

IS THE FED TOO TIMID? MONETARY POLICY IN AN UNCERTAIN WORLD

Glenn D. Rudebusch*

Abstract—Estimates of the Taylor rule using historical data from the past decade or two suggest that monetary policy in the U.S. can be characterized as having reacted in a moderate fashion to output and inflation gaps. In contrast, the parameters of optimal Taylor rules derived using empirical models of the economy often recommend much more vigorous policy responses. This paper attempts to match the historical policy rule with an optimal policy rule by incorporating uncertainty into the derivation of the optimal rule and by examining plausible variations in the policymaker's model and preferences.

I. Introduction

THE question that motivates this paper is: How close was recent monetary policy to the behavior recommended by an optimal policy rule? An optimal rule can be derived with a structural model and a loss function for policymakers. For example, Rudebusch and Svensson (1999) used a small empirical model of the U.S. economy and a loss function penalizing output, inflation, and interest-rate variability to derive the optimal coefficients of a Taylor rule. A notable feature of this optimal Taylor rule was the large size of the inflation and output-gap response coefficients, which suggest that the ideal monetary policy behavior by the Fed would be quite responsive to economic conditions.¹ A Taylor rule can also be used to model historical monetary policy, and empirical estimates of such a rule appear to capture recent Fed behavior fairly well. However, the historical Taylor rule estimated using recent data appears to have relatively low response coefficients for output and inflation. That is, this estimated rule demonstrates a more cautious adjustment of the monetary policy instrument than is recommended by the optimal rule. This paper attempts to reconcile historical and optimal policy rules.

Received for publication September 28, 1999. Revision accepted for publication June 21, 2000.

* Federal Reserve Bank of San Francisco.

For helpful comments, I thank participants at the 1998 NBER Universities Research Conference on the Formulation of Monetary Policy, seminar participants at the Federal Reserve Board, the European Central Bank, Stockholm University, and the Federal Reserve Bank of San Francisco, as well as Charlie Evans, Sharon Kozicki, Frank Smets, Carl Walsh, and Ken West. Heather Royer and Niloofar Badie provided excellent research assistance. The views expressed in this paper do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

¹ This result does not appear to be unique to the particular model used by Rudebusch and Svensson. For example, McCallum and Nelson (1999) also obtained large optimal Taylor rule parameters in a very different model.

Of course, one possibility is that historical monetary policy cannot be described as the outcome of an economic optimization problem. This resolution, besides cutting short the current paper, would seem unsatisfactory on several levels. First, although it is hard to fit a stable reaction function or rule to the entire postwar history of U.S. monetary policy (Rudebusch, 1998), as described below, some success has been achieved in this regard for the past decade or so. Second, recent Fed policy and the economic performance it has helped foster have garnered both academic and general acclaim, so it seems likely that some sort of optimum has been approximated. Finally, a long-standing principle of economics is that any economic behavior can be understood as a problem in constrained optimization, and this principle should apply to central banks as forcefully as to the representative firm or agent.

The obvious avenue for reconciling the historical policy rule and the optimal rule is to alter the macroeconomic model or objective function used in deriving the latter in order to obtain a better match with real-world policy. With regard to the objective function, this paper does not explore in any great detail possible variations in the goals postulated for the central bank. It maintains a fairly standard assumption that the Fed is concerned with minimizing (in a quadratic fashion) output variation around potential, inflation variation around a target, and interest-rate volatility. This paper instead focuses on the context for decision making and, in particular, on how the addition of uncertainty into the model may alter the calculation of optimal policy. I also consider the uncertainty about the model used by policymakers and examine plausible model variation.

Especially since Brainard (1967), it has been recognized that uncertainty about model parameters can produce smaller responses, or "stodginess" (Blinder, 1998), in optimal policy rules. Indeed, policymakers often note that typically little new information is obtained between policy meetings (or from quarter to quarter) to justify large changes in the stance of policy. In particular, uncertainty about the state of the economy (data uncertainty) and about the trajectory and responsiveness of the economy (model or parameter uncertainty) appear to weigh heavily on policymakers. For example, at the Federal Open Market Committee (FOMC) meeting on December 16, 1987, a Federal

Reserve Board research director, after summarizing the staff forecast, stated:

By depicting these two [forecast] scenarios, I certainly don't want to suggest that a wide range of other possibilities doesn't exist. However, I believe both scenarios are well within the range of plausible outcomes, and they point up what we perceive to be a dilemma for the Committee: namely, given the lags in the effect of policy action, an easing or tightening step might be appropriate now, but it isn't clear which. This, of course, isn't an unprecedented problem, . . .

Similarly, in discussing rules for policy, another Fed research director (Kohn, 1999) notes that members of the FOMC "are quite uncertain about the quantitative specifications of the most basic inputs required by most rules and model exercises. They have little confidence in estimates of the size of the output gap [or] the level of the natural or equilibrium real interest rate . . ." (p. 195).

This paper then is an attempt to reconcile recommendations about optimal policy rules with actual estimates of the historical policy rule. I largely focus on how much and what type of uncertainty must be added to the model so that the resulting calculated optimal policy rule matches the historical one. This reverse engineering is conducted in the context of a Taylor rule for policy. The next two sections set the stage by presenting actual historical policy—in the form of estimated Taylor rules—and the contrasting optimal Taylor rules derived without uncertainty. Sections IV, V, and VI introduce, in isolation, parameter uncertainty, model variation, and data uncertainty, respectively, into the derivation of optimal policy. Section VII combines various types of uncertainty, and section VIII concludes.

II. Historical Estimates of the Policy Rule

Taylor (1993) proposed a simple rule for monetary policy:

$$i_t = r^* - 0.5\pi^* + 1.5\bar{\pi}_t + 0.5y_t, \quad (1)$$

where i_t is the quarterly average federal funds rate at an annual rate in percent;

$\bar{\pi}_t$ is the four-quarter inflation rate in percent;

y_t is the percent gap between actual real GDP (Q_t) and potential real GDP (Q_t^*), that is, $100(Q_t - Q_t^*)/Q_t^*$;

r^* is the equilibrium real interest-rate in percent; and

π^* is the inflation target.

As a descriptive rule, Taylor (1993, 1999) argued that the rule (1), with $r^* = \pi^* = 2.0$, seemed to capture some important factors influencing monetary policy and the general stance of policy from the mid-1980s onward. (Also see Judd and Rudebusch (1998).)

Actual estimates of a generalized Taylor rule of the form

$$i_t = k + g_\pi \bar{\pi}_t + g_y y_t, \quad (2)$$

(with the constant term $k \equiv r^* - (g_\pi - 1)\pi^*$) seem to bear this out. At the most rudimentary level, a simple least-squares regression of equation (2) from 1987:Q4 to 1996:Q4 yields

$$i_t = 0.63 + 1.78 \bar{\pi}_t + 0.82 y_t + \varepsilon_t^i, \quad (3)$$

(0.86) (0.22) (0.10)

where inflation is defined using the GDP chain-weighted price index (denoted P_t so $\pi_t = 400(\ln P_t - \ln P_{t-1})$ and $\bar{\pi}_t = \frac{1}{4} \sum_{j=0}^3 \pi_{t-j}$), and the output gap is defined with potential output as estimated by the Congressional Budget Office (1995). In this regression, the values of the estimated rule parameters (namely, $g_\pi = 1.78$ for the inflation response and $g_y = 0.82$ for the output response) are just slightly higher than the 1.5 and 0.5 that Taylor (1993) originally proposed. (Robust standard errors are reported in parentheses.)

More careful econometric analysis also supports such moderate policy response parameters. A key feature of the various studies that estimate Taylor rules with the historical data is that they take account of the apparent slow adjustment of the actual rate to the level recommended by the Taylor rule; thus, lagged interest rates are added to the regression to account for the apparent serial correlation in the ε_t^i .² For example, Judd and Rudebusch (1998) estimate a Taylor rule like equation (2) in the context of an error-correction framework (from 1987:Q3 to 1997:Q4) and find an inflation response of $g_\pi = 1.54$ and an output response of $g_y = 0.99$ (with standard errors of 0.18 and 0.13, respectively). Similarly, with closely related dynamic Taylor rule specifications, Kozicki (1999) estimates $g_\pi = 1.42$ and $g_y = 0.49$ (from 1983 to 1997), and Clarida, Gali, and Gertler (2000) estimate $g_\pi = 2.02$ and $g_y = 0.99$ (from 1982:Q4 to 1996:Q4).

As a rough benchmark then, the historical estimates of U.S. monetary policy during the late 1980s and the 1990s suggest that policy can be broadly described by a Taylor rule with a g_π in the range of 1.4 to 2.0 and a g_y in the range of 0.5 to 1.0. This paper will attempt to find a control problem that produces an optimal Taylor rule with response coefficients in these ranges. The next section considers this issue under certainty.

III. Optimal Policy Under Certainty

A. An Empirical Model of Output and Inflation

The optimal policy rules in this paper are derived in a simple model of output and inflation:

² The size of the estimated coefficient on the lagged interest rate is often quite large—on the order of 0.8—however, as noted below—and discussed in detail in Rudebusch (2000b)—such estimated lagged dynamics is open to interpretation.

$$\pi_t = \alpha_0 + \alpha_{\pi 1}\pi_{t-1} + \alpha_{\pi 2}\pi_{t-2} + \alpha_{\pi 3}\pi_{t-3} + \alpha_{\pi 4}\pi_{t-4} + \alpha_y y_{t-1} + \varepsilon_t \quad (4)$$

$$y_t = \beta_0 + \beta_{y1}y_{t-1} + \beta_{y2}y_{t-2} - \beta_r(\bar{i}_{t-1} - \bar{\pi}_{t-1}) + \eta_t, \quad (5)$$

with \bar{i}_{t-1} equal to the four-quarter average federal funds rate ($\frac{1}{4}\sum_{j=1}^4 i_{t-j}$), and the other variables defined as above.

The first equation is a Phillips curve that relates inflation to a lagged output gap and to lags of inflation, which represent an autoregressive or adaptive form of inflation expectations. The second equation is an IS curve that relates the output gap to its own lags and to the difference between the average funds rate and average inflation over the previous four quarters—an approximate *ex post* real rate. As described by Rudebusch and Svensson (1999, 2000), the use of this model can be motivated by a variety of considerations. In particular, although its simple structure facilitates the production of benchmark results, this model also appears to roughly capture the views about the dynamics of the economy held by some monetary policymakers, including Federal Reserve Governor Meyer (1997) and former Federal Reserve Vice-Chairman Blinder (1998). This point is fundamental to my analysis, which is predicated on the assumption that policymakers acted in an optimal manner. If I find an optimal policy rule for a particular model that matches the historical policy rule, this result is surely undercut if policymakers believed that they were optimizing in a completely different model.

The empirical fit of the model is also quite good. The estimated equations, using the sample period 1961:1 to 1996:4, are shown below. (Coefficient standard errors are given in parentheses, and the standard error of the residuals and Durbin-Watson statistics are reported.)

$$\begin{aligned} \pi_t &= 0.08 + 0.67\pi_{t-1} - 0.08\pi_{t-2} \\ &\quad (0.09) \quad (0.08) \quad (0.10) \\ &+ 0.29\pi_{t-3} + 0.12\pi_{t-4} \\ &\quad (0.10) \quad (0.08) \\ &+ 0.15y_{t-1} + \varepsilon_t^i, \\ &\quad (0.03) \end{aligned} \quad (6)$$

$$\sigma_\varepsilon = 1.007, \quad DW = 1.99,$$

$$\begin{aligned} y_t &= 0.19 + 1.17y_{t-1} - 0.27y_{t-2} \\ &\quad (0.10) \quad (0.08) \quad (0.08) \\ &- 0.09(\hat{i}_{t-1} - \hat{\pi}_{t-1}) + \eta_t, \\ &\quad (0.03) \end{aligned} \quad (7)$$

$$\sigma_\eta = 0.822, \quad DW = 2.09.$$

These equations were estimated individually by OLS.³ The hypothesis that the sum of the lag coefficients of inflation equals 1 had a *p*-value of 0.48, so this restriction was imposed in estimation. Thus, this is an accelerationist form of the Phillips curve, which implies a long-run vertical Phillips curve. (This is reconsidered in Section VI.) The fit and dynamics of this model compare favorably to an unrestricted VAR. Indeed, the model can be interpreted as a restricted VAR, in which the restrictions imposed are not at odds with the data as judged, for example, with standard model information criteria. (See Rudebusch and Svensson (1999).)

In addition, the model appears to be stable over various subsamples, which is an important condition for drawing inference. With a backward-looking model, the Lucas critique may apply with particular force, so it is important to gauge its historical importance with econometric stability tests (Oliner, Rudebusch, and Sichel, 1996). For example, consider a stability test from Andrews (1993): the maximum value of the likelihood-ratio test statistic for structural stability over all possible breakpoints in the middle 70% of the sample. For our estimated inflation equation, the maximum likelihood-ratio test statistic is 10.89 (in 1972:3), and the 10% critical value is 14.31 (from table 1 in Andrews (1993)). Similarly, for the output equation, the maximum statistic is 11.51 (in 1982:4), and the 10% critical value is 12.27. However, it should be noted that the Lucas critique is arguably less of an issue for this paper, which analyzes a policy rule that tries to match the historical policy regime, than for the typical policy evaluation exercise (Rudebusch and Svensson, 1999), which analyzes policy rules that differ from the historical regime.

Finally, note that it is useful to rewrite the IS curve as

$$\begin{aligned} y_t &= \beta_{y1}y_{t-1} + \beta_{y2}y_{t-2} \\ &- \beta_r(\bar{i}_{t-1} - \bar{\pi}_{t-1} - r^*) + \eta_t, \end{aligned} \quad (8)$$

where $r^* = \beta_0/\beta_r$ is the equilibrium real rate relevant for the constant term (*k*) of the Taylor rule (2). Empirically, a (nonlinear) least-squares regression of equation (8) yields $r^* = 2.21$ with a coefficient standard error of 0.80. (Of course, nothing else about the regression equation (7) changes.) This point estimate of the equilibrium real rate is used as the constant term in the Taylor rule; furthermore, the standard error of this coefficient will be useful in analyzing the effect of uncertainty about the equilibrium real rate on the performance of the Taylor rule.

³ These coefficient estimates differ slightly from those in Rudebusch and Svensson (1999) because of a slightly longer sample and revised data. Also, almost identical parameter estimates were obtained by SUR and by system ML methods because the cross-correlation of the errors is essentially zero.

TABLE 1.—BASELINE RESULTS WITH NO UNCERTAINTY

Rule Parameters		Volatility Results			Expected Loss		
g_π	g_y	Std [$\bar{\pi}_t$]	Std [y_t]	Std [Δi_t]	λ	ν	Loss
Panel A. Historical Estimated Rule: $i_t = k + g_\pi \bar{\pi}_t + g_y y_t$							
1.78	0.82	2.96	2.20	0.97	1	0.5	14.05
Panel B. Optimal Taylor Rule: $i_t = k + g_\pi \bar{\pi}_t + g_y y_t$							
2.83	1.63	2.16	2.22	1.80	1	0.5	11.23
2.30	2.19	2.54	1.93	2.19	4	0.5	23.77
3.22	1.33	2.02	2.52	1.72	0.25	0.5	7.15
3.59	2.68	1.97	2.15	2.88	1	0.1	9.34
1.96	0.65	2.75	2.38	0.93	1	5	17.59
Panel C. Optimal Dynamic Rule: $i_t = (1 - \rho)(k + g_\pi \bar{\pi}_t + g_y y_t) + \rho i_{t-1}$							
2.87	1.80	2.17	2.24	1.74	1	0.5	11.21
Panel D. Optimal Unrestricted Rule							
—	—	2.13	2.23	1.72	1	0.5	11.00
Panel E. Optimal Lagged Rule: $i_t = k + g_\pi \bar{\pi}_{t-1} + g_y y_{t-1}$							
2.59	1.61	2.36	2.42	1.78	1	0.5	13.01

The optimal inflation- and output-response parameters are given in the first two columns. The associated standard deviations are given in the next three columns. The final column gives the associated loss, and λ and ν are the weights on output and interest-rate volatility in the loss function.

B. The Optimal Taylor Rule

Deriving the optimal monetary policy rule for equation (6) and (7) requires an objective function, and I use one that is fairly standard in the literature (Rudebusch and Svensson, 1999). The central bank is assumed to minimize the variation in inflation around its target π^* , in the output gap, and in changes in the interest rate. Specifically, expected loss equals the weighted sum of unconditional variances,

$$E[L_t] = \text{Var}[\bar{\pi}_t - \pi^*] + \lambda \text{Var}[y_t] + \nu \text{Var}[\Delta i_t], \quad (9)$$

where $\Delta i_t = i_t - i_{t-1}$, and the parameters $\lambda \geq 0$ and $\nu \geq 0$ are the relative weights on output stabilization and interest-rate smoothing, respectively, with respect to inflation stabilization.⁴

Table 1 provides a variety of results assuming no uncertainty.⁵ Panel A shows the results for the estimated Taylor rule (3) with $g_\pi = 1.78$ and $g_y = 0.82$ (which are taken as representative of the range of historical estimates). This rule when coupled with model (6) and (7) produces $\text{Std}[\bar{\pi}_t] = 2.96$, $\text{Std}[y_t] = 2.20$, and $\text{Std}[\Delta i_t] = 0.97$. (For these results and those below, $\pi^* = 0$; thus, $k = r^* = 2.21$, the

⁴ Alternative preferences are a possibility. For example, policymaker preferences may be asymmetric with regard to outcomes, perhaps leading to the opportunistic behavior described by Bomfim and Rudebusch (2000).

⁵ All the results in this paper are obtained by simulating the relevant model with $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ and $\eta_t \sim N(0, \sigma_\eta^2)$. The g_π and g_y parameters of the optimal Taylor rule are obtained via a grid search. Specifically, a simulation of 100,000 observations is formed for each candidate Taylor rule, and the optimal one is chosen. In the case of table 1, identical optimal parameter results (to within one- or two-hundredths) were obtained with the analytical methodology described by Rudebusch and Svensson (1999). A simulation method is employed in this paper to easily allow the incorporation of various (often nonstandard) types of uncertainty into the model.

estimate from equation (8).⁶ Assuming that $\lambda = 1$ and $\nu = 0.5$, the expected loss is evaluated at 14.05. This is substantially more than the optimal Taylor rule of form (2). As shown in the first line of panel B, the optimal Taylor rule when $\lambda = 1$ and $\nu = 0.5$ has $g_\pi = 2.83$ and $g_y = 1.63$. This much more vigorous rule provides substantially less expected loss through better inflation control. The essence of this paper is an attempt to reconcile these first two lines.

One way to reconcile the historical and optimal results is to consider different weightings of the goal variables in the loss function. The baseline loss function in this paper, with $\lambda = 1$ and $\nu = 0.5$, penalizes equally a 1% output gap, a one percentage point inflation gap, and a 1.4 percentage point quarterly change in the funds rate. These weights seem plausible, but the rest of panel B in table 1 presents optimal Taylor rules for four other sets of preferences over goals. These include two cases in which output variability is much more costly ($\lambda = 4$, $\nu = 0.5$) and much less costly ($\lambda = 0.25$, $\nu = 0.5$) than inflation variability. They also include two cases in which the variability of nominal interest rate changes is not very costly ($\lambda = 1$, $\nu = 0.1$) and is quite costly ($\lambda = 1$, $\nu = 5$). In the first three cases, the resulting optimal Taylor rule parameters are, taken together, no closer to those of the historical estimated Taylor rule than with the baseline loss function. Most importantly, varying the importance of output stabilization relative to inflation stabilization over a wide range does not provide a reconciliation between the optimal and historical rules.

In contrast, as shown in the final line of panel B, with a very strong interest-rate smoothing motive ($\lambda = 1$, $\nu = 5$), the optimal rule does fairly closely match the historical

⁶ The numerical value of π^* has no implications for the results, which are based on second moments in a linear model. On the related issue of a nonnegativity constraint on nominal interest rates, see Rudebusch and Svensson (1999).

one.⁷ This is one possible resolution to the problem addressed in this paper (namely, that the Fed responds modestly to output and inflation gaps because it prefers not to make large changes in the funds rate). This central-bank utility-function resolution is not, however, a completely satisfying one, because the desire for smooth interest rates is difficult to motivate (Rudebusch, 1995; Woodford, 1999). Thus, the subsequent analysis focuses on adding uncertainty to reconcile optimal and historical policy.

One might argue that conducting the analysis of this paper with a Taylor rule might impose a straitjacket that distorts the results. In this regard, it should be noted that the Taylor rule has captured the attention of some monetary policymakers as a useful simple summary rule for policy (Meyer, 1998). Still, the remainder of table 1 examines three different policy rules as a check on the robustness of the results.

The first alternative is the partial adjustment form of the Taylor rule

$$i_t = (1 - \rho)(k + g_\pi \bar{\pi}_t + g_y y_t) + \rho i_{t-1}, \quad (10)$$

which adds the lagged level of the interest rate as a third argument. As shown in panel C, the optimal form of this dynamic Taylor rule also displays relatively high output and inflation response coefficients (the optimal $g_\pi = 2.87$, $g_y = 1.80$, and $\rho = 0.18$). The values of the loss functions for the optimal Taylor rule and dynamic Taylor rule—which are 11.23 and 11.21, respectively—are essentially indistinguishable. Indeed, none of the empirical results in this table or in the subsequent tables in this paper change qualitatively with the use of a dynamic Taylor rule. (Some of these alternative results are given in appendix A.)

Panel D goes further to consider the fully optimal policy rule for this model, which for $\lambda = 1$ and $\nu = 0.5$ has the form

$$\begin{aligned} i_t = & 0.86\pi_t + 0.31\pi_{t-1} + 0.38\pi_{t-2} \\ & + 0.13\pi_{t-3} + 1.34y_t - 0.36y_{t-1} + 0.50i_{t-1} \\ & - 0.06i_{t-2} - 0.03i_{t-3}. \end{aligned}$$

As noted by Rudebusch and Svensson (1999), the Taylor rule (with two arguments) comes close to matching this unrestricted rule (with all nine state variables as arguments) by setting the first four parameters equal to 0.68 (that is, $g_\pi/4$), the y_t parameter equal to 1.57, and the other parameters equal to 0. The value of the loss function for the fully optimal rule is 11.00, only slightly less than for the optimal Taylor rule.

Finally, another common criticism of central-bank rules such as equation (2) is that they assume too much informa-

tion to be plausibly implemented in real time by policymakers. Thus, some have suggested that a policy rule with only lagged information is more appropriate (McCallum, 1998). However, as shown in panel E, the results for the optimal Taylor rule with lagged arguments are little different from those for contemporaneous arguments. Most significantly, the optimal lagged-rule coefficients are no nearer to the historical estimates than are those of the optimal contemporaneous rule. Still, the issue of data uncertainty is reconsidered in section VI.

IV. Optimal Policy with Parameter Uncertainty

Parameter uncertainty is probably the obvious element to include in an attempt to produce an optimal policy rule that matches a cautious historical one. Ever since the classic Brainard (1967) analysis, uncertainty about the quantitative impact of policy and the dynamics of the economy has been widely cited as a rationale for damped policy action. Svensson (1999) and Estrella and Mishkin (1999) provide recent illustrative theoretical analyses that demonstrate the “less activist” nature of optimal policy under model parameter uncertainty. Specifically, in the context of models similar to equation (4) and (5), they show that uncertainty about β_r , which gauges the interest rate sensitivity of the economy, will cause the policymaker to be more cautious. In the general case, though, as Chow (1975, ch. 10) makes clear, almost nothing can be said even qualitatively about how the optimal rule under model uncertainty changes relative to the optimal rule under certainty. For example, the optimal policy-response parameters are not necessarily reduced in the presence of uncertainty about several parameters. Thus, quantitative answers are required; however, until very recently, there has been a surprising lack of empirical analysis on this topic.⁸

A general approach to the calculation of optimal policy with model uncertainty is to calculate the optimal policy separately for each of the possible individual models and then obtain a globally optimal policy by combining these individual optimal policies using the relative model likelihoods. (See the discussion by Stock (1999).) As noted by Brainard (1967), this one calculation can be given two different interpretations. In the first interpretation, there is one true model, but the policymaker is uncertain what it is and has a probability distribution over possible models. In the other interpretation, the actual generating process for the data varies over time according to a probability distribution over possible models. The results in this section are naturally viewed from this second, time-varying model perspective, which has a long history in policy discussions (Fischer and Cooper, 1973). Specifically, in order to determine the optimal Taylor rule with model-parameter uncertainty, I

⁷The loss function with $\lambda = 1$ and $\nu = 5$ implies that the policymaker considers an output gap of one percentage point or an inflation gap of one percentage point as distasteful as a 44 basis point change in the quarterly average funds rate.

⁸Recent analyses available at the time of the first draft of this paper included Estrella and Mishkin (1999) and Sack (1998). The subsequent flood of papers is described below.

TABLE 2.—OPTIMAL TAYLOR RULES WITH PARAMETER UNCERTAINTY

Type of Uncertainty	Rule Parameters		Expected Loss
	g_π	g_y	
Panel A. No Uncertainty			
None	2.83	1.63	11.23
Panel B. Uncertainty about a Single Parameter			
$\sigma^2(\beta_r)$	2.81	1.60	11.42
$4 \cdot \sigma^2(\beta_r)$	2.73	1.57	11.82
$\sigma^2(r^*)$	2.83	1.63	11.28
$4 \cdot \sigma^2(r^*)$	2.83	1.63	12.16
$\sigma^2(\beta_r)$, 16Q	2.87	1.73	12.25
$\sigma^2(r^*)$, 16Q	2.83	1.69	11.44
Panel C. Uncertainty about All Parameters			
Var-Cov	2.77	1.60	11.60
$4 \cdot$ Var-Cov	2.63	1.52	13.33
Var-Cov, 16Q	2.91	1.82	14.20

All parameter uncertainty is normally distributed with a standard error equal to or double the estimated standard error. Parameters change every sixteen quarters in the "16Q" cases and every quarter otherwise.

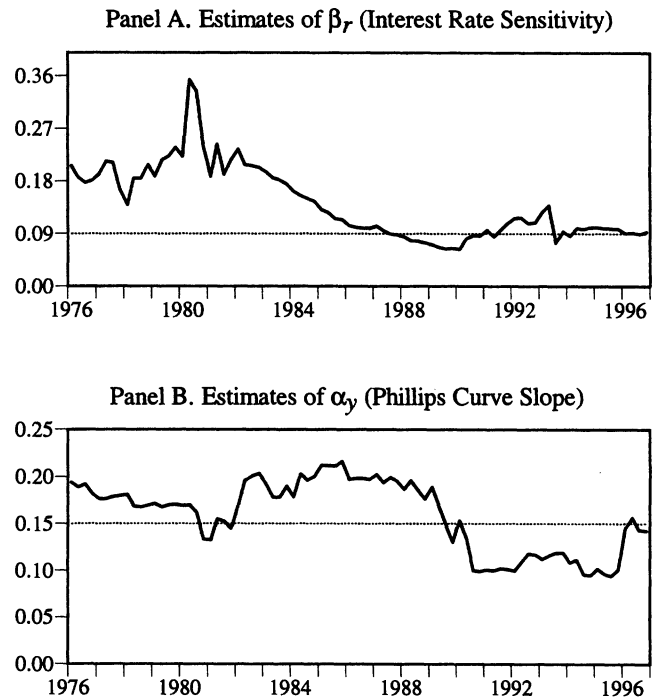
simulate the model with different types of parameter variation over time. I assume that the policymaker faces an economy like equation (6) and (7) on average but that, in any given quarter, the coefficients may take on a random value. The policymaker has to choose the g_π and g_y parameters of the Taylor rule (2) so that the loss function (9) is minimized in the presence of this parameter uncertainty.⁹

The results are shown in table 2. For comparison, the row in panel A displays the results under certainty. Then, panel B provides some initial results assuming uncertainty about only one parameter. The first line considers uncertainty about β_r (in equation (5)), which is a key parameter in the earlier theoretical work described above. I assume that the policymaker faces an economy like equation (6) and (7) except that each quarter the coefficient β_r takes on a different value, which is randomly drawn from a normal distribution with the least-squares estimated mean of 0.09 and estimated standard error of 0.03. (Under the first interpretation given above, this uncertainty can be considered estimation or sampling uncertainty.) Relative to the certainty case, the quantitative effect of this parameter uncertainty is negligible. As shown in the line labeled " $\sigma^2(\beta_r)$," the optimal inflation and output response parameters for this case, $g_\pi = 2.81$ and $g_y = 1.60$, are essentially the same as under certainty. Clearly, sampling uncertainty is not nearly enough, so in the second line of panel B (denoted " $4 \cdot \sigma^2(\beta_r)$ "), I add much more variation by drawing the β_r from a normal distribution with four times the estimated variance. Although in this case there is a more notable reduction in the rule coefficients, it is still economically insignificant.

Uncertainty about another single parameter has also attracted some attention in considering the implementation of the Taylor rule. As noted in the introduction, uncertainty

⁹For my basic simulation methodology, in each period t , one or more new coefficients are drawn to generate y_t and π_t on the basis of lagged data, and i_t is formed according to a candidate Taylor rule. Such a simulation (with 100,000 observations) is repeated over a grid of candidate Taylor rules, and the optimal one is chosen based on the variances.

FIGURE 1.—SEQUENTIAL PARAMETER ESTIMATES FROM ROLLING FIFTEEN-YEAR SAMPLES



Each parameter estimate is obtained from a fifteen-year regression subsample that ends in the quarter in which the parameter estimate is plotted.

about r^* —the equilibrium real rate of interest—seems important for policymakers. If r^* were overestimated by a percentage point, then the interest rate in equation (2) would be set too high by a similar amount (through the constant term). This type of model-parameter uncertainty, however, is benign. As shown in the middle of panel B, period-by-period uncertainty about r^* , at any level of variance, has no effect on the optimal Taylor rule parameters.¹⁰ This is not surprising because theoretically increased uncertainty about r^* in equation (8) is equivalent to increased uncertainty about β_0 in equation (5), which is equivalent to added variation in the error ϵ_t , and altering the variance of the residuals of the structural equations will not change the parameters of the optimal rule.

As another case, I also consider persistent parameter shifts. The most plausible form of parameter uncertainty may not be the quarter-by-quarter, i.i.d. random draw of coefficients, in which a policymaker may face a β_r that is unusually high one quarter and quite low the next. Instead, as a more plausible case, there may be changes in parameter values that endure for some time. In the extreme case, considered by Stock (1999), the policymaker faces one-shot uncertainty and must choose Taylor rule parameters once and for all in the face of an uncertain (and worst-case scenario) β_r that will then be fixed once and for all. A less extreme case is suggested by figure 1, which provides

¹⁰The simulation experiment conducted here assumes that the policymaker uses the estimated mean value of r^* in the Taylor rule but faces an IS curve (8) in which r^* is changed each period.

subsample estimates of β_r and α_y (from sequential estimates of equation (5) and (6)). These rolling coefficient estimates are based on successive fifteen-year subsamples of data, with the estimates shown on the last quarter of their associated estimation subsample. (For example, the α_y estimate of about 0.20 in 1984:Q4 is estimated over a sample from 1970:Q1 to 1984:Q4.) The full-sample estimates are shown as dotted lines. If policymakers use such restricted windows to discern the changing structure of the economy, they may uncover sluggish fluctuations in the parameter estimates. In this spirit, the last two lines of panel B consider a case in which the β_r and r^* are redrawn only every sixteen quarters. Such persistent shifts appear if anything to push the optimal Taylor rule coefficients higher, which is consistent with Stock's (1999) results. In this model, if β_r is quite low, the system borders on dynamic instability. This instability is highlighted when the parameters shift infrequently, because a sustained bad parameter draw could allow the economy to be driven quite far from its targets, so high response coefficients are required to mitigate the instability.¹¹

Finally, panel C in table 2 provides the results under uncertainty about all of the coefficients of the model. In this case, period-by-period parameter uncertainty again leads to some reduction in the optimal Taylor rule coefficients, whereas more-persistent parameter shifts tend to push the optimal rule parameters higher. Overall, however, the changes in the optimal rule in table 2 are quite modest, so this type of model uncertainty does not appear to be a likely rationale for lowering the optimal Taylor rule coefficients within this framework.

This conclusion accords with much, but not all, of the most recent research on parameter uncertainty. The differences among the various studies appear to stem from whether the models and rules employed are tightly or loosely parameterized. In particular, research that uses either a parsimoniously parameterized structural model or policy rule appears to find that parameter uncertainty is not an important source of policy-response attenuation. Notably, in the analysis above—as well as in Estrella and Mishkin (1999) and Peersman and Smets (1999)—with a parsimonious structural model and a simple rule, there is no significant attenuation of the rule parameters. Similarly, no attenuation results are obtained with a parsimonious model and a many-parameter unrestricted rule in Peersman and Smets (1999), Shuetrim and Thompson (1999), and Debelle and Cagliarini (1999). Finally, in other preliminary work, it appears that the converse combination of an unrestricted VAR model with a simple rule also provides no significant

change in the rule coefficients with parameter uncertainty. Indeed, the only research that reports finding some attenuation in policy-response parameter are those studies, such as Sack (2000), Salmon and Martin (1999), and Söderström (1999), that use an unrestricted VAR together with a many-parameter unrestricted policy rule. A plausible explanation for this last set of results is that an optimized unrestricted policy rule tends to have coefficients that are overly attuned to the random idiosyncracies in the estimated coefficients of a VAR. For example, a spuriously marginally significant estimated coefficient on some lag of a particular variable in the VAR likely induces some outlier among the optimal rule parameters. In this case, recognizing the VAR coefficient uncertainty helps damp such rule parameters. Such a result, however, may well be of limited interest to policymakers.

V. Model Perturbation

In the previous section, I examined optimal policy given uncertainty about the future values of model parameters in the context of known parameter first moments. Thus, in section IV, the policymaker knew that the model equations (6) and (7) were true on average on an ex ante basis, but the policymaker was uncertain about the values of the parameters in any particular quarter. This section reexamines optimal policy under various perturbations of the baseline model. Under one interpretation, this is the analog to the variation in preference parameters considered in table 1, and changes to the rule induced by perturbations to the model provide, in essence, an important robustness check on the results in section III. However, there may be good reasons to consider the case in which a policymaker believes that the parameters of the model, even on average, may be different from the least-squares estimates in equation (6) and (7). Such uncertainty would not be too surprising (Söderström, 1999). The model (4) and (5) is quite simple and certainly misspecified—for example, with omitted variables—and it is quite possible that the resulting least-squares estimates are biased. Alternatively, policymakers may have a very different idea of, say, the interest-rate sensitivity of the economy based on more detailed estimated models or other information. To examine the effects of such model uncertainty—or, more precisely, the uncertainty about the mean parameters of a given model—I explore the optimal Taylor rule over a range of parameters for the model (4) and (5).¹²

I perturb the model over four different dimensions: the “slope” of the Phillips curve (α_y in equation (4)), the interest-rate sensitivity of output (β_r in equation (5)), the sum of the own-lag dynamics of output ($\sum_{j=1}^2 \beta_{yj}$ in (5)), and the sum of the own-lag dynamics of inflation ($\sum_{j=1}^4 \alpha_{\pi j}$ in equation (4)). The results are given in table 3. In the first column, Ψ represents, in turn in each of four panels, the

¹¹ Various assumptions about the frequency of real-time variation in coefficients have been made in the literature. For example, Svensson (1999) and Söderström (1999) assume that coefficients shift every period, whereas Sack (2000) essentially redraws coefficients every fourteen years (the size of his sample), and Debelle and Cagliarini (1999) draw coefficients every twelve years. Furthermore, it should be noted that the issue of model estimation in the presence of potentially random coefficients has not been addressed in this literature.

¹² This is similar to the definition of Levin, Wieland, and Williams (1999) who use the term *model uncertainty* to “refer to lack of knowledge about which model among a given set of alternatives provides the best description of the economy” (p. 263).

TABLE 3.—OPTIMAL TAYLOR RULES WITH MEAN PARAMETER UNCERTAINTY

Coefficient Ψ	p-value	Rule Parameters		Expected Loss
		g_π	g_y	
Panel A. $\Psi = \alpha_y$; Phillips curve slope				
0.08	0.09	2.92	1.37	13.55
0.15	0.85	2.83	1.63	11.23
0.22	0.04	2.71	1.92	11.10
Panel B. $\Psi = \beta_r$; Interest rate sensitivity				
0.02	0.05	4.66	1.85	16.08
0.09	1.00	2.83	1.63	11.23
0.15	0.05	2.35	1.40	10.17
0.35	0.00	1.83	0.97	9.49
Panel C. $\Psi = \sum_{j=1}^2 \beta_{yj}$; Output dynamics				
0.85	0.05	2.97	1.29	10.34
0.90	0.90	2.82	1.63	11.18
0.95	0.09	2.68	2.00	12.55
Panel D. $\Psi = \sum_{j=1}^4 \alpha_{\pi j}$; Inflation dynamics				
0.85	0.00	1.37	1.30	6.60
0.90	0.05	1.67	1.37	7.38
0.95	0.53	2.12	1.49	8.78
1.00	0.49	2.83	1.63	11.23

quantity being varied. The p -values for the standard F -statistic restricting Ψ to equal these values are given in the second column; however, given the likely misspecifications noted above, these p -values should be taken only as a rough starting point for assessing the likelihood of various models. In each panel, although Ψ varies by row, the non- Ψ parameters are held at their original least-squares values.

Panel A presents the optimal Taylor rule results with α_y ($\equiv \Psi$), the Phillips curve slope coefficient, set equal to 0.08, 0.15, and 0.22. These values cover a two-standard-deviation range around the least-squares estimate of α_y and a considerable range of economic behavior. However, there is surprisingly little variation in the optimal rule parameters over this range, and, even more serious, the variation in α_y affects the optimal g_π and g_y in opposite directions. Thus, variation or uncertainty about the slope of the Phillips curve apparently will not help match the historical estimates.


Different results are obtained in panel B with uncertainty about the interest-rate sensitivity of output (so $\Psi \equiv \beta_r$). Over a range of plus-or-minus two standard deviations around the least-squares estimate of β_r , none of the optimal rules come close to matching the historical Taylor rule coefficients. However, in contrast to the results in panel A, the optimal g_π and g_y parameters do vary in the same direction. As the interest-rate sensitivity (β_r) increases, a more cautious Taylor rule is optimal. Still, even with $\beta_r = 0.15$, the optimal g_π and g_y are significantly above the historical estimates. Indeed, it would require a β_r at least as high as 0.35—as shown in the last line of panel B—to push the optimal g_π and g_y into the range of the historical estimates. Such a high β_r is not very plausible. The p -values in column 2 suggest that a β_r greater than 0.2 is quite unlikely. Similarly, the β_r coefficient estimates from rolling subsamples in panel A of figure 1 show that, in any fifteen-

year estimation period that ended after the mid-1980s, β_r was close to the full-sample estimate. This evidence is supported by recent empirical estimates of variants of the IS curve (5) by other authors. For example, Clark, Laxton, and Rose (1995, 1996) and McCallum and Nelson (1999) both estimate β_r to be equal to 0.16 (with fairly tight standard errors), although Smets (1999) estimates β_r to be equal to 0.06. Although central bankers and many economists often appear to believe that interest-rate elasticities are higher than the data might suggest (say for business investment), a β_r as high as 0.35 seems implausibly high. Thus, model structure uncertainty regarding the interest-rate sensitivity of output alone is likely not enough to reconcile the optimal and historical policies.

The other important coefficients in the model concern dynamics. Panel C examines varying the sum of the own-lag parameters for output, $\Psi \equiv \sum_{j=1}^2 \beta_{yj}$, over the values 0.85, 0.90, and 0.95. The least-squares estimate of this sum is close to 0.90, with lower (higher) values of this sum representing alternative less (more) persistent processes. This range about covers the sequential point estimates of the lag sum in equation (7)—shown in figure 2. As for α_y , variation in output dynamics does not help reconcile the optimal and historical Taylor rule coefficients. Although reduced persistence in the output process is able to lower the optimal g_y , the optimal g_π is pushed in the opposite direction.

Panel D examines variation in the sum of the own-lag coefficients for inflation. The sum of these coefficients, $\Psi \equiv \sum_{j=1}^4 \alpha_{\pi j}$, ranges over the values 0.85, 0.90, 0.95, and

FIGURE 2.—SEQUENTIAL PARAMETER SUM ESTIMATES FROM ROLLING FIFTEEN-YEAR SAMPLES



Each parameter estimate is obtained from a fifteen-year regression subsample that ends in the quarter in which the parameter estimate is plotted.

1.0. In the least-squares estimates in equation (6) and (7), this sum was constrained to equal 1.0, in part because the unconstrained point estimate is 0.97 and there may be downward bias to the least-squares estimate (Rudebusch, 1992, 1993). Furthermore, a unitary sum is clearly consistent with the Natural Rate Hypothesis and a long-run vertical Phillips curve. However, if agents are forward-looking in setting prices, then the sum of the lagged terms in an estimated Phillips curve with no long-run tradeoff may be less than 1 (Sargent, 1971). Smets (1999), for example, estimates a close variant of equation (4) and obtains a point estimate of the inflation lag sum of 0.84. Other empirical estimates are consistent with a forward-looking element to the inflation process (Fuhrer, 1997; Clark et al., 1995, 1996). If this is the case and agents think there is a constant inflation target, then the appropriate backward-looking representation of inflation does not contain a unit root. Historically, of course, over the estimation sample, it seems unlikely that agents in the U.S. had a constant inflation target in mind, so a slow reversion to a changing mean (or target) is more likely (as in Kozicki and Tinsley (1998)). Equation (6) may capture this persistence as a unit-root process. But looking forward, policymakers in the late 1980s and the 1990s may have believed that they had credibly conveyed their commitment to a constant inflation target to the populace. (Indeed, as shown in the lower panel of figure 2, the estimated sum of inflation lags for the fifteen-year subsample from 1982 to 1996 is only about 0.75.) In this case, a less persistent process than equation (6) would be appropriate for calculating optimal policy.¹³

In any case, lowering the sum of the inflation lags, even by a little, does reduce the optimal Taylor rule coefficients and particularly the g_π parameter. Indeed, the results are quite dramatic, as lowering the sum from 1.0 to 0.9 reduces the g_π parameter from 2.83 to 1.67 and the g_y parameter from 1.63 to 1.37. The intuition for this result is clear: Given some, even small, tendency for inflation to revert to a target, policymakers do not have to react as strongly to deviations from inflation or, given the less persistent Phillips curve implications, to output gaps. Still, modifying the dynamics of inflation alone does not appear to give a small enough output response parameter to match the historical Taylor rule.

VI. Optimal Policy with Data Uncertainty

The Taylor rule sets the interest rate in quarter t on the basis of output and inflation in quarter t ; however, as noted in the introduction, real-time uncertainty about the output

gap seems to loom large for policymakers.¹⁴ As noted in subsection IIIB, merely lagging the rule arguments to capture real-time data uncertainty was not successful. Furthermore, the lagged Taylor rule is likely no better an approximation to the real-time information set than the contemporaneous Taylor rule. The real-time data have two important features. First, in real time, the policymaker does have a large amount of information about the current-quarter state of the economy by the way of monthly, or even weekly, statistics on production, employment, spending, and prices. Second, the lagged data cannot be treated as flawless, for they are subject to extensive revisions even after several quarters (or years). Thus, it is unlikely that simply lagging the variables in the policy rule—but still using the final revised data available many years after the fact—will capture the real-time policy calculation.

Instead, it seems more appropriate to model data uncertainty directly. Namely, given the true model (6) and (7), the policymaker may use a real-time Taylor rule of the form

$$i_t = k + g_\pi \bar{\pi}_{t|t} + g_y y_{t|t}, \quad (11)$$

in which the arguments of the rule are explicitly the inflation and output gaps in time t that are estimated by the policymaker as of time t .¹⁵ These noisy data arguments are assumed to be related to the true series that are generated by the model (6) and (7) as

$$\bar{\pi}_{t|t} = \bar{\pi}_t + n_t^\pi \quad (12)$$

and

$$y_{t|t} = y_t + n_t^y, \quad (13)$$

where n_t^π and n_t^y are the contemporaneous measurement errors that plague the policymaker in real time, with standard errors σ_{n^π} and σ_{n^y} , respectively, and uncorrelated with $\bar{\pi}_t$ and y_t . The error n_t^π contains data revisions stemming from the use of initial and early estimates of recent quarterly inflation rates (that is, $\pi_{t-1|t}$, $\pi_{t-2|t}$, and $\pi_{t-3|t}$) in the four-quarter average, as well as the “forecast” or estimation error involved with $\pi_{t|t}$. The error n_t^y is based on estimation errors in assessing contemporaneous actual and potential output, $Q_{t|t}$, and $Q_{t|t}^*$.

The appropriate sizes of σ_{n^π} and σ_{n^y} (the standard deviations of the measurement errors) can be obtained by comparing historical estimates of the inflation rate and the output gap to the final estimates as they stand today. Figure 3, for example, plots the time series of actual estimates of $\bar{\pi}_{t|t}$ as they were made in real time by the Federal Reserve

¹⁴For a criticism of VAR reaction functions along these lines, see Rudebusch (1998).

¹⁵See Diebold and Rudebusch (1991a, 1991b), Rudebusch (1998), Judd and Rudebusch (1998), Clarida, Gali, and Gertler (2000), Orphanides (1998), and Wieland (1998) for discussions of real-time information sets.

¹³Also, note that policymakers are not trying to exploit the Phillips curve through changes in inflation but are committed to a constant inflation target.

FIGURE 3.—REAL-TIME AND FINAL ESTIMATES OF INFLATION





FIGURE 4.—REAL-TIME AND FINAL ESTIMATES OF OUTPUT GAP



Board staff as well as the current estimate $\bar{\pi}_t$.¹⁶ The associated time series of measurement errors has a standard

error, $\sigma_{n\pi}$, equal to 0.34. Similarly, figure 4 plots the available sample of output gap real-time estimates, $y_{t|t}$, and

¹⁶These real-time estimates are Federal Reserve staff Greenbook forecasts that were presented at mid-quarter FOMC meetings (typically in February, May, August, and November). These data are released publicly with a five-year lag, so the sample ends in 1993. For figure 3 only, inflation

rates were constructed with the output price deflator rather than the chain-price index, but this difference is negligible.

TABLE 4.—CORRELATIONS BETWEEN REVISIONS AND FINAL AND REAL-TIME DATA

Revision	Final Data	Real-Time Data
n_t^y	y_t	$y_{t t}$
	0.40	0.74
	(0.05)	(0.00)
	[0.09]	[0.00]
n_t^π	$\bar{\pi}_t$	$\bar{\pi}_{t t}$
	0.27	0.41
	(0.02)	(0.00)
	[0.09]	[0.01]

OLS p -values for the significance of the correlation coefficients are given in parentheses, whereas the p -values in brackets use a GMM autocorrelation correction.

final estimates, y_t .¹⁷ Not surprisingly, the output gap measurement errors have a relatively larger standard error, σ_{ny} , of 0.94.¹⁸

With such data uncertainty, the question remains as to whether the optimal g_π and g_y chosen for a system consisting of the model (6) and (7), the reaction function (2), and the measurement equation (12) and (13) can match the historical estimates of the Taylor rule. This would seem unlikely at first glance, for a general guideline that economic analysis gives to policymakers is the principle of certainty equivalence, which states that optimal policy requires the same response when there is only partial information about the state of the economy as when there is full information (Svensson and Woodford, 2000). However, there are two exceptions to the principle of certainty equivalence that operate here. First, as shown by Smets (1999) and Peersman and Smets (1999), the coefficients of simple policy rules that respond to only a subset of the state variables do change with partial information. In particular, the use of the Taylor rule above rather than the unrestricted fully optimal rule eliminates certainty equivalence. Second, in the usual certainty-equivalence case, the policymaker reacts in the same given fashion to an *efficient* estimate of the unknown output gap. The real-time estimates written above contain a element of inefficient noise relative to the final estimates, and, in this case, certainty equivalence will also fail (Swanson, 2000). Thus, an important empirical issue to be decided is whether the revisions n_t^π and n_t^y are correlated with the final estimates or with the real-time estimates. The first type of correlation implies that the initial estimate is an efficient forecast (revisions are “news”), whereas the second type implies that they are just noisy final estimates (revisions are “noise”). (See Mankiw and Shapiro (1986) and Rudebusch (1998).) As shown in table 4, the correlations in the data favor a substantial noise element,

¹⁷ The $y_{t|t}$ data, again shown with the five-year release lag, were supplied by Athanasios Orphanides, and y_t is defined as above.

¹⁸ The sample of real-time data that is available is fairly short; however, the size of the output gap measurement errors are broadly consistent with those from more mechanical estimates. For example, using the Kalman filter to roll through a sample of final revised data, Smets (1999) estimates a σ_{ny} of about 1.1.

TABLE 5.—OPTIMAL TAYLOR RULES WITH DATA UNCERTAINTY

Data Noise		Rule Parameters		Expected Loss
$\sigma_{n\pi}$	σ_{ny}	g_π	g_y	
Panel A. Period-by-period i.i.d. data uncertainty				
0.0	0.0	2.83	1.63	11.23
0.0	1.0	2.50	1.02	12.92
0.3	0.0	2.63	1.54	11.91
0.3	1.0	2.39	1.00	13.44
Panel B. Time series data uncertainty; $n_t^y \sim AR(1)$ and/or $n_t^\pi \sim MA(3)$				
0.0	1.27	2.59	1.16	12.66
0.39	0.0	2.64	1.55	11.87
0.39	1.27	2.47	1.13	13.15


and the revision process is modeled as equation (12) and (13) with i.i.d. revisions added to the final data.

Panel A in table 5 gives the coefficients for the optimal Taylor rules under various assumptions about the degree of data uncertainty. Assuming $n_t^\pi \sim N(0, \sigma_{n\pi}^2)$ and $n_t^y \sim N(0, \sigma_{ny}^2)$, four $(\sigma_{ny}, \sigma_{n\pi})$ pairs are considered. For the case with no uncertainty (the (0, 0) pair), the optimal Taylor rule coefficients match those in table 1. Increasing the amount of data uncertainty reduces the optimal Taylor rule coefficients in an intuitive fashion. Namely, greater output-gap (inflation) uncertainty especially reduces the output (inflation) response coefficient. However, a major shortcoming of the results in panel A is that they assume i.i.d. measurement errors quarter by quarter. For example, in quarter t , the contemporaneous estimate of the output gap may be a percentage point too high ($y_{t|t} - y_t = 1$), whereas, in the next quarter, the gap may be contemporaneously estimated to be a half of a percentage point too low ($y_{t|t} - y_t = -0.5$). This quarter-by-quarter volatility in the measurement errors is unlikely; instead, the data-measurement errors appear to be quite persistent over time.¹⁹ As shown in figure 4, this persistence is particularly evident for the output gap. This is not surprising, because the main source of uncertainty about the output gap is determining the trend growth of potential output, and errors in setting this trend will affect the level of the gap in adjacent quarters.²⁰ Another example of persistently noisy real-time output-gap estimates is provided by real-time and final revised capacity utilization data. Figure 5 shows a long sample of the Federal

¹⁹ Note that these data-measurement errors are not one-step-ahead forecast errors. In particular, the measurement errors are not being judged relative to the contemporaneous, quarter- t information set but to the end-of-sample information set. Thus, their serial correlation is not a violation of rationality. However, to the extent that such measurement errors are the basis for monetary policy shocks in a VAR context (as hypothesized by Rudebusch, 1998), then the monetary VAR approach, which assumes serially uncorrelated shocks, will be incorrect. In the same vein, much of the lagged dynamics in the historical dynamic Taylor rule estimates may reflect serially correlated data-measurement errors rather than sluggish Fed responses.

²⁰ Equivalently, errors tend to persist through time about the level of the unemployment gap. For example, in the mid-1990s, there was only a gradual change in the consensus estimate of the natural rate from roughly 6.5% to 5.5%, so the measurement error in contemporaneous unemployment gap estimates persisted for several years.

FIGURE 5.—REAL-TIME AND FINAL ESTIMATES OF CAPACITY UTILIZATION



Reserve Board's contemporaneous estimates of manufacturing capacity utilization, CU_{it} , and the final revised series, CU_t (as of 1998).²¹ Again, the measurement errors for this alternative measure of the output gap are quite persistent.

This persistence can be modeled with a time-series process. For example, the estimated first-order autoregression (AR(1)),

$$n_t^y = 0.75 n_{t-1}^y + \varepsilon_t^y, \quad \sigma_{\varepsilon^y} = 0.838, \quad (14)$$

gives a good fit to the historical errors in the output gap series in figure 4.²² (Although the output gap sample is quite short, an AR(1) model estimated with the capacity utilization measurement error series from figure 5 has a very similar estimated coefficient of 0.71.) This is a fairly persistent process. The results from using the AR(1) process (14) to generate output gap measurement errors and assuming no inflation measurement errors (so $\sigma_{n\pi} = 0.0$) are shown in the first line of panel B of table 5. Given equation (14), the unconditional volatility of the output measurement error (that is, $\text{Std}[n_t^y]$ or σ_{n^y}) is equal to 1.27. Accordingly, the AR(1) results are consistent with the results in the second line in panel A.²³

²¹ These data were kindly provided by Evan Koenig (1996). In this figure, the quarterly series CU_{it} is defined as the average of the initial release of the second month and the first revision of the first month of each quarter. These data were available in the third month of the same quarter.

²² For this regression and the following, Q -statistics suggested no residual serial correlation, and a t test suggested that a constant was insignificant.

²³ Persistence in output gap errors has two effects. First, a policymaker is likely to lower g_y because of the fear that the error may be magnified

The measurement error between the real-time and final estimates of four-quarter inflation, $\bar{\pi}_{it}$ and $\bar{\pi}_t$, is naturally modeled as a third-order moving average (MA(3)) process, which is estimated as

$$n_t^\pi = \varepsilon_t^\pi + 0.63 \varepsilon_{t-1}^\pi + 0.26 \varepsilon_{t-2}^\pi + 0.18 \varepsilon_{t-3}^\pi, \quad (15)$$

(0.12) (0.14) (0.13)

$$\sigma_{\varepsilon^\pi} = 0.320.$$

The results from using just this MA(3) model for inflation-measurement errors (with an implied $\text{Std}[n_t^\pi]$, or σ_{n^π} , equal to 0.39) are shown in the second line of panel B in table 5. The results are similar to those in panel A.

Finally, the last line in the table includes both the AR(1) and the MA(3) time-series representations. Comparing the top line (with no uncertainty) and the final line (the most plausible form of real-time measurement error), it appears that data uncertainty alone is not able to completely reconcile optimal and historical Taylor rule estimates. Accounting for data uncertainty is able to reduce the optimal Taylor rule parameters substantially, but plausible amounts of such uncertainty leave the response coefficients—especially g_π —above their historical estimates.

VII. Model Perturbation and Data Uncertainty

Although no single exercise above appears plausibly able to reconcile an optimal Taylor rule with the cautious histor-

through the dynamics of the model. (In contrast, a high-frequency measurement error may simply wash out.) Second, a persistent measurement error implies a much smaller variance for the change in the real-time output gap estimate; hence interest rates will be smoother.

TABLE 6.—OPTIMAL TAYLOR RULES WITH MODEL AND DATA UNCERTAINTY

Rule Parameters		Expected Loss
g_π	g_y	
Panel A. $\sum_{i=1}^4 \alpha_{\pi i} = 0.9$; $n_t^y \sim AR(1)$; $n_t^\pi \sim MA(3)$		
1.55	0.90	8.35
Panel B. $\beta_r = 0.25$; $n_t^y \sim AR(1)$; $n_t^\pi \sim MA(3)$		
1.80	0.83	11.94
Panel C. $\beta_r = 0.15$; $\sum_{i=1}^4 \alpha_{\pi i} = 0.95$; $n_t^y \sim AR(1)$; $n_t^\pi \sim MA(3)$		
1.66	0.88	9.44

ical one, combining the various types would appear to be a promising alternative. Perhaps the most obvious combination is data uncertainty, which delivered a low output-response parameter, along with less persistent inflation dynamics, which delivered a low inflation-response parameter; that is, $n_t^y \sim AR(1)$ and $n_t^\pi \sim MA(3)$ and $\sum_{j=1}^4 \alpha_{\pi j} = 0.90$. This combination, shown in panel A of table 6, delivers a $g_\pi = 1.55$ and $g_y = 0.90$, which are well within the range of historical estimates of the policy rule. Finally, then, this provides an optimal rule that matches actual Fed policy.

Table 6 also provides two other cases. In panel B, an elevated interest-rate sensitivity ($\beta_r = 0.25$)—and one that is still probably outside the bounds of what is likely—is examined with data uncertainty. This combination also provides a fairly cautious optimal Taylor rule, with $g_\pi = 1.80$ and $g_y = 0.83$. Finally, in panel C, data uncertainty is combined with a slightly elevated interest-rate sensitivity ($\beta_r = 0.15$) and a slowly mean-reverting inflation process ($\sum_{j=1}^4 \alpha_{\pi j} = 0.95$). This combination also provides an optimal Taylor rule well within the historical range.

VIII. Conclusion

This paper can be summarized with three statements. First, data uncertainty—particularly about the output gap—matters for policymakers. Second, the specification of the model structure and coefficients—particularly about the interest-rate sensitivity of output and the inflation dynamics—also matters. Finally, parameter uncertainty—even about the value of the equilibrium real rate—does not matter. The first two of these statements likely accord with the views of policymakers (see the introduction), while the third does not. Each of these results is important and needs to be examined in further work.

As a guide to further work, it is useful to recognize five distinct elements in the above analysis: a historical Fed reaction function, a Fed objective function, a structural model of the economy, data uncertainty, and parameter or model uncertainty. These five elements encompass a broad range of topics worthy of future investigation, and I consider each in turn.

There is of course a long history of research on the reaction function. (See Rudebusch (1998).) The Taylor rule is a workable approximation, but recent research has sug-

gested that the actual rule may be forward-looking and dynamic (for example, Clarida et al. (2000)). A forward-looking rule, however, can be rewritten in terms of a projection on lagged state variables, and the analysis above (and Rudebusch and Svensson (1999)) suggests that the Taylor rule does quite well in this regard. (Simple Taylor-type rules appear to perform well in a variety of models, too (Levin, Wieland, and Williams, 1999).) The apparent sluggish dynamics or inertia of Fed responses could reflect central-bank preferences for smooth rates or be an optimal response to the structure of the economy. (See Sack (1998) and Woodford (1999).) However, these explanations suggest the presence of predictable movements in the funds rate at multi-quarter-ahead horizons that may contradict the evidence in financial markets (Rudebusch, 1998, 2000b). Alternatively, the lagged interest rate may just be a proxy for serially correlated shocks (perhaps of the measurement-error variety as noted above) or serially correlated movements in r^* . Certainly, the historical monetary policy rule deserves further examination.

A related issue concerns the weight applied to interest-rate smoothing in the Fed objective function. As noted above, given a policy environment, it is possible to reverse-engineer the preferences of the policymaker. (See Favero and Rovelli (1999).) According to the above analysis, it may be that the Fed acts timidly because interest-rate smoothing is a very important direct consideration in the Fed's loss function. If this is so, then motivating and understanding such preferences should be a priority.

The structural model used above has certain advantages and has proved a popular vehicle for research (see Onatski and Stock (2000) and Smets (1999)), but, of course, it is only a starting point. A different model might produce different results, although, as noted in section V, the likely important two directions for model perturbations are altering the sensitivity of output to interest rates and the inflation dynamics. Such perturbations could be obtained in a forward-looking model (Woodford, 1999), or in a VAR, or with additional structural detail such as an exchange-rate channel (DeBelle and Cagliarini, 1999). The explanation for the Fed's behavior pursued in this paper is that uncertainty is the source of timidity. As a general statement, this is unsurprising; however, it should be noted that, in the analysis above, the typical conventional wisdom does not hold true. In particular, simple Brainard-style multiplicative parameter uncertainty is not an important source of cautious behavior. As noted above, this result accords with several other studies that have used small structural models (such as Estrella and Mishkin (1999)).²⁴ Still, the manner in which parameter or model uncertainty unfolds through time de-

²⁴In contrast, studies that employ unrestricted VARs and unrestricted policy rules suggest that parameter uncertainty has a larger effect on the optimal rule coefficients (Sack, 2000). As noted above, this result likely reflects the interaction of the small-sample estimates of the many economically superfluous variables included in the VAR with a many-parameter optimal unrestricted rule.

serves further analysis. In particular, as noted above, there is an important remaining issue as to whether the parameters are best viewed as changing every period or as fixed for some period of time between changes.

Finally, modeling the real-time data set of policymakers was a decisive factor in determining the optimal reaction coefficients. Smets (1999) and Orphanides (1998) provide similar analyses by examining data uncertainty in the context of a Taylor rule and (slight variants) of the Rudebusch-Svensson (1999) model.²⁵ Smets also finds policy-parameter attenuation for the Taylor rule with a different form of data revision (an efficient or "news" form), whereas Orphanides does not consider optimal rules. In related work, Rudebusch (2000a) provides a broader investigation of the optimal rule by comparing the Taylor rule against various nominal income rules; however, the nominal income rules are dominated by the Taylor rule even with substantial output-gap uncertainty. Also, Drew and Hunt (1999) and Lansing (2000) examine the evolution of real-time output gap information. Still, much more research is necessary in this area.

REFERENCES

- Andrews, Donald W. K., "Tests for Parameter Instability and Structural Change with Unknown Change Point," *Econometrica* 61 (1993), 821–856.
- Blinder, Alan S., *Central Banking in Theory and Practice* (Cambridge, MA: MIT Press, 1998).
- Bomfim, Antulio N., and Glenn D. Rudebusch, "Opportunistic and Deliberate Disinflation Under Imperfect Credibility," *Journal of Money, Credit and Banking* 32 (2000), 707–721.
- Brainard, William, "Uncertainty and the Effectiveness of Policy," *American Economic Review Papers and Proceedings* 57 (1967), 411–425.
- Chow, G. C., *Analysis and Control of Dynamic Economic Systems* (New York: John Wiley and Sons, 1975).
- Clarida, Richard, Jordi Gali, and Mark Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* 115 (2000), 147–180.
- Clark, Peter, Douglas Laxton, and David Rose, "Capacity Constraints, Inflation and the Transmission Mechanism: Forward-Looking Versus Myopic Policy Rules," International Monetary Fund working paper 95/75 (1995).
- , "Asymmetry in the U.S. Output-Inflation Nexus," *IMF Staff Papers* 43 (1996), 216–251.
- Congressional Budget Office, "CBO's Method for Estimating Potential Output," CBO Memorandum (October 1995).
- Debelle, Guy, and Adam Cagliarini, "The Effect on Uncertainty on Monetary Policy: How Good are the Brakes?" Reserve Bank of Australia manuscript (1999).
- Diebold, Francis X., and Glenn D. Rudebusch, "Forecasting Output With the Composite Leading Index: A Real-Time Analysis," *Journal of the American Statistical Association* 86 (1991a), 603–610.
- , "Turning Point Prediction with the Composite Leading Index: An Ex Ante Analysis" (pp. 231–256), in K. Lahiri and G. H. Moore (Eds.), *Leading Economic Indicators: New Approaches and Forecasting Records* (Cambridge, MA: Cambridge University Press, 1991b).
- Drew, Aaron, and Benjamin Hunt, "The Effects of Potential Output Uncertainty on the Performance of Simple Policy Rules" (pp. 73–102), in Benjamin Hunt and Adrian Orr (Eds.), *Monetary Policy Under Uncertainty* (Wellington, New Zealand: Reserve Bank of New Zealand, 1999).
- Estrella, Arturo, and Frederic S. Mishkin, "Rethinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty" (pp. 405–430), in John B. Taylor (Ed.), *Monetary Policy Rules* (Chicago: Chicago University Press, 1999).
- Favero, Carlo, and Riccardo Rovelli, "Modeling and Identifying Central Banks' Preferences," University of Bocconi manuscript (1999).
- Fischer, Stanley, and J. Phillip Cooper, "Stabilization Policy and Lags," *Journal of Political Economy* (1973), 847–877.
- Fuhrer, Jeffrey C., "The (Un)Importance of Forward-Looking Behavior in Price Specifications," *Journal of Money, Credit, and Banking* (August 1997), 338–350.
- Judd, John P., and Glenn D. Rudebusch, "Taylor's Rule and the Fed: 1970–1997," *Economic Review* 3 (1998), Federal Reserve Bank of San Francisco (1998), 3–16.
- Kohn, Donald, "Comments on 'Forward-Looking Rules for Monetary Policy'" (pp. 192–199), in John B. Taylor (Ed.), *Monetary Policy Rules* (Chicago: Chicago University Press, 1999).
- Koenig, Evan, "Capacity Utilization As a Real-Time Predictor of Manufacturing Output," *Economic Review*, Federal Reserve Bank of Dallas (third quarter 1996), 16–23.
- Kozicki, Sharon, "How Useful Are Taylor Rules for Monetary Policy?," *Economic Review* Federal Reserve Bank of Kansas City (second quarter 1999), 5–33.
- Kozicki, Sharon, and Peter A. Tinsley, "Moving Endpoints and the Internal Consistency of Agents' Ex Ante Forecasts," *Computational Economics* 11 (1998), 21–40.
- Lansing, Kevin, "Learning about a Shift in Trend Output: Implications for Monetary Policy and Inflation," Federal Reserve Bank of San Francisco manuscript (2000).
- Levin, Andrew, Volker Wieland, and John C. Williams, "Robustness of Simple Monetary Policy Rules under Model Uncertainty" (pp. 263–299), in John B. Taylor (Ed.), *Monetary Policy Rules* (Chicago: Chicago University Press, 1999).
- Mankiw, Greg, and Matthew Shapiro, "News or Noise: An Analysis of GNP Revisions," *Survey of Current Business* 66 (1986), 20–25.
- McCallum, Bennett T., "Issues in the Design of Monetary Policy Rules" (pp. 1483–1530), in John B. Taylor and Michael Woodford (Eds.), *Handbook of Macroeconomics* (New York: North-Holland, 1998).
- McCallum, Bennett T., and Edward Nelson, "Performance of Operational Policy Rules in an Estimated Semi-Classical Structural Model" (pp. 15–45), in John B. Taylor (Ed.), *Monetary Policy Rules* (Chicago: Chicago University Press, 1999).
- Meyer, Laurence H., "The Economic Outlook and Challenges for Monetary Policy," remarks at the Charlotte Economics Club, Charlotte, North Carolina (January 16, 1997).
- , "The Strategy of Monetary Policy," The Alan R. Holmes Lecture, Middlebury College, Middlebury, Vermont (March 16, 1998).
- Onatski, Alexei, and James H. Stock, "Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy," NBER working paper 7490 (2000).
- Oliner, Stephen D., Glenn D. Rudebusch, and Daniel Sichel, "The Lucas Critique Revisited: Assessing the Stability of Empirical Euler Equations for Investment," *Journal of Econometrics* 70 (1996), 291–316.
- Orphanides, Athanasio, "Monetary Policy Evaluation with Noisy Information," FEDS working paper 1998-50, Federal Reserve Board (1998).
- Peersman, Gert, and Frank Smets, "The Taylor Rule: A Useful Monetary Policy Benchmark for the Euro Area?," *International Finance* 2 (1999), 85–116.
- Rudebusch, Glenn D., "Trends and Random Walks in Macroeconomic Time Series: A Reexamination," *International Economic Review* 33 (1992), 661–680.
- , "The Uncertain Unit Root in Real GNP," *American Economic Review* 83 (1993), 264–272.
- , "Federal Reserve Interest Rate Targeting, Rational Expectations, and the Term Structure," *Journal of Monetary Economics* 35 (1995), 245–274.
- , "Do Measures of Monetary Policy in a VAR Make Sense?," *International Economic Review* 39:4 (1998), 907–931.
- , "Assessing Nominal Income Rules for Monetary Policy with Model and Data Uncertainty," Federal Reserve Bank of San Francisco working paper 2000-03 (2000a).

²⁵ Initial drafts of Smets (1999), Orphanides (1998), and the analysis above were all independently circulated during the summer of 1998.

- , "Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia," Federal Reserve Bank of San Francisco manuscript (2000b).
- Rudebusch, Glenn D., and Lars E. O. Svensson, "Policy Rules for Inflation Targeting" (pp. 203–246), in John B. Taylor (Ed.), *Monetary Policy Rules* (Chicago: Chicago University Press, 1999).
- , "Eurosystem Monetary Targeting: Lessons from U.S. Data," Federal Reserve Bank of San Francisco manuscript (2000).
- Sack, Brian, "Does the Fed Act Gradually? A VAR Analysis," *Journal of Monetary Economics* 46 (2000), 229–256.
- Salmon, Chris, and Ben Martin, "Should Uncertain Monetary Policymakers Do Less?" (pp. 226–256), in Benjamin Hunt and Adrian Orr (Eds.), *Monetary Policy Under Uncertainty* (Wellington, New Zealand: Reserve Bank of New Zealand, 1999).
- Sargent, Thomas J., "A Note on the Accelerationist Controversy," *Journal of Money, Credit, and Banking* 8:3 (1971), 50–60.
- Shuetrim, Geoffrey, and Christopher Thompson, "The Implications of Uncertainty for Monetary Policy" (pp. 259–292), in Benjamin Hunt and Adrian Orr (Eds.), *Monetary Policy Under Uncertainty*, (Wellington, New Zealand: Reserve Bank of New Zealand, 1999).
- Smets, Frank, "Output Gap Uncertainty: Does It Matter for the Taylor Rule?" (pp. 10–29), in Benjamin Hunt and Adrian Orr (Eds.), *Monetary Policy Under Uncertainty* (Wellington, New Zealand: Reserve Bank of New Zealand, 1999).
- Söderström, Ulf, "Should Central Banks Be More Aggressive?," Central Bank of Sweden manuscript (1999).
- Stock, James H., "Comment on 'Policy Rules for Inflation Targeting'" (pp. 253–259), in John B. Taylor (Ed.), *Monetary Policy Rules* (Chicago: Chicago University Press, 1999).
- Svensson, Lars E. O., "Inflation Targeting: Some Extensions," *Scandinavian Journal of Economics* 101 (1999), 337–361.
- Svensson, Lars, and Michael Woodford, "Indicator Variables for Monetary Policy," Princeton University manuscript (2000).
- Swanson, Eric T., "On Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rules," Federal Reserve Board manuscript (2000).
- Taylor, John B., "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39 (1993), 195–214.
- , "An Historical Analysis of Monetary Policy Rules" (pp. 319–341), in John B. Taylor (Ed.), *Monetary Policy Rules* (Chicago: Chicago University Press, 1999).
- Wieland, Volker, "Monetary Policy and Uncertainty About the Natural Unemployment Rate," FEDS working paper 1998-22, Federal Reserve Board (1998).
- Woodford, Michael, "Optimal Monetary Policy Inertia," NBER working paper 7261 (1999).

TABLE A1.—OPTIMAL PARTIAL ADJUSTMENT TAYLOR RULES

Type of Uncertainty	Rule Parameters			Expected Loss
	g_{π}	g_y	ρ	
None	2.87	1.80	0.18	11.21
A. Uncertainty About a Single Parameter				
$\sigma^2(\beta_r)$	2.85	1.74	0.17	11.29
$4 \cdot \sigma^2(\beta_r)$	2.77	1.67	0.14	11.78
B. Uncertainty About All Parameters				
Var-Cov	2.82	1.72	0.16	11.54
$4 \cdot \text{Var-Cov}$	2.65	1.59	0.11	13.30

APPENDIX A:

Results with a Partial-Adjustment Taylor Rule

As noted in the text, a popular alternative rule is the partial-adjustment form of the Taylor rule:

$$i_t = (1 - \rho)(k + g_{\pi}\bar{\pi}_t + g_y y_t) + \rho i_{t-1}. \quad (16)$$

This form provides a natural metric through the size of the ρ parameter to measure "interest-rate smoothing." Sack (2000) argues that such smoothing increases with parameter uncertainty. As shown in table A1, this is not the case here. With no uncertainty, the optimal ρ is 0.18. The second row adds the estimated uncertainty about a single coefficient, β_r , with little effect. The first row in panel B adds the estimated uncertainty about all the model coefficients, and again, as with the simple Taylor rule, there is only a modest attenuation of the optimal rule parameters; furthermore, there is no indication of increased interest-rate smoothing. In this framework, adding coefficient uncertainty does not significantly change the optimal value of the lagged interest-rate parameter or the other rule parameters.