Optimal Interest-Rate Rules: I. General Theory *

Marc P. Giannoni
Columbia University

Michael Woodford
Princeton University

September 9, 2002

Abstract

This paper proposes a general method for deriving an optimal monetary policy rule in the context of dynamic linear(ized) rational-expectations models with a quadratic objective function. A commitment to a policy rule of the type proposed here has several desirable properties. It results in a determinate rational-expectations equilibrium, in which the responses to shocks are the same as in an optimal state-contingent plan chosen subject to the constraints implied by rational-expectations equilibrium. Furthermore, the proposed policy rule is completely independent of the specification of the disturbance processes in one's model of the economy, and it is formulated entirely in terms of the observed and projected paths of variables that enter the policymaker's objective function. Finally, the proposed rules can be justified from the "timeless perspective" proposed by Woodford (1999b), so that commitment to such a rule need not imply time-inconsistent decision-making.

We show that under quite general conditions, optimal policy can be represented by a form of generalized Taylor rule, in which however the relation between the interest-rate instrument and the paths of the other target variables is not purely contemporaneous, as in Taylor's specification. We also offer general conditions under which optimal policy can be represented by a "super-inertial" interest-rate rule (in the sense of Rotemberg and Woodford, 1999), and under which it can be represented by a pure "targeting rule" that makes no explicit reference to the path of the instrument.

*An earlier version of this work was delivered by the second author as the Jacob Marschak Lecture at the 2001 Far Eastern Meeting of the Econometric Society, Kobe, Japan, July 21, 2001. We thank Ed Nelson, Julio Rotemberg and Lars Svensson for helpful discussions, and the National Science Foundation, through a grant to the NBER, for research support.

Optimal Interest-Rate Rules: II. Applications *

Marc P. Giannoni
Columbia University

Michael Woodford
Princeton University

August 25, 2002

Abstract

We calculate optimal monetary policy rules for several variants of a simple optimizing model of the monetary transmission mechanism with sticky prices and/or wages. In accordance with the general method expounded in Giannoni and Woodford (2002), our rules have the property that a commitment to ensure that the rule is satisfied at all times results in a determinate rational-expectations equilibrium; this equilibrium involves optimal dynamic responses to shocks as well as optimal long-run average values of inflation and output; and the same rule results in optimal responses regardless of the assumed statistical properties of the (additive) disturbances. We show that robustly optimal rules of this kind can be represented by interest-rate feedback rules that generalize the celebrated proposal of Taylor (1993). We also show that optimal rules can be represented by history-dependent inflation targets, generalizing the sort of "targeting rule" advocated by Svensson (1997, 1999, 2001). Optimal rules are, however, more complex than these simple proposals, even in the simplest optimizing model that we consider; in particular, they require that policy be history-dependent in ways not contemplated by the well-known proposals. We furthermore find that a robustly optimal policy rule is almost inevitably an *implicit* rule, that requires the central bank to use a structural model to project the economy's evolution under the contemplated policy action. However, our calibrated examples suggest that optimal rules do not place nearly as much weight on projections of inflation or output many quarters in the future as is the current practice of inflation-forecast targeting central banks.

*An earlier version of this work was delivered by the second author as the Jacob Marschak Lecture at the 2001 Far Eastern Meeting of the Econometric Society, Kobe, Japan, July 21, 2001. We thank Julio Rotemberg and Lars Svensson for helpful discussions, and the National Science Foundation, through a grant to the NBER, for research support.

In a companion paper (Giannoni and Woodford, 2002), we have expounded a general approach to the design of an optimal criterion on the basis of which a central bank should determine its operating target for the level of overnight interest rates. Our discussion there was framed in the context of a fairly general linear-quadratic policy problem. Here we consider the implications of our approach in the context of a particular (admittedly stylized) model of the monetary transmission mechanism, or rather a group of related variant models.[1] This allows us to address a number of questions raised by recent characterizations of actual central-bank policies in terms of "Taylor rules" or "flexible inflation targets".

One basic question is whether it makes sense for a central bank's policy commitment to be formulated in terms of a relationship between the bank's interest-rate instrument, some measure of inflation, and some measure of an "output gap" that the bank seeks to ensure will hold — as is true both of the interest-rate rule recommended by Taylor (1993) and the target criterion recommended by Svensson (1999). Can a desirable policy rule be expressed without reference to a monetary aggregate? Can the rule be optimal despite a lack of any explicit dependence of policy upon the nature of the exogenous disturbances affecting the economy? If a desirable rule can be expressed in terms of a relation between these variables, which inflation measure, and which conception of the output gap should it involve? And how do the optimal coefficients of the respective variables depend on quantitative features of one's model of the monetary transmission mechanism?

A particular concern here will be with the optimal *dynamic* specification of a monetary policy rule. Taylor's well-known proposal prescribes a purely contemporaneous relation between the federal funds rate target, an inflation measure, and an output-gap measure; but estimated central-bank reaction functions (*e.g.,* Judd and Rudebusch, 1998) always involve additional partial-adjustment dynamics for the funds rate, and sometimes other sorts of lagged responses as well. To what extent are such lags in the rule used to set interest

[1]While the general approach to the construction of robustly optimal policy rules used here is the same as that discussed in the companion paper, the derivations presented here are self-contained and do not rely upon any of the results for the general linear-quadratic problem presented in the earlier paper. It is our hope that a self-contained exposition of the relevant calculations for these simple models will serve to increase insight into the method, in addition to delivering results of interest with regard to these particular models.

rate desirable? Some empirical studies (*e.g.,* Clarida *et al.,* 2000) imply that central banks respond to forecasts of the future levels of inflation and/or output rather than to current values. Is this preferable, and if so, how far in the future should these forecasts look?

We address these questions here along the lines proposed in Giannoni and Woodford (2002). Rather than optimizing over some parametric family of policy rules, we consider the design of a rule in order to bring about the optimal equilibrium pattern of responses to disturbances — more precisely, to determine an equilibrium that is optimal from the "timeless perspective" explained in our previous paper. In addition to requiring the rule to be consistent with this optimal equilibrium, we ask that the rule imply that rational expectations equilibrium should be *determinate,* so that commitment to the rule can be relied upon to bring about the desired equilibrium rather than some other, less desirable one. We also construct rules that can expressed purely in terms of the "target variables" that the central bank seeks to stabilize, and that are optimal regardless of the nature of the (additive) exogenous disturbances to which the economy is subject, and regardless of the statistical properties of those disturbances. The requirement that our rule simultaneously satisfy each of these desiderata allows us to narrow the class of optimal rules to a fairly small set; among these, we give primary attention to those rules that are simplest in form. Even so, we are typically left with more than one possible representation of optimal policy. In particular, in most of the cases considered here, optimal policy can be represented *either* by a generalized Taylor rule or by a history-dependent inflation target, and we discuss both of these formulations.

1 Optimal Rules for a Simple Forward-Looking Model

We first illustrate our method of constructing robustly optimal policy rules in the context of the basic optimizing model of the monetary transmission mechanism expounded in Woodford (2002, chap. 4), and used as the basis for the discussion of the optimal responses to real disturbances in Woodford (1999a) and Giannoni (2001). The model may be reduced to two structural equations

$$x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n) \tag{1.1}$$

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t \tag{1.2}$$

for the determination of the inflation rate π_t and the output gap x_t, given the central bank's control of its short-term nominal interest-rate instrument i_t, and the evolution of the composite exogenous disturbances r_t^n and u_t. Here the output gap is defined relative to an exogenously varying natural rate of output, chosen to correspond to the gap that belongs among the target variables in the central bank's loss function. The "cost-push shock" u_t then represents exogenous variation in the gap between the flexible-price equilibrium level of output and this natural rate, due for example to time-varying distortions that alter the degree of inefficiency of the flexible-price equilibrium.[2] The microfoundations for this model imply that $\sigma, \kappa > 0$, and that $0 < \beta < 1$. The unconditional expectation of the natural rate of interest process is given by $E(r^n) = \bar{r} \equiv -\log\beta > 0$, while the cost-push disturbance is normalized to have an unconditional expectation $E(u) = 0$. Otherwise, our theoretical assumptions place no *a priori* restrictions upon the statistical properties of the disturbance processes, and we shall be interested in policy rules that are optimal in the case of a general specification of the additive disturbance processes of the form discussed in Giannoni and Woodford (2002, sec. 4).

The assumed objective of monetary policy is to minimize the expected value of a loss criterion of the form

$$W = E_0 \left\{ \sum_{t=0}^{\infty} \beta^t L_t \right\}, \tag{1.3}$$

where the discount factor β is the same as in (1.2), and the loss each period is given by

$$L_t = \pi_t^2 + \lambda_x (x_t - x^*)^2 + \lambda_i (i_t - i^*)^2, \tag{1.4}$$

for certain optimal levels $x^*, i^* \geq 0$ of the output gap and the nominal interest rate, and certain weights $\lambda_x, \lambda_i > 0$. A welfare-theoretic justification is given for this form of loss

[2] See Woodford (2002, chap. 6) for discussion of the welfare-relevant output gap and of the nature of "cost-push shocks".

function in Woodford (2002, chap. 6), where the parameters are related to those of the model of the structural model. However, our conclusions below are presented in terms of the parameters of the loss function (1.4), and are applicable in the case of any loss function of this general form, whether the weights and target values are the ones that can be justified on welfare-theoretic grounds or not. In the numerical results presented below, the model parameters are calibrated as in Table 1 of Woodford (1999a). (For convenience, the parameters are reported in Table 1 below.)

1.1 The Optimal Taylor Rule

Before turning to the question of fully optimal policy in this model, it may be of interest to briefly consider the optimal choice of a rule within a restricted class that has been widely discussed, which is to say, the class of "simple Taylor rules",[3]

$$i_t = \bar{\imath} + \phi_\pi(\pi_t - \bar{\pi}) + \phi_x(x_t - \bar{x})/4, \tag{1.5}$$

involving only contemporaneous feedback from the inflation rate and the output gap, and no direct responses to real disturbances.[4] A rule of this form reflects the intuitive notion that it may be desirable to adjust the bank's instrument in response to deviations of its target variables (other than the instrument itself) from certain desired levels. The conditions for such a rule to imply a determinate equilibrium in this model have already been treated in Woodford (2002, chap. 4).

A rule of the form (1.5) represents an example of a *purely forward-looking* rule, so if it implies a determinate equilibrium, that equilibrium is one in which all three target variables

[3] In the rule proposed by Taylor (1993), the inflation variable is actually the most recent four-quarter change in the GDP deflator, whereas we here consider rules that respond to the change from the previous to the current quarter only. However, Taylor's intention seems to have been to assume feedback from contemporaneous measures of the Fed's (implicit) target variables. We here assume that the central bank seeks to stabilize the one-period inflation rate π_t rather than some average of inflation over a longer time span, because this is the objective that we have been able to justify on welfare-theoretic grounds, in Woodford (2002, chap. 6). Our analysis is also simplest in this case, though similar methods could be used to analyze optimal policy in the case of an alternative inflation-stabilization objective, that might reflect the true goal of a particular central bank.

[4] The coefficient on the output gap is denoted $\phi_x/4$ rather than ϕ_x, so that ϕ_x corresponds to Taylor's output-gap coefficient, writing the rule in terms of annualized data. Here we assume that "periods" of our model correspond to quarters.

will be functions solely of the current and expected future values of the real disturbances. Hence the best pattern of responses to disturbances that could possibly be implemented by a rule in this family is the one that we have called the *optimal non-inertial plan* in Woodford (1999a). In general, even the optimal non-inertial plan can only be implemented by a rule more complex than (1.5). One case in which a rule of this form suffices, however — at least for an open set of possible parameter values — is that in which both the r_t^n and u_t disturbances are Markovian (*i.e.*, first-order autoregressive processes), as assumed in the numerical examples presented in Woodford (1999a) and Giannoni (2001). In such a case, the rule that implements the optimal non-inertial plan is clearly the optimal member of the family, and this makes calculation of the optimal Taylor rule quite straightforward.

Thus we assume once again disturbances of the form

$$\hat{r}_t^n = \rho_r \hat{r}_{t-1}^n + \epsilon_{rt}, \tag{1.6}$$

$$u_t = \rho_u u_{t-1} + \epsilon_{ut}, \tag{1.7}$$

where $\hat{r}_t^n \equiv r_t^n - \bar{r}$, ϵ_{rt} and ϵ_{ut} are i.i.d. mean-zero exogenous shocks, and $0 \leq \rho_r, \rho_u < 1$. In the case, the constraints upon the feasible evolution of the target variables $\{\pi_t, x_t, i_t\}$ from date t onward depend only upon the vector of current disturbances $e_t \equiv [\hat{r}_t^n \; u_t]'$, and the optimal non-inertial plan is given by linear functions of the form

$$z_t = \bar{z} + F e_t, \qquad\qquad i_t = \bar{\imath} + f_i e_t,$$

where $z_t \equiv [\pi_t \; x_t]'$ is the vector of endogenous variables other than the policy instrument. The long-run average values $\bar{z}, \bar{\imath}$ and response coefficients F, f_i are given in the Appendix (section A.3).

In the (generic) case that the matrix F is invertible, an instrument rule consistent with this pattern of responses to shocks is given by

$$i_t = \bar{\imath} + f_i F^{-1}(z_t - \bar{z}), \tag{1.8}$$

which takes the form of a simple Taylor rule (1.5). Note that while we have here written the rule in terms of deviations from implicit targets for each of the variables that correspond to the optimal long-run average values of these variables, the only thing that matters for the constrained optimality of (1.8) is the value of the total intercept term

$$\bar{\imath} - f_i F^{-1} \bar{z}.$$

Of course, the above derivation guarantees only that the suggested rule is *consistent* with the equilibrium responses to shocks that constitute the optimal non-inertial plan. In order to implement the plan, we also need for the rule to imply a determinate equilibrium. The conditions under which this will be true have been discussed in Woodford (2002, chap. 4). The following result states conditions under which the coefficients of the rule just proposed satisfy this additional requirement .

PROPOSITION 4. Suppose the disturbances are of the form (1.6) – (1.7), with autocorrelation coefficients satisfying the bounds

$$\begin{aligned} 0 &< (1-\rho_r)(1-\beta\rho_r) - \rho_r \kappa\sigma \\ &\leq (1-\rho_u)(1-\beta\rho_u) - \rho_u \kappa\sigma < \frac{\kappa\sigma}{\lambda_i}. \end{aligned}$$

Then (1.8) defines a Taylor rule of the form (1.5) with coefficients $\phi_\pi > 1, \phi_x > 0$. Furthermore, commitment to this rule implies a determinate rational expectations equilibrium, which implements the optimal non-inertial plan.

The proof is given in the Appendix. Note that the inequalities assumed in the proposition may equivalently be written

$$\underline{\rho} < \rho_u \leq \rho_r < \bar{\rho},$$

where $\underline{\rho} < \bar{\rho}$ and the bounds are functions of the model parameters β, κ, σ and λ_i. Thus there is an open set of values of ρ_r and ρ_u for which the conditions are satisfied, and these are not obviously unreasonable; for example, the calibrated values reported in Table 1 below satisfy these conditions.

Under the conditions assumed in this last proposition, we thus obtain theoretical justification for crucial aspects of Taylor's recommendation. In particular, we provide support for his recommendation that the operating target for the federal funds rate should respond positively to fluctuations in both the current inflation rate and the current output gap. The rule proposed here also satisfies the "Taylor Principle", according to which an increase in inflation above the target rate results in an even greater increase in the nominal interest rate.

The need for non-zero response coefficients for both target variables follows from a desire to implement the optimal non-inertial responses to two distinct types of real disturbances — disturbances to the natural rate r_t^n and "cost-push shocks" u_t — or more precisely, to respond optimally to any of a range of real disturbances, which shift the model's structural equations in these two different ways to differing extents. If instead we assume that there are no cost-push shocks — not that there are no "supply" disturbances, but that all real disturbances shift the natural rate of output and the efficient rate of output to the same extent — then the requirement that our rule implement the optimal non-inertial response to disturbances to the natural rate of interest imposes only a single linear restriction upon the coefficients ϕ_π and ϕ_x,[5] and it is possible to find a rule that implements the optimal non-inertial plan with $\phi_x = 0$. On the other hand, adding the requirement that the rule also implement the optimal non-inertial response to cost-push shocks, should they ever occur, has no cost in terms of a less desirable response to disturbances to the natural rate of interest,[6] and thus robustness concerns make it advisable that policy respond to variations in the output gap as well. Interestingly the optimal degree of response to variations in the output gap is independent of the assumed importance of cost-push shocks (*i.e.*, the assumed variance of the u_t disturbance); all that matters for the recommendation (1.8) is the assumed degree

[5] This case is analyzed in Woodford (1999a).

[6] We assume here that there is no difficulty in measuring and hence responding to either of the target variables, inflation and the output gap. In practice, measurement of the output gap is likely to be more problematic, and for this reason implementation of the optimal non-inertial responses to variations in r_t^n through a rule that involves a large coefficient ϕ_x may result in some deterioration in the ability of policy to successfully respond to those disturbances. The problem of the optimal conduct of policy when measurement problems are taken into account is considered below in section xx.

of serial correlation of such disturbances when they occur.

Our analysis also provides at least partial support for Taylor's recommendation of what we have called a *direct* policy rule: one that specifies adjustment of the instrument purely in terms of feedback from the observed (or projected) behavior of the target variables. Of course we have not shown that the rule (1.8) cannot be improved upon; but we have shown that it is optimal within the class of purely forward-looking rules. This means that if we consider only possible dependence of the central bank's instrument upon various state variables that are relevant to the determination of current of future values of the target variables, there is no possible gain from introducing dependence upon variables other than those already allowed for in (1.5).

In particular, our analysis justifies Taylor's neglect of any response to projections of *future* inflation or output gaps, as opposed to projections for the current quarter. If we were to introduce additional terms representing feedback from $E_t\pi_{t+j}$ or $E_t x_{t+k}$ for some horizons $j, k > 0$, the optimal rule within that broader family would achieve no better an outcome. For such a rule would continue to be purely forward-looking, and so could at best implement the optimal non-inertial plan, and this is already achieved by the optimal Taylor rule, under the assumptions of Proposition 4.

Nor is it even entirely correct to say that a forecast-based rule would be an equally useful way of achieving the same outcome. Under the assumption that the central bank has access to perfectly accurate forecasts (so that a forecast-based rule can, in principle, involve exactly the same equilibrium adjustment of interest rates as under a Taylor rule), then the optimal response coefficients become *larger* the longer the horizon of the forecasts that are used to implement policy. For example, suppose we consider forward-looking Taylor rules of the form

$$i_t = \bar{\imath} + \phi_\pi(E_t\pi_{t+k} - \bar{\pi}) + \phi_x(E_t x_{t+k} - \bar{x})/4, \tag{1.9}$$

for a given forecast horizon $k > 0$, and for simplicity assume that $\rho_r = \rho_u = \rho$, for some $0 < \rho < 1$.[7] Then the unique rule within this family that is consistent with the optimal

[7] Note that in the case that either ρ_r or ρ_u is equal to zero, it will be impossible for a purely forecast-based

non-inertial plan is given by

$$i_t = \bar{\imath} + \rho^{-k} f_i F^{-1}(E_t z_{t+k} - \bar{z}); \tag{1.10}$$

both response coefficients must be multiplied by the factor $\rho^{-k} > 1$, which may be quite large in the case of a horizon several quarters in the future.[8] But such an alternative rule has the unpalatable feature that it involves a commitment to extremely strong responses to something that, in practice, is likely to be estimated with considerable error.

Furthermore, even when highly accurate conditional forecasts are available, a commitment to strong response to them by the central bank makes it likely that equilibrium will be indeterminate.[9] For the class of forward-looking rules just considered, we can establish the following.

PROPOSITION 5. For all forecast horizons k longer than some critical value, the rule of the form (1.9) that is consistent with the optimal non-inertial plan implies indeterminacy of rational-expectations equilibrium.

The proof is in the Appendix. Thus if the forecast horizon k is sufficiently long, it is not possible to implement the optimal non-inertial plan using a rule of the form (1.9).[10] It follows that, at least when the parameters satisfy the conditions of Proposition 4, the best rule in this family is not as desirable as the best simple Taylor rule.

In the case of the calibrated parameter values proposed in Table 1 below, including the values $\rho_r = \rho_u = .35$ for the serial correlation of the disturbance processes, the optimal

---

rule such as (1.9) to implement the optimal non-inertial plan, because under that pattern of responses to disturbances, the forecasts will not reveal information about the current value of the transitory disturbance.

[8] For example, if we assume a serial correlation coefficient of $\rho = .35$, as in the baseline calibration in Woodford (1999a), and a forecast horizon k of 8 quarters, a fairly typical horizon for inflation-targeting central banks, this factor is greater than 4000.

[9] Recall the discussion of this defect of forward-looking rules in Woodford (2002, chap. 4, sec. xx). The possibility that too strong a response to forecasts can lead to indeterminacy was first shown by Bernanke and Woodford (1997), while the possibility that too long a forecast horizon can lead to indeterminacy is illustrated by Levin *et al.*(2001).

[10] For example, in the case of the calibrated parameter values given in Table 1 below, the rule (1.10) implies indeterminacy for all $k \geq 1$.

Taylor rule is given by

$$i_t^{ann} = .03 + 1.72\pi_t^{ann} + .57x_t, \tag{1.11}$$

where we now (for comparability with Taylor's prescription) report the rule in terms of an annualized interest rate and inflation rate ($i_t^{ann} = 4i_t$, $\pi_t^{ann} = 4\pi_t$).[11] These parameter values are quite similar to those recommended by Taylor. Particularly worthy of note is the substantial response coefficient ϕ_x for variations in the output gap; thus the low assumed value for λ_x (relative to *ad hoc* loss functions often assumed in the literature on monetary policy evaluation) does not imply a low Taylor-rule response coefficient, relative to conventional recommendations.[12] If, instead, one believes that a proper weight on output-gap stabilization as a policy goal requires that λ_x be much higher than the value assumed here, the optimal value of ϕ_x should be correspondingly higher; see equation (1.12) below.

Probably the most important difference between this constrained-optimal rule and Taylor's is that the implicit target inflation rate here is near zero, whereas Taylor assumes a target rate of 2 percent per year. It is important also to note that the value $\phi_x = .57$ refers to the optimal response to fluctuations in a theoretical concept of the output gap ($x_t \equiv \hat{Y}_t - \hat{Y}_t^e$) that may not correspond too closely to conventional "output gap" measures, which are often simply real GDP relative to some smooth trend.[13] Instead, the microeconomic foundations of our model imply that the efficient level of output $\hat{Y}_t^e$ should be affected by real disturbances of all sorts (including some that would conventionally be classified as "demand" disturbances), and these disturbances may include some high or medium-frequency compo-

[11] While the value of β reported in Table 1 is equal to .99, rounding to only two significant digits, and this value would imply that $\bar{\imath}$ should equal approximately .01 per quarter, 4 percent per year, the estimates of Rotemberg and Woodford (1997) actually imply a long-run average real federal funds rate closer to 3 per cent per year, so this is the value reported for $4\bar{\imath}$ here. Because the Rotemberg-Woodford estimates imply an optimal inflation target only slightly above zero, the value of the constant term in this rule is essentially the value of the annualized interest rate consistent with zero average inflation.

[12] In particular, this result shows that the reason for the extremely low optimal output-response coefficients obtained in the study of Rotemberg and Woodford (1999) is not the low value of λ_x assumed in that analysis; it is rather the fact that in their various families of simple policy rules, the coefficients in question indicate response to a conventional output-gap measure rather than to the welfare-relevant gap, as discussed below.

[13] In the case of Taylor's (1993) discussion of the degree to which his proposed rule could account for actual US policy under Greenspan's chairmanship of the Fed, the linearly detrended log of real GDP is used an empirical proxy for x_t.

nents. If one were instead to ask what the constrained-optimal rule would be within the simple family (1.5), but with x_t replaced by detrended output $\hat{Y}_t$, the optimal value of the output coefficient may be quite different — it need not even be positive! For example, Gali (2000) considers this question in the context of a calibrated model similar to our baseline model, in which the real disturbances are technology shocks, and concludes that the optimal output response coefficient is zero when detrended output is used in the Taylor rule instead of the theoretically correct gap measure. Rotemberg and Woodford (1999) reach a similar conclusion (an optimal output-gap coefficient of only 0.02) in the context of their related but more complex model, with disturbance processes inferred from US time series. In the case that there are substantial deviations of the efficient level of output from a smooth trend, as both of these analyses imply, the conventional gap measure is not at all closely related to variations in the welfare-relevant gap, and a substantial positive response to it — stabilizing the conventional gap but thereby *destabilizing* the welfare-relevant gap, as well as inflation — can have undesirable consequences from the point of view of the welfare-theoretic stabilization goals assumed here.[14]

But even if the rule incorporates the correct implicit inflation target and is implemented using a correct measure of the output gap, there remain disadvantages of the Taylor rule as a policy prescription. For one, the constrained-optimality of the coefficients in (1.11) is demonstrated only in the case of a particular specification of the real disturbance processes — only two disturbances, each an AR(1) process with a serial correlation coefficient of exactly .35. The optimal coefficients are in fact quite sensitive to the assumed degree of persistence of the disturbances; for example, in the special case that $\rho_r = \rho_u = \rho$, they are given by

$$\phi_\pi = \frac{\kappa\sigma}{\lambda_i[(1-\rho)(1-\beta\rho) - \rho\kappa\sigma]}, \quad \phi_x = \frac{4\lambda_x\sigma(1-\beta\rho)}{\lambda_i[(1-\rho)(1-\beta\rho) - \rho\kappa\sigma]}. \tag{1.12}$$

In the case of our calibrated values for the other parameters, this implies that the optimal coefficients take the values $\phi_\pi = .96, \phi_x = .41$ if the common ρ is assumed to be as low

[14]See also McCallum and Nelson (2001) and Woodford (2001) for related discussions, with additional evidence suggesting that conventional and welfare-relevant gap measures may not be at all closely related in historical time series for the US.

as .17, while they instead take unboundedly large values if it is assumed to be as high as .68.[15] Thus this constrained-optimal rule is not robustly optimal in the sense discussed above. This makes it an unappealing policy prescription, for in practice policymakers are not simply likely to doubt whether the assumed values of ρ_r and ρ_u correctly represent the typical degree of persistence of disturbances of the two types; they are instead likely to deny that all such disturbances possess *any* single degree of persistence, and thus to remain skeptical about the wisdom of commitment to a rule that is optimal only if all disturbances are linear combinations of only two types.

Furthermore, even if it is literally true that only two types of disturbances ever occur and they are correctly described by (1.6) – (1.7), the Taylor rule (1.11) is not a fully optimal rule. In the best case, it implements the optimal non-inertial plan, but as is shown in Woodford (1999a), this is not generally the optimal plan. It is possible to do better by committing to a rule that incorporates an appropriate form of history-dependence. As we shall see, introducing history-dependence of the right kind can eliminate both of these defects of the simple Taylor rule.

1.2 A Robustly Optimal Instrument Rule

We turn now to the search for a rule that can instead implement the optimal pattern of responses to real disturbances. We recall that the state-contingent plan that minimizes the objective (1.3) – (1.4) subject to the constraints (1.1) – (1.2) satisfies the first-order conditions

$$\pi_t - \beta^{-1}\sigma\Xi_{1t-1} + \Xi_{2t} - \Xi_{2t-1} = 0, \tag{1.13}$$

$$\lambda_x(x_t - x^*) + \Xi_{1t} - \beta^{-1}\Xi_{1t-1} - \kappa\Xi_{2t} = 0, \tag{1.14}$$

[15] In the case of $\rho < .17$, the coefficients given by (1.12) cease to imply a determinate equilibrium, as the "Taylor Principle" ceases to be satisfied. In the case of $\rho > .68$, the denominators of both expressions in (1.12) become negative, implying $\phi_\pi, \phi_x < 0$. While these are possible rules in our discrete-time model, and even imply a determinate equilibrium as long as $\rho < .79$, the analysis for this range of parameter values takes too literally the assumption that all economic decisions are made only at discrete (quarterly) intervals, and so we choose not to emphasize the possibility of using a Taylor rule to implement the optimal non-inertial plan in this case.

$$\lambda_i(i_t - i^*) + \sigma\Xi_{1t} = 0, \tag{1.15}$$

for each date $t \geq 0$,[16] together with the initial conditions

$$\Xi_{1,-1} = \Xi_{2,-1} = 0. \tag{1.16}$$

(Here Ξ_{1t} and Ξ_{2t} are the Lagrange multipliers associated with constraints (1.1) and (1.2) respectively.) In the case that a bounded optimal plan exists, we have seen that it can be described by equations for $\pi_t, x_t, i_t, \Xi_{1t}$ and Ξ_{2t} as linear functions of $\Xi_{1,t-1}$ and $\Xi_{2,t-1}$ together with the current and expected future values of the exogenous disturbances; these linear equations with constant coefficients apply in all periods $t \geq 0$, starting from the initial conditions (1.16).

It follows from these first-order conditions that in the case of an optimal commitment that has been in force since at least period $t-2$, it is possible to infer the values of $\Xi_{1,t-1}$ and $\Xi_{2,t-1}$ from the values that have been observed for x_{t-1}, i_{t-1}, and i_{t-2}. Specifically, one can infer the value of $\Xi_{1,t-1}$ from the value of i_{t-1} using (1.15), and similarly the value of $\Xi_{1,t-2}$ from the value of i_{t-2}. Then substituting these values into (1.14) for period $t-1$, one can also infer the value of $\Xi_{2,t-1}$ from the value of x_{t-1}. One can, of course, similarly solve for the period t Lagrange multipliers as functions of x_t, i_t, and i_{t-1}. Using these expressions to substitute out the Lagrange multipliers in (1.13), one obtains a linear relation among the endogenous variables $\pi_t, x_t, x_{t-1}, i_t, i_{t-1}$ and i_{t-2} that must hold in any period $t \geq 2$. This thus provides a candidate policy rule that is consistent with the optimal state-contingent plan.

Because the relation in question involves a non-zero coefficient on i_t, it can be expressed as an implicit instrument rule of the form

$$i_t = (1-\rho_1)i^* + \rho_1 i_{t-1} + \rho_2 \Delta i_{t-1} + \phi_\pi \pi_t + \phi_x \Delta x_t/4, \tag{1.17}$$

where

$$\rho_1 = 1 + \frac{\kappa\sigma}{\beta} > 1, \qquad \rho_2 = \beta^{-1} > 1, \tag{1.18}$$

[16] In terms of the notation of section 1.2, we here assume that $t_0 = 0$.

$$\phi_\pi = \frac{\kappa\sigma}{\lambda_i} > 0, \qquad \phi_x = \frac{4\sigma\lambda_x}{\lambda_i} > 0. \tag{1.19}$$

We can furthermore show (see Appendix for proof) that commitment to this rule implies a determinate equilibrium.

PROPOSITION 6. Suppose that a bounded optimal state-contingent plan exists. Then in the case of any parameter values $\sigma, \kappa, \lambda_x, \lambda_i > 0$ and $0 < \beta < 1$, a commitment to the rule described by (1.17) – (1.19) implies a determinate rational-expectations equilibrium.

The equilibrium determined by commitment to this rule from date $t = 0$ onward corresponds to the unique bounded solution to equations (1.13) – (1.15) when the initial conditions (1.16) are replaced by the values of $\Xi_{1,-1}$ and $\Xi_{2,-1}$ that would be inferred from the historical values of x_{-1}, i_{-1}, and i_{-2} under the reasoning described above.

It follows that the equilibrium determined by commitment to the time-invariant instrument rule (1.17) involves the same responses to random shocks in periods $t \geq 0$ as under the optimal commitment. This is thus an example of an instrument rule that is optimal from a timeless perspective, in the sense defined in section 1.2. Note that we could instead implement precisely the optimal once-and-for-all commitment from date $t = 0$ onward (the bounded solution to (1.13) – (1.15) with initial conditions (1.16)) by committing to (1.17) in all periods $t \geq 2$, but to a modified version of the rule in periods $t = 0$ and 1. But this would be a non-time-invariant rule (policy would depend upon the date relative to the date at which the commitment had been made), and the preferability of this alternative equilibrium, from the standpoint of expected welfare looking forward from date $t = 0$, would result from the alternative policy's optimal exploitation of prior expectations that are already given in that period. Choice of a rule that is optimal from a timeless perspective requires us to instead commit to set the interest rate according to the time-invariant rule (1.17) in all periods.

The rule (1.17) has the additional advantage of being robustly optimal, in the sense defined in section 1.3. We note that our derivation of the optimal rule has required no hypotheses about the nature of the disturbance processes $\{r_t^n, u_t\}$, except that they are exogenously given and that they are bounded. In fact, the rule is optimal regardless of their nature; commitment to this rule implies the optimal impulse responses displayed in Woodford (1999a) in the case of the particular disturbance processes assumed in the numerical illustrations there, but it equally implies optimal responses in the case of any other types of disturbances to the natural rate of interest and/or "cost-push shocks" — disturbances that may be anticipated some quarters in advance, disturbances the effects of which do not die out monotonically with time, and so on.[17] Indeed, one may assume that both of the disturbances r_t^n and u_t in equations (1.1) – (1.2) are composite disturbances of the general form discussed in Giannoni and Woodford (2002, sec. 4), and (1.17) remains an optimal rule. This robustness of the rule is a strong advantage from the point of view of its adoption as a practical guide to the conduct of monetary policy.

It is important to note that (1.17) is not a *uniquely* optimal instrument rule; it is not even the only rule that is robustly optimal in the sense just discussed. For example, other rules that are equally consistent with the optimal responses to disturbances, regardless of the nature of the disturbance processes, may be obtained by substituting for variables in (1.17) using one or the other of the structural equations (1.1) – (1.2).[18] However, alternative optimal rules derived in this way will not be *direct* rules, insofar as they will involve feedback from past, current, or expected future real disturbances as well as from the paths of the target variables. (One might arrange for the disturbance terms to cancel, under a particular hypothesis about the statistical properties of the disturbances, but the version of the rule that omitted reference to the disturbances would not be robustly optimal.)

[17] This is a substantial advantage of this instrument rule over the one proposed in Woodford (1999a), which expresses the federal funds rate as a function of the lagged funds rate, the lagged rate of increase in the funds rate, the current inflation rate, and the previous quarter's inflation rate. That rule would also be consistent with optimal responses to real disturbances, but only if (as assumed in the earlier calculation) all disturbances perturb the natural rate of interest in a way that can be described by an AR(1) process (1.6) with a single specified coefficient of serial correlation, and have no effect on the natural rate of output that is different than the effect on the efficient rate of output (*i.e.,* there are no cost-push shocks). In this special case, however, the rule discussed earlier has the advantage that its implementation requires no information on the part of the central bank other than an accurate measure of inflation (including an accurate projection of period t inflation at the time that the period t funds rate is set).

[18] A specific example: one might use (1.2) to substitute for π_t in (1.17), and obtain a rule for setting i_t as a function of $i_{t-1}, i_{t-2}, x_t, x_{t-1}, E_t\pi_{t+1}$, and u_t.

A robustly optimal direct rule must be an implication of the first-order conditions (1.13) – (1.15) only, in order for it not to refer to the structural disturbances; and in order for it not to refer to the Lagrange multipliers, either, it must in fact be an implication of (1.17). This still does not make (1.17) the unique such rule. For example, if (1.17) holds in all periods, it follows that

$$\begin{aligned} i_t &= (1-\rho_1)(1-\rho_3)i^* + [\rho_1(1-\rho_3)+\rho_3]i_{t-1} + (\rho_2+\rho_1\rho_3)\Delta i_{t-1} \\ &\quad - \rho_2\rho_3\Delta i_{t-2} + q_t - \rho_3 q_{t-1} \end{aligned} \tag{1.20}$$

must also hold in all periods, where

$$q_t \equiv \phi_\pi \pi_t + (\phi_x/4)\Delta x_t \tag{1.21}$$

and ρ_3 is an arbitrary coefficient. (This relation is obtained from (1.17) by adding to the right-hand side ρ_3 times i_{t-1} minus the right-hand side at date $t-1$.) Condition (1.20) can also be interpreted as a direct implicit instrument rule, and it too is consistent with the optimal responses to all real disturbances, regardless of the statistical properties of those disturbances. Since we know that the rule implies a determinate equilibrium when $\rho_3 = 0$, it follows by continuity that it will also imply a determinate equilibrium for all small enough $\rho_3 \neq 0$. Hence there exist rules of this form that are also robustly optimal direct instrument rules. But the additional history-dependence introduced into (1.20) is unnecessary; (1.17) is unambiguously a simpler rule. The same objection may be raised against the rules with additional lead terms that can be derived from (1.17) by substituting for some terms using the conditional expectation at date t of both sides of (1.17) at some future date.

Another relation implied by (1.17) that does not involve a larger number of terms is

$$i_t = (\rho_1+\rho_2)^{-1}[(\rho_1-1)i^* + E_t i_{t+1} + \rho_2 i_{t-1} - E_t q_{t+1}]. \tag{1.22}$$

(This relation is equivalent to the statement that (1.17) holds at date $t+1$ only in expectation conditional upon public information at date t.) This too might conceivably be interpreted as an implicit instrument rule for setting i_t at date t, though in this case a forecast-based rule.

However, while this relation is consistent with the optimal responses to disturbances, imposition of (1.22) as a monetary policy rule does not determine a unique rational-expectations equilibrium.[19] Thus (1.22) — which is implied by but does not imply (1.17) — does not represent a completely specified monetary policy rule under the criterion proposed in section 1.1, for it does not imply a determinate state-contingent path for the central bank's policy instrument. The same is true *a fortiori* of the relation that would be obtained from the conditional expectation at t of (1.17) at $t+2$. Thus we conclude that (1.17) is of unique interest as the simplest possible robustly optimal direct instrument rule, in the case of our basic neo-Wicksellian model.

The optimal rule (1.17) has a number of important similarities to the Taylor rule. Like the Taylor rule, (1.17) is an example of a direct, implicit instrument rule. The rule is also similar to Taylor's recommendation in that the contemporaneous effect of an increase in either inflation or the output gap upon the federal funds rate operating target is positive ($\phi_\pi, \phi_x > 0$); and the rule satisfies the "Taylor principle," given that $\phi_\pi > 0$ and $\rho_1 > 1$.[20] However, this optimal rule involves additional history-dependence, owing to the non-zero weights on the lagged funds rate, the lagged rate of increase in the funds rate, and the lagged output gap. And the optimal degree of history-dependence is non-trivial: the optimal values of ρ_1 and ρ_2 are both necessarily greater than one, while the optimal coefficient on x_{t-1} is as large (in absolute value) as the coefficient on x_t. It is particularly worth noting that the optimal rule implies not only intrinsic inertia in the dynamics of the funds rate — a transitory deviation of the inflation rate from its average value increases the funds rate not only in the current quarter, but in subsequent quarters as well — but is actually *super-inertial:* the implied dynamics for the funds rate are explosive,[21] if the initial overshooting of the long-

[19] We can easily see this by noting that (1.22) makes i_t a function of no predetermined state variables other than i_{t-1}. Hence if this rule *did* imply a determinate equilibrium, in that equilibrium, π_t, x_t and i_t would all be linear functions of i_{t-1} and the exogenous states that suffice to forecast the real disturbances from period t onward. Yet we know that the optimal responses to shocks generally involve more complex dependence upon history than can be summarized by a single predetermined variable such as i_{t-1}; for one cannot generally infer the values of *both* $\Xi_{1,t-1}$ and $\Xi_{2,t-1}$ from the value of i_{t-1} alone. We thus show by contradiction that (1.22) cannot imply a determinate equilibrium.

[20] Recall the discussion in Woodford (2002, chap. 4, sec. xx) of the generalization of this principle to the case of policy rules with interest-rate inertia.

run average inflation rate is not offset by a subsequent undershooting (as actually always happens, in equilibrium). In this respect this optimal rule is similar to those found to be optimal in the numerical analysis by Rotemberg and Woodford (1999) of a more complicated empirical version of the model.

In the case of the calibrated parameter values in Table 1 below, the coefficients of the optimal instrument rule are given by $\rho_1 = 1.15, \rho_2 = 1.01, \phi_\pi = .64$, and $\phi_x = .33$. These may be compared with the coefficients of the Fed reaction function of similar form estimated by Judd and Rudebusch (1998) for the Greenspan period: $\rho_1 = .73, \rho_2 = .43, \phi_\pi = .42$, and $\phi_x = .30$, except in this empirical reaction function ϕ_x represents the reaction to the current quarter's *level* of the output gap, rather than its first difference.[22] (Interestingly, they find that an equation with feedback from the first difference of the output gap, rather than its level, fits best during an earlier period of Fed policy, under Paul Volcker's chairmanship.) The signs of the coefficients of the optimal rule agree with those characterizing actual policy; in particular, the estimated reaction function includes substantial positive coefficients ρ_1 and ρ_2, though these are still not as large as the optimal values. Thus the way in which actual Fed policy is more complex than adherence to a simple Taylor rule can largely be justified as movement in the direction of optimal policy, according to the simple model of the transmission mechanism assumed here.

We find that in the case of this simple model at least, it is not necessary for the central bank's operating target for the overnight interest rate to respond to forecasts of the future evolution of inflation or of the output gap in order for policy to be fully optimal — and not just optimal in the case of particular assumed stochastic processes for the disturbances, but robustly optimal. Thus the mere fact that the central bank may sometimes have information

[21] Technically, this corresponds to the observation that in the equivalent representation (1.23) of the policy rule given below, there exists a root $\lambda_2 > 1$. A sufficient condition for this is that $\rho_1 > 1$, in which case exactly one of the roots is greater than 1.

[22] It should also be noted that the output gap measure used in Judd and Rudebusch's empirical analysis, while a plausible measure of what the Fed is likely to have responded to, may not correspond to the welfare-relevant output gap indicated by the variable x_t in the optimal rule (1.17). In addition, ϕ_π indicates response to the most recent four-quarter growth in the GDP deflator, rather than an annualized inflation rate over the past quarter alone.

about future disturbances, that are not in any way disturbing demand or supply conditions yet, is not a reason for feedback from current and past values of the target variables to be insufficient as a basis for optimal policy. This does not mean that it may not be desirable for monetary policy to restrain spending and/or price increases even before the anticipated real disturbances actually take effect. But in the context of a forward-looking model of private-sector behavior, a commitment to respond to fluctuations in the target variables only contemporaneously and later does not preclude effective pre-emptive constraint of that kind. First of all, such a policy may well mean that the central bank *does* adjust its policy instrument immediately in response to the news, insofar as forward-looking private-sector behavior may result in an immediate effect of the news upon current inflation and output.[23] And more importantly, in the presence of forward-looking private-sector behavior, the central bank mainly affects the economy through changes in expectations about the *future* path of its instrument in any event; a predictable adjustment of interest rates *later,* once the disturbances substantially affect inflation and output, should be just as effective in restraining private-sector spending and pricing decisions as a preemptive increase in overnight interest rates immediately.

At the same time, it is important to note that the optimal rule (1.17), while not "forecast-based" in the sense in which this term is usually understood, *does* depend upon projections of inflation and output in the same quarter as the one for which the operating target is being set. Thus the rule is not an explicit instrument rule in the sense of Svensson and Woodford (1999). And this implicit character (a feature that it shares with the Taylor rule) is crucial to the optimality of the rule, at least if we wish to find an optimal rule that is also a direct rule (specifying feedback only from the target variables). For optimal policy must generally involve an immediate adjustment of the short-term nominal interest rate in response to shocks, as shown in Woodford (1999a);[24] and so unless the rule is to be specified in terms of

[23]This is obviously not the case if, as more realistic models often assume, there are delays in the effect of any new information on prices and spending. But in this case, it is probably not desirable for overnight interest rates to respond immediately to news, either; see section xx below.

[24]This is not true if there are delays in the effects of shocks upon inflation and output, as discussed in section xx below. But in that case, even the *delayed* effect upon the central bank's instrument that is required

the central bank's response to particular shocks, it will have to specify a contemporaneous response to fluctuations in the target variables, and not simply a lagged response. Thus implementation of such a rule will involve judgment of some sophistication about current conditions; it cannot be implemented mechanically on the basis of a small number of publicly available statistics.

1.3 An Optimal Targeting Rule

We have just seen that policy need not be forecast-based in order to be optimal, and even robustly optimal. At the same time, this does not mean that a forecast-based decision procedure cannot have equally desirable properties. In fact, the policy rule just discussed is *equivalent*, in a certain sense, to a forecast-based rule. This alternative representation of optimal policy is also of interest as an example of a robustly optimal policy rule that takes the form of a *targeting* rule, rather than as an expression that presents, even implicitly, a formula for the bank's interest-rate operating target.

We can write the implicit instrument rule (1.17) in the form

$$(1 - \lambda_1 L)(1 - \lambda_2 L)\hat{\imath}_t = \hat{q}_t, \tag{1.23}$$

where $\hat{\imath}_t \equiv i_t - \bar{\imath}$, and $\hat{q}_t$ similarly denotes the deviation of q_t (the function of the target variables to which the central bank responds, defined in (1.21)) from its long-run average value, $\phi_\pi \bar{\pi}$. Because the optimal coefficients (1.18) are such that $\rho_1 > 1, \rho_2 > 0$, the roots in the factorization (1.23) necessarily satisfy $0 < \lambda_1 < 1 < \lambda_2$. It then follows that relation (1.17) is equivalent to the relation

$$(1 - \lambda_1 L)\hat{\imath}_{t-1} = -\lambda_2^{-1} E_t[(1 - \lambda_2^{-1} L^{-1})^{-1} \hat{q}_t], \tag{1.24}$$

in the following sense.

Proposition 7. Two bounded stochastic processes $\{\hat{\imath}_t, \hat{q}_t\}$ satisfy (1.23) for all $t \geq 0$ if and only if they satisfy (1.24) for all $t \geq 0$.

by optimal policy cannot be implemented on the basis only of lagged observations of the target variables, because of the delay with which shocks affect these variables.

(See proof in the Appendix.)

Thus there is no difference between the way in which a central bank must adjust its interest-rate instrument to ensure that (1.17) holds in all periods and the way that it would adjust it to ensure that (1.24) holds in all periods, for the two conditions imply one another. (This does not mean that arranging for (1.24) to hold in a *single* period t is equivalent to arranging for (1.17) to hold in that single period, regardless of how policy is expected to be conducted thereafter; but a *permanent commitment* to either rule from some date t_0 onward has identical consequences.) This equivalence does not apply only in the case of processes that are possible equilibria of the model consisting of structural equations (1.1) – (1.2); thus the rules are equivalent regardless of whether that model is correctly specified, and regardless of whether the central bank expects the economy to actually evolve according to a rational-expectations equilibrium of that model or not (for example, regardless of whether the private sector is believed to correctly understand the bank's policy rule or not).

It follows from this equivalence that a commitment to ensure that (1.24) holds in all periods from some date onward represents a coherent complete specification of a monetary policy rule, at least in the context of the model described by equations (1.1) – (1.2). Hence this represents a well-defined *targeting rule,* even though the criterion (1.22) cannot be solved by itself to yield even an implicit expression for the period t instrument setting: the left-hand side involves only lagged interest rates, while the right-hand side refers only to the evolution of inflation and the output gap. A model of the monetary transmission mechanism must be used in order to determine the instrument setting that is consistent with a projection that satisfies the target criterion.[25]

The target criterion (1.24) can be expressed in the form

$$F_t(\pi) + \frac{\phi_x}{4} F_t(x) = \frac{\theta_x}{4} x_{t-1} - \theta_i (i_{t-1} - i^*) - \theta_\Delta \Delta i_{t-1}, \tag{1.25}$$

[25]Note, however, that the situation is not really different in the case of a commitment to ensure that (1.17) is satisfied: a model is still needed to determine the instrument setting that should result in current period inflation and output that imply that the implicit instrument rule is satisfied.

where for each of the variables $z = \pi, x$ we use the notation $F_t(z)$ for a conditional forecast

$$F_t(z) \equiv \sum_{j=0}^{\infty} \alpha_{z,j} E_t z_{t+j},$$

involving weights $\{\alpha_{z,j}\}$ that sum to one. Thus the criterion specifies a time-varying target value for a weighted average of an inflation forecast and an output-gap forecast, where each of these forecasts is in fact a weighted average of forecasts at various horizons, rather than a projection for a specific future date. The rule represents a variant of what Svensson (1999) calls "flexible inflation targeting," though this rule differs from the examples that he discusses in the history-dependence of the inflation-forecast target (indicated by the non-constant terms on the right-hand side).

In representation (1.25) of this policy rule, there is no constant term, indicating an inflation-forecast target of zero except insofar as this is corrected in response to deviations (past or projected) of the output gap and/or the nominal interest rate from their target values.[26] The optimal coefficients indicating the degree to which the inflation-forecast target is adjusted are given by

$$\phi_x = \theta_x = 4(1 - \lambda_2^{-1})\frac{\lambda_x}{\kappa} > 0,$$

$$\theta_i = \lambda_2(1 - \lambda_1)(1 - \lambda_2^{-1})\frac{\lambda_i}{\kappa\sigma} > 0,$$

$$\theta_\Delta = \lambda_1\lambda_2(1 - \lambda_2^{-1})\frac{\lambda_i}{\kappa\sigma} > 0,$$

while the optimal weights in the conditional forecasts are

$$\alpha_{\pi,j} = \alpha_{x,j} = (1 - \lambda_2^{-1})\lambda_2^{-j}.$$

Thus the optimal conditional forecast is one that places positive weight on the projection for each future period, beginning with the current period, with weights that decline exponentially

[26] Note, however, that this does not mean that the rule sets the inflation forecast equal to zero on average. This is because the target interest rate i^* is in general not consistent with an average inflation rate of zero.

as the horizon increases. The mean distance in the future of the projections that are relevant to the target criterion is equal to

$$\sum_{j=0}^{\infty} \alpha_{z,j} j = (\lambda_2 - 1)^{-1}$$

for both the inflation and output-gap forecasts.

In the case of the calibrated parameter values reported in Table 1, the rate at which these weights decay per quarter is $\lambda_2^{-1} = .68$, so that the mean forecast horizon in the optimal target criterion is 2.1 quarters. Thus while our optimal targeting rule can be expressed in terms of a target for inflation and output-gap forecasts, the forecast horizon involved is short compared to those typically considered in the recent literature, or those typical of the actual practice of inflation forecast-targeting central banks. For these same parameter values, the optimal relative weight on the output-gap forecast is $\phi_x = .15$, indicating that the target criterion is essentially an inflation-forecast target, albeit a modified one. The direction of modification is the one suggested by Svensson: a forecast of a lower output gap than normal should cause the central bank to tolerate a higher than average inflation forecast. Finally, the remaining optimal coefficients are $\theta_x = .15, \theta_i = .24$, and $\theta_\Delta = .51$, indicating a substantial degree of history-dependence of the optimal modified inflation-forecast target. The fact that $\theta_x = \phi_x$ indicates that it is really the forecasted increase in the output gap relative to the previous quarter's level, rather than the absolute level of the gap, that should modify the inflation-forecast target. The signs of θ_i and θ_Δ imply that policy will be made tighter (in the sense of demanding a lower modified inflation forecast) when interest rates have been high and/or increasing in the recent past; this is another way of committing to interest-rate inertia of the kind discussed above.

The equivalence expressed in Proposition 7 implies that commitment to a history-dependent modified inflation-forecast target of this kind is a robustly optimal policy rule in exactly the same sense as the instrument rule (1.17). Thus commitment to a targeting rule can be a sound approach to policy. This alternative representation of optimal policy has the possible advantage (from the point of view of successfully steering private-sector expectations) of emphasizing the way in which the outlook for inflation and the output gap are adjusted at each point in time (at least as far as the intentions of the central bank are concerned) in response to variations in the recent evolution of the target variables. While this is *implied* by a commitment to implement the instrument rule (1.17) from now on, it might not be clear to the private sector — for example, because the central bank's commitment to continue to implement the instrument rule in the future might not be clear. Hence communication with the public about current policy decisions in terms of their implications for inflation and output-gap forecasts might be a superior way of conveying the central bank's commitments with regard to subsequent developments.

The representation of optimal policy in terms of a targeting rule also has the advantage of continuing to be possible even in the limiting case that $\lambda_i = 0$, *i.e.*, even when reducing the variability of interest rates is not an independent concern.[27] In that limit, the weights ϕ_π and ϕ_x in (1.17) become unboundedly large, so that a representation of optimal policy in terms of a direct instrument rule ceases to be possible. Instead, the coefficients of (1.25) remain well-defined: θ_i and θ_Δ become equal to zero, while ϕ_x, θ_x, and the weights $\{\alpha_{z,j}\}$ continue to take well-defined positive values. Thus in this limiting case, the optimal targeting rule is one in which the inflation-forecast target must be modified in proportion to the projected change in the output gap, but it is no longer also dependent on lagged interest rates.

In fact, the optimal target criterion in this case can be written more simply as

$$\pi_t + \frac{\phi_x}{4}\Delta x_t = 0, \tag{1.26}$$

or even in terms of a flexible price-level target

$$\log P_t = \log P^* - \frac{\phi_x}{4} x_t,$$

where P^* is a target price level.[28] Thus optimal policy in this can be implemented by commitment to a target of the sort proposed by Hall (1984), though the optimal weight ϕ_x

[27] Svensson and Woodford (1999) consider a model closely related to this one, but assume a stabilization objective in which $\lambda_i = 0$. It is for this reason that they find that it is possible to formulate a robustly optimal targeting rule, a forward-looking variant of (1.26) below, but not a robustly optimal instrument rule.

[28] See Woodford (1999b) for further discussion. Even when $\lambda_i > 0$, "Wicksellian" rules, in which the

implied by our theory is considerably smaller (for any plausible calibration of the model's parameters) than the values suggested by Hall. Furthermore, such a policy may be implemented through a forecast-targeting procedure of the kind practiced by inflation-targeting central banks, rather than necessarily requiring institution of the sort of automatic mechanism proposed by Hall.

2 Optimal Rules for a Model with Inflation Inertia

The basic model considered above is often criticized as being excessively forward-looking, particularly in its neglect of any sources of intrinsic inertia in the dynamics of inflation. It might be suspected that this feature of the model is responsible for our strong conclusion above, according to which a robustly optimal policy rule need involve no dependence upon forecasts of the target variables beyond the current period. In Svensson's (1997) classic argument for the optimality of inflation-forecast targeting, it is the existence of lags in the effect of monetary policy on inflation that causes the optimal rule to involve a target criterion for a forecast, with the optimal forecast horizon coinciding with the length of the policy transmission lag. It might reasonably be suspected that forecasts are not necessary in our analysis above because our simple model includes no lags in the effects of policy.

Here we take up this question by extending our analysis to the case of the model of inflation inertia developed in Woodford (2002, chap. 3). In this extension of our basic model, prices are not held constant between the dates at which they are re-optimized, but instead are automatically adjusted on the basis of the most recent quarter's increase in the aggregate price index, by a percentage that is a fraction γ of the percentage increase in the index. As shown in Woodford (2002, chap. 3), the aggregate supply relation (1.2) then takes the more general form

$$\pi_t - \gamma\pi_{t-1} = \kappa x_t + \beta E_t(\pi_{t+1} - \gamma\pi_t) + u_t, \tag{2.1}$$

nominal interest rate is adjusted in response to deviations of the price level from a deterministic target path, rather than to deviations of inflation from a target rate as in the Taylor rule, maybe be desirable by comparison with other equally simple rules, as shown by Giannoni (2001).

where the coefficient κ and the disturbance u_t are defined as before. For γ substantially greater than zero, this makes past inflation an important determinant of current inflation, along with current and expected future output gaps and cost-push shocks; if γ is close enough to one, even a monetary disturbance that has only a transitory effect on real activity can have a much longer-lasting effect on inflation.

The aggregate-demand side of our model remains as before, and our model can accordingly be summarized by the two structural equations (1.1) and (2.1), together with exogenous stochastic processes for the disturbances $\{r_t^n, u_t\}$. As shown in Woodford (2002, chap. 6), the change in our assumptions about pricing behavior implies a corresponding change in the appropriate welfare-theoretic stabilization objective for monetary policy. This is once again a discounted criterion of the form (1.3), but the period loss function becomes

$$L_t = (\pi_t - \gamma\pi_{t-1})^2 + \lambda_x(x_t - x^*)^2 + \lambda_i(i_t - i^*)^2. \tag{2.2}$$

We wish to consider policies that minimize the criterion defined by (1.3) and (2.2), subject to the constraints imposed by the structural equations (1.1) and (2.1), for arbitrary values of the indexation parameter $0 \leq \gamma \leq 1$.[29]

2.1 Optimal Instrument Rules

In the case of this generalization of our policy problem, the first-order condition (1.13) becomes instead

$$\pi_t^{qd} - \beta\gamma E_t\pi_{t+1}^{qd} - \beta^{-1}\sigma\Xi_{1t-1} - \beta\gamma E_t\Xi_{2,t+1} + (1+\beta\gamma)\Xi_{2t} - \Xi_{2t-1} = 0, \tag{2.3}$$

where

$$\pi_t^{qd} \equiv \pi_t - \gamma\pi_{t-1} \tag{2.4}$$

[29]An alternative way of modeling inflation inertia would be to assume the existence of backward-looking "rule of thumb" price-setters, as in Gali and Gertler (1999). This leads to a modification of the aggregate supply relation that is similar, though not quite identical, to (2.1). Steinsson (2000) and Amato and Laubach (2001b) derive welfare-theoretic loss functions for this model, and find that the loss each period is a quadratic function of both π_t and π_{t-1} that is similar, though again not identical, to our loss function (2.2). Hence we conjecture that similar conclusions as to the degree to which optimal policy is forward-looking would be obtained using the Gali-Gertler model, though we do not take this up here.

is the quasi-differenced inflation rate that appears in both the aggregate supply relation (2.1) and the loss function (2.2). Conditions (1.14) – (1.15) remain as before, and this system of three equations, together with initial conditions (1.16) and an initial condition for π_{-1}, continues to define the optimal once-and-for-all commitment to apply from date $t = 0$ onward.

As above, we can use conditions (1.14) – (1.15) to substitute for Ξ_{1t} and Ξ_{2t} in (2.3), obtaining an Euler equation of the form

$$E_t[A(L)(i_{t+1} - i^*)] = -f_t \tag{2.5}$$

for the optimal evolution of the target variables. Here $A(L)$ is a cubic lag polynomial

$$A(L) \equiv \beta\gamma - (1 + \gamma + \beta\gamma)L + (1 + \gamma + \beta^{-1}(1 + \kappa\sigma))L^2 - \beta^{-1}L^3, \tag{2.6}$$

while the term f_t is a function of the observed and expected future paths of the target variables, defined by

$$f_t \equiv \frac{\kappa\sigma}{\lambda_i}[\tilde{q}_t - \beta\gamma E_t\tilde{q}_{t+1}], \tag{2.7}$$

$$\tilde{q}_t \equiv \pi_t^{qd} + \frac{\lambda_x}{\kappa}\Delta x_t. \tag{2.8}$$

(Note that the above definition generalizes the earlier (1.21), and that in the limit where $\gamma = 0$, f_t is equal to $\tilde{q}_t$, which equals q_t.)

By an argument directly analogous to the proof of Proposition 6, we can show that if a bounded optimal state-contingent plan exists, the system obtained by adjoining (2.5) to the structural equations (1.1) and (2.1) implies a determinate rational-expectations equilibrium, in which the responses to exogenous disturbances are the same as under the optimal commitment. (The only difference between this equilibrium and the optimal once-and-for-all commitment just defined relates to the initial conditions, as in our earlier discussion, and once again this difference is irrelevant to the design of a policy rule that is optimal from a timeless perspective.) Hence we could regard (2.5) as implicitly defining a policy rule, and the rule would once again be robustly optimal. In the limiting case that $\gamma = 0$, (2.5) ceases to involve any dependence upon $E_t i_{t+1}$, and the proposed rule would coincide with the optimal instrument rule (1.17) discussed above.

However, (2.5) is an even less explicit expression for the central bank's interest-rate policy than the implicit instrument rules considered earlier, for (when $\gamma > 0$) it defines i_t only as a function of $E_t i_{t+1}$. This means that the central bank defines the way in which it is committed to set its instrument only as a function of the way that it expects to act further in the future. This failure to express the rule in "closed form" is especially undesirable from the point of view of our question about the optimal forecast horizon for a monetary policy rule. Expression (2.5) involves no conditional expectations for variables at dates more than one period in the future. However, this does not really mean that the central bank's forecasts for later dates are irrelevant when setting i_t. For this "rule" directs the bank to set i_t as a function of its forecast of i_{t+1}, and (if the same rule is expected to be used to set i_{t+1}) the bank's forecast at t of i_{t+1} should involve its forecast at t of $\tilde{q}_{t+2}$. It should also involve its forecast of i_{t+2}, and hence (by similar reasoning) its forecast of $\tilde{q}_{t+3}$, and so on. Hence it is more revealing to describe the proposed policy rule in a form that eliminates any reference to the future path of interest rates themselves, and instead refers only to the bank's projections of the future paths of inflation and the output gap.[30]

To obtain an equivalent policy rule of the desired form, we need to partially "solve forward" equation (2.5). This requires factorization of the lag polynomial as

$$A(L) \equiv \beta\gamma(1 - \lambda_1 L)(1 - \lambda_2 L)(1 - \lambda_3 L). \tag{2.9}$$

[30] A rule expressed in this way will also conform better to the evident preference of central bank's to justify their monetary policy decisions to the public in terms of their projections for the future paths of inflation and output, rather than in terms of their assumptions about the future path of interest rates. Public communications such as the Bank of England's *Inflation Report* put projections for both inflation and output at center stage, while being careful not to express any opinion whatsoever about the likely path of interest rates over the period under discussion. The forecast-based rules proposed below still refer to forecast paths conditional upon intended policy, rather than upon "constant-interest-rate" forecasts, and so it will not be possible to implement these rules without taking a stand (at least for internal purposes) on the likely future path of interest rates. But the rules make it possible to discuss the way in which the current instrument setting is required by the bank's inflation and output projections, without also discussing the interest-rate path that is implicit in those projections, and to this extent they require a less radical modification of current procedures.

We note the following properties of the roots of the associated characteristic equation.

PROPOSITION 8. Suppose that $\sigma, \kappa > 0, 0 < \beta < 1$, and $0 < \gamma \leq 1$. Then in the factorization (2.9) of the polynomial defined in (2.6), there is necessarily one real root $0 < \lambda_1 < 1$, and two roots outside the unit circle. The latter two roots are either two real roots $\lambda_3 \geq \lambda_2 > 1$, or a complex pair λ_2, λ_3 of roots with real part greater than 1. Three real roots necessarily exist for all small enough $\gamma > 0$, while a complex pair necessarily exists for all γ close enough to 1.

(See proof in the Appendix.) We use the conventions in the statement of this proposition in referring to the distinct roots in what follows. It is also useful to rewrite (2.5) as

$$E_t[A(L)\hat{\imath}_{t+1}] = -\hat{f}_t, \tag{2.10}$$

where once again hats denote the deviations of the original variables from the long-run average values implied by the policy rule (2.5), or equivalently, by the optimal commitment.

In the case that three real roots exist, the existence of two distinct roots greater than one allows us two distinct ways of "solving forward", resulting in two alternative relations,

$$(1 - \lambda_1 L)(1 - \lambda_2 L)\hat{\imath}_t = (\beta\gamma\lambda_3)^{-1} E_t[(1 - \lambda_3^{-1} L^{-1})^{-1} \hat{f}_t], \tag{2.11}$$

or

$$(1 - \lambda_1 L)(1 - \lambda_3 L)\hat{\imath}_t = (\beta\gamma\lambda_2)^{-1} E_t[(1 - \lambda_2^{-1} L^{-1})^{-1} \hat{f}_t]. \tag{2.12}$$

We can also derive other relations of the same form by taking linear combinations of these ones. Of special interest is the relation

$$\begin{aligned} (1 - \lambda_1 L)\left(1 - \frac{\lambda_2 + \lambda_3}{2} L\right)\hat{\imath}_t &= \frac{1}{2}(\beta\gamma\lambda_3)^{-1} E_t[(1 - \lambda_3^{-1} L^{-1})^{-1} \hat{f}_t] \\ &+ \frac{1}{2}(\beta\gamma\lambda_2)^{-1} E_t[(1 - \lambda_2^{-1} L^{-1})^{-1} \hat{f}_t]. \end{aligned} \tag{2.13}$$

Here relations (2.11) and (2.12) are defined (with real-valued coefficients) only in the case that three real roots exist, while relation (2.13) can also be derived (and has real coefficients on all leads and lags) in the case that λ_2, λ_3 are a complex pair. Because $|\lambda_2|, |\lambda_3| > 1$, the right-hand side of each of these expressions is well-defined, and describes a bounded stochastic process in the case of any bounded process $\{\hat{f}_t\}$. (In what follows, we shall refer to the three possible expressions for an optimal instrument rule presented in (2.11) – (2.13) as Rule I, Rule II, and Rule III respectively.)

Each of the relations (2.11) – (2.13) can be solved for $\hat{\imath}_t$ as a function of two of its own lags and expectations at date t regarding current and future values of $\hat{f}_t$. These can thus be interpreted as implicit instrument rules, each of which now avoids any direct reference to the planned future path of the central bank's instrument (though assumptions about future monetary policy will be implicit in the inflation and output-gap forecasts). Each of these policy rules is equivalent to (2.5), and they are accordingly equivalent to one another, in the following sense.

Proposition 9. Under the assumptions of Proposition 8, and in the case that the factorization (2.9) involves three real roots, a pair of bounded processes $\{\hat{\imath}_t, \hat{f}_t\}$ satisfy any of the equations (2.11), (2.12) or (2.13) at all dates $t \geq t_0$ if and only if they satisfy (2.10) at all of those same dates. In the case that a complex pair exists, (2.13) is again equivalent to (2.10), in the same sense.

(See proof in Appendix.) Each of the rules thus represents a feasible specification of monetary policy in the case that its coefficients are real-valued, and when this is true it implies equilibrium responses to real disturbances that are those associated with an optimal commitment. Accordingly, each represents an optimal policy rule from a timeless perspective. (Note that although the coefficients differ, these are not really different policies. Proposition 9 implies that they involve identical actions, if the bank expects to follow one of them indefinitely, regardless of the model of the economy used to form the conditional forecasts.)

In the case that three real roots exist, we have a choice of representations of optimal policy in terms of an instrument rule, and this time we cannot choose among them on grounds of simplicity. But rule I seems particularly appealing in this case. This is the rule (among our three possibilities, or any other linear combinations of these) that puts the least weight on forecasts far in the future. It is proper to ask at what rate the weights on forecasts shrink with the forecast horizon, under the assumption that these shrink as fast as possible consistent with robust optimality of the policy rule, if we wish to determine how much forecast-dependence is *necessary* for robust optimality. This choice is also uniquely desirable in the sense that it remains well-defined in the limit as γ approaches zero. In this limit, rule 1 reduces to

$$(1 - \lambda_1 L)(1 - \lambda_2 L)\hat{\imath}_t = \hat{f}_t,$$

which is the optimal instrument rule (1.23) derived earlier.[31] Instead, in the case of any of the other rules, the coefficients on lagged interest rates become unboundedly large as γ approaches zero. Thus rule 1 is clearly the preferable specification of policy in the case of small γ. The desire for a rule that varies continuously with γ, so that uncertainty about the precise value of γ will not imply any great uncertainty about how to proceed, then make rule I an appealing choice over the entire range of γ for which it is defined.

One might think that the same continuity argument could instead be used to argue for the choice of rule III in all cases, since this is the only one of our optimal instrument rules that continues to be defined for high values of γ. Yet the instruction to follow rule I if three real roots exist, but rule III if there is a complex pair, is also a specification that makes all coefficients of the policy rule continuous functions of γ. The reason is that as γ passes through a critical value $\bar{\gamma}$ at which the real roots of the characteristic equation bifurcate, the two larger real roots, λ_2 and λ_3, come to exactly equal one another. When $\bar{\gamma}$ is approached from the other direction, the imaginary parts of the complex roots λ_2 and λ_3 approach zero;

[31] Note that as $\gamma \to 0$, $\lambda_3 \to +\infty$, while $\gamma\lambda_3 \to \beta^{-1}$. Recall also that in this limiting case, $\hat{f}_t = \hat{q}_t$. One can show furthermore that the two smaller roots λ_1, λ_2 in the factorization (2.9) approach the two roots in the factorization (1.23) of our earlier quadratic lag polynomial.

Table 1: Calibrated parameter values for the basic neo-Wicksellian model.

Structural parameters	
β	0.99
σ^{-1}	0.16
κ	.024
Shock processes	
$\rho(\hat{r}^n), \rho(u)$	0.35
Loss function	
λ_x	.048
λ_i	.077

at the bifurcation point their common real value is the repeated real root obtained as the common limit of the two real roots from the other direction. Hence when $\gamma = \bar{\gamma}$, rules I, II, and III are all identical. There is thus no ambiguity about whether rule I or rule III should be applied in this case, and no discontinuity in the coefficients of the recommended rule as γ approaches $\bar{\gamma}$ from either direction. At the same time, this proposal results in a rule that remains well-defined as γ approaches zero, and for small $\gamma > 0$ results in a rule that is very close to the one previously recommended for an economy with no inflation inertia.

Each of rules I, II, and III can be written in the form

$$i_t = (1 - \rho_1)i^* + \rho_1 i_{t-1} + \rho_2 \Delta i_{t-1} + \phi_\pi F_t(\pi) + \frac{\phi_x}{4} F_t(x) - \theta_\pi \pi_{t-1} - \frac{\theta_x}{4} x_{t-1}, \tag{2.14}$$

where here we have added the constant terms again to indicate the desired level of interest rates (and not just the interest rate relative to its long-run average level), and where $F_t(z)$ again denotes a linear combination of forecasts of the variable z at various future horizons, with weights normalized to sum to one. This form of rule generalizes the specification (1.17) that suffices in the case $\gamma = 0$ in two respects: the interest-rate operating target i_t now depends upon lagged inflation in addition to the lagged variables that mattered before, and it now depends upon forecasts of inflation and the output gap in future periods, and not simply upon the projections of those variables for the current period.

Except in these respects, the coefficients are qualitatively similar to those in (1.17), as indicated by the following proposition.

Proposition 10. Under the assumptions of Proposition 8, and a loss function with $\lambda_x, \lambda_i > 0$, each of rules I, II, and III has a representation of the form (2.14) for all values of γ for which the rule is well-defined, and in this representation,

$$\rho_1 > 1, \qquad \rho_2 > 0,$$

$$0 < \theta_\pi \leq \phi_\pi,$$

and

$$0 < \theta_x = \phi_x.$$

Furthermore, for given values of the other parameters, as $\gamma \to 0$ (for rule I) the coefficient θ_π approaches zero, though ϕ_π approaches a positive limit; while as $\gamma \to 1$ (for rule III) the coefficients θ_π and ϕ_π approach the same positive limit.

(The proof is again in the Appendix.) It is especially noteworthy that once again the optimal instrument rule is superinertial. We also note that once again what should matter is the projected output gap relative to the previous quarter's output gap, rather than the absolute level of the projected gap; and once again interest rates should be increased if the gap is projected to rise. Once again a higher projected inflation rate implies that the interest rate should be increased; but now the degree to which this is true is lower if recent inflation has been high, and in the extreme case $\gamma = 1$, it is only the projected inflation rate relative to the previous quarter's rate that should matter.

The numerical values of these coefficients are plotted, for alternative values of γ ranging between zero and one, in the various panels of Figure 1, where the assumed values for the other parameters are as in Table 1. For all values $\gamma < \bar{\gamma} = .35$, there are three real roots, and for each value of γ the three values corresponding to rules I, II, and III are each plotted; for $\gamma > \bar{\gamma}$, only rule III is defined. An interesting feature of these plots, is that if one considers the coefficients associated with rule I for $\gamma \leq \bar{\gamma}$ and rule III for $\gamma \geq \bar{\gamma}$, one observes that the

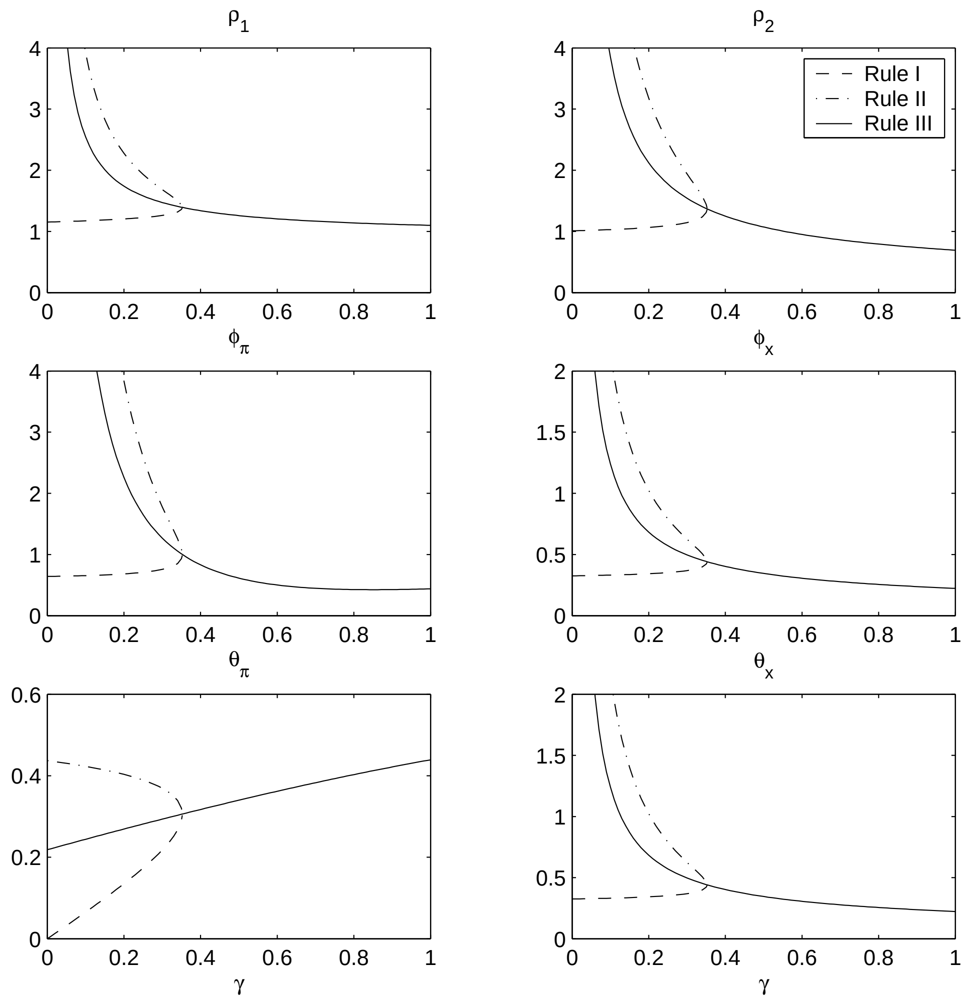

Figure 1: Coefficients of the optimal instrument rule (2.14) as functions of γ.

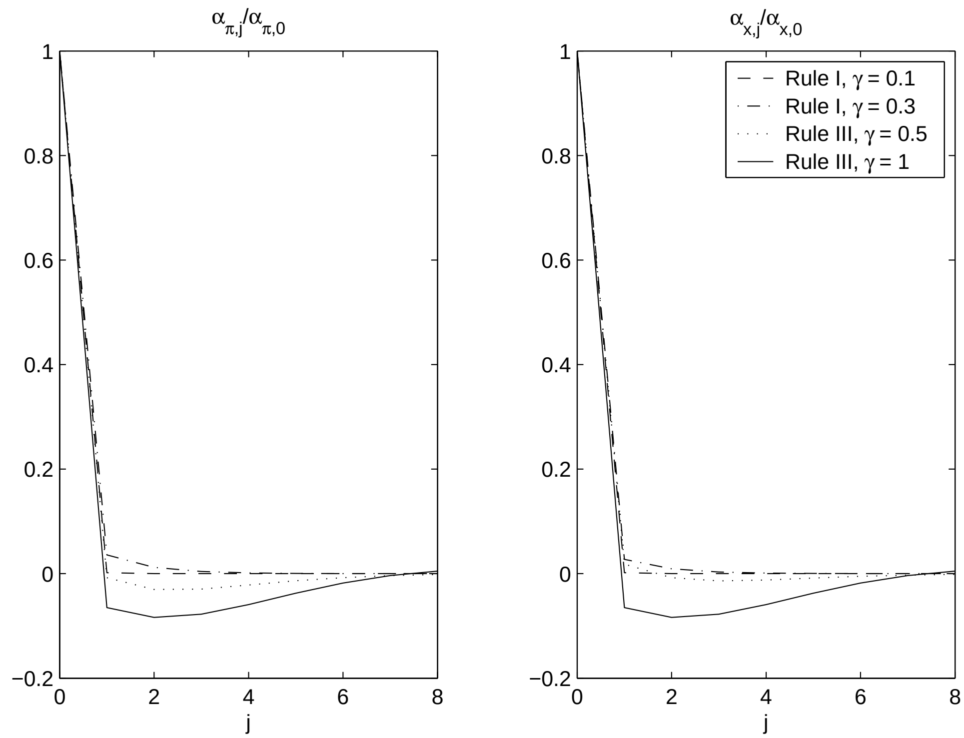

Figure 2: Relative weights on forecasts at different horizons in the optimal rule (2.14).

magnitude of each of the coefficients remains roughly the same, regardless of the assumed value of γ. (The exception is θ_π, which approaches zero for small γ, but becomes a substantial positive coefficient for large γ, as indicated by Proposition 10.)

The panels of Figure 2 similarly plot the relative weights $\alpha_{z,j}/\alpha_{z,0}$ for different horizons j of the inflation and output-gap forecasts to which the optimal instrument rule refers,[32] for each of several different possible values of γ. (The weights associated with rule I are plotted in the case of values $\gamma < \bar{\gamma}$, and those associated with rule III in the case of values $\gamma > \bar{\gamma}$.) Here we observe that in this case the forecasts $F_t(z)$ are not actually weighted averages of forecasts at different horizons, because the weights are not all non-negative. Thus while in the presence of inflation inertia, the optimal instrument rule is to some extent forecast-

[32]Here we plot the relative weights, rather than the absolute weights, because this makes visual comparison between the degree of forecast-dependence of optimal policy in the different cases easier. The absolute weights can be recovered by integrating the plots shown here, since the relative weights in each case must sum to $1/\alpha_{z,0}$.

based, the optimal responses to forecasts of future inflation and output gaps are not of the sort generally assumed in forward-looking variants of the Taylor rule. In the case of high γ, a higher forecasted inflation rate (or output gap) in any of the next several quarters implies, for given past and projected current conditions, that a *lower* current interest rate is appropriate. According to the optimal rule, a higher current inflation rate should be tolerated in the case that high inflation is forecast for the next several quarters. This is because (in an economy with γ near one) it is sudden *changes* in the inflation rate that creates the greatest distortions in the economy, by making automatic adjustment of prices in response to lagged inflation a poor rule of thumb.

In addition to this difference from the conventional wisdom with respect to the sign with which forecasts should affect policy, one notes that under the optimal rule it is only forecasts regarding the near future that matter much at all. Even if we consider only the weights put on forecasts for $j \geq 1$ quarters in the future, the mean future horizon of these forecasts, defined by

$$\sum_{j \geq 1} \alpha_{z,j} j / \sum_{j \geq 1} \alpha_{z,j},$$

is equal to only 2.2 quarters in the case of our calibrated example with $\gamma = 1$. Thus forecasts other than for the first year following the current quarter matter little under the optimal policy. Even more notably, none of the projections beyond the current quarter should receive too great a weight; in our example, the sum of the relative weights on *all* future quarters,

$$\sum_{j>0} |\alpha_{z,j}| / \alpha_{z,0},$$

is equal to only 0.39 even in the extreme case $\gamma = 1$, while this fraction falls to zero for small γ. Thus while a robustly optimal direct instrument rule does have to be forecast-based in the presence of inflation inertia, the degree to which forecasts matter under the optimal policy rule is still relatively small. Instead, a strong response to projections of inflation and the output gap for the current period, as called for by the Taylor rule, continues to be the crucial element of optimal policy.

the notion of "certainty equivalence" must be applied with care.

Nonetheless, optimal policy under imperfect information is in many ways similar to optimal policy under full information. For example, it is no more forward-looking than our previous results indicated; in the present example, the optimal interest-rate operating target for period t depends on the projections $\pi_{t|t}$ and $x_{t|t}$, but not on projections of inflation or the output gap for any dates farther in the future. And once again, optimal policy will imply substantial persistence in interest-rate fluctuations. Taking the expectation of (3.20) conditional upon the central bank's information set at date $t-2$, one obtains

$$i_{t|t-2} = (1-\rho_1)i^* + (\rho_1+\rho_2)i_{t-1|t-2} - \rho_2 i_{t-2} + \phi_\pi \pi_{t|t-2} + \frac{\phi_x}{4}(x_{t|t-2} - x_{t-1|t-2}).$$

This in turn implies that

$$\rho_{k+2}(i) = (\rho_1+\rho_2)\rho_{k+1}(i) - \rho_2\rho_k(i) + \phi_\pi \beta_{k+2}(\pi, i) + +\frac{\phi_x}{4}(\beta_{k+2}(x,i) - \beta_{k+1}(x,i)), \quad (3.22)$$

for each $k \geq 0$, where we use the notation

$$\rho_k(z) \equiv corr(z_{t+k}, z_t), \qquad \beta_k(y,z) \equiv \mathrm{cov}(y_{t+k}, z_t)/\mathrm{var}(z_t).$$

Equation (3.22) is not enough by itself to allow one to solve for the autocorrelation function of the equilibrium interest rate process $\{\rho_k(i)\}$; this depends on the regression coefficients $\{\beta_k(\pi, i)\}$ and $\{\beta_k(x, i)\}$, which cannot be determined from the form of the monetary policy rule alone. But (3.22) is the same restriction on the autocorrelation function as is implied by (1.17) in the full-information case; thus the fact that $\rho_1, \rho_2 > 1$ in (3.20) makes a high degree of serial correlation of interest rates likely in the case of imperfect information as well.

4 Conclusions

We have shown that robustly optimal policy rules can be constructed for each of a variety of simple forward-looking models of the monetary transmission mechanism. We have seen that these rules may take the form either of an implicit instrument rule — an interest-rate feedback rule that generalizes the one proposed by Taylor (1993) — or of a pure targeting rule — a history-dependent inflation target that generalizes the "flexible inflation targeting" rule proposed by Svensson (1999). In the cases where both representations of optimal policy are possible, these are actually equivalent policy rules, at least as far as their implications for rational-expectations equilibrium are concerned, though the policy commitment is described in apparently different ways.

Our examples offer insights into several questions posed in the introduction. First, we have seen that, at least in the case of the simple models considered here,[45] optimal policy rules can be expressed in terms of a commitment to bring about a certain (time-invariant) linear relationship between the paths of the short-term nominal interest-rate instrument of the central bank, an inflation measure, and a measure of the output gap, as proposed by both Taylor and Svensson. And this is not only a possible representation of optimal policy in these cases, but one with uniquely desirable properties, discussed in detail in Giannoni and Woodford (2002). Of course, in more complex (but more realistic) models, optimal rules are likely to involve additional state variables besides these three; we have already seen an illustration of this in section 3.1, where, in general, the presence of wage as well as price stickiness implies that an optimal rule will involve responses to more than one inflation measure (wage as well as price inflation). Nonetheless, the general form of rules that have been widely discussed in the recent literature are not found to be fundamentally misguided; there is no reason why an optimal rule must pay attention to monetary aggregates, for example, or why it must explicitly respond in different ways to different types of disturbances.

An issue that is much debated by monetary economists is whether it is desirable for monetary policy to respond to a measure of the output gap, as both the "Taylor rule" and Svensson's "flexible inflation targeting" prescribe. Here we have found that our optimal policy rules all prescribe interest-rate adjustments or modification of the inflation target in response to changes in the projected path of the (correctly defined) output gap, and thus

[45] While the models are simple, it is perhaps worth recalling that they do represent log-linearizations of completely specified intertemporal general-equilibrium models of the monetary transmission mechanism, and at least those considered in section 3 are already sophisticated enough to match a number of salient features of the econometric evidence on the effects of monetary policy shocks, as discussed in Woodford (2002, chaps. 3-4).

our results provide some justification for the emphasis upon this variable by both Taylor and Svensson. While it might also be possible to formulate policy rules consistent with an optimal equilibrium that would *not* involve explicit reference to this variable, such alternative representations of optimal policy would either not be *robustly* optimal like the rules derived here — the coefficients of the optimal rule would depend upon precise details of the assumed statistical character of the disturbances[46] — or they would not be *direct* rules — they would involve explicit reference to variables other than the target variables, such as specific exogenous disturbances.[47] Hence, insofar as a robustly optimal, direct rule is desirable, there is an important advantage to expressing the central bank's policy commitment in terms of a rule that involves the bank's estimate of (or projection of) the path of the output gap.

There are certainly substantial difficulties involved in accurate measurement of the output gap in practice. But many of these can be well-represented by additive measurement error of the kind considered by Svensson and Woodford (2001), so that certainty-equivalence applies, as discussed in section 3.3. In this case, it is still optimal to commit to a policy rule with the same coefficient on the central bank's estimate of the output gap as would be optimal under full information; the measurement error affects only the way in which it is optimal for the central bank to form its estimate of current and past output gaps. For example, an optimal estimate of the output gap will generally make use of information about wages and prices,

[46]For example, Woodford (1999a) derives an optimal policy rule for the model of section 1 that involves only the short-term nominal interest rate and an inflation measure, and argues that this representation of optimal policy is desirable because it can be implemented without requiring the central bank to measure the output gap. However, the rule discussed there is optimal only under a very specific assumption about the disturbances: there are no cost-push shocks, and the natural rate of interest is an AR(1) process, with innovations that are revealed in the same period that they affect the natural rate. Furthermore, the numerical coefficients of the optimal rule depend on the coefficient of serial correlation of the natural-rate disturbances. Instead, the rules discussed in section 1 are equally optimal under this specification of the disturbances and a vast number of other possibilities.

[47]For example, an optimal policy rule could be constructed for the model of section 1, under a particular specification of the statistical properties of the disturbances, by solving for the optimal state-contingent paths $\{\bar{\imath}_t, \bar{\pi}_t\}$ of the nominal interest rate and inflation, and then choosing a policy rule of the form $i_t = \bar{\imath}_t + \phi(\pi_t - \bar{\pi}_t)$, where $\phi > 1$ in conformity with the "Taylor Principle". It follows from Proposition xx of Woodford (2002, chap. 4) that a rule of this kind implies a determinate rational-expectations equilibrium, and it is obvious that the rule is consistent with the optimal paths of the variables. However, such a rule involves explicit reference to the state of the world as defined by the history of exogenous disturbances, and the way in which the terms $\bar{\imath}_t$ and $\bar{\pi}_t$ vary with the history of disturbances also depends on the details of the assumed statistical properties of the disturbances.

and not simply available quantity measures. But this does not mean that it is not useful for the central bank to describe its policy commitment in terms of a relationship between the output gap and other variables. For *this* description of policy will be much more robust than an explicit description of the way that the central bank should respond to specific indicator variables, which will depend on the bank's current beliefs about the statistical properties of the various disturbances (including the ones responsible for the measurement problems).

It is also worth noting that the errors that have been observed historically in real-time estimates of the output gap (documented by Orphanides, 2000) have been much greater in the case of estimates of the absolute level of the output gap than in the case of estimates of the quarter-to-quarter changes in the gap.[48] (Errors in the recognition of shifts in the trend rate of growth of potential output until years later have caused substantial, highly persistent mis-estimates of the absolute gap; but this particular source of measurement error has little effect on the higher-frequency components of the output gap estimate.) But the optimal rules exhibited above all involve only the projected path of quarter-to-quarter *changes* in the output gap, and are independent of the absolute *level* of the gap, even though it is the absolute size of the gap that one wishes to stabilize. Because of this, it is less obvious that output-gap mismeasurement should be a serious problem in the case of the optimal rules derived here than in the case of the simpler rules proposed by Taylor (1993) and Svensson (1999), both of which make policy depend on the absolute level of the current or projected future output gap.

Our analysis also offers insights into the question of which inflation measure policy should respond to or target. The answer given here is the particular measure or measures that appear as target variables in the welfare-theoretic loss function derived according to the principles set out in Woodford (2002, chap. 6), which is to say, the inflation measures that are directly related to measures of the relative-price distortions that result from imperfect

[48]This assumes, as does Orphanides, that current conventional estimates of past levels of potential output are in fact correct. Of course, the conception of potential output upon which such estimates are based may not be the same one as in the "output gap" to which an optimal policy rule would respond, as suggested by Woodford (2001). But this sort of error is not an inevitable one, resulting from data limitations, but rather one that can be eliminated through clarification of the optimal rule.

synchronization of wage and price changes. The inflation measure that is correct will thus depend on the nature of the nominal rigidities associated with wage and price-setting, which is ultimately an empirical question. In our baseline model, with flexible wages and the same degree of stickiness of all goods prices, the relevant inflation rate is the change in a uniformly weighted index of goods prices, which conforms fairly closely to the kind of price index actually targeted by the central banks with inflation targets. But under other assumptions, the correct inflation measure will differ. For example, we have shown in section 3.1 that if wages as well as prices are sticky, the optimal rule must involve wage inflation as well as price inflation. Similarly, if some goods prices are sticky while others are not, the correct inflation measure will be an index of "core inflation," — an index of the changes in the prices only of the sticky-price goods, as discussed in Woodford (2002, chap. 6, sec. xx).

We also obtain some tentative conclusions about the degree of history-dependence of optimal policy rules. Even in our baseline model, which posits an extremely simple dynamic structure, our optimal policy rules involve substantial history-dependence of a kind not present in proposals such as those of Taylor (1993) and Svensson (1999). In addition to the fact that policy should respond to the projected *change* in the output gap rather than its level, which makes the recent *past* level of the output gap relevant for current policy, we find that past nominal interest rates should affect the current policy setting. Specifically, both our optimal instrument rule (1.17) and our optimal targeting rule (1.25) have the feature that, for any given inflation and output-gap projections, interest rates should be higher than they otherwise would be if (i) interest rates have recently been higher than average, or (ii) interest rates have recently been rising.[49] Thus the optimal rules incorporate both the interest-rate *persistence* (a positive effect of i_{t-1} on the choice of i_t) and interest-rate *momentum* (a positive effect of Δi_{t-1} on the choice of Δi_t) that characterize the actual Fed reaction functions estimated by Judd and Rudebusch (1998).

Finally, we have also explored the degree to which optimal rules should make policy a

[49]In the case of (1.25), one observes that either a high value of i_{t-1} or a high value of Δi_{t-1} require a *lower* value for the output-gap-adjusted inflation forecast — that is, these conditions require policy to be tightened, though the rule itself does not specify the interest-rate setting that this involves.

function of projections of inflation and/or output many quarters in the future. In our baseline model, it is possible to formulate a robustly optimal policy rule (the implicit instrument rule (1.17) that involves no projections farther in the future than the period for which the nominal interest-rate operating target is being set. Perhaps surprisingly, this rule is optimal regardless of what we may assume about the availability of advance information about future disturbances. Of course, this strong result depends on the purely forward-looking character of that simple model of inflation and output determination. But even when we allow for a high degree of inflation inertia, in section 2, we find that an optimal policy rule depends much more on the projected inflation rate and output gap in the quarter for which policy is being set that on the projections for any later horizons. And while projections for later quarters do matter to some extent if the degree of inflation inertia is sufficiently great, projections farther than a year in the future matter little even in this case. Thus we find little justification for a policy that gives primary attention to the inflation forecast at a horizon two years in the future, as is true of the inflation-forecast targeting currently practiced at the Bank of England.

It is important nonetheless to stress that our results do not justify a purely backward-looking approach to the conduct of policy. In all of the cases considered, our optimal rules are *implicit* rules, which is to say that they specify a criterion that must be satisfied by the central bank's projections of inflation and output given its policy. The criterion in question involves variables the values of which depend on the current policy action that is chosen; hence they must be *projected* using a model of the monetary transmission mechanism, rather than simply being *measured.* It is true that optimal policy could also be described by an explicit (purely backward-looking) instrument rule, specifying the instrument setting as a function of current exogenous disturbances and past (or at any rate predetermined) state variables, that need simply be measured. But such a representation of optimal policy would not be robust to changes in the assumed character of the disturbance processes, unlike the implicit rules derived here. Hence we would argue that the use of a quantitative model, that can be used to project the effects of prospective policy settings, is essential to the optimal conduct of monetary policy. And in a model that takes account of forward-looking private sector behavior, projections for the current quarter cannot generally be made without forecasting the economy's subsequent evolution as well.

Furthermore, in the case that spending and pricing decisions are predetermined, as assumed in empirical models such as those of Rotemberg and Woodford (1997) or Christiano *et al.* (2001), the optimal policy is one under which the interest operating target is chosen d periods in advance, on the basis of projections of inflation and output for the period for which the interest rate is being chosen (if not projections farther in the future as well). In this case, policy decisions necessarily will depend crucially on projections of conditions at least d periods in the future. However, the lag d by which spending or pricing decisions are predetermined is not plausibly longer than one or two quarters. And even in this case, no justification is provided for basing the interest-rate operating target for a given period on forecasts regarding points in time that are much more distant than the period for which the interest-rate decision is being made. Hence our results provide little support for the desirability of basing interest-rate decisions primarily on forecasts of conditions as long as two years in the future.

References

Amato, Jeffery D., and Thomas Laubach, "Estimation and Control of an Optimization-Based Model with Sticky Wages and Prices," unpublished, October 2001a. (Forthcoming, *Journal of Economic Dynamics and Control.*)

—- —- and —- —-, "Rule-of-Thumb Behavior and Monetary Policy," unpublished, Bank for International Settlements, December 2001b.

Christiano, Lawrence J., Martin S. Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," NBER Working Paper no. 8403, July 2001.

Clarida, Richard, Jordi Gali, and Mark Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* 115: 147-180 (2000).

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics* 46: 281-313 (2000).

Gali, Jordi, "The Conduct of Monetary Policy in the Face of Technological Change: Theory and Postwar U.S. Evidence," unpublished, Universitat Pompeu Fabra, October 2000.

—- —-, and Mark Gertler, "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics* 44: 195-222 (1999).

Giannoni, Marc P., "Optimal Interest-Rate Rules in a Forward-Looking Model, and Inflation Stabilization versus Price-Level Stabilization," unpublished, Princeton University, 2001.

—- —-, and Michael Woodford, "Optimal Interest-Rate Rules: I. General Theory," unpublished, Princeton University, August 2002.

Judd, John F., and Glenn D. Rudebusch, "Taylor's Rule and the Fed: 1970-1997," *Federal Reserve Bank of San Francisco Economic Review,* 1998(3), pp. 3-16.

Levin, Andrew, Volker Wieland, and John C. Williams, "The Performance of Forecast-Based Monetary Policy Rules under Model Uncertainty," FEDS paper no. 2001-39, Federal Reserve Board, August 2001.

Orphanides, Athanasios, "Monetary Policy Evaluation with Noisy Information," FEDS paper no. 1998-50, Federal Reserve Board, October 1998.

—- —-, "The Quest for Prosperity without Inflation," Working Paper no. 15, European Central Bank, March 2000.

Rotemberg, Julio J., and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual* 12: 297-346 (1997).

—- —- and —- —-, "Interest-Rate Rules in an Estimated Sticky-Price Model," in J.B. Taylor, ed., *Monetary Policy Rules,* Chicago: University of Chicago Press, 1999.

Steinsson, Jon, "Optimal Monetary Policy in an Economy with Inflation Persistence," Working Paper no. 11, Central Bank of Iceland, December 2000.

Svensson, Lars E.O., "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," *European Economic Review* 41: 1111-1146 (1997).

—- —-, "Inflation Targeting as a Monetary Policy Rule," *Journal of Monetary Economics* 43: 607-654 (1999).

—- —-, and Michael Woodford, "Implementing Optimal Policy through Inflation-Forecast Targeting," unpublished, Princeton University, November 1999.

—- —- and —- —-, "Indicator Variables for Optimal Policy under Asymmetric Information," NBER Working Paper no. 8255, April 2001. (Forthcoming, *Journal of Economic Dynamics and Control.*)

Taylor, John B., "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214 (1993).

Woodford, Michael, "Optimal Monetary Policy Inertia," NBER Working Paper no. 7261, August 1999a.

—- —-, "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?" in *New Challenges for Monetary Policy*, Kansas City: Federal Reserve Bank of Kansas City, 1999b.

—- —-, "The Taylor Rule and Optimal Monetary Policy," *American Economic Review* 91(2): 232-237 (2001).

—- —-, *Interest and Prices: Foundations of a Theory of Monetary Policy*, unpublished, Princeton University, July 2002.