0} = \alpha_{ST-C,0} = 0.01/2$, $\alpha_{LS,0} = 0.01$, and $\alpha_{NK,0} = 0.01$.

Appendix 3

Given the Phillips curve equations (1) – (4) and the policy equation (12), each of the models included in the model space can be written in state-space representation as:

$$S_{i,t} = A_i S_{i,t-1} + B_i x_{t|t-1} + C_i \eta_t^i \quad (17)$$

For the state vectors, $S_{SS-K,t} = S_{SS-C,t} = S_{SS,t}$ and $S_{ST-K,t} = S_{ST-C,t} = S_{ST,t}$. The state vectors are defined as:

$$\begin{aligned} S_{SS,t} &= [u_t \quad u_{t-1} \quad \pi_t \quad \pi_{t-1} \quad \pi_{t-2} \quad \pi_{t-3} \quad 1]' \\ S_{ST,t} &= [(u_t - u_t^*) \quad (u_{t-1} - u_{t-1}^*) \quad \pi_t \quad \pi_{t-1} \quad \pi_{t-2} \quad \pi_{t-3} \quad 1]' \\ S_{LS,t} &= [(u_t - u_t^*) \quad (u_{t-1} - u_{t-1}^*) \quad \pi_t \quad 1]' \\ S_{NK,t} &= [(u_t - u_t^*) \quad \pi_t \quad \pi_{t-1} \quad 1]' \end{aligned}$$

The matrices A_i and C_i , and the vectors B_i , contain the model parameters and can be written from the equations of each model.

The loss function (13) can be written using (17) as:

$$\mathcal{L}_i = E_t \sum_{j=0}^{\infty} \beta^j \left(S'_{i,t+j} M_i' Q M_i S_{i,t+j} + x'_{t+j|t-1} R x_{t+j|t-1} \right) \quad (18)$$

where R is the cost of using the policy instrument, which is set equal to 0.001.²² The matrix M_i selects u_t and π_t from the state vector and introduces the target values of the variables, while the

²²It is standard to include this variable, mostly for computational reasons. Its inclusion does not affect the results of the optimization problem.

matrix Q attaches the weight λ to inflation. More specifically:

$$\begin{aligned}
 M_{SS-K} = M_{SS-C} &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & -u^* \\ 0 & 0 & 1 & 0 & 0 & 0 & -\pi^* \end{bmatrix} \\
 M_{ST-K} = M_{ST-C} &= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & -\pi^* \end{bmatrix} \\
 M_{LS} &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & -\pi^* \end{bmatrix} \\
 M_{NK} &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & -\pi^* \end{bmatrix} \\
 Q &= \begin{bmatrix} 1 & 0 \\ 0 & \lambda \end{bmatrix}
 \end{aligned}$$

The optimal policy rule for each specification of the Phillips curve is computed by minimizing (18) with respect to $\{x_{t+j|t-1}\}_{j=0}^{\infty}$. This is a standard optimal control problem, and the solution to this problem can be obtained using well known techniques. Under standard conditions, the solution will take the form: $x_{t+j|t-1} = -F^i S_{i,t+j-1}$, where F^i is a vector of constant policy parameters. As the estimated values of the parameters of each model will change over time based on the updating procedure described in Appendix 2, the vector F^i will be computed again in each period using the new set of parameters for each model.

The robust policy under model uncertainty is obtained by implementing the same approach as in Cogley and Sargent (2005). The policymaker is assumed to consider a composite (“encompassing”) problem constructed from the models of Phillips curve included in the model space, each of them weighted based on its normalized posterior probability. The loss function for the composite problem is:

$$\mathcal{L}_E = E_t \sum_{j=0}^{\infty} \beta^j \left(S'_{E,t+j} Q_{E,t} S_{E,t+j} + x'_{t+j|t-1} R x_{t+j|t-1} \right) \quad (19)$$

The composite state space representation is

$$S_{E,t} = A_E S_{E,t-1} + B_E x_{t|t-1} + C_E \eta_{E,t} \quad (20)$$

with:

$$S_{E,t} = [S'_{SS-K,t} \ S'_{ST-K,t} \ S'_{SS-C,t} \ S'_{ST-C,t} \ S'_{LS,t}; S'_{NK,t}]'$$

and

$$\eta_{E,t} = [\eta_t^{SS-K} \ \eta_t^{ST-K} \ \eta_t^{SS-C} \ \eta_t^{ST-C} \ \eta_t^{LS}; \eta_t^{NK}]'$$

The matrix A_E is a block diagonal matrix with blocks composed by the individual matrices A_i , and similarly C_E is a block diagonal matrix composed of the individual matrices C_i . The vector B_E is defined as:

$$B_E = \begin{bmatrix} B_{SS-K} \\ B_{ST-K} \\ B_{SS-C} \\ B_{ST-C} \\ B_{LS} \\ B_{NK} \end{bmatrix} \quad (21)$$

Finally, the matrix $Q_{E,t}$ is a block diagonal matrix composed of the elements $\alpha_{i,t} M_i' Q M_i$, where M_i and Q are as previously defined, and $\alpha_{i,t}$ are the models' normalized posteriors computed from (8) and updated as described in Appendix 2. The posteriors $\alpha_{i,t}$ and the matrices A_E , B_E and C_E will change from period to period as the policymaker updates his estimates using the new data that becomes available.

References

- [1] Blanchard, O., 2016. The Phillips Curve: Back to the '60s? *American Economic Review*, vol. 106(5), pp. 31-34.
- [2] Blanchard, O., Cerutti, E., Summers, L., 2015. Inflation and Activity – Two Explorations and their Monetary Policy Implications. NBER Working Papers n. 21726.
- [3] Brock, W. A., Durlauf, S. N, West, K. D., 2003. Policy Evaluation in Uncertain Economic Environments. *Brookings Papers on Economic Activity*, pp.235-301.
- [4] Brock, W. A., Durlauf, S. N, West, K. D., 2007. Model Uncertainty and Policy Evaluation: Some Theory and Empirics. *Journal of Econometrics*, vol. 136, no. 2, February, pp. 629-64.
- [5] Cogley, T., Colacito, R., Sargent, T. J., 2007. Benefits from U.S. Monetary Policy Experimentation in the Days of Samuelson and Solow and Lucas. *Journal of Money, Credit, and Banking*, vol. 39, Supplement February, pp. 67-99.
- [6] Cogley, T.; De Paoli, B.; Matthes, C.; Nikolov, K.; Yates, T.; 2011. A Bayesian approach to optimal monetary policy with parameter and model uncertainty. *Journal of Economic Dynamics and Control*, vol. 35(12), pp. 2186-2212.
- [7] Cogley, T., Sargent, T.J., 2001. Evolving post world war II US inflation dynamics. In: *NBER Macroeconomics Annual*, vol. 16, pp. 331–373.
- [8] Cogley, T., Sargent, T.J., 2005. The Conquest of U.S. inflation: Learning and Robustness to Model Uncertainty. *Review of Economic Dynamics*, vol. 8, no. 2, April, pp. 262-302.
- [9] Coibion, O., Gorodnichenko, Y., 2015. Is the Phillips Curve Alive and Well after All? *Inflation Expectations and the Missing Disinflation*. *American Economic Journal: Macroeconomics*, vol. 7(1), pp. 197-232.
- [10] Conway, R. K.; Gill, G., 1991. Is the Phillips curve stable? A time-varying parameter approach. *Journal of Policy Modeling*, vol. 13(1), pp. 141-151.
- [11] DeLong, J.B., 1997. America's only peacetime inflation: the 1970s. In: Romer, C.D., Romer, D.H. (Eds.), *Reducing Inflation NBER Studies in Business Cycles*, vol. 30. Univ. of Chicago Press, Chicago.
- [12] Evans, G.W., Honkapohja, S., 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton.
- [13] Friedrich, C., 2014. Global Inflation Dynamics in the Post-Crisis Period: What Explains the Twin Puzzle? *Staff Working Papers 14-36*, Bank of Canada.
- [14] Gali, J., Gertler, M., 1999. Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, vol. 44(2), pp. 195-222.
- [15] Hansen, L., Sargent, T., 2008. *Robustness*. Princeton University Press, Princeton, New Jersey.
- [16] King, R.G., Stock, J.H., Watson, M.W., 1995. Temporal instability of the unemployment-inflation relationship. *Economic Perspectives*, Federal Reserve Bank of Chicago, May, pp. 2-12.

- [17] King, R.G., Watson, M.W., 1994. The post-war US Phillips curve: A revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy* 41, pp. 157–219.
- [18] Kreps, D., 1998. Anticipated Utility and Dynamic Choice, in D.P. Jacobs, E. Kalai, and M. Kamien, eds., *Frontiers of Research in Economic Theory*, (Cambridge: Cambridge University Press), pp. 242-74.
- [19] Lucas Jr., R.E., 1976. Econometric Policy Evaluation: A Critique. In Brunner, K.; Meltzer, A. *The Phillips Curve and Labor Markets*. *Carnegie-Rochester Conference Series on Public Policy*. 1. New York: American Elsevier. pp. 19–46.
- [20] Lucas Jr., R.E., 1972. Expectations and the neutrality of money. *Journal of Economic Theory* 4, pp. 103–124.
- [21] Matheson, T., Stavrev, E., 2013. The Great Recession and the inflation puzzle. *Economics Letters*, vol. 120(3), pp. 468-472.
- [22] Nason, J.M., Smith, G.W., 2008a. Identifying the new Keynesian Phillips curve. *Journal of Applied Econometrics*, vol. 23(5), pp. 525-551.
- [23] Nason, J.M., Smith, G.W., 2008b. The New Keynesian Phillips curve : lessons from single-equation econometric estimation. *Federal Reserve Bank of Richmond*, issue Fall, pp. 361-395.
- [24] Phillips, A.W.H., 1958. The Relation Between Unemployment and the Rate of Change of Money Wages in the United Kingdom, 1861-1957. *Economica*, ns. 25, pp. 283-299.
- [25] Primiceri, G.E., 2006. Why inflation rose and fell: Policymakers’ beliefs and US postwar stabilization policy. *The Quarterly Journal of Economics*, vol. 121, August, pp. 867-901.
- [26] Rudebusch, G., Svensson, L.E.O. 1999. Policy Rules for Inflation Targeting. *NBER Chapters*, in: *Monetary Policy Rules*, pp. 203-262.
- [27] Samuelson, P.A., Solow, R.M., 1960. Analytical aspects of anti-inflation policy. *American Economic Review*, *Papers and Proceedings* 50, pp. 177–184.
- [28] Sargent, T.J., 1973. Rational expectations, the real rate of interest, and the natural rate of unemployment. *Brookings Papers on Economic Activity*, pp. 429–472.
- [29] Sargent, T.J., 1976. The Observational Equivalence of Natural and Unnatural Rate Theories of Macroeconomics. *Journal of Political Economy*, vol. 84(3), pp. 631-640.
- [30] Sargent, T.J., 2001. *The Conquest of American Inflation*. Princeton Univ. Press, Princeton, NJ.
- [31] Sargent, T. J., Williams, N., 2005. Impacts of Priors on Convergence and Escapes from Nash Inflation. *Review of Economic Dynamics*, vol. 8, n. 2, pp. 360–91.
- [32] Sargent, T.J, Williams, N., Zha, T., 2006. Shocks and Government Beliefs: The Rise and Fall of American Inflation. *American Economic Review*, vol. 96(4), pp. 1193-1224.
- [33] Sims, C.A., 2001. Comment on Sargent and Cogley’s ‘Evolving post World War II US inflation dynamics’. In: *NBER Macroeconomics Annual*, vol. 16, pp. 373–379.
- [34] Solow, R.M., 1968. Recent controversy on the theory of inflation: an eclectic view. In: Rousseaus, S. (Ed.), *Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control*. New York University, New York.

- [35] Stock, J.H., Watson, M.W., 2010. Modeling Inflation After the Crisis. NBER Working Papers n. 16488.
- [36] Svensson, L.E.O., Williams, N.M. 2007. Bayesian and Adaptive Optimal Policy under Model Uncertainty. NBER Working Papers 13414.
- [37] Tobin, J., 1968. Discussion. In: Rouseaus, S. (Ed.), Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control. New York University, New York.

FIGURES

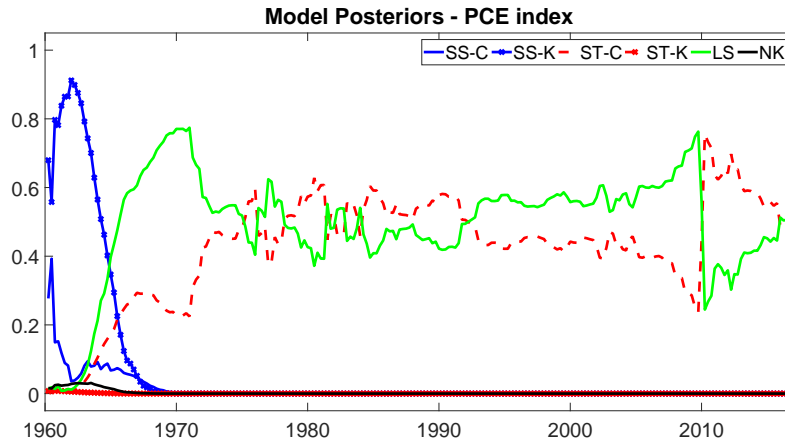


Figure 1 - Posterior probabilities over the period 1960 : *I* – 2017 : *IV* for the 6 model specifications of the Phillips curve included in the model space. The posterior probabilities are normalized so that they add up to one in each quarter.

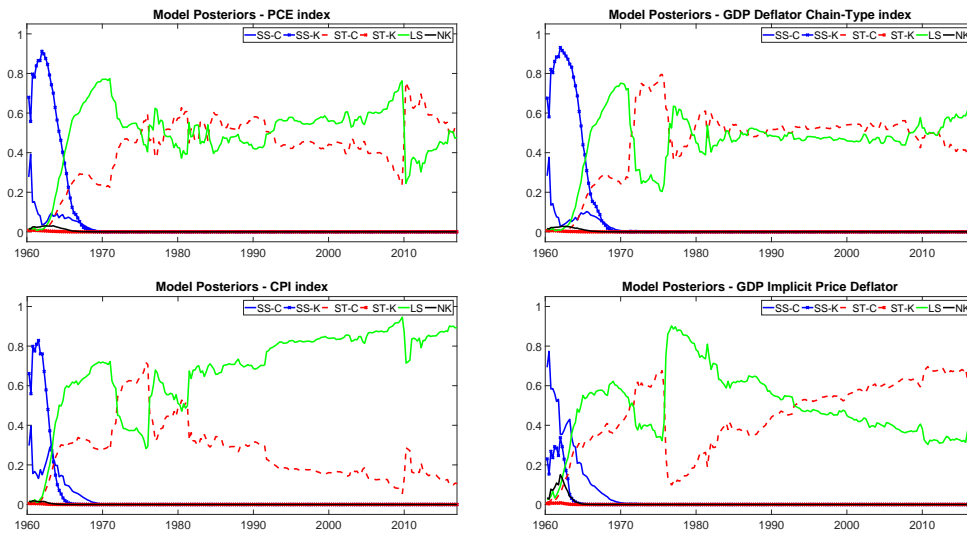


Figure 2 - Posterior probabilities over the period 1960 : *I* – 2017 : *IV* for the 5 specifications of the Phillips curve included in the model space. The posterior probabilities are normalized so that they add up to one in each quarter. The top-right panel is obtained using the GDP deflator chain-type index. The bottom-left panel is obtained using the CPI for all items. Finally, the bottom-right panel is obtained using the GDP implicit price deflator. For comparison, the top-left panel reports the posteriors using the PCE chain-type index, as in Figure 1.

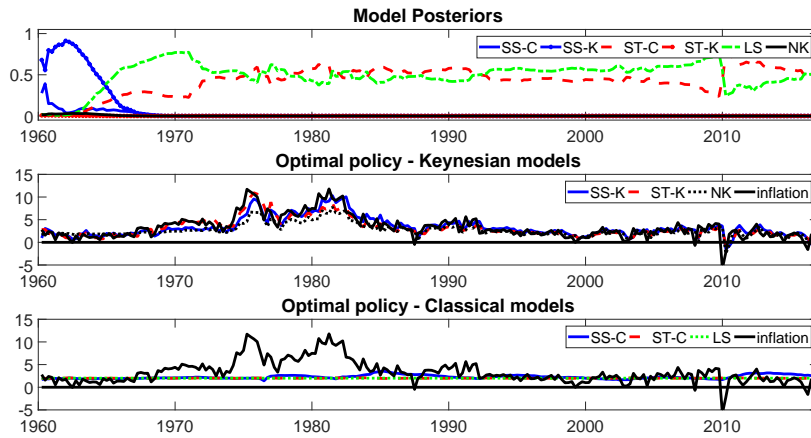


Figure 3 - The top panel reports the posterior probabilities already shown in Figure 1. The middle and bottom panels report the optimal policy recommendations for the Keynesian and Classical models respectively, obtained using the approach described in Appendix 3. The actual inflation rate during the same period is also shown for comparison.

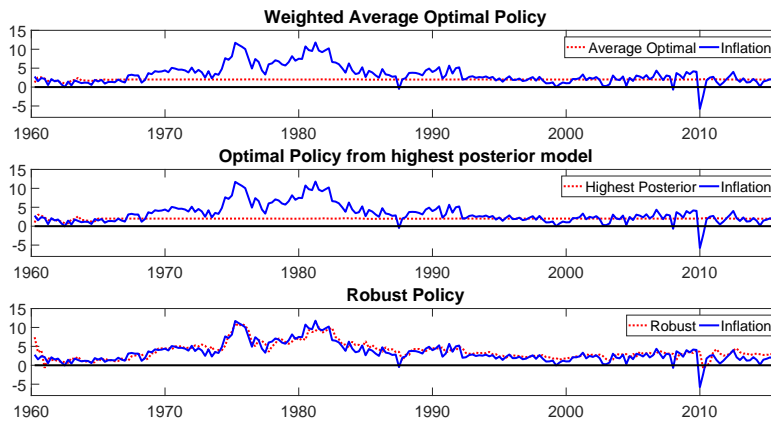


Figure 4 - The top panel shows the weighted average optimal policy computed as the sum of the optimal policies weighted by the models' posteriors. The middle panel reports the optimal policy for the model with the highest posterior in each quarter. The bottom panel reports the robust policy under model uncertainty obtained using the approach proposed by Cogley and Sargent (2005).

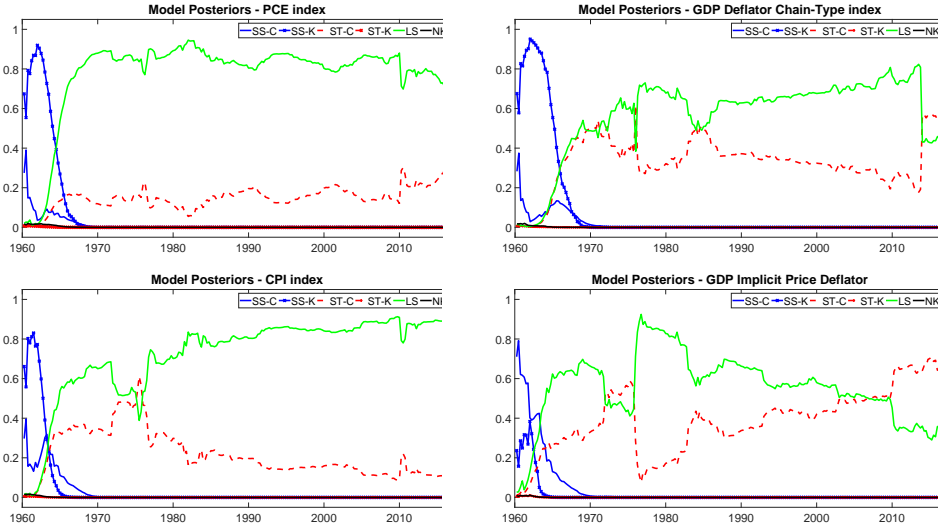


Figure 5 - Posterior probabilities over the period 1960 : I – 2017 : IV for the 5 model specifications of the Phillips curve included in the model space, under the alternative assumption $E_{t-1}(\pi_t) = x_{t|t-1}^E$. The posterior probabilities are normalized so that they add up to one in each quarter. The top-left panel is obtained using the PCE chain-type index. The top-right panel is obtained from the GDP deflator chain-type index. The bottom-left panel is obtained using the CPI. Finally, the bottom-right panel is obtained using the GDP implicit price deflator.