
KEVIN LEE
JAMES MORLEY
KALVINDER SHIELDS

The Meta Taylor Rule

We characterize U.S. monetary policy within a generalized Taylor rule framework that accommodates uncertainties about the duration of policy regimes and the specification of the rule, in addition to the standard parameter and stochastic uncertainties inherent in traditional Taylor rule analysis. Our approach involves estimation and inference based on Taylor rules obtained through standard linear regression methods, but combined using Bayesian model averaging techniques. Employing data that were available in real time, the estimated version of the “meta” Taylor rule provides a flexible but compelling characterization of monetary policy in the United States over the last 40 years.

JEL codes: C32, D84, E32

Keywords: Taylor rule, real-time policy, model uncertainty, U.S. interest rates.

DISCUSSIONS OF MONETARY POLICY in recent years invariably make reference to the “Taylor rule.” This is a simple rule suggested by Taylor (1993) in which the federal funds rate is set with reference to a linear function of U.S. inflation and the output gap and that accurately described Federal Reserve policy during the period considered by Taylor, that is, 1987q1–1992q3. Its simplicity has meant that it has since been widely used as a benchmark against which policy decisions have been judged and its properties as a rule for policymaking have been thoroughly investigated in the context of various macroeconomic models (notably in Woodford 2003).¹ There have also been numerous empirical exercises investigating

The authors are indebted to participants at seminars at the universities of Cambridge, Hitotsubashi, Montreal, Western Australia, and Wellington for helpful comments. Financial support from the ARC (Discovery Grant DP0988112) is gratefully acknowledged.

KEVIN LEE is at the University of Nottingham (E-mail: Kevin.lee@nottingham.ac.uk). JAMES MORLEY is at the University of New South Wales (E-mail: james.morley@unsw.edu.au). KALVINDER SHIELDS is at the University of Melbourne (E-mail: k.shields@unimelb.edu.au).

Received October 20, 2011; and accepted in revised form August 20, 2013.

1. The work of Woodford (2003) provides a justification for the use of the Taylor rule framework by relating the rule to the underlying payoffs (utility and welfare), choices and beliefs that might hold in the economy.

the extent to which Taylor's original finding that the rule describes Federal Reserve behavior extends to other data periods. Orphanides (2003) in particular considers whether the Taylor rule can be used as an organizing device with which to characterize U.S. monetary policy, concluding that policy since the early 1950s, and indeed before, can be readily interpreted within this framework.

Interest in estimating Taylor rules does not necessarily arise from a desire to expose an actual rule that was used in formulating policy. Rather the Taylor rule framework can be used to characterize past decision making and to impose a useful structure for drawing inferences about changes in the systematic reaction of monetary policy to economic conditions. But there remains considerable structural uncertainty even within a Taylor rule framework. Specifically, there is *specification uncertainty* relating to the precise form of the Taylor rule to be estimated. For example, the choice of model can vary according to the horizon over which policymakers consider their decisions since they might focus on recently experienced inflation and business cycle fluctuations or on expected future outcomes. The choice of model can also vary according to the degree of interest rate smoothing allowed, the chosen measure of inflation (including or excluding asset prices, say), and the chosen measure of the output gap.² There is also uncertainty about the stability of the policy responses to economic conditions, that is, *regime uncertainty*. There have been considerable changes over the decades in policymakers' understanding of the operation of the macroeconomy and in the perceived payoffs from policy interventions. This translates into changes in policy regime, sometimes occurring abruptly with the appointment of a new Federal Reserve chairman and sometimes involving an evolution of policy as priorities and beliefs change.³ This potential for structural instability generates uncertainty about the relevance of past interest rate decisions to current decisions and about the choice of the sampling window in empirical work. An analysis that accommodates and characterizes these two forms of structural uncertainty would extend traditional Taylor rule analysis, which is typically only concerned with inferences based on the estimated responsiveness of the interest rate to inflation and the gap within a particular model (i.e., relating to "parameter uncertainty") and on the fit of the model (i.e., relating to the residual "stochastic uncertainty").

This paper provides a characterization of U.S. monetary policy based on a novel and pragmatic modeling approach that accommodates specification and regime uncertainties as well as the parameter and stochastic uncertainties in traditional Taylor rule analysis. This approach involves estimation and inference based on a set of specific Taylor rules obtained through linear regression methods, but combined using

2. See Kozicki (1999) for discussion of various forms of specification uncertainty in monetary policy decisions, and Orphanides and van Norden (2003) and Garratt et al. (2008, 2009) for detailed discussion of appropriate measures of the output gap. As discussed below, when the data are published only with a delay or are subject to revision, these measurement problems are compounded by the need to consider the data that are available in real time.

3. See, for example, Cogley and Sargent's (2005) study of the extent to which postwar U.S. inflation can be explained by changes in policymakers' understanding of the processes determining inflation and the likely consequences of policy actions.

Bayesian model averaging techniques. The Taylor rule is a vehicle for characterizing past interest rate decisions and the weights employed in combining individual Taylor rules to obtain the “meta” rule are determined according to the ability of the individual rules to explain past interest rate movements.⁴ The weights can change over time so that the approach is very flexible, even compared to more computationally demanding time-varying parameter models of Taylor rules (e.g., Boivin 2005, Kim and Nelson 2006, Kim, Kishor, and Nelson 2006, McCulloch 2007, Alcidi, Flamini, and Fracasso 2011), and could be usefully applied to investigate many behavioral relations in economics. As we see, the estimated meta Taylor rule is able to capture many of the changes in the reaction of monetary policy to economic conditions over the last 40 years, while still providing a compelling characterization of monetary policy decision making in a single coherent and simple modeling framework. Section 1 of the paper formalizes the modeling approach taken to estimate the meta Taylor rule and relates the approach to the relevant model selection and Bayesian literature. Section 2 presents the results of the estimation of the U.S. meta Taylor rule over the period 1972q1–2008q4, highlighting phases of policy in which anti-inflationary policy was pursued more or less aggressively, when fear of recession or overheating dominated decisions, and when policy was more or less interventionist. Section 3 concludes.

1. STRUCTURAL UNCERTAINTY AND THE TAYLOR RULE

1.1 Taylor Rule Specifications

The rule reported in Taylor (1993) relates the federal funds rate in time t , r_t , to the rate of change of the implicit output deflator over the previous four quarters, π_t , and a measure of the output gap at t , x_t , as follows:

$$\begin{aligned} r_t &= \gamma_0 + \gamma_\pi \pi_t + \gamma_x {}_{T+1}x_t + \varepsilon_t, \\ &: \quad t = 1987q1, \dots, T, \quad \text{and } T = 1992q3. \end{aligned} \tag{1}$$

In Taylor’s paper, the parameters of interest were taken to be $\gamma_0 = 1$, $\gamma_\pi = 1.5$, $\gamma_x = 0.5$ but many subsequent papers have estimated rules of the form (1) where ε_t represents the deviations from the rule characterizing policy in a given quarter assumed to have (unconditional) mean zero and variance σ^2 . Note that here, ${}_{T+1}x_t$ is a measure of the output gap in time t as made available at time $T + 1$ and introduces the distinction between the measure of the gap that was available to Taylor when he undertook the analysis in 1992q4, that is, ${}_{T+1}x_t$, and the measure that would have been

4. Typically, “meta” analysis averages estimates of interest across different studies. However, the settings are too disparate in past studies of U.S. monetary policy to make this practicable. Thus, we recreate “meta” analysis by averaging estimates across a range of model specifications and sample periods considered in past studies, but otherwise controlling the setting in terms of variables, data, and general structure.

available in real time, that is, ${}_t x_t$. Orphanides (2001) provides a detailed analysis of the Taylor rule when estimated over the sample originally used by Taylor but paying attention to this distinction between the real-time and end-of-sample measures of the gap. He demonstrates that the performance of the Taylor rule in capturing interest rate movements is considerably reduced when the real-time measures are used and urges policymakers to take this informational issue into account when using simple rules in decision making. We find this to be a persuasive argument and make the real-time dimension of the analysis explicit in all that follows.

Orphanides' (2001) analysis also raises the issue of whether, in practice, policymakers are as myopic as is suggested by the rule of (1) where only contemporaneous measures of inflation and the gap are assumed to influence interest rate decisions. He considers the possibility that policymakers are forward looking and estimates alternative Taylor rules using direct measures of expected future inflation and the expected future gap, ${}_t \pi_{t+i}$ and ${}_t x_{t+i}$, $i = 1, \dots, 4$, in place of the contemporaneous values. The direct measures are the expected inflation data and the Federal Reserve staff estimates of the gap based on potential output as published in the Greenbook. Again focusing on Taylor's original sample (and also using the slightly extended sample 1987q1–1993q4 used in Taylor 1994), Orphanides shows that very different parameter estimates are obtained in the alternative rules based on these different policy horizons, establishing that the uncertainty surrounding the policy horizon is empirically important.

Analysis of monetary policy should accommodate the model uncertainty surrounding the policy horizon but it also should address the possibility of changes in policy regime if the analysis is to span a reasonably long data period. There has been considerable debate on the different approaches to monetary policy formulation taken by successive Federal Reserve chairmen over the years (see, e.g., Romer and Romer 2004) and the extent to which these pursued more or less active counter-inflationary policies. To the extent that there has been variation in policy approaches, it should be reflected by different values for the γ_0 , γ_π , γ_x , and σ parameters in Taylor rule models estimated at different times. However, unless there is a clear-cut break in regime, at precisely the time of a change of the Fed chairmanship, for example, there will be uncertainty on the sample periods relevant for estimating the different rules that describe policymaking over a protracted period. Indeed, the regime uncertainty surrounding the choice of sample might interact with the specification uncertainty surrounding the choice of policy horizon if, for example, new regimes behave cautiously at first, focusing on contemporaneous or short-horizon outcomes, and become more forward looking over time if the policy is seen to succeed and credibility is established.⁵

The structural uncertainties discussed above can be accommodated within a set of Taylor rule models M_{ijT} , each distinguished according to the policy horizon, i , and the sample period for which the model is relevant ($T - j, \dots, T$). Taking a

5. See Pesaran and Timmerman (1995) and Elliott and Timmerman (2008) for discussion of uncertainty over model instability in the context of forecasting.

real-time perspective, T here denotes the final observation in the sample available at the time an interest rate decision is made. The set of models characterizing interest rate determination over the period T_1, \dots, T_n is given by

$$\begin{aligned}
 M_{ijT} : \quad & r_t = \rho_{ijT} r_{t-1} + (1 - \rho_{ijT}) (\gamma_{0ijT} + \gamma_{\pi ijT} \pi_{t+i} + \gamma_{xijT} x_{t+i}) + \varepsilon_{ijT,t}, \\
 & : \text{ where } i = -1, \dots, 4, \quad j = j_{\min}, \dots, j_{\max}, \\
 & : \quad \quad \quad t = T - j, \dots, T, \text{ and } T = T_1, \dots, T_n,
 \end{aligned} \tag{2}$$

and $\varepsilon_{ijT,t}$ are i.i.d. innovations with mean zero and standard deviation σ_{ijT} . All of the models take the Taylor rule form of (1) extended to allow for interest rate smoothing. In any model, the policy horizon considered by the decision maker is assumed to look back one quarter or to look forward for up to 1 year ($i = -1, \dots, 4$). The models are also distinguished by the time span over which a rule is assumed to have operated, considered here to be in operation for j periods ending in period T ; that is, j is a duration measure describing the age of the regime. Of course, when there is a break, the regime period starts afresh so that $j_{\min} = 1$ although, in practice, the choice of minimum regime length will be driven by the need to have enough observations for estimation purposes (so that we might choose $j_{\min} = 16$, say). The maximum period for the survival of an unchanged policy stance is, in principle, unlimited. In practice in the United States, though, there have been six Federal Reserve chairmen since the mid-1960s so that, in the absence of any other information, one might anticipate that there would be breaks every 6 or 7 years and that a given policy rule would not last longer than 10 years, that is, $j_{\max} = 40$, say.

1.2 The Meta Taylor Rule and Model Averaging

The considerable structural uncertainty surrounding interest rate determination is reflected by the idea that the interest rate observed at a particular moment T could be explained by any of $6 \times 25 = 150$ different models according to (2) if we set $j_{\min} = 16$ and $j_{\max} = 40$.⁶ The meta Taylor rule proposed here accommodates this uncertainty by using a weighted average of the models in (2). Model averaging is now in widespread use in forecasting but is much less widely employed in structural modeling even though the statistical arguments to support the approach are equally valid in inference and prediction.

The foundation of the approach is the Bayesian modeling average (BMA) formula (see Draper 1995, Hoeting et al. 1999):

$$\Pr(\theta_T | \mathbf{Z}_T) = \sum_{i=-1}^4 \sum_{j=16}^{40} \Pr(\theta_T | M_{ijT}, \mathbf{Z}_T) \times \Pr(M_{ijT} | \mathbf{Z}_T), \tag{3}$$

6. Of course, estimation of these separate models would also expose the parameter uncertainty surrounding the estimates of ρ_{ijT} , γ_{0ijT} , $\gamma_{\pi ijT}$, and γ_{xijT} , and the stochastic uncertainty surrounding the estimated $\varepsilon_{ijT,t}$ as considered in standard Taylor rule estimation exercises.

where θ_T represents the unknown responsiveness of interest rates in time T to inflation and the output gap, reflected by the parameters $(\rho, \gamma_0, \gamma_\pi, \gamma_x)$ in the Taylor rule form; $\mathbf{Z}_T = (\mathbf{z}_1, \dots, \mathbf{z}_T)$ represents the data available at T with $\mathbf{z}_t = (r_t, {}_t\pi_{t+i}, {}_t x_{t+i}, \forall i)$; and $\Pr(\theta_T | \mathbf{Z}_T)$ is the inferential distribution that describes our understanding of the parameters of interest. The BMA formula deals with the structural uncertainty accommodated within $\Pr(\theta_T | \mathbf{Z}_T)$ by decomposing it into a weighted average of the conditional distributions (i.e., conditional on a specific model), $\Pr(\theta_T | M_{ijT}, \mathbf{Z}_T)$, using as weights the posterior model probabilities $\Pr(M_{ijT} | \mathbf{Z}_T)$.

The conditional distributions. A typical Taylor rule analysis considers the first element on the right-hand side of (3) only, working with a specific model (say M^*) and making inference that takes into account the stochastic and parameter uncertainties surrounding this specific model, noting that

$$\Pr(\theta_T | M^*, \mathbf{Z}_T) = \int \Pr(\theta_T | M^*, \theta, \mathbf{Z}_T) \Pr(\theta | M^*, \mathbf{Z}_T) d\theta.$$

As $\Pr(\theta | M^*, \mathbf{Z}_T) \propto \Pr(\mathbf{Z}_T | M^*, \theta) / \Pr(\theta | M^*)$, a strict Bayesian approach to evaluating this distribution requires a prior position to be taken on the the likely value of the parameters in the specified model. Alternatively, the conditional inferential distribution can be approximated using the maximum likelihood estimator of the parameters in M^* and its associated density. Specifically, in the case of a standard linear regression model, we have $(\widehat{\theta}_T^* - \theta_T | M^*, \mathbf{Z}_T) \sim N(\mathbf{0}, \widehat{\mathbf{V}}_T^*)$ where $\widehat{\theta}_T^*$ and $\widehat{\mathbf{V}}_T^*$ denote the ML estimator and its estimation variance, respectively. Although θ_T is taken as fixed at the estimation stage, it can be viewed as a random variable at the inference stage, so that $\Pr(\theta_T | M^*, \mathbf{Z}_T)$ is approximated by $N(\widehat{\theta}_T^*, \widehat{\mathbf{V}}_T^*)$ and standard inference carried out. Moreover, this simplification can be made for any model so that we can look at all 150 of our models of interest and base $\Pr(\theta_T | M_{ijT}, \mathbf{Z}_T)$ on the models' maximum likelihood estimates.

The model weights. The meta Taylor rule accommodates specification and instability uncertainty in (3) through the use of the model weights. Further application of the Bayes rule provides

$$\Pr(M_{ijT} | \mathbf{Z}_T) = \frac{\Pr(M_{ijT}) \Pr(\mathbf{Z}_T | M_{ijT})}{\sum_p \sum_q \Pr(M_{pqT}) \Pr(\mathbf{Z}_T | M_{pqT})},$$

where $\Pr(M_{ijT})$ is the prior probability of model M_{ijT} and $\Pr(\mathbf{Z}_T | M_{ijT})$ is the integrated likelihood, which can itself be decomposed into elements involving the prior probability on θ_{ijT} given the model M_{ijT} . Again, a strict Bayesian approach to estimation of the meta rule therefore involves the specification of meaningful prior probabilities on the models of interest and on the associated parameters. Alternatively, in the forecasting literature, simple averaging (using equal weights) or weights based on the models' likelihoods or their information criteria have been proposed as a way of accommodating the structural uncertainty and have been shown to improve forecasting performance considerably (see, e.g., Garratt et al. 2003, Elliott and Timmermann 2008, for discussion).

Given our modeling context, where there is uncertainty on the period over which any model is relevant due to the possibility of regime change, another sensible set of weights might be chosen by allowing these to evolve over time, recursively updating a model's weight to reflect the extent to which it remains useful for sample periods up to $T = T_1, \dots, T_n$. Specifically, we can write

$$\begin{aligned} \Pr(M_{ijT} | \mathbf{Z}_T) &= \Pr(M_{ijT} | \mathbf{Z}_{T-1}, \mathbf{z}_T) \propto \Pr(\mathbf{z}_T | M_{ijT}, \mathbf{Z}_{T-1}) * \Pr(M_{ijT} | \mathbf{Z}_{T-1}) \\ &= \Pr(\mathbf{z}_T | M_{ijT}, \mathbf{Z}_{T-1}) * \sum_{k=-1}^4 \sum_{l=16}^{40} \Pr(M_{ijT} | M_{klT-1}, \mathbf{Z}_{T-1}) \\ &\quad * \Pr(M_{klT-1} | \mathbf{Z}_{T-1}) \end{aligned} \quad (4)$$

so that a model's weight in period T depends on two things: the probability or density of the final observation in the sample conditional on the model, and the likelihood of the model given data up to $T - 1$. For the first of these, we note that under standard normality assumptions on the residuals for a model, the log density of the final observation is proportional to the value of the squared residual in period T . For the second element, we note that the likelihood of the model given data up to $T - 1$ depends, in turn, on the weights for all models in $T - 1$ and on the transition probabilities $\Pr(M_{ijT} | M_{klT-1}, \mathbf{Z}_{T-1})$.⁷ A simple structure for the transition probabilities is that, irrespective of the prevailing policy regime in the previous period, there is a constant probability of a break in regime in the next period, λ , and that, if there is a break, the new regime uses a Taylor rule with policy horizon i with equal probability for each of the possible policy horizons. That is,

$$\Pr(M_{i,j,T} | M_{klT-1}, \mathbf{Z}_{T-1}) = \begin{cases} 1 - \lambda & \text{if } i = k, j = l + 1 & \text{i.e., if no break} \\ \lambda/6 & \text{if } i = -1, \dots, 4, j = 16, & \text{i.e., if break to policy horizon } i \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

Taken together then, (4) and (5) provide a straightforward means for producing a set of weights for each $T = T_1, \dots, T_n$.⁸ This weighting scheme allows new regimes to be "born" in each period and otherwise updates the weights on different models recursively from one period to the next to reflect the likelihood that the models remain relevant for the updated sample.

7. It is worth noting that, given the same number of parameters for each model and the same initial model weights prior to the first sample ending in T_1 , a weighting scheme based only on squared residuals would be equivalent to using the SIC approximation of log marginal likelihoods for BMA under the assumption of a fixed policy regime (since a fixed policy regime implies transition probabilities equal to one for a given model and the same sample periods for all models in each T). The factorization of the model weights into the different elements in (4) addresses the different relevant sample periods based on j and different transition probabilities for each model given the possibility of policy regime changes.

8. We also need initial weights for the models prior to the first sample period that ends in T_1 . For simplicity, we assume these initial weights are equal, as is standard in BMA.

The use of sensibly chosen model weights for $\Pr(M_{ijt}|\mathbf{Z}_T)$ in (3), along with the use of ML estimation of individual rules, represents a pragmatic approach to accommodating structural uncertainty in discussing inference in an estimated Taylor rule, which could be applied more widely in modeling behavioral relationships in economics. Estimation of the individual Taylor rules is based on standard linear regressions of the form found throughout the literature. But the combination of these into a meta model accommodates specification uncertainty and can capture the effects of complicated structural change. The formula in (4) constrains the weights to evolve over time according to the models' historical fit. This corresponds with the idea discussed earlier, and by Orphanides, that the Taylor rule provides a framework for characterizing decision making according to its ability to capture past policy outcomes. Moreover, the evolution of the weights itself provides useful information with which to interpret the changing policy regime. Also, the weighting scheme allows for considerable flexibility in the ways in which the sensitivity of interest rates to inflation and the gap can develop; for example, as we shall see in the empirical section below, the evolving weights can accommodate periods in which the responsiveness of policy changes slowly over time and periods when policy changes abruptly. The approach is more flexible than a standard time-varying parameter (TVP) model, for example, in which the form of the instability is defined at the outset, while estimation of a more elaborate TVP model that allows for more complex forms of instability would be computationally more demanding than the meta approach proposed here.⁹ The "meta" approach provides a pragmatic, easy-to-implement, and easy-to-interpret means of accommodating structural uncertainty therefore. The approach also clearly addresses some of the dangers implicit in many model selection algorithms that use the data \mathbf{Z}_t to identify a single preferred choice of M^* and then proceed to make inferences as if M^* was known to be correct.¹⁰

2. THE U.S. META TAYLOR RULE

In this section, we describe the meta Taylor rule, obtained as a weighted average of the various models described in (2), estimated using U.S. data for the period 1969q1–2008q4.¹¹ Our primary data set consists of the federal funds rate plus real-time data on ${}_t\pi_{t+i}$ and ${}_ty_{t+i}$, $i = -1, \dots, 4$. These represent the first-release measures of inflation and output (released with a one-period delay) when $i = -1$, the nowcast of

9. The use of $N(\hat{\theta}_T^*, \hat{\mathbf{V}}_T^*)$ as an approximation for $\Pr(\theta_T|M^*, \mathbf{Z}_T)$ is akin to a Bayesian approach with noninformative priors for θ_T . The model averaging allows for a specification of diffuse priors over different types of models and parameters. A Bayesian TVP model with comparable flexibility would require more informative/restrictive priors and would be computationally much more complicated to estimate.

10. See Draper (1995) for further discussion.

11. In unreported analysis, we have extended the sample period to include the period after 2008q4 where interest rates have been close to their zero lower bound. Not surprisingly given the lack of variation in the dependent variable, the estimated responses to inflation and the output gap do not change much after 2008q4, but the confidence bands on these estimates get wider. These results are available from the authors upon request.

current inflation and output as provided by the Survey of Professional Forecasters (SPF) when $i = 0$, and their one-, two-, three-, and four-quarter-ahead forecasts when $i = 1, \dots, 4$. In some of our analysis below, we also make use of the Federal Reserve staff estimates of the gap between actual and potential output as published in the Greenbook and the GDP gap constructed by the Congressional Budget Office (CBO). For each model, we consider OLS estimation because the right-hand-side variables were available to policymakers when setting the federal funds rate and, therefore, predetermined.

Our primary aim is to accommodate in our “meta” Taylor rule the uncertainty arising from the choice of policy horizon and the uncertainty arising from changing policy regimes. For most of the analysis, therefore, we abstract from the uncertainties arising from the choice of inflation and output gap measures by using simple and readily available measures of these key variables.¹² Specifically, we use the GDP deflator for inflation and we use a measure of the gap based on Taylor’s original exercise but constructed using real-time data only. Taylor’s gap measure was calculated as the difference between (log) output and a linear trend running through the observations of (log) output series between 1984q1 and 1992q3, where the output data used was the 1992q4 vintage. Clearly this measure would not have been available within the sample period and so could not have served as a basis of policy decisions in real time.¹³

Our measure of the gap in each period uses only data available in real time, considering the historical output series available up to one quarter earlier (since there is a one-quarter delay in publication of output data) plus the output data available from the SPF giving direct measures of expected output contemporaneously and up to four quarters ahead. This allows the construction of a gap measure for our different policy horizons $i = -1, \dots, 4$. The real-time output trend values are constructed using a linear trend through a rolling sample of 40 quarters of the real-time data series (including the 35 historical and 5 expected observations) and the gap is measured as the difference between the expected contemporaneous output value and the value of this trend. The idea is to consider a gap measure that is as close as possible to that originally considered by Taylor to focus attention on the uncertainties surrounding the policy horizon and regime. Later, however, we do use the Federal Reserve’s gap measures in an extension to the main analysis to gauge the impact of accommodating this further element of uncertainty on the estimation of the Taylor rule.

12. See Garratt, Lee, and Shields (2014) for a discussion of a measure of the natural output gap that is calculable in real time and that has an explicit economic motivation. See also Garratt et al. (2008, 2009) for a more comprehensive discussion of the characterization of the output gap when there is uncertainty on how the concept is best measured.

13. Orphanides (2001) shows that the gap measure used by Taylor is, by coincidence, relatively close to measures that were produced by Federal Reserve staff in real time over this particular time frame so that the original Taylor characterization is robust to the real-time measurement issues for his particular sample of data.




Fig. 1. Contemporaneous Real Time (solid line) and Greenbook-CBO (dashed line) Gap Measures with NBER Peak-Trough Dates (shaded areas).

Figure 1 plots our real-time measure of the output gap from 1969q1 to 2008q4 and the Greenbook/CBO gap that we use in our later analysis for its shorter available sample period of 1987q3–2008q4. There is a close correspondence of the two output gap measures and of our measure with the National Bureau of Economic Research (NBER) peak/trough dates, also displayed in the figure.

2.1 Taylor Rules for the Taylor Sample, 1987q1–1993q4

Tables 1 and 2 describe a series of Taylor rules estimated over the period 1987q1–1993q4. These illustrate some of the empirical issues involved in estimating Taylor rules and provide a point of contact with some estimated rules in the previous literature. The tables correspond to tables 5 and 6 in Orphanides (2001) that also consider this (extended Taylor 1994) sample period using the 1994q4 vintage data and using real-time data. The difference between Tables 1 and 2 here and Orphanides' tables is in the measure of the gap; Orphanides uses a measure based on the Federal Reserve staff's estimates of potential output as reported in the Greenbook while we use the measure based on the linear trend described above. The results of Tables 1 and 2 show that the choice of gap measure is not the dominating feature of this analysis since the results are qualitatively similar to those of Orphanides.

TABLE 1
TAYLOR RULES ESTIMATED WITH 1994q4 DATA: 1987q1–1993q4

	Horizon relative to decision period (in quarters)					
	-1	0	1	2	3	4
$(1 - \rho)\gamma_0$	1.507 (0.446)	0.288 (0.326)	0.061 (0.547)	-0.881 (0.545)	-2.257 (0.629)	-3.276 (0.574)
$(1 - \rho)\gamma_\pi$	-0.306 (0.281)	0.613 (0.230)	0.351 (0.224)	0.733 (0.220)	1.241 (0.256)	1.530 (0.222)
$(1 - \rho)\gamma_x$	0.332 (0.103)	0.395 (0.055)	0.208 (0.073)	0.177 (0.063)	0.155 (0.056)	0.092 (0.048)
ρ	0.894 (0.144)	0.575 (0.101)	0.781 (0.088)	0.715 (0.073)	0.631 (0.071)	0.619 (0.056)
γ_π	-2.882 (3.673)	1.442 (0.261)	1.600 (0.723)	2.570 (0.545)	3.366 (0.430)	4.018 (0.385)
γ_x	3.126 (3.478)	0.930 (0.177)	0.949 (0.358)	0.621 (0.218)	0.419 (0.150)	0.243 (0.124)
γ_π^{MR}	-1.570 (2.036)	1.412 (0.279)	1.318 (0.660)	2.325 (0.491)	3.231 (0.404)	3.878 (0.351)
γ_x^{MR}	1.702 (0.400)	0.911 (0.154)	0.782 (0.241)	0.562 (0.187)	0.402 (0.141)	0.234 (0.119)
\bar{R}^2	0.971	0.982	0.961	0.971	0.978	0.985
SEE	0.384	0.30	0.447	0.384	0.336	0.281
LL	-10.31	-3.84	-14.43	-10.33	-6.69	-1.84
$SC(1)$	3.11 [0.92]	4.18 [0.96]	6.85 [0.99]	3.153 [0.92]	0.00 [0.04]	0.23 [0.37]

NOTES: The table presents least squares estimates of the regression equations

$$M_t : r_t = \rho r_{t-1} + (1 - \rho)(\gamma_0 + \gamma_\pi \pi_{t+i} + \gamma_x x_{t+i}) + \eta_t, \quad \text{for } t = 1987q1, \dots, 1993q4$$

based on information available in 1994q4 and using a linear trend to obtain a measure of the output gap. The columns correspond to different values for i . For forward-looking variants of the Taylor rule, survey forecasts from the Survey of Professional Forecasters are used and the output gap is obtained as the difference between the relevant forward-looking forecast series and the corresponding forecast values of the linear trend. \bar{R}^2 is the squared multiple correlation coefficient, SEE the standard error of the regression. LL is the log likelihood value and $SE(1)$ gives an LM test statistic of residual first-order serial correlation. Standard errors are given in parentheses and p -values are given in brackets.

The estimated Taylor rule obtained using 1994q4 data with no interest rate smoothing and contemporaneous inflation and gap measures used as regressors takes the form¹⁴

$$r_t = -0.091 + 1.765\pi_t + 0.583_{T+4}x_t + \hat{\varepsilon}_t, \tag{6}$$

(0.541) (0.158) (0.078)

: $t = 1987q1 - T, \quad T = 1993q4,$

$$\bar{R}^2 = 0.947, \quad SEE = 0.535, \quad LL = -20.613, \quad SC(1) = 10.608$$

matching closely the Taylor rule of (1). Table 1 shows the corresponding partial adjustment Taylor rules estimated for policy horizons ranging from $i = -1$ (backward

14. See the footnote to Table 1 for an explanation of diagnostic statistics.

TABLE 2
TAYLOR RULES ESTIMATED WITH REAL-TIME DATA: 1987q1–1993q4

	Horizon relative to decision period (in quarters)					
	-1	0	1	2	3	4
$(1 - \rho) \gamma_0$	(0.123) (0.317)	-0.377 (0.304)	-0.832 (0.280)	-1.367 (0.315)	-2.390 (0.541)	-3.360 (0.584)
$(1 - \rho) \gamma_\pi$	0.239 (0.209)	0.465 (0.183)	0.541 (0.142)	0.665 (0.150)	1.004 (0.251)	1.386 (0.272)
$(1 - \rho) \gamma_x$	0.172 (0.037)	0.179 (0.029)	0.157 (0.023)	0.132 (0.022)	0.097 (0.027)	0.048 (0.031)
ρ	0.860 (0.079)	0.824 (0.066)	0.835 (0.050)	0.833 (0.048)	0.781 (0.067)	0.707 (0.070)
γ_π	1.708 (0.754)	2.638 (0.481)	3.274 (0.532)	3.975 (0.614)	4.574 (0.594)	4.738 (0.435)
γ_x	1.227 (0.526)	1.018 (0.324)	0.952 (0.283)	0.790 (0.263)	0.440 (0.214)	0.165 (0.134)
γ_π^{MR}	1.113 (0.758)	1.958 (0.485)	2.349 (0.386)	2.872 (0.386)	3.766 (0.432)	4.318 (0.314)
γ_x^{MR}	0.799 (0.123)	0.756 (0.123)	0.683 (0.117)	0.571 (0.122)	0.363 (0.139)	0.150 (0.114)
\bar{R}^2	0.975	0.979	0.983	0.984	0.981	0.984
SEE	0.355	0.326	0.295	0.285	0.308	0.282
LL	-8.54	-6.17	-3.41	-2.45	-4.563	-2.15
$SC(1)$	1.48 [0.78]	0.53 [0.53]	0.34 [0.44]	0.024 [0.12]	0.00 [0.05]	0.08 [0.22]

NOTES: The table presents least squares estimates of the regression equations

$$M_t : r_t = \rho r_{t-1} + (1 - \rho)(\gamma_0 + \gamma_\pi r_{\pi t+i} + \gamma_x r_{x t+i}) + \eta_t, \quad \text{for } t = 1987q1, \dots, 1993q4$$

where the output gap is obtained by detrending a rolling sample of 40 quarters of data using the historical time series available during that quarter. The columns correspond to different values for i . For forward-looking variants of the Taylor rule, survey forecasts from the Survey of Professional Forecasters are used. \bar{R}^2 is the squared multiple correlation coefficient, SEE the standard error of the regression. LL is the log likelihood value and $SE(1)$ gives an LM test statistic of residual first-order serial correlation. Standard errors are given in parentheses and p -values are given in brackets.

looking) to $i = 4$ (four quarters ahead) all based on the 1994q4 vintage of data. The column headed $i = 0$ provides a straight point of comparison with the model in (6). This demonstrates the empirical importance of including the lagged dependent variable to deal with residual serial correlation in (6) and to distinguish between the impact responses of interest rates and the long-run responses (with the impact effect $(1 - \rho)\gamma_\pi = 0.613$ and the long-run effect $\gamma_\pi = 1.442$ for inflation, for example). The other columns of Table 1 show the sensitivity of the results to the inclusion of inflation and gap measures at the different policy horizons $i = -1, \dots, 4$. As in Orphanides (2001), the estimated coefficient on inflation gets larger and the estimated coefficient on the gap falls as longer policy horizons are considered. For example, the estimated long-run coefficient on inflation is actually negative for $i = -1$, although statistically insignificantly different from zero, but rises from a statistically significant value of 4.018 for $i = 4$. It is also worth noting at this point the sensitivity of the estimated long-run responses to the estimated value of ρ : the precision of the estimated long-run response declines rapidly as ρ approaches unity so that, for example, the standard errors of the estimated long run inflation and output gap responses are very high in

column (1), where $\rho = 0.894$, compared to the remaining columns where ρ takes values of 0.8 or below. For this reason, in much of what follows, we report also the “medium-run” inflation and output gap response observed over a six-quarter period, denoted γ_π^{MR} and γ_x^{MR} , which provides a more precisely estimated indication of the interest rate response over the medium term even in models with very high degrees of interest rate smoothing.¹⁵

Table 2 shows the corresponding results where the real-time output gap measure is employed. The extent of interest rate smoothing is typically estimated to be larger here than in Table 1 and the inflation and gap coefficients are typically smaller. All of the coefficients are more precisely estimated in Table 2 compared to Table 1; the fit of the equations, reflected by the standard errors and R^2 statistics, are generally improved; and the problems of residual serial correlation observed in Table 1 are resolved in Table 2. The results obtained using real-time data are more satisfactory in a statistical sense than but, more importantly, they are quite different from those obtained using the end-of-sample data in some columns, confirming Orphanides’ (2001) point on the importance of using real-time data in the study of Taylor rules.

Focusing on the results in Table 2, we note that there is more consistency in parameter estimates across the policy horizons than in Table 1, particularly for the long-run coefficients. There remain some considerable differences in the short-run coefficients and implicit dynamics though, illustrating the specification uncertainty discussed in the previous section. However, it is straightforward to provide a meta Taylor rule for interest rate determination during this period by averaging across the separate models of Table 2. Specifically, a reasonable set of weights for the six models, denoted M_i , $i = -1, \dots, 4$, to reflect the dependence on the policy horizon, might be given by

$$w_i = \frac{RSS_i^{-1}}{\sum_{j=-1}^4 RSS_j^{-1}},$$

where $RSS_i = \sum_{t=87q1}^{93q4} \widehat{\varepsilon}_{i,t}^2$ so that the weight is (inversely) proportionate to the sum of squared residuals for the individual regressions. For Table 2, this would give weights of 0.1238, 0.1466, 0.1786, 0.1912, 0.1644, and 0.1954 for models M_{-1} to M_4 , respectively, reflecting the relatively good performance of the longer-horizon regressions. A summary of the model-averaged results can then be written in the form of a Taylor rule

$$\begin{aligned} r_t &= 0.802r_{t-1} + (1 - 0.802) \left(\begin{array}{ccc} -6.701 & + & 3.621\pi_{at} & + & 0.727_t x_{at} \end{array} \right) + \widehat{\varepsilon}_{at}, \\ &\quad (0.045) \qquad\qquad (1.597) \qquad (0.439) \qquad (0.203) \qquad (7) \\ &: \quad t = 1987q1 - T, \quad T = 1993q4, \end{aligned}$$

15. So, for example, $\gamma_\pi^{MR} = \gamma_\pi \sum_{s=0}^6 \rho^s$.

where the constructed coefficients on the lagged interest rate, inflation and gap variables are simply the weighted averages of the corresponding coefficients from the individual models, and where the π_{at} and x_{at} are notional variables denoting the inflation and gap pressures averaged across the various policy horizons. The standard errors of the constructed coefficients in (7) are readily calculated using the formulae in Lee, Pesaran, and Pierser (1990), taking the weights for each model as fixed.¹⁶

The Taylor rule of (7) accommodates the model uncertainty raised by the ambiguity on the policy horizon used by decision makers as well as the parameter and stochastic uncertainty that is more usual in estimated Taylor rule models. It shows that, for Taylor's sample period at least, the parameters are broadly consistent with the sort of policy rule advocated by Taylor, with a reasonably high degree of smoothing but with positive and statistically significant feedback from inflation and the gap to the interest rate with coefficients 3.621 and 0.727 over the long run.

2.2 Recursive Estimation of Taylor Rules, 1978q4–2008q4

We now turn to the main empirical exercise of the paper, broadening the analysis beyond the Taylor sample to use the data for the whole period 1969q1–2008q4. The beginning of the sample period is constrained by the availability of direct measures of expectations: expectations of output and inflation are available from the SPF at the one-, two-, three-, and four-quarter-ahead forecast horizon only from 1968q4. In the first instance, the set of models considered is exactly that described by (2), using (4) and (5) to construct model weights.¹⁷ To ensure sufficient degrees of freedom in estimating our Taylor rules, we assume that regimes last a minimum of 4 years (i.e., $j_{\min} = 16$). We also assume regimes do not last longer than 10 years (i.e., $j_{\max} = 40$) although, in the event, models of this duration get very small or zero weight. Our choice of j_{\max} is innocuous in this sense, although it does constrain the estimation period for the meta Taylor rule to run from 1978q4 to 2008q4. This issue is considered further in the extended analysis below.

Given our setup, the first set of 150 Taylor rules that were estimated relate to the sample window of 40 observations from 1969q1 to 1978q4, estimating six rules over the whole period with $i = -1, \dots, 4$, then six over the period 1969q2–1978q4, and so on, finishing with six models estimated over 1975q1–1978q4. For this initial period, we assume equal weights across all 150 models. A second set of 150 rules was then estimated relating to the 40 observations from 1969q2 to 1979q1. Weights were

16. Writing model $M_i : r_t = \mathbf{z}_t \boldsymbol{\theta}_i + \mathbf{u}_t$ for $i = 1, \dots, m$, and taking weights w_i as fixed, the covariance matrix of $\hat{\boldsymbol{\theta}}_a = \sum_{i=1}^m w_i \hat{\boldsymbol{\theta}}_i$ is given by $\sum_{i,j=1}^m w_i w_j \text{cov}(\hat{\boldsymbol{\theta}}_i, \hat{\boldsymbol{\theta}}_j)$ where $\text{cov}(\hat{\boldsymbol{\theta}}_i, \hat{\boldsymbol{\theta}}_j) = \hat{\sigma}_{ij} (\mathbf{z}'_i \mathbf{z}_i)^{-1} \mathbf{z}'_i \mathbf{z}_j (\mathbf{z}'_j \mathbf{z}_j)^{-1}$ and $\hat{\sigma}_{ij} = \hat{\mathbf{u}}'_i \hat{\mathbf{u}}_j$ under the assumptions on the error structure described in Lee, Pesaran, and Pierser (1990). Also, since $r_t = \boldsymbol{\theta}_a \mathbf{z}_t + (\boldsymbol{\theta}_i - \boldsymbol{\theta}_a) \mathbf{z}_t + \mathbf{u}_t$ in model i , averaging across i gives $r_t = \boldsymbol{\theta}_a \mathbf{z}_t + \mathbf{v}_t$, where $\mathbf{z}_t = \mathbf{n}^{-1} \sum_i \mathbf{z}_t$ and $\mathbf{v}_t = \mathbf{n}^{-1} \sum_i ((\boldsymbol{\theta}_i - \boldsymbol{\theta}_a) \mathbf{z}_t + \mathbf{u}_t)$. Hence, $\hat{\boldsymbol{\theta}}_a$ is the estimated responsiveness of r_t to \mathbf{z}_t assuming the $\boldsymbol{\theta}_i$ and \mathbf{z}_t are distributed independently across i so $E(\mathbf{z}_t \mathbf{v}_t) = 0$.

17. Unstable estimated rules, in which $\hat{\rho}_{iJT}$ exceeds unity, were excluded from the meta rule and replaced by models explaining Δr_t ; in other words, $\hat{\rho}_{iJT}$ was capped at unity. This only impacts on results up to 1981q2. Before this time, the proportion of capped models averaged around 30%, but very few unstable models were obtained afterward.




FIG. 2. Recursive Estimation of the Sample Horizon.

calculated for each of these models based on the weights for the models in the previous period but updated according to their ability to explain the final observation in 1979q1 as in (5). Here we assume a value for the break probability of $\lambda = 6/163 = 0.037$. This was based on a series of tests (one-step-ahead predictive failure tests, CUSUM tests, etc.) applied to a Taylor rule estimated (using contemporaneous inflation and gaps) over the 163 observations of our full sample, all of which suggested the presence of five or six breaks when working at the 5% level of significance.¹⁸ This procedure was then repeated as we moved recursively through the data set.

Figures 2–6 summarize the results of estimating the meta Taylor rule in this way, with the vertical lines denoting the start of the terms of office of Paul Volcker (1979q3), Alan Greenspan (1987q3), and Ben Bernanke (2006q1) as chairman of the Federal Reserve. Figure 2 plots the probability-weighted average sample length across the 150 models at each point in time, $\bar{j}_T = \sum_{i=-1}^4 \sum_{j=16}^{40} w_{ijT} \times j$, while Figure 3 plots the probability-weighted average policy horizon, $\bar{i}_T = \sum_{i=-1}^4 \sum_{j=16}^{40} w_{ijT} \times i$, to provide a sense of the relative importance of the 150 alternative models in each period. The corresponding confidence bands are plotted to show the precision of the estimated statistics and are obtained through stochastic simulation.¹⁹ Figures 4–6

18. Working at the 1% or 10% level of significance, the tests suggested as few as 4 or as many as 10 breaks, respectively. The meta rules obtained with corresponding values for λ are qualitatively unchanged from those reported in what follows, although the average duration of regimes grows a little more (less) rapidly with $\lambda = 4/163$ (10/163).

19. Specifically, the estimated meta Taylor rule describes a data-generating process for interest rates for each point in the sample given the history of inflation and the output gap and based on estimates of the 150 individual models and of the weights and transition probabilities. This data-generating process was used to provide 10,000 alternative simulated “histories” for interest rates. For each simulated series,




FIG. 3. Recursive Estimation of the Policy Horizon.




FIG. 4. Recursive Estimation of the Partial Adjustment Coefficient.

show the probability-weighted averages of the partial adjustment coefficients, $\bar{\rho}_T$, and the inflation and output gap elasticities over the medium run, $\bar{\gamma}_{\pi T}^{MR}$ and $\bar{\gamma}_{xT}^{MR}$.

Regimes and policy horizons. Figure 2 suggests that policy over the period can be usefully grouped into four broad phases: the Volcker/early-Greenspan phase (1979q4–1993q4), the mid-Greenspan phase (1994q1–1999q2), the late-Greenspan

the procedure described to estimate the meta Taylor rule was implemented and distributions of average sample lengths and average policy horizons obtained. The confidence intervals illustrate the range covered by two standard deviations of these distributions.




FIG. 5. Recursive Estimation of the Medium-Run Inflation Coefficient.




FIG. 6. Recursive Estimation of the Medium-Run Output Gap Coefficient.

phase (1999q3–2005q4), and the (pre-zero-lower-bound) Bernanke phase (2006q1–2008q4). The first phase starts at the beginning of Volcker’s term of office where the probability-weighted average sample length drops to a very low level, indicating that monetary policy at that time was implemented very differently than previously. This coincides with the well-publicized move in October 1979 by the Federal Reserve to reorientate policy toward price stability by targeting nonborrowed reserves to control monetary growth instead of the federal funds rate (see Lindsey, Orphanides, and Rasche 2005). The average sample length rises only slowly throughout the early 1980s, reflecting the ongoing challenges in targeting nonborrowed reserves instead

of the federal funds rate, which was restored as the primary policy instrument by 1982 (see Axilrod 2005).²⁰

The protracted rise in the probability-weighted average sample length to the end of 1993, evident in Figure 2, implies some continuity in policy regime that included the rest of the Volcker years, but especially took hold in the early years of the Greenspan years. The continuity appears to end in 1994 when successive rises in the interest rate, and the Federal Open Market Committee (FOMC) decision to announce its policy actions immediately upon making them, herald the beginning of the mid-Greenspan phase. The rise in average sample length over the subsequent 5 years implies a further period of policy stability. But the drop in sample length, and the start of the late-Greenspan phase, in 1999q3 also coincides with an important change in policy operation as the FOMC started releasing press statements including “intended federal funds rate” and “policy bias” from May 1999, having including a numeric value of the “intended federal funds rate” in each policy directive since August 1997 (see Poole, Rasche, and Thornton 2002 for a thorough discussion of changes in policy operations over this time).

The decline in the probability-weighted average sample length observed in Figure 2 around the beginning of Bernanke’s term of office corresponds with a fourth phase of policy. The especially sharp drop in August 2007 coincides with the Federal Reserve’s change of discount window policies in order to “promote the restoration of orderly conditions in financial markets” (see August 17, 2007 press release of the Federal Reserve Board of Governors).

The evolution of policy within and across phases reflects changes in the responsiveness of interest rates to inflation and business cycle fluctuations, sometimes occurring abruptly and sometimes more gradually, that we discuss in detail below. But the flexibility of the meta-modeling approach also allows us to capture changes over time in decision makers’ policy horizons. Specifically, Figure 3 indicates that the policy horizon was generally forward looking, with the average sample horizon ranging between one and two periods ahead when considered over the whole sample period. However, the estimation of the policy horizon varies considerably over the sample and, in particular, there is evidence that the policy horizon shortens during times of recession. This is apparent as the troughs in the probability-weighted policy horizon occur in 1980q3, 1982q3, 1991q3, 2000q4, and 2008q4, which correspond closely to the troughs in activity in the five recessions identified by the NBER during this period. This ability of the meta-modeling approach to capture changes in the policy horizon is a major advantage of it over a standard model-based approach, even when considering models that are flexible enough to capture time-varying parameters.

20. The fact that the probability-weighted average sample length never reaches the upper bound of 40 quarters confirms that there is genuine model averaging in the sense of nontrivial weights distributed across multiple models. If one model were to have dominated, the average sample length should have increased to the upper bound as the sample period progressed. The result also directly suggests that that our particular choice of 40 quarters for the upper bound on a policy regime is sufficiently large when considering U.S. monetary policy for this sample period.

Smoothing, inflation and gap effects. Figures 4–6 show the probability-weighted averages $\bar{\rho}_T$, $\bar{\gamma}_{\pi T}^{MR}$, and $\bar{\gamma}_{xT}^{MR}$ to provide further insights into the nature of changing policy regimes implied by Figure 2. The bands on these diagrams are 95% confidence intervals based on standard errors obtained analytically taking the weights as given and using the approach of Lee, Pesaran, and Pierse (1990) although bands obtained through simulation were qualitatively similar. The Volcker/early-Greenspan phase starts with a high, albeit imprecisely estimated, coefficient on inflation that remains high, and indeed slowly rises, throughout the period to peak at 3.30 in 1993q4.²¹ The phase is also characterized by increasingly smooth interest rate decisions and a growing influence from the output gap measure. The latter has a low and statistically insignificant influence through to the late 80s but begins to show significantly during the recessionary period of 1990–91 and stays at the new level of around 0.5 to 1993.

The second, mid-Greenspan, phase starting at the beginning of 1994 reflects a shift to a policy to preempt inflation (see Goodfriend 2003 for discussion). Policy continues to involve strong responses to inflation but shows a marked shift of focus toward the output gap with its coefficient doubling in size over this phase. This was a period when the output gap became positive for the first time after the early 1990s recession and remained high throughout the boom years to mid-1999. The increasing influence of the output gap influence will have kept interest rates higher than they would have been if earlier versions of the rule had been implemented and reflects a desire to manage the growth in demand and to avoid overheating.

The third, late-Greenspan, phase running 1999q3–2005q4 saw the gap continuing to exert a relatively high level of influence on policy, but introduced a more agile responsiveness as evidenced by a noticeable reduction in the partial adjustment coefficient. Importantly, this period also saw the inflation coefficient falling and becoming insignificantly different to zero during 2004. The decline in the response to inflation could be related to a worry at the time that, with rates having been lowered very rapidly through the recession of 2001, a strong reaction to expected inflation might actually trigger deflation and the federal funds rate might eventually hit the zero lower bound.²² In any case, the decline in the influence of inflation was reversed during the last quarters of Greenspan’s term of office and through the first quarters of Bernanke’s term to 2007q4, matched by a slight shift in emphasis away from the output gap. But these trends were halted completely by the financial crisis, which appears to have shifted attention solely to the gap and eliminated the influence of inflation.

21. The imprecision in the early years is associated with the instability of estimated rules, and their replacement with models explaining Δr_t , during this time. This is as might be expected while nonborrowed reserves were targeted rather than the federal funds rate itself.

22. See Alcidi, Flamini, and Fracasso (2011) for a nonlinear Taylor rule that captures this period as a regime in which policymakers worry about the zero lower bound problem.

This commentary illustrates the usefulness of the meta Taylor rule in providing a single framework with which to interpret monetary policy since the late 70s, at least prior to the fed funds rate hitting the zero lower bound. It captures the continuity and strong anti-inflationary stance of the Volcker/early-Greenspan years, the fear of overheating in the mid-Greenspan years, the easing on inflation during most of the late-Greenspan years, the reassertion of anti-inflationary policies in the late-Greenspan and early-Bernanke years, and the dramatic impact of the financial crisis on policy. It accommodates the shifts to more myopic decision making during periods of recession and changes in the extent to which the Fed chooses to smooth its policy responses. The meta rule is able to capture this complexity and these nuances in policy in a very straightforward way, without recourse to complex nonlinearities or strategies to deal with structural breaks because of the simplicity of the individual linear Taylor rules that underlie it and the flexibility of the model averaging framework.

2.3 Extensions of the Analysis

The flexibility of meta rule approach can be further exploited to deal with two data limitations that were mentioned earlier and to extend the Taylor rule analysis. The first limitation concerns the estimation period which runs 1978q4–2008q4. The start date in 1978q4 is defined by the start date of the sample, 1969q1, and by our use of $j_{\max} = 40$ in 2 reflecting our wish to allow for Taylor rules that are unchanged for up to 10 years (even if this turns out to be unlikely in practice). This approach insists on considering the same set of 150 potential models at each point. A more pragmatic approach might be to consider all 150 models when data allows at 1978q4, but to allow for a maximum length of 39 observations in 1978q3, 38 in 1978q2, and so on. If we continue to assume that we need at least 16 observations to be able to reliably estimate a Taylor rule, this means that we could extend our analysis to run from 1972q4 to 2008q4. Only six alternative models will be considered to explain interest rate determination in 1972q4, relating to the six policy horizons, and the meta model is initiated with equal weights on each of the six. The meta model will extend to consider 12 models in 1973q1, rising to 150 models by 1978q4 as uncertainty on sample length is introduced progressively and with the weights evolving as in (4) and (5) as before.

The second data limitation discussed in the analysis above relates to the measure of the output gap and the fact that the Greenbook measure of the gap—which describes policymakers’ stated views on the size of the gap in real time—is available publicly only for the period 1987q3–2004q4. One reaction to the absence of a complete run of data is simply to use an alternative series, as we did above. But in reality, new sources of information do become available over time and it is interesting to see how the availability of a new data series might have impacted on monetary policy decisions in real time. The model-averaging approach is able to accommodate this sort of break by including in the meta rule an extra set of models that uses the new variable. So, here, an extra set of 6 models that incorporate the Greenbook measure of the gap can




FIG. 7. Recursive Estimation of the Sample Horizon with Expanded Model Space.

be considered in explaining the interest rate decision in 1991q2, in addition to the 150 used previously, using the first 16 observations of the new gap series in place of Taylor's linear trend-based measure. The number of extra models increases to 12 in 1991q3 and grows to 150 models by 1997q2 taking into account the uncertainty on sample length in these extra Greenbook gap models. At this point, the meta Taylor Rule includes as many models incorporating the Greenbook gap measure as models that incorporate the gap measure based on the linear trend. If the new information source becomes influential, then this would be reflected by a shift of weights toward these alternative models.²³

Figures 7–11 summarize the meta Taylor rule obtained to accommodate these two extensions. The estimates run from the earlier starting date of 1972q4 using just 6 models estimated on 16 observations between 1969q1 and 1972q4 and using the estimated transition probabilities to build up to 150 models at 1978q4. The estimates also accommodate additional models that make use of the Greenbook/CBO output gap data from 1991q2. The extended estimates illustrate two further distinct monetary policy regimes associated with the chairmanships of Arthur Burns to 1978q1 and William Miller 1978q2–1979q3. Both periods are characterized as having a high degree of inertia and in fact the number of models in which the partial adjustment coefficient ρ_{ijT} is capped at unity averages around 50% during these phases so that

23. The policymakers' gap measure is extended to the end of the sample by splicing the Greenbook series that is available upto 2004q4 with the GDP gap measures provided by the CBO. See Poole (2007) for a discussion of the properties of such a spliced series. The CBO figures for ${}_t x_{t+i}$, $i = -1, \dots, 4$, are published biannually, in January and August, and the q2 and q4 figures are obtained through linear interpolation of the real time data.




FIG. 8. Recursive Estimation of the Policy Horizon with Expanded Model Space.




FIG. 9. Recursive Estimation of the Partial Adjustment Coefficient with Expanded Model Space.

the meta rule incorporates many models effectively explaining Δr_t . The responses to inflation and the output gap are estimated imprecisely, but there is little evidence of any feedback from inflation to interest rates during the Burns period, consistent with the much highlighted finding in Clarida, Gali, and Gertler (2000). Inflation becomes more significant for policy during the Miller term of office and the fall in the average sample length over this period reflects the redirection of policy that was then pursued during the Volcker/early-Greenspan phase as identified in the previous subsections.

Notably, the estimated meta rule in the extended analysis looks very similar to that described above for most of the remaining sample. This is despite the differences in




FIG. 10. Recursive Estimation of the Medium-Run Inflation Coefficient with Expanded Model Space.




FIG. 11. Recursive Estimation of the Medium-Run Output Gap Coefficient with Expanded Model Space.

the starting point of the analysis and the potential impact of a different set of estimated transition probabilities during the early part of the sample and despite the different measures of the gap used in the later part of the sample. As it turns out, the sum of the weights on the models including the Greenbook output gap measure are generally high, rising from zero when they are first introduced in 1991q2 to around 0.8 by 1995q4. The weights are more evenly distributed during the periods when the gap has less influence on interest rates (1998–99 and 2004 onward) but they emphasise the Greenbook measure throughout 2000–04 (with an average weight of 0.77) when the influence of the gap on interest rates was at its greatest. In any case, the distinct phases of policy discussed in the previous subsections are again recognizable in the data, with the same characterization of policy regimes appearing to hold. The main change in interpretation relates to the timing of the reassertion of the anti-inflationary

policies of the late-Greenspan/early-Bernanke years, which show a more abrupt shift to an aggressively anti-inflation policy from 2005q2, not long before Ben Bernanke became chairman.

3. CONCLUSION

The use of model averaging is now widespread in the forecasting literature. The analysis of this paper demonstrates that the approach is also useful in the context of behavioral modeling and inference, providing an extremely flexible tool with which to model and characterize economic decision making. The modeling approach can accommodate the uncertainties surrounding parameter estimates and random shocks to relationships as usual, but can also accommodate a modeler's uncertainty over the period during which relations hold and on the measures of variables used in decision making. Our results show that a "meta" Taylor rule provides a flexible but compelling characterization of monetary policy in the United States over the last 40 years, with no single model of monetary policy dominating over the sample period or even at any point of time. The estimated rule highlights the lack of feedback from inflation to monetary policy during the early to mid 70s the change in direction in the Miller years and the continuity in a policy stance based on appropriately strong feedback from inflation and output gaps to policy during the Volcker/early-Greenspan years. The rule also provides evidence of changes in the emphasis on inflation and on the output gap in the rule subsequently, illustrating the successive effects of the fear of overheating, an easing on inflation, a reassertion of anti-inflationary policies and the dramatic impact of the financial crisis on policy. We also find that the Federal Reserve cares more about the immediate future than longer horizons during periods of recession. The meta-modeling approach draws these inferences out in a straightforward way that is more flexible than even highly complicated time-varying parameter models.

As a caveat, our analysis is predicated on the Taylor-rule assumption that U.S. monetary policy responds to inflation and the output gap in a linear fashion and that we have accurately measured inflation and the gap as they are perceived by monetary policy authorities. To the extent that policymakers respond to other variables or respond nonlinearly or have different perceptions about inflation and the gap, we may be overestimating the instability of policy regimes. For example, if the perceived output gap is strongly asymmetric, being more negative in recessions than positive in expansions, as found in Morley and Piger (2012), we would expect an increase in the estimated response to our linear measure of the output gap during recessions and a decrease in expansions, even if the true response to the gap is stable. Notably, the estimates in Figure 11 for the response to the output gap behave this way, at least in the 2000s. However, a full investigation of more complicated models and additional measures of inflation and the output gap are left for future research. Meanwhile, at least within the linear Taylor-rule framework and assuming our measures of inflation and the output gap do provide a reasonable approximation of the perceptions of

policymakers, we find strong support for sizable and interesting changes in the systematic elements of U.S monetary policy over the last 40 years.

LITERATURE CITED

- Alcidi, Cinzia, Alessandro Flamini, and Andrea Fracasso. (2011) "Policy Regime Changes, Judgment and Taylor Rules in the Greenspan Era." *Economica*, 78, 89–107.
- Axilrod, Stephen H. (2005) Commentary, *Federal Reserve Bank of St. Louis Review*, 87, 237–42.
- Boivin, Jean. (2005) "Has US Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data." *Journal of Money, Credit and Banking*, 38, 1149–73.
- Clarida, Richard, Jordi Gali, and Meric Gertler. (2000) "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics*, 115, 147–80.
- Cogley, Timothy, and Thomas J. Sargent. (2005) "The Conquest of US Inflation: Learning and Robustness to Model Uncertainty." *Review of Economic Dynamics*, 8, 528–63.
- Elliott, Graham, and Allan Timmermann. (2008) "Economic Forecasting." *Journal of Economic Literature*, 46, 3–56.
- Garratt, Anthony, Kevin C. Lee, Emi Mise, and Kalvinder Shields. (2008) "Real Time Representations of the Output Gap." *Review of Economics and Statistics*, 90, 792–804.
- Garratt, Anthony, Kevin C. Lee, Emi Mise, and Kalvinder Shields. (2009) "Real Time Representations of the UK Output Gap in the Presence of Model Uncertainty." *International Journal of Forecasting*, 25, 81–102.
- Garratt, Anthony, Kevin C. Lee, M. Hashem Pesaran, and Yongcheol Shin. (2003) "Forecast Uncertainty in Macroeconometric Modeling: An Application to the UK Economy." *Journal of the American Statistical Association*, 98, 829–38.
- Garratt, Anthony, Kevin C. Lee, M. Hashem Pesaran, and Yongcheol Shin. (2006) *National and Global Macroeconometric Modeling: A Long-Run Structural Modeling Approach*. Oxford, UK: Oxford University Press.
- Garratt, Anthony, Kevin C. Lee, and Kalvinder Shields. (2014) "Information Rigidities and the News-Adjusted Output Gap." University of Melbourne Working Paper, Melbourne, Australia.
- Goodfriend, Marvin. (2003) "The Phases of U.S. Monetary Policy: 1987–2001." In *Central Banking, Monetary Theory and Practice*, Vol. 1, edited by P. Mizen, pp. 173–89. Cheltenham, UK: Edward Elgar Publishing.
- Kim, Chang-Jin, and Charles R. Nelson. (2006) "Estimation of a Forward-Looking Monetary Policy Rule: A Time-Varying Parameter Model Using Ex-Post Data." *Journal of Monetary Economics*, 53, 1949–66.
- Kozicki, Sharon. (1999) "How Useful Are Taylor Rules for Monetary Policy?" *Federal Reserve Bank of Kansas Economic Review*, Second Quarter, 5–33.
- Lee, Kevin C., M. Hashem Pesaran, and Richard G. Pierse. (1990) "Testing for Aggregation Bias in Linear Models." *Economic Journal*, 100, 137–50.
- Lindsey, David E., Athanasios Orphanides, and Robert H. Rasche. (2005) "The Reform of October 1979: How It Happened and Why." *Federal Reserve Bank of St. Louis Review*, 87, 187–235.

- McCulloch, J. Huston (2007) "Adaptive Least Squares Estimation of the Time-Varying Taylor Rule." Working Paper, University of Ohio, Columbus.
- Morley, James, and Jeremy Piger. (2012) "The Asymmetric Business Cycle." *Review of Economics and Statistics*, 94, 204–21.
- Orphanides, Athanasios. (2001) "Monetary Policy Rules Based on Real-Time Data." *American Economic Review*, 91, 964–85.
- Orphanides, Athanasios. (2003) "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics*, 50, 983–1022.
- Orphanides, Athanasios, and Simon van Norden. (2002) "The Unreliability of Output Gap Estimates in Real Time." *Review of Economics and Statistics*, 84, 569–83.
- Orphanides, Athanasios, Richard D. Porter, David Refschneider, Robert Tetlow, and Frederico Finan. (2000) "Errors in the Measurement of the Output Gap and the Design of Monetary Policy." *Journal of Economics and Business*, 52, 117–41.
- Pesaran, M. Hashem, and Allan Timmerman. (1995) "Small Sample Properties of Forecasts from Autoregressive Models under Structural Breaks." *Journal of Econometrics*, 129, 183–217.
- Poole, William. (2007) "Understanding the Fed." *Federal Reserve Bank of St. Louis Review*, 89, 3–13.
- Poole, William, Robert H. Rasche, and Daniel L. Thornton. (2002) "Market Anticipations of Monetary Policy Actions." *Federal Reserve Bank of St. Louis Review*, 84, 65–94.
- Romer, Christina, and David Romer. (2004) "Choosing the Federal Reserve Chair: Lessons from History." *Journal of Economic Perspectives*, 18, 129–62.
- Taylor, John B. (1993) "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, 39, 195–214.
- Woodford, Michael. (2001) "The Taylor Rule and Optimal Monetary Policy." *American Economic Review*, 91, 232–37.
- Woodford, Michael. (2003) *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.