

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Monetary Policy Rules

Volume Author/Editor: John B. Taylor, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-79124-6

Volume URL: <http://www.nber.org/books/tayl99-1>

Publication Date: January 1999

Chapter Title: Interest Rate Rules in an Estimated Sticky Price Model

Chapter Author: Julio J. Rotemberg, Michael Woodford

Chapter URL: <http://www.nber.org/chapters/c7414>

Chapter pages in book: (p. 57 - 126)

Interest Rate Rules in an Estimated Sticky Price Model

Julio J. Rotemberg and Michael Woodford

This paper seeks to evaluate monetary policy rules that generalize the rule proposed by Taylor (1993). In particular, we consider rules in which the Fed sets the federal funds rate as a function of the history of inflation, output, and the federal funds rate itself. Even though this is not part of Taylor's original formulation, we introduce the possibility that the federal funds rate depends on the history of the funds rate itself in order to allow for interest rate smoothing of the kind that appears to be an important feature of current Fed policy. We also consider the character of optimal policy, that is, the policy that maximizes the utility of the representative agent, assuming unlimited information about the exogenous disturbances to the economy. We then compare optimal policy in this unrestricted sense with the best rule of the generalized Taylor family.

We evaluate these rules under the assumption that interest rate, inflation, and output determination in the U.S. economy can be compactly represented by the small structural model whose parameters we estimate in Rotemberg and Woodford (1997). This is a rational expectations model derived from explicit intertemporal optimization, in which firms are unable to change their prices every period, and in which purchases are determined somewhat in advance of when they actually take place. In evaluating different monetary rules we use two approaches. The first approach is simply to compute the welfare of the

Julio J. Rotemberg is the MBA Class of 1942 Professor of Business Administration at the Harvard Business School and a research associate of the National Bureau of Economic Research. Michael Woodford is the Harold Helm '20 Professor of Economics and Banking at Princeton University and a research associate of the National Bureau of Economic Research.

The authors thank Marc Giannoni for excellent research assistance and Larry Christiano, Martin Feldstein, Hans Genberg, Mark Gertler, Marvin Goodfriend, Huw Pill, Argia Sbordone, Lars Svensson, and John Taylor for comments. They also thank the Harvard Business School Division of Research and the National Science Foundation, through a grant to the NBER, for research support.

representative household according to our model of the U.S. economy. Because this places great strain on the assumptions that the model contains accurate descriptions of the preferences of American residents, and that we have correctly identified the nature of the real disturbances to which monetary stabilization policy must respond, we also study separately the variability of output, inflation, and interest rates induced by different policy rules. This latter way of characterizing economic performance under alternative rules is less dependent on the “deep structural” interpretation of the residuals in our structural equations, although it is, of course, still dependent on the specification of those structural equations and on the statistical properties of their disturbance terms.

We proceed as follows. In section 2.1, we describe the structure of the model, which is discussed more thoroughly in Rotemberg and Woodford (1997, 1998). Section 2.2 is devoted to the analysis of simple policy rules that represent variations on the rule proposed by Taylor (1993), while section 2.3 considers optimal policy. Section 2.4 concludes.

2.1 A Framework for Analysis

We begin by reviewing the structure of the estimated sticky price model developed in Rotemberg and Woodford (1997). This also allows us to derive the utility-based measure of deadweight loss due to price level instability that is the basis for our subsequent discussion of optimal policy.

2.1.1 A Small, Structural Model of the U.S. Economy

We suppose that there is a continuum of households indexed by i , where i runs between 0 and 1. Each of these households produces a single good while it consumes the composite good. The utility of household i at t is given by

$$(1) \quad E_i \sum_{T=t}^{\infty} \beta^{T-t} [u(C_T^i; \xi_T) - v(y_T^i; \xi_T)],$$

where β is a discount rate, y_t^i is the household's production of its own good, and ξ_t is a vector of preference (or technological) disturbances. The argument C_t^i represents an index of the household's purchases of the continuum of differentiated goods produced in the economy. Following Dixit and Stiglitz (1977), this index is given by

$$(2) \quad C_T^i \equiv \left[\int_0^1 c_T^i(z)^{(\theta-1)/\theta} dz \right]^{\theta/(\theta-1)},$$

where $c_t^i(z)$ is the quantity purchased of good z and the constant elasticity of substitution θ is assumed to be greater than one. We assume that all purchasers, including the government, care only about an aggregate of the form (2). As usual, this implies that the total demand $y_t(z)$ for differentiated good z is given by a constant-elasticity demand function

$$(3) \quad y_t^z = Y_t \left(\frac{p_t(z)}{P_t} \right)^{-\theta},$$

where $p_t(z)$ is the period t price of good z , P_t is the price index defined by

$$(4) \quad P_t \equiv \left[\int_0^1 p_t(z)^{1-\theta} dz \right]^{1/(1-\theta)},$$

and Y_t measures aggregate demand for the composite good defined by equation (2).

One of the delays we assume is that households must choose their index of purchases C_t^i at date $t - 2$. As we show in Rotemberg and Woodford (1997) this, or an assumption like it, seems necessary if one wishes to explain the response of U.S. GDP to monetary disturbances, because this response is itself delayed by about two quarters. This delay implies that the standard Euler equation for optimal intertemporal allocation of consumption spending need not hold, except (approximately) conditional on information available two periods in advance. Because C_t^i is chosen in advance, household optimization requires only

$$(5) \quad E_t[u_C(C_{t+2}^i; \xi_{t+2})] = E_t(\lambda_{t+2}^i P_{t+2}),$$

where λ_t^i is a Lagrange multiplier indicating the marginal utility for household i of additional nominal income in period t . Assuming borrowing limits that never bind in equilibrium, these marginal utilities of income must satisfy

$$(6) \quad \lambda_t^i = \beta R_t E_t \lambda_{t+1}^i,$$

where R_t is the gross return on a riskless nominal one-period asset in which the household invests at t . We assume the existence of complete insurance markets, so that all households consume the same amount at any time and have the same marginal utility of income. Then equations (5) and (6) also hold when we drop the i superscripts and interpret them as equations relating aggregate consumption C_t to the marginal utility of income λ_t of the representative household. However, because of the conditional expectations in equation (5), these two equations still do not imply the standard Euler equation relating aggregate consumption spending in two consecutive periods to the real rate of return between those two periods. Finally, substituting into equation (5) the equilibrium requirement that $C_t = Y_t - G_t$, where G_t represents exogenous variation in government purchases of the composite good, we obtain an equilibrium relation between the index of aggregate demand Y_t and variations in the marginal utility of income, which provides the crucial link in our model between interest rate variations and aggregate demand.

For our numerical work, we rely on log-linear approximations to the model's structural equations. We assume an equilibrium in which the economy always

stays near a steady state path, which represents a stationary, deterministic equilibrium in the case of no exogenous disturbances ($\xi_t = 0$ and $G_t = \bar{G}$ at all times) and a monetary policy consistent with stable prices. In this steady state, output is constant at a level \bar{Y} (defined below), and consumption is constant at the level $\bar{C} \equiv \bar{Y} - \bar{G}$.¹ It follows that the marginal utility of *real* income, $\lambda_t P_t$, is also constant, at the value $\bar{\lambda} \equiv u_c(\bar{C}; 0)$. We log-linearize the structural equations of the model around these steady state values. Percentage deviations in the marginal utility of consumption $u_c(C_t; \xi_t)$ around the steady state value $u_c(\bar{C}; 0)$ can be written as $-\tilde{\sigma}(\hat{C}_t - \bar{C}_t)$, where $\hat{C}_t \equiv \log(C_t/\bar{C})$, \bar{C}_t is an exogenous shift variable (a certain linear combination of the elements of ξ_t),² and $\tilde{\sigma} \equiv -u_{cc}\bar{C}/u_c$, where the partial derivatives are evaluated at the steady state level of consumption. With this substitution, the log-linear approximations to equations (5) and (6) are given by

$$(7) \quad -\tilde{\sigma}E_{t-2}(\hat{C}_t - \bar{C}_t) = E_{t-2}\hat{\lambda}_t,$$

$$(8) \quad \hat{\lambda}_t = \hat{R}_t - \pi_{t+1} + E_t\hat{\lambda}_{t+1},$$

where $\hat{R}_t \equiv \log(R_t/R^*) = \log(\beta R_t)$ is the percentage deviation of the short-run nominal interest rate from its steady state value, $\pi_t \equiv \log(P_t/P_{t-1})$ is the inflation rate, and $\hat{\lambda}_t \equiv \log(\lambda_t P_t/\bar{\lambda})$ measures the percentage deviation of the marginal utility of real income from its steady state value. (Eq. [8] refers to actual rather than expected inflation because inflation π_{t+1} is known with certainty at date t in our model.)

A similar log-linear approximation to the market-clearing condition allows us to replace $\hat{C}_t - E_{t-2}\bar{C}_t$ with $s_c^{-1}(\hat{Y}_t - \hat{G}_t)$, where $s_c \equiv \bar{C}/\bar{Y}$, $\hat{Y}_t \equiv \log(Y_t/\bar{Y})$, and \hat{G}_t collects the exogenous disturbance terms that shift the relation between aggregate demand and the marginal utility of consumption. Substituting this into equation (7) yields

$$(9) \quad E_{t-2}\hat{\lambda}_t = -\sigma E_{t-2}(\hat{Y}_t - \hat{G}_t),$$

where $\sigma \equiv s_c^{-1}\tilde{\sigma}$. Then taking the conditional expectation of equation (8) two periods earlier and substituting equation (9), we obtain

$$(10) \quad E_{t-2}(\hat{Y}_t - \hat{G}_t) = -\sigma^{-1}E_{t-2}(\hat{R}_t - \pi_{t+1}) + E_{t-2}(\hat{Y}_{t+1} - \hat{G}_{t+1}).$$

(Thus, in our log-linear approximation, the standard Euler equation does hold, but only conditional on lagged information.) Solving forward, we may equivalently write

1. Throughout these derivations, we assume an economy with zero growth for simplicity. More properly, we assume a deterministic trend for real activity, and variables such as Y_t refer to detrended values, which take constant values in the steady state. In our estimation of the model, we use a series for \hat{Y}_t obtained by removing a linear trend from the log of real GDP.

2. Details of this and other aspects of our Taylor series expansions are presented more fully in the appendix.

$$(11) \quad \hat{Y}_t = \hat{G}_t - \sigma^{-1} E_{t-2} \sum_{T=t}^{\infty} (\hat{R}_T - \pi_{T+1}).$$

Equation (11) plays a role analogous to the “IS equation” of traditional Keynesian models but is consistent with intertemporal optimization.³ It relates output to the long-run real interest rate (with a negative sign) and to autonomous spending disturbances. The latter include disturbances both to private impatience to consume resources and to government spending, summarized in the composite disturbance term \hat{G}_t . We assume that \hat{G}_t is determined at $t - 1$, so that it is determined after C_t has already been chosen, but in time for the central bank to adjust the period t interest rate R_t in response to it. Letting \hat{G}_t be determined after \hat{C}_t ensures that output is not predetermined as of $t - 2$ (i.e., it allows us an interpretation for the output innovations in our vector autoregression—VAR—model of the U.S. data), even though output responds with a two-period delay to exogenous disturbances to monetary policy.

The source of the real effects of monetary policy in our model is that prices do not adjust immediately to shocks. Following Calvo (1983), we assume that prices are changed at exogenous random intervals.⁴ Specifically, a fraction $1 - \alpha$ of sellers get to choose new prices at the end of any given period, whereas the others must continue using their old prices. Of those who get to choose new prices, a fraction γ start charging the new price at the beginning of the next period. The remaining fraction $1 - \gamma$ must wait until the following period to charge the new price, or put differently, they must post their prices one quarter in advance. These assumed delays explain why no prices respond in the quarter of the monetary disturbance and why the largest response of inflation to a monetary shock takes place only two quarters after the shock.

Let p_t^1 denote the price set by sellers that decide at date $t - 1$ on a new price to take effect at date t , and p_t^2 the price set by sellers that decide at date $t - 2$ on a new price to take effect only two periods later. These prices are chosen to maximize the contributions to expected utility resulting from sales revenues on the one hand, and the disutility of output supply on the other, at each of the future dates and in each of the future states in which the price commitment still applies. This means that p_t^1 is chosen to maximize

$$(12) \quad \Phi_{t-1}(p) \equiv E_{t-1} \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [\lambda_T (1 - \tau) p Y_T (p/P_T)^{-\theta} - v(Y_T (p/P_T)^{-\theta}; \xi_T)]$$

over p . Here we have substituted the demand function (3) into the household's objective function, written λ_T for the marginal utility (in units of period T util-

3. Except for our introduction of the two-period delay in the determination of interest-sensitive purchases, our derivation of this “expectational IS equation” follows the earlier work of authors such as Koenig (1987), Kerr and King (1996), McCallum and Nelson (forthcoming), and Woodford (1996).

4. This general approach to modeling the dynamics of price adjustment is adopted in a large number of recent quantitative equilibrium business cycle studies, beginning with Yun (1996) and King and Watson (1996).

ity flow) of additional nominal income during period T , and assumed that revenues each period are taxed at the constant rate τ .⁵ The factor α^{T-t} appears as the probability that the price that is charged beginning in period t is still in effect in period $T \geq t$ (where we assume that this contingency is independent of all aggregate disturbances). Note that our assumption of complete contingent claims markets (including full opportunities for households to insure one another against idiosyncratic risk associated with different timing of their price changes) implies that the marginal utility of income process $\{\lambda_t\}$ is the same for all households and can be treated as an exogenous stochastic process by an individual household (whose pricing decisions will have only a negligible effect on aggregate prices, aggregate incomes, and aggregate spending decisions). Similarly, an individual household treats the processes $\{P_t, Y_t\}$ as exogenous in choosing its desired price. The optimizing choice of p_t^1 then must satisfy the first-order condition

$$(13) \quad \Phi'_{t-1}(p_t^1) = 0,$$

where the prime denotes the derivative with respect to p in the explicit expression given in equation (12).

As before, we wish to log-linearize this equilibrium condition around a steady state in which $Y_t = \bar{Y}$, $P_t/P_{t-1} = 1$, $p_t^1/P_t = 1$, and $\lambda_t P_t = \bar{\lambda}$ at all times. (The requirement that these constant values satisfy equation [13] when $\xi_t = 0$ at all times determines the steady state value \bar{Y} .)⁶ Percentage deviations of $v_y(y_t^j; \xi_t)$ from its steady state value can be written as $\omega(\hat{y}_t^j - \bar{Y}_t)$, where $\omega \equiv v_{yy}\bar{Y}/v_y$, with partial derivatives evaluated at the steady state, $\hat{y}_t^j \equiv \log(y_t^j/\bar{Y})$, and \bar{Y}_t is a certain linear function of ξ_t . Using this notation, the log-linear approximation to equation (13) takes the form

$$(14) \quad E_{t-1} \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [(\hat{\lambda}_T + \hat{Y}_T - (\theta - 1)\hat{p}_{t,T}^1) - \omega(\hat{Y}_T - \theta\hat{p}_{t,T}^1 - \bar{Y}_T) - (\hat{Y}_T - \theta\hat{p}_{t,T}^1)] = 0,$$

where in addition $\hat{p}_{t,T}^1 \equiv \log(p_t^1/P_T)$. Introducing the notation⁷

$$\hat{X}_t \equiv \frac{1 - \alpha}{\alpha} \log\left(\frac{p_t^1}{P_t}\right),$$

so that

5. The allowance for nonzero τ is primarily so that we can linearize around a steady state in which the constant level of output is efficient. The convenience of this for our purposes is discussed in the appendix.

6. Under the assumption that $\tau = -(\theta - 1)^{-1}$, this requires that \bar{Y} satisfy the equation $u_c(\bar{Y} - \bar{G}; 0) = v_y(\bar{Y}; 0)$, which also defines the efficient level of output.

7. The factor $(1 - \alpha)/\alpha$ turns out to be convenient in giving a simpler form to equations such as eq. (20) below.

$$\hat{p}_{i,T}^1 = \frac{\alpha}{1 - \alpha} \hat{X}_t - \sum_{s=t+1}^T \pi_s,$$

we can solve equation (14) to obtain

$$(15) \quad \hat{X}_t = \frac{1 - \alpha}{\alpha} \frac{1 - \alpha\beta}{1 + \omega\theta} E_{t-1} \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [-\hat{\lambda}_T + \omega(\hat{Y}_T - \bar{Y}_T) + (1 + \omega\theta) \sum_{s=t+1}^T \pi_s]$$

as the optimizing choice of the relative price in period t of goods with new prices chosen just the period before.

We can use equation (9) to eliminate the $E_{t-1}\hat{\lambda}_T$ terms in equation (15), for all $T > t$. Taking the conditional expectation of equation (8) at $t - 1$ and using equation (9), we see that we can also write

$$(16) \quad E_{t-1}\hat{\lambda}_t = \phi_{t-1} - \sigma E_{t-1}(\hat{Y}_t - \hat{G}_t),$$

where

$$(17) \quad \begin{aligned} \phi_t &\equiv E_t[\hat{R}_{t+1} - \pi_{t+2} - \sigma(\hat{Y}_{t+2} - \hat{G}_{t+2} - \hat{Y}_{t+1} + \hat{G}_{t+1})] \\ &= E_t \sum_{T=t+1}^{\infty} (\hat{R}_T - \pi_{T+1}) - E_{t-1} \sum_{T=t+1}^{\infty} (\hat{R}_T - \pi_{T+1}). \end{aligned}$$

Note that the final equality in expression (17) follows from substitution of (11). Then, substituting equations (9) and (16) into (15), we obtain

$$(18) \quad \begin{aligned} \hat{X}_t &= \frac{1 - \alpha}{\alpha} \frac{1 - \alpha\beta}{1 + \omega\theta} \left\{ E_{t-1} \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [(\sigma + \omega)(\hat{Y}_T - \hat{Y}_T^s) + (1 + \omega\theta) \sum_{s=t+1}^T \pi_s] - \phi_{t-1} \right\} \\ &= \frac{1 - \alpha}{\alpha} \frac{1 - \alpha\beta}{1 + \omega\theta} \left\{ E_{t-1} \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} [(\sigma + \omega)(\hat{Y}_T - \hat{Y}_T^s) + (1 + \omega\theta) \frac{\alpha\beta}{1 - \alpha\beta} \pi_{T+1}] - \phi_{t-1} \right\}, \end{aligned}$$

where

$$\hat{Y}_t^s \equiv \frac{\omega}{\omega + \sigma} E_{t-1} \bar{Y}_t + \frac{\sigma}{\omega + \sigma} \hat{G}_t$$

is a composite exogenous disturbance. We can think of \hat{Y}_t^s as representing variation in the “natural” or “potential” level of output, since it is expected deviations $\hat{Y} - \hat{Y}^s$, rather than deviations in the level of output relative to trend, that results in a desire by price setters to increase the *relative* price of their goods,

which in equilibrium requires inflation of the average level of prices. (An equilibrium in which no prices are ever changed is consistent with eq. [18] as long as $\hat{Y}_t = \hat{Y}_t^S$ at all times and interest rates vary so as to ensure that $\phi_t = 0$ at all times. Note that the latter condition ensures that eq. [10] and hence [11] is also satisfied at all times.)

We turn next to the price-setting decision of sellers that choose a new price p_t^2 at $t - 2$ to apply beginning in period t . Because such a price is expected to apply in periods $t + j$ with exactly the same probabilities as for the price p_t^1 , the objective of these sellers is simply $E_{t-2}\Phi_{t-1}(p)$, and the first-order condition that determines p_t^2 is given by $E_{t-2}\Phi_{t-1}(p_t^2) = 0$. Comparison with equation (13) implies that, in our log-linear approximation,

$$(19) \quad \log p_t^2 = E_{t-2} \log p_t^1.$$

Finally, our definition of the price index (4) implies that this index evolves according to

$$P_t = [\alpha P_{t-1}^{1-\theta} + (1 - \alpha)\gamma(p_t^1)^{1-\theta} + (1 - \alpha)(1 - \gamma)(p_t^2)^{1-\theta}]^{1/(1-\theta)}.$$

Dividing both sides by P_t , log-linearizing, and substituting equation (19), we obtain

$$\pi_t = \gamma \hat{X}_t + (1 - \gamma) \left[E_{t-2} \hat{X}_t - \frac{1 - \alpha}{\alpha} (\pi_t - E_{t-2} \pi_t) \right].$$

Taking the conditional expectation of both sides at $t - 2$, one observes that $E_{t-2} \pi_t = E_{t-2} \hat{X}_t$. Substitution of this then allows the equation to be written in the form

$$(20) \quad \pi_t = \psi \hat{X}_t + (1 - \psi) E_{t-2} \hat{X}_t,$$

where $\psi \equiv \gamma\alpha/[1 - \gamma(1 - \alpha)]$. This indicates how aggregate inflation results from the incentives of individual price setters to choose a higher relative price.

These results may be collected in the form of an implied aggregate supply relation between inflation variation and deviations of output from potential. Equation (18) may be expressed in quasi-differenced form as

$$(21) \quad \begin{aligned} \hat{X}_t &= \kappa(\hat{Y}_t - \hat{Y}_t^S) + (1 - \alpha)\beta E_{t-1} \pi_{t+1} + \alpha\beta E_{t-1} \hat{X}_{t+1} - \frac{\kappa}{\sigma + \omega} \phi_{t-1} \\ &= \kappa(\hat{Y}_t - \hat{Y}_t^S) + \beta E_{t-1} \hat{X}_{t+1} - \frac{\kappa}{\sigma + \omega} \phi_{t-1}, \end{aligned}$$

where $\kappa \equiv (1 - \alpha)(1 - \alpha\beta)(\omega + \sigma)/\alpha(1 + \omega\theta)$, and the second line follows from the fact that equation (20) implies that $E_{t-2} \pi_t = E_{t-2} \hat{X}_t$. Solving this forward, we obtain

$$\hat{X}_t = \kappa E_{t-1} \left[\sum_{T=t}^{\infty} \beta^{T-t} (\hat{Y}_T - \hat{Y}_T^S) \right] - \frac{\kappa}{\sigma + \omega} \phi_{t-1},$$

where we have used the fact that equation (17) implies that $E_{t-1}\phi_t = 0$. Substitution of this into equation (20) then yields

$$(22) \quad \begin{aligned} \pi_t = & (1 - \psi)E_{t-2}\pi_t + \psi \left\{ \kappa E_{t-1} \sum_{T=t}^{\infty} \beta^{T-t} (\hat{Y}_T - \hat{Y}_T^s) \right. \\ & \left. - \frac{\kappa}{\omega + \sigma} \left[E_{t-1} \sum_{T=t}^{\infty} (\hat{R}_T - \pi_{T+1}) - E_{t-2} \sum_{T=t}^{\infty} (\hat{R}_T - \pi_{T+1}) \right] \right\}. \end{aligned}$$

This is our aggregate supply (AS) equation, relating inflation variation to deviations of output from potential. Because prices are set in advance, expectations of future increases in output relative to Y^s also raise prices. In addition, inflation declines when the long-term real interest rate at t is higher than had been expected at $t - 1$. The reason for this is that such upward revisions raise the returns households can expect to earn from their revenues at t . As a result, they are inclined to raise these revenues by cutting their prices. Only surprise variations in the long rate contribute to this term because only those variations result in changes in the current marginal utility of income that are not reflected in the current level of aggregate consumption demand, and hence in the output gap.

To complete our model specification, we posit that interest rates are set according to a feedback rule of the form

$$(23) \quad r_t - r^* = \sum_{j=0}^{m_\pi} a_j (\pi_{t-j} - \pi^*) + \sum_{j=0}^{m_Y} b_j \hat{Y}_{t-j} + \sum_{j=1}^{m_R} c_j (r_{t-j} - r^*).$$

Here r_t is the continuously compounded nominal interest rate (identified with $\log R_t$ in terms of our theoretical model, and with the federal funds rate in our empirical implementation of the model), r^* is the steady state value of r implied by the policy rule, and π^* is the steady state inflation rate implied by the rule. In equilibrium, the steady state nominal interest rate r^* must equal the sum of the equilibrium steady state real interest rate ρ and the steady state inflation rate π^* . Thus, if ρ is independent of the monetary policy rule (as our model implies),⁸ the monetary authority's choice of π^* implies a value for r^* . Thus the pair of values π^* and r^* represent only a single free parameter in the specification of the policy rule, which we shall treat in the subsequent discussion as the choice of π^* .⁹

The aim of our paper is to discuss the effects of alternative rules of the form

8. Up to the log-linear approximation used in all of our computations of the equilibria associated with alternative policy rules, the steady state real interest rate is given by $\rho = -\log \beta$, as a consequence of eq. (6).

9. We need not assume that the monetary authority actually knows the true value of ρ . The rule (23) involves a single constant term $K = (1 - \sum_j c_j)r^* - \sum_j a_j \pi^*$, and it is this that the authority must know in order to implement the rule. However, according to our model, a given value of K implies (generically) a unique value of π^* (which may or may not be correctly estimated by the authority). We choose to parameterize alternative policy rules in terms of π^* rather than in terms of K because of the simpler interpretation of the latter parameter.

(23). In our discussion, we will generally treat separately the effects of the parameters a_j , b_j , and c_j , which indicate how the interest rate reacts to the history of the economy, and the effects of the choice of π^* . This is because, in our log-linear approximation to the model's equilibrium conditions, the parameter π^* has no effect on the implied responses to shocks (and hence on the equilibrium variability of the various state variables), while the parameters a_j , b_j , and c_j have no effect on the implied steady state (and hence on the average equilibrium values of the state variables). We may thus study separately the determination of the steady state and the determination of fluctuations around the steady state, and different parameters of the policy rule matter for each of these investigations. Our overall welfare criterion (discussed in the next subsection) depends, however, on both aspects of equilibrium, and so on both sets of policy parameters.

Our complete model of the economy consists of the IS and AS equations (11) and (22) together with the monetary policy rule (23). To evaluate the effect of changing the monetary rule we need to know both the parameters of the model as well as the stochastic process for the two structural disturbances \hat{G}_t and \hat{Y}_t^S , the first of which affects only our IS equation while the second affects only our AS equation. In Rotemberg and Woodford (1997, 1998) we describe both our method for estimating and calibrating the behavioral parameters and our approach to reconstructing the structural disturbances and their stochastic process. Here we give an outline of this approach.

We start with a recursive VAR model of the state vector¹⁰

$$(24) \quad Z_t = [\hat{r}_t, \hat{\pi}_{t+1}, \hat{Y}_{t+1}]',$$

where $\hat{r}_t \equiv r_t - r^*$ and $\hat{\pi}_t \equiv \pi_t - \pi^*$.¹¹ We estimate a system of the form

$$(25) \quad \bar{Z}_t = B\bar{Z}_{t-1} + U\bar{e}_t,$$

where the vector \bar{Z}_t is the transpose of $[Z_t', Z_{t-1}', Z_{t-2}']$ and U is a lower triangular matrix with ones on the diagonal and nonzero off-diagonal elements only in the first three rows, the off-diagonal elements of which are estimated so as to make the residuals in \bar{e}_t orthogonal to one another. The first three rows of the vectors \bar{e}_t contain the VAR residuals $e_{1,t}$, $e_{2,t}$, and $e_{3,t}$ while the other elements are zero. The number of lags included in our VAR is sufficient to eliminate nearly all evidence of serial correlation in the disturbances.

The first equation in this VAR is our estimate of the monetary policy rule.

10. In fact, we estimate a VAR model of r_t and π_t that includes constant terms in the equations and use these constant terms to obtain our econometric estimates of r^* and π^* . The latter estimates then imply our estimated value for the model parameter ρ (and hence β). In the exposition here, we drop the constant terms for simplicity.

11. Note that, with these definitions, $\hat{r}_t = \hat{R}_t - \pi^*$. The difference in the definitions follows from the difference in the rate of inflation in the steady state with respect to which deviations are calculated under the two definitions. It is also worth noting that \hat{r}_t bears the same relation to \hat{R}_t as $\hat{\pi}_t$ bears to π_t .

This estimated rule has the same structure as rule (23), except that it also includes a white noise residual $e_{1,t}$. Note that while the interest rate comes first in the casual ordering, the timing of the variables ensures that the interest rate in period t responds to inflation and output in period t , while these variables only react to lagged interest rates. We suppose that $e_{1,t}$ is independent of the two “real” disturbances \hat{Y}_t^s and \hat{G}_t so that it is exclusively a monetary policy disturbance. (Note that these identifying assumptions are ones that are implied by the decision lags assumed in our theoretical model.) Under these assumptions, we can estimate not only the coefficients of the historical monetary policy rule but also the impulse responses of output, inflation, and the interest rate to a monetary policy disturbance. We can then recover most of the structural parameters of our model by minimizing the discrepancy between the estimated responses of these variables to the monetary disturbance $e_{1,t}$ and the responses predicted by our theoretical model when the systematic part of the monetary policy rule is given by the estimated coefficients in equation (23).¹² By calibrating the remaining parameters on the basis of other evidence, we obtain numerical values for the model parameters α , β , γ , σ , θ , and ω . These are, respectively, 0.66, 0.99, 0.63, 0.16, 7.88, and 0.47, so that κ equals 0.024 and ψ equals 0.53.

Armed with our parameter values and the VAR, we can reconstruct the stochastic processes for the structural disturbances as follows.¹³ Equation (11) gives \hat{G}_t as the sum of \hat{Y}_t and σ times the expected long-term real rate. Given that the VAR allows us to forecast both inflation and interest rates, this expected long-term real rate is a function of \bar{Z}_t . Similarly, solving equation (20) for \hat{X}_t as a function of inflation and expected inflation, and substituting this into (21), we find that \hat{Y}_t^s must be given by

$$\hat{Y}_t^s = \hat{Y}_t - \frac{1}{\kappa\psi}\pi_t + \frac{1-\psi}{\kappa\psi}E_{t-2}\pi_t + \frac{\beta}{\kappa}E_{t-1}\pi_{t+1} - \frac{1}{\omega+\sigma}\phi_{t-1}.$$

Furthermore, using equation (17), the last term in this equation can be written as a function of expectations of future interest rates and inflation rates, as of periods $t-1$ and $t-2$. Using the VAR to forecast future variables, the right-hand side of the equation then depends just on \bar{Z}_{t-1} , on \bar{Z}_{t-2} , and on the model parameters. It is thus straightforward to use the structural parameters as well as the matrices B and U to compute matrices C and D such that

$$(26) \quad s_t \equiv [\hat{G}_{t+1}, \hat{Y}_{t+1}^s]^\top = C\bar{Z}_{t+1} + D\bar{e}_t.$$

The resulting historical time series for the two disturbances s_t could then be identified with the residuals of the model’s structural equations.

If the model fit the properties of the U.S. time series perfectly, the vector s_t

12. Our estimation strategy is discussed in more detail in appendix 1 of Rotemberg and Woodford (1998).

13. Rotemberg and Woodford (1998) provides more details about both this method of construction and the properties of the constructed series.

constructed in this way would be orthogonal to the identified monetary policy disturbance $e_{1,t}$. In practice, the right-hand side of equation (26) does depend on the first element of the vector of VAR residuals \bar{e}_t , which we identify as the monetary policy disturbance. Perhaps more troubling is the observation that if the real disturbances s_t are generated by a law of motion of the kind implied by conjoining equation (26) with equation (25) for the evolution of \bar{Z}_t , then we should not expect all three of the independent structural disturbances \bar{e}_t that matter for the evolution of s_t to be revealed by data on the three variables in Z_t alone. (This is because one of the VAR innovations corresponds to the monetary policy shock, so that only the other two orthogonal innovations can reveal information about the real disturbances.) But this would mean that forecasts of the future values of the variables in Z_t using the VAR should not correspond, in principle, to the expectations of these variables conditional on the public's information set (assuming that the public has complete information about the structural disturbances); and thus our method for identifying the historical series for our structural equation residuals would not be internally consistent.

We prefer instead to work with a theoretical model not subject to this last problem, that is, one in which the evolution of the real disturbances s_t depends only on two orthogonal disturbances each period, which then should in principle correspond to the two VAR residuals $\bar{e}_{2,t}$ and $\bar{e}_{3,t}$. The structural disturbances s_t of our theoretical model then have moments that do not correspond precisely to those of the residuals of our model equations; but this discrepancy will exist only insofar as our model (quite apart from the law of motion chosen for the structural disturbances) is in fact inconsistent with the estimated VAR (and in particular, with the estimated impulse responses to a monetary policy shock). We accordingly consider a law of motion

$$(27) \quad s_t = CZ_{t-1}^\dagger + D^\dagger e_t^\dagger$$

for the structural disturbances, where the matrix C is the one referred to in equation (26), D^\dagger corresponds to D with the first column deleted, and e_t^\dagger is a vector of two orthogonal white noise disturbances (which correspond to $\bar{e}_{2,t}$ and $\bar{e}_{3,t}$). Here Z_t^\dagger is a vector of exogenous state variables that evolve according to

$$(28) \quad Z_t^\dagger = BZ_{t-1}^\dagger + U^\dagger e_t^\dagger,$$

where B is the same as in equation (25) and U^\dagger corresponds to U with the first column deleted.

Note that because the elements of Z_t^\dagger refer to exogenous states (underlying states for the dynamics of the real disturbances s_t), unlike the elements of Z_t (which correspond to endogenous variables of our model), this specification does *not* imply the existence of any feedback from the evolution of the endogenous variables to the exogenous disturbance processes s_t . What this construc-

tion *does* guarantee is that the empirical impulse response functions of inflation, output, and interest rates to the two VAR disturbances orthogonal to the monetary policy shock are identical to the impulse responses predicted by our theoretical model. This property of the predicted impulse responses is independent of the structural parameters assumed in the model. Thus, given this method for constructing the laws of motion for the real disturbances, only the estimated responses to the monetary policy shock contain any information that can be used to help identify the structural parameters. This is our justification for the strategy that we use for parameter estimation, mentioned above.

It is worth noting that the stochastic processes for the real disturbances that we obtain with this method imply a great deal of variability for both \hat{G}_t and \hat{Y}_t^s . For example, the standard deviations of these two series are 29.5 and 13.7 percent, respectively.¹⁴ This extreme volatility is consistent with the fact that the literature reports many “failures” in fitting equations very similar to our IS and AS curves by either ordinary least squares or by using lags as instruments. Our interpretation of these “failures” is that they say simply that these equations are subject to disturbances whose variance is large and whose serial correlation pattern is rich (so that they are correlated with the lags that are used as instruments).

In this paper, we evaluate monetary rules by evaluating how well they perform when the economy is buffeted by these shocks to \hat{G} and \hat{Y}^s . In other words, we are asking how the U.S. economy would perform if it were subject to structural disturbances whose properties are the same as those that have affected it in the past while, at the same time, the way interest rates are set by the central bank is different. Because the structural equations (11) and (22) follow simply from the Euler equations for optimal intertemporal behavior on the part of households, and so can be derived without reference to any particular specification of the monetary policy rule, they should remain invariant under contemplated changes in that rule. Thus our stochastic simulation methodology responds to the Lucas (1976) critique of more traditional methods of econometric policy evaluation.¹⁵

14. This compares, e.g., with a standard deviation of only 2.1 percent for our detrended output process \hat{Y}_t . The high volatility of the constructed \hat{G}_t process is mainly due to its high serial correlation (serial correlation coefficient of .92), rather than to extraordinary volatility of the \hat{G}_t innovations, which correspond in fact to the \hat{Y}_t innovations in our VAR model. It is possible that the data would be better described by a model in which C_t , \bar{C}_t , and G_t are not required to have a common deterministic trend. The volatility of the constructed \hat{Y}_t^s process, instead, is largely due to the presence of a very volatile transitory component. E.g., the standard deviation of $E_{t-2}\hat{Y}_t^s$ is only 4.3 percent.

15. It should go without saying that this does not imply that the model is necessarily correct. If our model is incorrectly specified, changes in the monetary policy rule will have effects other than those implied by our analysis. What makes the model preferable to purely backward-looking models is that, as stressed by Lucas (1976), it is highly implausible that purely backward-looking specifications of IS and AS curves will remain invariant with respect to changes in the monetary rule.

2.1.2 The Welfare Loss from Price Level Instability

One of the primary advantages of our derivation of our structural equations from explicit optimizing foundations is that we are able to evaluate alternative monetary policy rules in terms of their welfare effects. Specifically, we consider the effects on the average level of welfare

$$(29) \quad W = E\left(u(C_t; \xi_t) - \int_0^1 v(y_t(z); \xi_t) dz\right)$$

in the stationary equilibrium associated with one or another policy rule within the class that we consider.

Here the expectation is over alternative possible histories of the preference shocks ξ_t , (which include the effects of technology shocks, since technological possibilities are implicit in our assumed disutility of supplying output). We only consider the welfare associated with alternative stationary rational expectations equilibria, in which all relative quantities are stationary and all quantities are trend stationary. Thus we can evaluate an unconditional expectation in expression (29) for each of the equilibria that we consider. We also restrict our attention to monetary policy rules that result in unique stationary rational expectations equilibria (in terms of inflation, all relative prices, detrended output, and all relative quantities); we thus obtain a unique welfare measure for each policy rule in the admissible set. Given that we evaluate the unconditional expectation, rather than conditioning on the current state of the economy at some particular date at which the policy choice is to be made, the criterion (29) is equivalent to comparing equilibria on the basis of the average level of expected utility of the households in our model (for the unconditional expectation of the latter quantity is simply $(1 - \beta)^{-1}W$). We evaluate the unconditional expectation in order to obtain a policy evaluation criterion that is not subject to any problem of time consistency.¹⁶

Following Rotemberg and Woodford (1998), we take a second-order Taylor series approximation of this welfare measure around the steady state values of the stationary variables that affect utility. The “steady state values” represent the constant equilibrium values of these variables in the absence of real disturbances, and in the case of a deterministic monetary policy consistent with zero inflation.¹⁷ The steady state considered for this purpose also involves a tax rate τ that is set so that the steady state level of output is efficient. (This involves a small output *subsidy*, in order to counteract the distortion caused by monopoly power.) Consideration of a Taylor series expansion around these values means that our approximate welfare measure will accurately rank alternative policy rules insofar as they result in only a small degree of variability of the relevant

16. For an alternative approach, cf. King and Wolman (1998).

17. The rate of inflation matters for the evaluation of eq. (29) in such a steady state because it determines the dispersion of relative prices and hence the dispersion of the relative quantities produced of the various goods z .

state variables *and* they result in average values of the state variables that are close to the assumed steady state values. Thus our analysis should be most reliable in the case of rules that imply an average rate of inflation not too different from zero and an average level of output near the optimal steady state level \bar{Y} , and in which the fluctuations in both inflation and output are small. In fact, the policies that we characterize as optimal within various families of possible policy rules all imply low inflation rates, and also low variability of inflation and output, in the case that the variability of the real disturbances (represented by ξ_t) is small enough.

Linearization around this particular (optimal) steady state is extremely convenient since our approximate measure of W takes an especially simple form in that case. In particular, in this case our second-order approximation for W depends only on a first-order approximation to the equilibrium responses of inflation and output to the exogenous shocks. This means that we can solve a log-linear approximation to the model's equilibrium conditions using standard linear methods, as sketched in the previous subsection, and obtain an approximation to W that neglects only terms of third order and higher in the deviations from the steady state. This result depends on the absence of any first-order contribution to our welfare measure from changes in the average level of output under alternative rules (as a result of the optimality of the level \bar{Y} relative to which we consider deviations); for if W contained a term of first order in the average level of output, then *second-order* terms in the equations determining output would matter for a second-order approximation to W .

In fact, in the calculations reported here, we furthermore assume that the tax rate τ actually varies depending on the monetary policy rule, so as to ensure that $E \log Y_t = \log \bar{Y}$ in any event. This allows us to obtain a measure of the deadweight loss associated with price level instability that abstracts from any effects of alternative monetary policies on the long-run average level of output. While many analyses of the welfare effects of monetary policy have emphasized exactly such effects,¹⁸ we think there is good reason to abstract from them. Our primary reason is that there exist other policy instruments, such as the general level of and structure of taxation, which allow the government to influence the average level of output while, at the same time, being much less well suited for the achievement of stabilization objectives since they cannot be adjusted quickly and precisely in response to shocks. It thus makes sense to assume that, in an optimal policy regime, the other instruments are chosen to achieve the desired average level of output for a given monetary policy, while the monetary policy rule is chosen to minimize those contributions to deadweight loss that are independent of the economy's average level of output. We do this by choosing the monetary policy rule that maximizes W under the assumption that the other instruments are adjusted in the manner stated in response to any change in the monetary policy rule.

18. See, e.g., Bénabou and Konieczny (1994), King and Wolman (1996), and Feldstein (1997).

Abstracting from these effects also has the advantage of making our results independent of a feature of our model about which we are especially uncomfortable, namely, its predictions about the effects of sustained inflation on the long-run level of output.¹⁹ One might think that sustained inflation should result in adaptations that eliminate any effects of the average inflation rate on average output. One such adaptation would be price commitments that specify a constant rate of price increase of π^* between the occasions on which the commitments are modified, as assumed in Yun (1996). With this modification, our model would come to satisfy the “natural rate hypothesis.” In the modified model, the correct second-order approximation to W would be exactly the one that we report here, but then it would apply to small fluctuations in the rate of inflation around *any* average value π^* .²⁰

We show in the appendix that, under these assumptions, a second-order approximation for W is given by

$$(30) \quad \begin{aligned} W = & -\frac{1}{2}u_c\bar{Y}(\sigma + \omega) \text{var}\{E_{t-2}(\hat{Y}_t - \hat{Y}_t^s)\} \\ & - \frac{1}{2}u_c\bar{Y}(\theta^{-1} + \omega)E \text{var}_z\{\log y_t(z)\} \\ & + \text{terms independent of policy} + \mathcal{O}(\|\xi\|^3), \end{aligned}$$

where the suppressed final terms are either independent of the evolution of the endogenous variables or of third order or smaller in the size of the exogenous disturbances. Note that this welfare measure depends solely on the allocation of real resources, summarized by the pattern of levels of production $\{y_t(z)\}$ at each point in time. However, equation (30) indicates that welfare depends not only on the degree to which aggregate output deviates from the natural level of output Y^s but also on the degree of (inefficient) dispersion of output levels across the different varieties of goods being produced at each point in time.

The dispersion of output levels directly corresponds, in equilibrium, to the degree of dispersion of output prices. Prices differ across goods, in turn, only because of variation in the overall price level (together with the fact that different suppliers adjust their prices at different times). The $E \text{var}_z\{\log y_t(z)\}$ term in equation (30) can accordingly be expressed as a function of the aggregate inflation process, as shown in the appendix. With this substitution, we obtain

19. E.g., McCallum and Nelson (chap. 1 of this volume) criticize the Calvo model of price setting on the ground that its failure to conform to the natural rate hypothesis is unrealistic. Our closely related model has exactly the feature that they criticize.

20. King and Wolman (1998) also argue, on alternative grounds, that optimal policy should involve a steady state inflation rate of $\pi^* = 0$, despite the fact that a small positive inflation rate can raise steady state output by lowering average markups, and that (if one assumes $\tau = 0$, as they do) this would raise the steady state value of the period utility flow $u(C) - v(y)$. Their argument involves calculation of the optimal time-dependent policy (under commitment) to maximize the average level of discounted utility over the infinite horizon. They show that this optimal time-dependent (and time-inconsistent) policy involves a commitment to an inflation rate that converges to zero asymptotically, even though the optimal *stationary* rate of inflation would be positive.

$$\begin{aligned}
(31) \quad W = & -\frac{1}{2}u_c \bar{Y} \left\{ (\omega + \sigma) \text{var}\{E_{t-2}(\hat{Y}_t - \hat{Y}_t^s)\} \right. \\
& + \frac{\theta(1 + \theta\omega)}{(1 - \alpha)^2} \left[\alpha \text{var}\{E_{t-2}\pi_t\} + \alpha\pi^{*2} \right. \\
& \left. \left. + \left(\alpha + \frac{1 - \gamma}{\gamma} \right) \text{var}\{\pi_t - E_{t-2}\pi_t\} \right] \right\} + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3).
\end{aligned}$$

Here π^* again denotes the steady state rate of inflation associated with a given policy rule (the rate of inflation when the shocks $\xi_t = 0$ for all time); it corresponds, neglecting terms of second order or higher, to the average rate of inflation (or to the unconditional mean of π) in the stationary equilibrium. The notation “t.i.p.” refers to the terms that are independent of policy.

Expression (31) can be written more compactly as

$$(32) \quad W = -\Omega(L + \pi^{*2}) + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3),$$

where

$$(33) \quad L = \text{var}\{\pi_t\} + (\psi^{-1} - 1) \text{var}\{\pi_t - E_{t-2}\pi_t\} + \Lambda \text{var}\{E_{t-2}(\hat{Y}_t - \hat{Y}_t^s)\}$$

and $\Omega, \Lambda > 0$. The quantity $L + \pi^{*2}$ represents the measure of deadweight loss due to price level instability that we shall use to evaluate alternative monetary policies. Here the loss measure L collects the terms that depend solely on the degree of variability of inflation and the output gap, while π^{*2} is proportional to the deadweight loss due to nonzero inflation, even when it is perfectly steady.²¹

Note that our loss measure L is similar in form to a type of ad hoc loss function,

$$\text{var}\{\pi_t\} + \lambda \text{var}\{\hat{Y}_t - \hat{Y}_t^s\},$$

for some $\lambda > 0$, assumed in many analyses of optimal monetary policy (e.g., Taylor 1979; Bean 1983). Our utility-based derivation, however, allows us to assign a specific numerical weight to the relative importance of stabilization of output around Y^s , as opposed to inflation stabilization. It also clarifies the kinds of stabilization that are important. Because of the lags involved in price-

21. Note that the latter measure considers only the welfare costs of steady inflation that result from the relative price distortions that follow from the lack of continuous price adjustment. As noted earlier, we abstract from any effects of steady inflation on the steady state level of aggregate output. We also abstract from other welfare costs of inflation, such as the costs of economizing on real money balances, that are emphasized in many discussions of this issue. It seems likely that the effects that we neglect should, if anything, make it even more desirable that average inflation be low. Since many of our results consider the trade-off between stabilization objectives and the objective of a low average rate of inflation, and since our results, when we consider the overall minimization of $L + \pi^{*2}$, recommend a low average rate of inflation in any event, we do not feel that an attempt to quantify such additional considerations is likely to change our conclusions dramatically.

ing, it turns out to be desirable to reduce the variability of both expected inflation and unexpected inflation. Moreover, the variability of unexpected inflation deserves somewhat greater weight, unlike what the ad hoc loss function above would imply. The analysis also makes it clear that it is the variability of quarter-to-quarter inflation, rather than some longer horizon average rate of inflation, or the deviation of the price level from some deterministic or stochastic trend path, that is most closely related to the welfare losses due to price level instability. Finally, it makes it clear that it is the variability of $\hat{Y} - \hat{Y}^s$, rather than the variability of deviations of output from trend or the variability of output growth, that matters for welfare. Specifically, it is the variability of the part of $\hat{Y} - \hat{Y}^s$ that is forecastable two quarters earlier that policy should seek to minimize.

It is worth noting that all three of the terms in equation (33) are directly related, in different ways, to inflation variability. For the analysis of optimal policy below, it is helpful to rewrite L so that it depends only on the stochastic process for the relative price variable \hat{X} . We show in the appendix that the model's structural equations imply that equation (33) may be rewritten in the form

$$(34) \quad L = \text{var}\{E_{t-2}\hat{X}_t\} + \psi \text{var}\{\hat{X}_t - E_{t-2}\hat{X}_t\} + \frac{\Lambda}{\kappa^2} \text{var}\{E_{t-2}(\hat{X}_t - \beta\hat{X}_{t+1})\}.$$

This shows that the deadweight losses measured by L are zero if variations in \hat{X} are eliminated (as we show below to be possible in principle). Thus a constant rate of inflation is both necessary and sufficient for achievement of the minimum value of $L = 0$. This means that, even though our proposed welfare criterion (30) assigns ultimate importance *only* to the efficiency of the level of real activity in each sector of the economy, it in fact justifies giving complete priority to inflation stabilization as opposed to output stabilization.

Given the model, one can compute the value of L as well as that of its components for any rule that sets the interest rate as a function of the history of inflation and output in such a way that there is a unique stationary equilibrium. But this still leaves open the question of whether there is a trade-off between stabilizing the economy by reducing L and keeping a low steady state level of inflation. As suggested by Summers (1991), the requirement that nominal interest rates must always be positive implies that a low average rate of inflation is inconsistent with a great deal of stabilization. The reason is that a low average rate of inflation implies that the average interest rate is low, and this means that the interest rate cannot be too variable. At the same time, keeping the variability of interest rates low weakens the government's ability to reduce L by having the interest rate respond to shocks. To see this, it is worth displaying the relation between interest rates and \hat{X} implied by our model.

This relationship can easily be derived from the equilibrium conditions (17), (20), and (21), together with the requirement that

$$(35) \quad E_t(\hat{Y}_{t+2} - \hat{G}_{t+2}) = \hat{Y}_{t+2} - \hat{G}_{t+2},$$

which is implied by the fact that interest-sensitive purchases in period $t + 2$ are determined at t . We first take the difference between equation (21) and the expectation of this equation at $t - 2$, and use equations (17) and (35) to obtain

$$\begin{aligned} & (\hat{X}_t - \beta E_{t-1} \hat{X}_{t+1}) - E_{t-2}(\hat{X}_t - \beta \hat{X}_{t+1}) \\ & = \kappa[(\hat{G}_t - \hat{Y}_t^s) - E_{t-2}(\hat{G}_t - \hat{Y}_t^s)] - \frac{\kappa}{\omega + \sigma} \phi_{t-1}. \end{aligned}$$

Using this expression to substitute for ϕ_{t-1} in equation (21), we obtain

$$(36) \quad \hat{Y}_t = \hat{Y}_t^s + \kappa^{-1} E_{t-2}(\hat{X}_t - \beta \hat{X}_{t+1}) + [(\hat{G}_t - \hat{Y}_t^s) - E_{t-2}(\hat{G}_t - \hat{Y}_t^s)].$$

We now rewrite equation (17) using (36) to substitute for \hat{Y}_t and the expression just above to substitute for ϕ_{t-1} . This yields

$$\begin{aligned} (37) \quad E_t \hat{R}_{t+1} & = E_t \hat{X}_{t+2} + \hat{\rho}_t + \frac{\sigma}{\kappa} E_t [(\hat{X}_{t+2} - \beta \hat{X}_{t+3}) - (\hat{X}_{t+1} - \beta \hat{X}_{t+2})] \\ & - \frac{\omega}{\kappa} [E_t(\hat{X}_{t+1} - \beta \hat{X}_{t+2}) - E_{t-1}(\hat{X}_{t+1} - \beta \hat{X}_{t+2})], \end{aligned}$$

where we have used the fact that equation (20) implies that $E_{t-2} \pi_t = E_{t-2} \hat{X}_t$, and where

$$\begin{aligned} (38) \quad \hat{\rho}_t & \equiv \omega[(\hat{G}_{t+1} - \hat{Y}_{t+1}^s) - E_{t-1}(\hat{G}_{t+1} - \hat{Y}_{t+1}^s)] \\ & - \sigma[E_t(\hat{G}_{t+2} - \hat{Y}_{t+2}^s) - (\hat{G}_{t+1} - \hat{Y}_{t+1}^s)]. \end{aligned}$$

Note that $\hat{\rho}_t$ is an exogenous stochastic process that can be expressed as a function of the history of the shocks \bar{e}_t .

Equation (37) represents the only restriction implied by our model on the behavior of \hat{R}_t given the evolution of \hat{X}_t . For any given process for \hat{X}_t , the variance of \hat{R}_t is obviously minimized by setting \hat{R}_{t+1} equal to the right-hand side of equation (37). In the case where one wishes to stabilize prices completely, this means that \hat{R}_{t+1} is given by $\hat{\rho}_t$, as discussed in Rotemberg and Woodford (1997). This means that the interest rate at $t + 1$ must rise whenever $\hat{G}_{t+1} - \hat{Y}_{t+1}^s$ increases unexpectedly at t . If, instead, upward revisions in $\hat{G}_{t+1} - \hat{Y}_{t+1}^s$ are matched by upward revisions in \hat{X}_{t+1} , \hat{R}_{t+1} need not rise as much. In other words, if inflation is allowed to respond to these shocks, the interest rate does not have to respond as much to them.

We propose a simple representation of the quantitative connection between average inflation and the variability of interest rates as in Rotemberg and Woodford (1997). In particular, we suppose that, along any equilibrium path, the lowest possible value of r^* (and π^*) consistent with a given degree of interest rate variability is given by

$$(39) \quad r^* = \rho + \pi^* = k\sigma(\hat{R}),$$

where $\sigma(\hat{R})$ refers to the standard deviation of the unconditional distribution for \hat{R}_t in the stationary equilibrium associated with a given policy rule. We let the factor k equal 2.26, which is the ratio of the mean funds rate to its standard deviation under the historical regime, so that, in effect, we are assuming that this is the minimum possible value for this ratio.²² For any monetary policy rule we consider, we thus compute the variance of the nominal funds rate, and then use equation (39) to determine the associated value of π^* . We then compare policy rules according to how low a value they imply for the overall dead-weight loss measure $L + \pi^{*2}$.²³

While minimizing the welfare losses of the agents in the economy is a rather obvious objective for policy, it is worth looking more generally at the effect of different monetary policy rules on the variances of output, inflation, and interest rates. This analysis has several benefits. First, it provides intuition for our results concerning the effects of different rules on $L + \pi^{*2}$. Second, because this analysis is not as dependent on the subset of parameters that we calibrate, it remains valid even if some our calibrations are inappropriate.

Finally, the model may be incorrect in ways that maintain the validity of our estimates of the structural parameters but vitiate our welfare analysis. We do not know the precise range of variations on the model for which this would be true. One simple example would be if there are changes over time in the elasticity of substitution of different goods for each other. This would imply that the Dixit-Stiglitz aggregator varies over time. The resulting changes in the elasticity of demand faced by each firm would lead firms to desire changes in the ratio of price to marginal cost. As far as the algebra of the model is concerned, such changes in the desired markup have the same effect as changes in Y_t^S . The difference is that, under this alternative interpretation, it is no longer socially desirable for output to track the time variation in Y_t^S . In particular, variation in desired markups would justify an objective of reducing the variance of output relative to trend more than is implied by our minimization of $L + \pi^{*2}$ below. For this reason, as well as for comparability of our results with those of other studies, we look at a relatively wide range of consequences of the monetary rules we study.

22. Note that our definitions imply that $\hat{R}_t = r_t - \rho$, so that $\sigma(\hat{R}) = \sigma(r)$. We refer to $\sigma(\hat{R})$ in eq. (39) because the structural equations of our model are written in terms of the variable \hat{R}_t , and so we solve for the equilibrium fluctuations in that variable.

23. The advantage of this substitute for the more rigorous approach of imposing the requirement that $R_t \geq 0$ at all times, given estimated shock distributions with bounded supports, is a considerable saving in computational effort. First, imposing a constraint of the form (39), our optimization problem continues to be a linear-quadratic one (if we use approximation [32] to the objective and a log-linear approximation to the model structural equations), and as a result the optimal policy is described by a linear rule, which we can obtain using linear methods. Second, under this form of constraint, the optimal policy does not depend on any more detailed description of the distribution of the exogenous shocks e_t^i than their means and variances. This means that we do not need to estimate more detailed properties of these distributions and that our conclusions do not depend on properties of such distributions that are likely to be very poorly estimated in a sample of our size.

2.2 Consequences of Simple Policy Rules

As noted earlier, we wish to compare a variety of types of monetary policy rules that make the interest rate r_t depend on the history of output, inflation, and the interest rate itself. In this section, we explore the effects of varying the parameter in some very simple rules of this kind. These simple rules, which are variants of the rule proposed by Taylor (1993), have some practical advantages. Their simplicity makes them easy to understand so that a central bank that adopted them ought to find it easy to explain what it is doing. As a result, the public ought to find it easy to monitor the central bank's compliance with its rule. Finally, the use of similar rules in the other papers in this volume makes our results concerning the desirability of these rules directly comparable to theirs.

When we study rules that can be described by only a small number of parameters, we study the consequences of parameter variation for two sorts of issues. First we analyze the range of parameter values that ensures that a determinate rational expectations equilibrium exists; as an extensive prior literature has stressed, determinacy of equilibrium cannot be taken for granted in rational expectations models, especially in the case of a monetary policy defined by an interest rate rule. (See, e.g., Bernanke and Woodford 1997 for general discussion of this issue, and illustrations in the context of a model similar to the one that we use here.) Next we study the effect of parameter variation within the range of parameter values for which equilibrium is determinate.

2.2.1 Performance Measures for Alternative Rules

For each of the rules we consider, we compute a number of statistics relating to the variability of inflation, output, and interest rates in the unique stationary rational expectations equilibrium associated with that rule. These statistics are reported in table 2.1 for a number of rules of particular interest. The significance of the parameters a , b , and c that define these rules is explained below.²⁴

Among the specific rules included in the table are several that are also considered in other papers in this volume. These are labeled A_i through D_i , with i equal to 0 in the case of rules where the interest rate responds to contemporaneous output and inflation, and i equal to 1 in the case where it responds with a lag. The table also reports the effects of setting the parameters at the values that represent the best rule (in the sense of minimization of our utility-based loss measure $L + \pi^{*2}$) within each of several families of simple rules discussed below (these are labeled E_0 , F_0 , and G_0 , and E_1 and G_1). Finally, we also report the statistics associated with our estimate of actual U.S. policy during the period 1979–95 (rule H), and for the unconstrained optimal policy according to our model, discussed in section 2.3 (rule I).

24. Briefly, in each case, a measures the extent to which the funds rate responds to deviations of inflation or the price level from its target value, b measures the extent to which the funds rate responds to deviations of output from trend, and c measures the extent to which the funds rate responds to deviations in its own lagged value.

Table 2.1 **Statistics for Several Policy Rules**

	a	b	c	$\text{var}\{R\}$	$\text{var}\{Y\}$	$\text{var}\{\pi\}$	$\text{var}\{\Delta p^e\}$	β^e	$\text{var}\{\pi - E\pi\}$	$\text{var}\{E(Y - Y^s)\}$	L	π^*	$L + \pi^{*2}$
<i>Contemporaneous-data Taylor rules</i>													
A ₀	3.00	0.80	1.00	6.24	7.24	0.66	4.00	3.82	0.27	8.77	1.32	2.66	8.38
B ₀	1.20	1.00	1.00	6.41	2.73	1.82	16.05	5.44	0.48	11.69	2.80	2.74	10.29
C ₀	1.50	0.50	0.00	17.14	3.87	7.34	121.14	11.14	0.81	13.86	8.72	6.37	49.29
D ₀	1.50	1.00	0.00	16.26	1.79	6.76	100.77	9.80	0.83	14.63	8.18	6.13	45.73
E ₀	1.22	0.06	1.28	1.88	13.47	0.38	1.47	-1.86	0.19	11.37	1.09	0.11	1.10
F ₀	2.88	0.02	0.00	8.65	15.42	1.04	16.77	6.38	0.37	10.35	1.86	3.66	15.26
<i>Lagged-data Taylor rules</i>													
A ₁	3.00	0.80	1.00	5.89	10.32	0.66	4.20	3.84	0.28	6.56	1.23	2.50	7.47
B ₁	1.20	1.00	1.00	5.99	3.30	1.83	16.30	5.43	0.49	9.58	2.71	2.54	9.18
C ₁	1.50	0.50	0.00	13.83	4.63	5.97	94.12	10.04	0.78	11.17	7.19	5.42	36.55
D ₁	1.50	1.00	0.00	14.86	3.32	5.85	83.99	9.08	0.81	11.69	7.12	5.73	39.91
E ₁	1.27	0.08	1.28	1.87	13.92	0.37	1.12	-1.35	0.19	12.18	1.12	0.10	1.13
<i>Price level targeting rules</i>													
G ₀	0.26	0.07	1.03	2.00	11.14	0.34	0.00	0.00	0.23	11.93	1.11	0.21	1.16
G ₁	0.38	0.10	0.85	2.04	11.35	0.34	0.00	0.00	0.25	12.57	1.15	0.24	1.21
<i>Estimated historical U.S. policy</i>													
H	-	-	-	7.64	4.79	2.28	26.19	5.60	0.66	12.14	3.43	3.26	14.06
<i>Optimal policy</i>													
I	-	-	-	1.93	11.30	0.39	1.77	-2.19	0.20	7.57	0.93	0.15	0.95

The statistics reported in table 2.1 include the variance of output around trend, the variance of inflation, and the variance of the federal funds rate. In addition to these conventional statistics, we also report the variance of quarterly innovations in the rational forecast of the long-run price level. This is the variance of changes in the variable

$$\rho_t^\infty \equiv E_t \lim_{T \rightarrow \infty} (\log P_T - T\pi^*),$$

which is just the stochastic trend in the price level in the sense of Beveridge and Nelson (1981). (Note that it follows from this definition that the first difference of p^∞ is also the innovation in this variable.) We include this statistic as an alternative index of the degree of price stability associated with different equilibria. The advantage of this statistic is that it reflects the extent to which agents make capital gains and losses on long-term nominal contracts and some analysts have expressed concern over these (e.g., Hall and Mankiw 1994). Finally, we also report the coefficient β^∞ of a regression of the innovation at t in the forecast of the long-run price level p^∞ on the quarter t innovation in the (log) price level at $t + 1$. (Recall that, according to our model, the price level P_{t+1} is determined at date t .) This coefficient tells us whether inflation innovations in quarter t eventually lead to a higher price level or whether instantaneous increases in the price level are later offset by subsequent expected reductions in prices. In the case of a random walk in the log price level, we should find $\beta^\infty = 1$, while if temporary price level increases are eventually completely offset, we should find $\beta^\infty = 0$.

The first column of table 2.1 serves as a key for figures 2.1, 2.2, and 2.3, where the consequences of these rules for the variability of output, inflation, interest rates, and long-run price-level forecasts are plotted. Figure 2.1 has a certain similarity to the policy frontier shown in Taylor (1979), in that rules that have smaller standard deviations of inflation tend to involve larger standard deviations of output and vice versa. The only rules that appear to be “dominated” in this plot are the rules with labels in the series C_i and D_i . These are simple “Taylor rules” that make the funds rate a function only of current inflation and output, and they respond much more strongly to output fluctuations than does our optimal rule in that family (labeled F_0). The rules in families C_i and D_i are worse than the B_i rules because they induce a higher standard deviation of inflation without reducing the standard deviation of output. Interestingly, the rule F_0 , which is the best rule of this type in terms of minimizing our utility-based loss measure, is something of an outlier as well in that it involves more variability of both inflation and output relative to other rules in the set. From the point of view solely of the criteria plotted in this figure, historical policy seems to be slightly worse than the rules described by B_i , but not significantly so.

Figure 2.2 paints a different picture, one that involves pure dominance relations and no trade-offs. Once again, the rules C_i and D_i are particularly bad in




Fig. 2.1 Selected rules: standard deviations of output and inflation

that they now also involve a high standard deviation of the funds rate. Among the remaining rules, those with a lower standard deviation of inflation tend to have a lower standard deviation of the funds rate so they allow average inflation to be lower as well. Thus the best rules in this plot are the rules E_i , which, as we shall see below, also minimize $L + \pi^2$ among rules that are as simple as these. These involve low standard deviations of both inflation and interest rates, while the other rules perform worse on both dimensions. When coupled with the results of figure 2.1, we see that—leaving aside C_i and D_i —the rules we consider here have the property that those that reduce the standard deviation of output tend to raise the standard deviations of inflation and interest rates simultaneously.

Figure 2.3 shows the implications of these rules for the variance of inflation and the variance in the innovation of the forecast of the long-run price level. We see in this figure that the specific rules we consider rank equally along these two dimensions. The price level rules G_i and the E_i rules have both the lowest variance of inflation and the smallest innovations in the long-run price level. That the price level rules have low variances in the long-run price level is not surprising, since they ensure the price level is stationary. What is perhaps more surprising is that the best of the rules that respond to deviations of the inflation rate from target have this property as well.




Fig. 2.2 Selected rules: standard deviations of interest rate and inflation




Fig. 2.3 Selected rules: standard deviations of inflation and innovation in long-run price level

The regression coefficients β^∞ of the innovation in the long-run price level on the current price level innovation reported in table 2.1 help to explain this finding. This coefficient is obviously zero for the price level rules, since these equilibria involve no change in the forecast of the long-run price level at any time. For the E, rules as well as for the rule marked I, which is the rule that minimizes $L + \pi^{*2}$ among all possible rules, this coefficient is actually smaller than -1 . This means that increases in the contemporaneous price level eventually lead to a lower price level, and, indeed, to a lower price level by an amount that is even greater than the size of the initial price level innovation (but with an opposite sign). Thus, while the long-run price level is not being stabilized, expected reductions in future inflation more than offset the initial increase in the price level. This stands in sharp contrast to the other rules reported in the table. For these rules, this coefficient exceeds one so that increases in the current price level lead to even larger increases in the long-run price level. This means that, on average, increases in inflation are followed by further inflation. This clearly destabilizes the long-run price level. In addition, because expected future inflation leads price setters who can change their prices at t to raise their prices by more, it also means that inflation at t is increased by policies that follow inflation at t with further inflation. For this reason, policies with high values of β^∞ have both variable inflation and large variances in the innovation of the long-run price level.

The remaining columns of table 2.1 report statistics that measure various components of the utility-based measure of deadweight loss derived in the previous section. The columns labeled $\text{var}\{\pi\}$, $\text{var}\{\pi - E\pi\}$, and $\text{var}\{E(Y - Y^S)\}$ report the values of the three unconditional variances that receive positive weights in expression (33) for the loss measure L . The third column from the right then reports the implied value for L . This is our summary measure of the deadweight losses due to variability of inflation and output, in units of the variance of inflation. We scale inflation so that $\pi = 1$ corresponds to a 1 percent inflation per year. Hence, $L = 1$ indicates the same degree of deadweight loss as results from this inflation rate. The next to last column reports the minimum value of π^* consistent with the degree of funds rate variability required by the policy rule, using equation (39) to derive this. Finally, the last column reports the implied value of $L + \pi^{*2}$, our total measure of deadweight loss.

One interesting fact about the table is that the ranking of alternative rules according to their implications for the variability of $\hat{Y} - \hat{Y}^S$ is quite different from their ranking according to their implications for the variability of output relative to its deterministic trend path. The rule D_0 , which minimizes $\text{var}\{\hat{Y}\}$ among those considered in the table, implies the *highest* degree of variability of output relative to the natural level \hat{Y}^S . This indicates that responding to deviations of output from a deterministic trend, while perhaps successful as a way of stabilizing output around that trend, may well be counterproductive if one is interested in keeping output close to its natural level. (Cf. figs. 2.6 and 2.8 below, for further illustration of this point.)

Another fact that is apparent from the table is that the ranking of different rules according to the value achieved for L is essentially the same as their ranking in terms of the variability of inflation. Thus our utility-based welfare criterion $L + \pi^{*2}$ leads to conclusions that are similar to those that would be reached by giving some weight to the reduction of both the variability of inflation and the variability of the funds rate. In both these respects, the rules labeled E_p , G_p , and I are better than the others. We turn now to a more systematic exploration of the consequences of parameter variations, in order to clarify why this is so.

2.2.2 Simple “Taylor Rules”

We first consider the consequences of varying a and b in simple “Taylor rules” of the form

$$(40) \quad \hat{r}_t = a\hat{\pi}_t + b\hat{Y}_t,$$

where once again $\hat{r}_t \equiv r_t - r^*$ and $\hat{\pi}_t \equiv \pi_t - \pi^*$. Note that both the rule C_0 proposed by Taylor (1993) and the related rule considered by Henderson and McKibbin (1993) belong to this family. Our aim here is to highlight the trade-offs involved in the choice between having interest rates respond to output and having interest rates respond to inflation.

In the case of simple Taylor rules of the form (40) with a constrained to be positive, our loss criterion $L + \pi^{*2}$ reaches a minimum when a equals 2.88 and b equals 0.02. The consequences of this rule for our loss measures is displayed in table 2.1, where the rule is designated F_0 . As one might guess, this rule (which places essentially all of the weight on inflation variations rather than output variations) allows much greater variations in output relative to trend than do rules C_0 and D_0 . However, according to our model, it leads to *less* variability of output relative to its natural level, which is what matters for our loss measure. It also results in significantly less variability of inflation, and noticeably less variability of the funds rate. (It is actually the latter difference that is most significant for our loss measure, because of the reduction in the average inflation rate π^* that it allows.) The ultimate result is a reduction in deadweight loss by a factor of three, relative to the other proposals. However, our model and our loss measure imply that this rule would not represent an improvement on historical U.S. policy in the Volcker-Greenspan period. To do better we must not simply vary the weights on inflation and output but must consider at least slightly more sophisticated rules.

Before turning to other families of rules, it is worth noting that the welfare criterion $L + \pi^{*2}$ reaches an even lower value, according to our model, if we allow a and b to be negative in equation (40). The optimum then involves a equal to -1 and b equal to -1.3 . The idea that negative values of a and b are acceptable may be surprising. For this reason, figure 2.4 displays both the region where equilibrium is determinate as well as a contour plot of $L + \pi^{*2}$ as we vary a and b . The equilibrium is not unstable for any of these parameter




Fig. 2.4 Simple Taylor rules: $L + \pi^{*2}$ as a function of a and b

values (i.e., a stationary equilibrium always exists), but equilibrium is indeterminate in the region labelled “Indet.” Indeterminacy arises, for example, when b is zero and a is small and positive. This indeterminacy implies, among other things, that inflation can vary simply as a result of changes in expectations. A “sunspot” can lead inflation at t to rise, for example. The real interest rate would then fall (because the nominal interest rate responds little) and the resulting increase in output means that expected future inflation is lower than current inflation. Thus the change in the expected future path of inflation that is required to justify the initial change in inflation is consistent with expected future inflation converging back to the target inflation rate π^* . In this case, a stationary rational expectations equilibrium is possible in which such fluctuations occur simply because they are expected to.

If, instead, a is large and positive, no such equilibrium is possible. Any increase in inflation above its unique saddle-path value is matched by increases in real interest rates that imply that output must fall. This, in turn, implies that expected future inflation rates must be higher than current inflation, given the nature of our AS curve. Thus inflation must be expected to explode, and since this is not consistent with stationarity, inflation must equal its saddle-path value in the unique stationary equilibrium. Similarly, as mentioned above, the equilibrium is determinate when a and b are both negative.




Fig. 2.5 Simple Taylor rules: $\text{var}\{\hat{\pi}\}$ as a function of a and b

Figure 2.4 presents contour lines for the value of our loss measure $L + \pi^{*2}$ in the regions where equilibrium is determinate. Policy F_0 appears as a star on this figure, at the point of a local minimum of the loss measure. However, the region of determinate equilibria with negative a and b also contains a local minimum. This point, which is shown with a star inside a circle, is actually the global minimum value. Nonetheless, we have chosen to present the local minimum F_0 in table 2.1, on the ground that restricting attention to values $a > 0$ corresponds to rules that are more similar to the Taylor and Henderson-McKibbin proposals. In addition, once we consider more general families of rules, we do find that the best rules involve tightening monetary policy (i.e., raising the funds rate) in response to inflation increases, as conventional wisdom (at least since the work of Wicksell [1907]) would indicate.

Similar contour plots for other statistics reported in table 2.1 provide further insight into why our loss measure varies with a and b as it does. Figure 2.5 shows the contour plots of the variance of inflation, while figure 2.6 shows the contour plots for the variance of $\hat{Y} - \hat{Y}^s$. These figures are essentially identical to each other, and they are both similar to the contour plot for L itself. There is thus no trade-off between stabilizing inflation and stabilizing $\hat{Y} - \hat{Y}^s$; the same parameters stabilize both. This follows immediately from our AS curve, which




Fig. 2.6 Simple Taylor rules: $\text{var}\{\hat{Y} - \hat{Y}^s\}$ as a function of a and b

relates inflation to departures of \hat{Y} from \hat{Y}^s . For the ranges considered in our figures, a wheel marks the global optimum for the performance criterion being considered. Thus the figures show that these variances become as small as possible when a is at its maximum possible value of 20 while b is set to a small negative number. Making a big contributes to stabilization because it ensures that interest rates rise a lot when either \hat{G} rises or \hat{Y}^s falls. This ensures that inflation does not rise much in either case and that, at least after the demand for output adjusts to changes in real rates, output does not rise in the former case while it declines substantially in the latter.

As figure 2.7 indicates, the rule that minimizes L by setting a equal to 20 leads to very variable interest rates. This is in part due to the delays in the response of output to interest rates. These delays imply that changes in \hat{G}_t that become known at $t - 1$ inevitably change output at t since C_t is predetermined. This leads firms to raise their prices at t unless long-term real interest rates rise unexpectedly. With c equal to zero, this means that prices can only be stabilized if the nominal interest rate at t rises a great deal. The resulting variability of interest rates then requires a high average inflation rate for interest rates never to be negative. This high inflation is so costly, at least relative to the benefits of the additional stabilization that is possible with a high value of a , that the contour plots for the variance of the interest rate are essentially identical to the contour plots for $L + \pi^{*2}$. The point that minimizes the variance of




Fig. 2.7 Simple Taylor rules: $\text{var}\{\hat{R}\}$ as a function of a and b

interest rates has a sufficiently stable inflation to be quite desirable as far as total welfare is concerned.

It is interesting to note that the stabilization of output requires a quite different set of parameters. This is demonstrated in figure 2.8, which gives the contour plots for the variance of output. This variance is reduced by keeping a small and positive while making b very large. Not surprisingly, output is stabilized if the real interest rate is raised significantly by the central bank whenever output rises, while it is lowered when output declines. What is interesting here is that the effects of the policy parameters on the variance of $\hat{Y} - \hat{Y}^s$, which are essentially the same as the effects on L , are very different from the effects on the variance of Y . The reason is that the VAR of Rotemberg and Woodford (1997) identifies large short-run fluctuations in \hat{Y}^s . As long as these are treated as variations in the welfare-maximizing level of output, setting b large is not desirable, and indeed, stabilization of $\hat{Y} - \hat{Y}^s$ requires that b be negative at least when a is 20. Even higher values of a reduce the variance of $\hat{Y} - \hat{Y}^s$ still further. Obviously, the result that the stabilization of \hat{Y} relative to \hat{Y}^s requires very different policies from those that stabilize output relative to trend is very sensitive to the assumption that our estimate of \hat{Y}^s is indeed the welfare-maximizing level of output. This conclusion would presumably change dramatically if movements in \hat{Y}^s were viewed as resulting from changes in distortions such as changes in desired markups. From an empirical point of view




Fig. 2.8 Simple Taylor rules: $\text{var}\{\hat{Y}\}$ as a function of a and b

these two interpretations may be difficult to disentangle because we identify \hat{Y}^s by measuring shifts in the empirically estimated AS equation given by (22). Unfortunately, changes in desired markups will shift this equation just as much as changes in technology or other changes in the welfare-maximizing level of output.

2.2.3 Rules That Involve a Lagged Interest Rate

We achieve improvements in household welfare if we generalize the family of simple Taylor rules to allow the funds rate to respond also to lagged values of itself. We thus consider generalized Taylor rules of the form

$$(41) \quad \hat{r}_t = a\hat{\pi}_t + b\hat{Y}_t + c\hat{r}_{t-1},$$

where we now allow c to be greater than zero. This allows for interest rate smoothing, so that sustained changes in output and inflation lead to only gradual changes in interest rates. Actual policy in the United States and elsewhere seems to involve some degree of interest rate smoothing, though academic commentators have often questioned why this should be so.²⁵ Nor is there any reason to restrict attention to the case $0 \leq c < 1$, though only in that case can

25. See, e.g., Goodfriend (1991), Rudebusch (1995), Goodhart (1997), and Sack (1998).




Fig. 2.9 Generalized Taylor rules: $\text{var}\{\hat{\pi}\}$ as a function of a and c

the policy rule be described as involving partial adjustment toward a “target” funds rate that depends on current output and inflation, as assumed in Clarida, Galí, and Gertler (1998), for example. An alternative is to follow Fuhrer and Moore (1995) and model U.S. interest rates by supposing that c is equal to one, so that it is *changes* in the funds rate, rather than the level of the funds rate, that respond to deviations of inflation and output from their typical levels. Policy proposals of this kind are considered elsewhere in this volume with rules A_0 and B_0 in table 2.1 being examples of such rules. We find that policies that involve values of c even greater than one often result in determinate rational expectations equilibria in our model, and so we consider arbitrary positive values of c . In fact, rules with $c > 1$ turn out to possess an important advantage, and this is one of our most important findings.

To gain some insight into the consequences of varying c , we set b equal to zero and discuss contour plots in the $\{a, c\}$ -plane for various measures of economic performance. Our motivation for starting with plots that set b equal to zero is that, as we show below, the welfare optimum obtains near this point. Moreover, the resulting family of rules has a very simple interpretation as the family in which interest rates depend only on inflation and lagged interest rates. Figure 2.9 displays the resulting contour plots for $\text{var}\{\hat{\pi}\}$, which, once again, are essentially identical to those for both the variance of $\hat{Y} - \hat{Y}^s$ and for

L itself. One interesting aspect of this figure is that it shows that determinacy obtains with c greater than one even if a is negative so that the Fed reacts perversely to inflation by cutting rates when inflation rises. The reason is that, as in the earlier case with negative values of a , these rules also induce explosions in response to deviations of inflation and output from saddle-point paths.

One surprising aspect of the figure is that it shows that “explosive” monetary rules in which c exceeds one do not produce explosive equilibria. In a way, this potential explosiveness of interest rates is effective at keeping the economy on track in this model. It means that, unless the price level reacts properly, the real interest rate falls or increases exponentially. An exponential increase in real rates represents a rather substantial reduction in expected future aggregate demand and thus leads firms to cut prices. The result is that the economy stays on a nonexplosive path in which increases in inflation are matched by subsequent reductions in inflation that ensure that the interest rate does not explode. In fact, higher values of c actually increase the range of values of a for which a determinate equilibrium exists, by helping to solve the problem of indeterminacy discussed above.

The figure also shows that, within the range being considered, the goal of inflation stabilization is furthered by setting a as large as possible. The variance of inflation reaches its minimum value (over the range of rules shown in the figure) when a equals 20 and c takes a positive value less than one. If the range of the figure were extended, the optimum would involve even higher values of a . Thus the key to inflation stabilization remains making sure that the interest rate reacts vigorously to inflation.

Interestingly, a higher value of c turns out to be better if one seeks to stabilize the long-run price level. This can be seen in figure 2.10, which shows that, for any given value of a , the variance of Δp^∞ reaches a minimum of zero for c equal to one. Further insight into this behavior of the variance of Δp^∞ can be obtained from figure 2.11, which shows β^∞ as a function of a and c . This figure shows that, when c is zero, β^∞ is greater than one so that initial increases in inflation are followed by further inflation. The reason for this is that an increase in \hat{G}_t raises the price level at t somewhat in spite of the increase in interest rates that takes place at t . But unless the price level continues rising, interest rates would immediately return to their steady state level. The result is that, in equilibrium, prices do keep rising because the initial increase in prices means that marginal cost has gone up for the firms that did not raise their prices at t . Consequently, increases in the price level at t are followed by further increases in prices, which, admittedly, are kept somewhat in check by the fact that the interest rate remains somewhat above the steady state for some time.

If, instead, c is made higher, the interest rate tends to stay high after an increase in \hat{G} even if the price level ceases to rise. This means that firms can be induced not to change their prices in the aftermath of an increase in \hat{G} . The result is that initial increases in prices are followed by smaller increases so that β^∞ is smaller and the variance of Δp^∞ falls. Setting c equal to one as suggested




Fig. 2.10 Generalized Taylor rules: $\text{var}\{\Delta p^\infty\}$ as a function of a and c




Fig. 2.11 Generalized Taylor rules: β^∞ as a function of a and c




Fig. 2.12 Generalized Taylor rules: $\text{var}\{\hat{R}\}$ as a function of a and c

by Fuhrer and Moore (1995) makes β^∞ equal to zero so that the shocks have no effect on the long-run price level. Even higher values of c imply that initial increases in inflation are followed by such high real rates that the expected long-run price level is lower than the initial price level so that β^∞ is negative.

For a given initial increase in inflation and interest rates, higher values of c imply that the long-run real rate rises more both because future short rates are expected to be higher and because future inflation is expected to be lower. Since unexpected increases in the long-term real rate prevent prices from rising this means that, for given a , increases in \hat{G} (and reductions in \hat{Y}^s) lead to smaller immediate price and interest rate increases the higher is c . This is reflected in figure 2.12, which shows that, for each a , the interest rate is less variable the higher is c . It also shows that, not surprisingly, the variance of interest rates increases with a .

While we have focused on stabilizing the variability of interest rates because of their implication for average inflation, the Fed also seems to be concerned with stabilizing the *change* in the funds rate from one week or month to the next. This would explain Rudebusch's (1995) finding that changes in the target rate are followed by further changes in the same direction. Figure 2.13 thus displays the variance of the change in interest rates in the $\{a, c\}$ -plane. Interestingly, this figure is nearly identical to the figure for the variance of the interest




Fig. 2.13 Generalized Taylor rules: $\text{var}\{\hat{R}_t - \hat{R}_{t-1}\}$ as a function of a and c

rate itself. Thus, in our model, stabilization of the short-term nominal rate is achieved in the same way as stabilization of the quarterly change in this rate.

As we saw in figure 2.9, setting c to a very high value destabilizes the inflation rate. In part this is because sufficiently high values of c imply that increases in inflation at t must be matched by reductions in inflation in the future. These predictable movements in inflation both raise the variance of $\hat{\pi}$ and increase the loss L . For that reason, figure 2.14 shows that $L + \pi^{*2}$ reaches its lowest value for a low value of a and a moderate value of c . This minimum is very close to the point that minimizes $L + \pi^{*2}$ within the family (41) since this minimum obtains when a , b , and c equal 1.22, 0.06, and 1.28, respectively. This is the rule labeled E_0 in table 2.1.

2.2.4 Rules Using Only Lagged Data

One criticism sometimes leveled (see, e.g., McCallum 1997) against all rules of the kind considered thus far is that they require the Fed to make use of data about current output and inflation that it does not actually have when it sets the current interest rate. There are two reasons why such variables may simply be unobservable by the central bank. These are that some important economic data are collected retrospectively and that even the data that are collected concurrently need to be processed before their message about the econ-




Fig. 2.14 Generalized Taylor rules: $L + \pi^{*2}$ as a function of a and c

omy as a whole can be distilled. A further difficulty with responding to contemporaneous variables may be that, even if these are observable immediately, the political process of responding to them takes time.

None of this denies that the central bank continually updates its estimate of the current state of the economy. And it should be recalled that our model of the delays in the response of output and inflation implies that the relevant data exist in principle in the quarter *prior* to the one in which the data must be used under rules (40) and (41). However, it is reasonable to suppose that the central bank's estimate of the state of the economy generally differs from the economy's actual state. In this case, responding to the current estimate of the current state differs from rules (40) and (41). If rules of the form (40) and (41) are applied to the error-ridden current estimates, the interest rate is affected by the measurement error, and a thorough evaluation of these rules would require an analysis of these effects.

Thus we now suppose instead that the Federal Reserve does not respond to output and inflation variations except with a one-quarter lag. In this class of rules,

$$(42) \quad \hat{r}_t = a\hat{\pi}_{t-1} + b\hat{Y}_{t-1} + c\hat{r}_{t-1}.$$




Fig. 2.15 Lagged response rules: $\text{var}\{\hat{\pi}\}$ as a function of a and b

Considering the effect of such a lag also allows us to compare our results with other papers in this volume since some of these also include the rules we label A_1 through D_1 in table 2.1.

Even if the Fed had a reasonably accurate estimate of the current state of the economy, there would be good reasons to be interested in lagged-data rules of this form. In particular, the use of such rules would make Fed operations more transparent to the public at large if the public only had this lagged information. By avoiding the use of information that the public does not have, it becomes both easier to describe Fed operations and easier for people to detect when the Fed has departed from the rule. An alternative, of course, might be to respond to internal estimates and publish these estimates of the state of the economy as they become available. The study of this alternative, and its effects on transparency given that this estimate will at least sometimes be wrong, is clearly beyond the scope of this paper.

We start in figure 2.15 by displaying how the variance of inflation varies with a and b when c is set equal to zero. This figure is quite different from figure 2.5, which involves the same parameters and performance criterion in the case of contemporaneous Taylor rules. Unlike what occurs with rules where the interest rate responds contemporaneously, large values of a and b lead to




Fig. 2.16 Lagged response rules: $\text{var}\{\hat{\pi}\}$ as a function of a and c

unstable equilibria in the case where the interest rate responds only to lagged output and inflation. Ignoring b , this can be understood as follows. Inflationary shocks now lead to delayed increases in interest rates, which imply delayed reductions in inflation. The rule then requires that subsequent interest rates fall so that inflation rises once again. For a sufficiently strong reaction of interest rates to lagged inflation, that is, a high value of a , the resulting oscillations are explosive. Thus the parameters that minimize the variance of $\hat{\pi}$ in the case of a contemporaneous rule no longer do so when the government can only react with a delay. In particular, this minimization now requires that a be equal to about 15.

Figure 2.16, which gives the contour plot for the variance of $\hat{\pi}$ when b is set to zero while a and c are allowed to vary, tells a similar story. Again, high values of a lead to explosive equilibria. By contrast, high values of c with low values of a do not. Note that high values of c coupled with moderate values of a mean that the eventual reaction of interest rates to increased inflation is extremely large. Nonetheless, these rules are less destabilizing than having the interest rate respond strongly to inflation after a delay of one quarter.

Even in the case of rules that react with a lag, the stabilization of interest rates continues to require high values of c together with small values of a . The result is that figure 2.17 shows that $L + \pi^{*2}$ achieves a minimum for a combi-




Fig. 2.17 Lagged response rules: $L + \pi^{*2}$ as a function of a and c

nation of a and c that is quite similar to the combination that was optimal in the case where the interest rate reacted contemporaneously. Moreover, the minimum value of $L + \pi^{*2}$ within the family (42) is obtained for very similar parameters. In particular, it requires that a , b , and c be equal to 1.27, 0.08, and 1.28, respectively.

This is the rule called E_1 in table 2.1. Clearly, these parameters are very similar to those (the rule E_0) that minimize $L + \pi^{*2}$ when contemporaneous data are used. What is more surprising, however, is that table 2.1 indicates that the minimized value of $L + \pi^{*2}$ is very similar in the two cases. In other words, this welfare criterion equals 1.10 when the best contemporaneous rule is used, while it equals 1.13 when the best of the rules that respond to lagged values is used. Recall that the units of this welfare criterion are squares of percentage *yearly* inflation rates. Thus the difference in loss is equivalent to the difference between having a completely stable annual inflation rate of 1.06 percent per year and having a completely stable annual inflation rate of 1.05 percent per year. The difference is trivial.

This similarity is not surprising once one recognizes that the optimal contemporaneous rule involves a high value of c . This fact means that, even in the case of contemporaneous rules, most of the reaction of interest rates to an inflationary shock such as an increase in \hat{G} or a reduction in \hat{Y}^s takes place

with a delay. Given this, it is not surprising that the further delay that comes about from responding to inflation and output with a lag has trivial welfare consequences. From an economic perspective, what is important is that delayed responses still allow for substantial revisions in long-term real interest rates, and it is these that help stabilize inflation.

2.2.5 Price Level Targeting Rules

In this subsection we consider the possibility of making the funds rate respond to deviations of the price *level* from some target path (assumed to be a deterministic trend with growth rate π^*), rather than responding to inflation. In particular, we consider rules of the form

$$(43) \quad \hat{r}_t = a\hat{P}_t + b\hat{Y}_t + c\hat{r}_{t-1}.$$

The rule given by (43) has the advantage that (if $a \neq 0$) it makes the price level stationary around the target (deterministic trend) path. Such rules thus reduce $\text{var}\{\Delta p^\infty\}$ to the maximum possible extent by ensuring this variance is zero. This may be considered a desirable goal of policy; for example, Hall and Mankiw (1994) discuss the advantages of a price level targeting rule in this regard.²⁶ Such rules also address the desire expressed by the 90 percent of the respondents to Shiller's (1997) survey, that any change in the price level be subsequently reversed. We wish to consider whether rules of this kind are also desirable in terms of the other measures of performance that we treat here, or to what degree one might have to sacrifice other goals for the sake of stability of the long-run price level forecast.

Figure 2.18 displays the contour plots of $L + \pi^{*2}$, once again setting b equal to zero. As the figure shows, price level rules tend to be unstable when a is negative and c is large; lower values of c with negative values of a lead to indeterminate equilibria instead. Within the positive orthant, these rules do lead to determinate equilibria, however. In particular, points with positive a and c equal to zero lead to unique determinate equilibria. Since the same is true for rules in the family (41) with b equal to zero, c equal to one, and a positive, the corresponding equilibria must be the same. To see this, note that, when b is zero, c is one, and a is positive, rules in family (41) take the form

$$(44) \quad \Delta\hat{r}_t = a\Delta\hat{P}_t.$$

Price level rules in the family (43) must also satisfy this equation when b and c are zero since, in this case, (44) is just the first difference of (43). Thus, if the equilibrium with the first-differenced rule (44) is unique, it must be the same as that of the corresponding price level rule. This explains why we found that rules in the family (41) with b equal to zero and c equal to one had the

26. But for a contrary view of the relative desirability of inflation stabilization and price level stabilization, see, e.g., Gertler (1996).




Fig. 2.18 Price level rules: $L + \pi^{*2}$ as a function of a and c

dual property that the long-run price level was stable and that β^∞ was equal to zero. These rules were in fact equivalent to price level rules.

However, the optimal price level rule is not a member of the family (41) because the optimal b and c within the class (43) are not equal to zero. In particular, the lowest value for $L + \pi^{*2}$ within the family of price level rules (43) obtains when a , b , and c equal 0.26, 0.07, and 1.03, respectively. Because the optimal b is zero, this point is close to the optimum depicted in figure 2.18. Once again, the desire to stabilize interest rates leads to a high value of c , though this parameter has a somewhat different meaning in the context of price level rules than it does in the context of the family (41).

Perhaps the most interesting aspect of these price level rules is that even the best such rule is somewhat worse from the perspective of $L + \pi^{*2}$ than the best rule that responds to contemporaneous inflation. Indeed, household welfare is slightly lower than it would be if the central bank followed the best rule that responds only to the lagged levels of inflation and output. The best among the rules that respond to lagged output and the lagged price level is even worse. Admittedly, the resulting differences in our welfare criterion are small, but it is worth knowing that price level rules are not particularly attractive in this context.

One could argue that these results do not really say whether it is worth stabi-

lizing the price level, because we are only looking at a very narrow class of rules. To see whether one can obtain some incremental improvement in our criterion function by responding to *both* the price level and the rate of inflation, we analyze hybrid rules of the form

$$(45) \quad \hat{r}_t = a_0 \hat{P}_t + a_1 \hat{P}_{t-1} + b \hat{Y}_t + c \hat{r}_{t-1}.$$

When we choose parameters a_0 , a_1 , b , and c to minimize $L + \pi^{*2}$, we obtain the values $a_0 = 1.22$, $a_1 = -1.22$, $b = 0.06$, and $c = 1.28$. Since $a_1 = -a_0$, the optimal member of this family is a member of the more restricted family (41) and, in fact, it is once again the rule labeled E_0 in table 2.1. There is thus nothing to be gained, from the point of view of our utility-based welfare criterion, by generalizing this family to add a term that ensures that the interest rate reacts to the price level, even though adding even a small term of that kind would serve to stabilize the long-run price level.

The reason is that the best rule within the class (41) involves some base drift, and this base drift is optimal. Interestingly, this base drift is very different from, and in some ways exactly opposite to, the base drift that people usually worry about. In particular, it is *not* optimal to respond to shocks that temporarily raise inflation by allowing the price level to be higher forever—that is, to choose a rule that implies $\beta^\infty > 0$. On the contrary, as discussed earlier, what is optimal is to have such shocks be followed by price declines that are sufficiently large that, eventually, the price level ends up *below* its initial value (when corrected for the average rate of inflation π^*). This is advantageous because the expectation of future price declines, by itself, dampens the initial inflationary effect of increases in \hat{G} and reductions in \hat{Y}^s . It is then possible to obtain the same degree of current price stabilization without having to raise interest rates so much. For this reason, the variability of interest rates is lower in the E_t rules than in the best price-level-targeting rules. This additional stability of interest rates is what makes the E_t rules more attractive from the point of view of the loss criterion $L + \pi^{*2}$ as well.

2.3 Optimal Policy

In this section we consider the best policy rule, from the point of view of minimization of our deadweight loss measure $L + \pi^{*2}$. We start by analyzing not monetary rules per se but allocations. In particular, we ask what (conditional) paths of output, inflation, and interest rates achieve the lowest value of $L + \pi^{*2}$ while being consistent with our IS and AS curves as well as with the stochastic process for \hat{G}_t and \hat{Y}_t^s given in equations (27) and (28). In other words, we compute the optimal response of the whole economy to these structural disturbances.

We then show that this optimal response of the economy is the unique equi-

librium that emerges when the interest rate is set according to a rule belonging to the general class (23). This means that one cannot do better from the point of view of minimizing $L + \pi^{*2}$ than using a rule within this class. Moreover, it should be obvious that the member of this class that induces the optimal allocation is also the optimal rule within the class (23).

2.3.1 Optimal Responses of the Economy to Real Disturbances

In this subsection we compute the optimal allocation and characterize it as a response of the economy to the innovations in \hat{G}_t and \hat{Y}_t^S . For this purpose we start by constructing a moving average representation of the stochastic process for the real disturbances \hat{G} and \hat{Y}^S . From equations (28) and (27) it follows that these variables can be written as functions of the history of the two independent identically distributed (i.i.d.) shocks in e_t^+ . Since these two shocks consist of $e_{2,t}$ and $e_{3,t}$, we can rewrite the stochastic process of the structural disturbances as

$$(46) \quad \hat{G}_{t+1} = \sum_{i=2}^3 \sum_{j=0}^{\infty} \Phi_{Gj}^i e_{i,t-j}, \quad \hat{Y}_{t+1}^S = \sum_{i=2}^3 \sum_{j=0}^{\infty} \Phi_{Sj}^i e_{i,t-j}.$$

The exact decomposition of the two shocks in (46) is irrelevant for present purposes; each is allowed to affect the evolution of both structural disturbances.

We now consider how the endogenous variables ought to evolve. Because we can write our loss measure in the form (34), it suffices to consider the evolution of \hat{X} . It should be obvious from equation (34) that there is no advantage to any random movements in \hat{X} apart from those needed for \hat{X}_t to respond to the shocks that contain information about the evolution of the real disturbances. Thus we may restrict attention to processes \hat{X} that may be written in the form

$$(47) \quad \hat{X}_{t+1} = \sum_{i=2}^3 \sum_{j=0}^{\infty} \Phi_{Xj}^i e_{i,t-j}.$$

Substituting this into equation (34), we find that L equals

$$(48) \quad L = \sum_{i=2}^3 \sigma_i^2 \left[\sum_{j=1}^{\infty} (\Phi_{Xj}^i)^2 + \psi(\Phi_{X0}^i)^2 + \frac{\Lambda}{\kappa^2} \sum_{j=1}^{\infty} (\Phi_{Xj}^i - \beta\Phi_{Xj+1}^i)^2 \right],$$

where σ_i^2 is the variance of $e_{i,t}$. We seek to obtain parameters Φ_{Xj}^i that make L as low as possible for a given variance of the funds rate, subject also to any constraints on the joint evolution of \hat{X} and \hat{R} implied by our structural model.

Using equations (38) and (46) we can write $\hat{\rho}_t$ as a function of the lagged e s. This means that using equation (47) in (37) we obtain an expression for the funds rate as a function of the history of the shocks,

$$(49) \quad \hat{R}_{t+1} = \sum_{i=2}^3 \sum_{j=0}^{\infty} \Phi_{Rj}^i e_{it-j},$$

where the coefficients Φ_{Rj}^i can be written as functions of the coefficients Φ_{Gj}^i , Φ_{Sj}^i , and Φ_{Xj}^i . This in turn allows us to write

$$(50) \quad \text{var}\{\hat{R}\} = \sum_{i=2}^3 \sum_{j=0}^{\infty} (\Phi_{Rj}^i)^2 \sigma_i^2.$$

The characterization of the optimal process \hat{X} then reduces to the choice of coefficients Φ_{Xj}^i to minimize the Lagrangian $L + \lambda \text{var}\{\hat{R}\}$, where we substitute expression (48) for L and (50) for $\text{var}\{\hat{R}\}$. Here $\lambda \geq 0$ is a multiplier indicating the weight placed on the variance of the funds rate. By minimizing the Lagrangian for different choices of $\lambda \geq 0$, we obtain the family of constrained optimal equilibria. This family corresponds to the frontier of minimum possible values of L for any given level of $\text{var}\{\hat{R}\}$ (and hence of π^*) that we report in Rotemberg and Woodford (1997).

There exists a particular value of λ such that the marginal reduction in π^{*2} from raising λ further (using eq. [39] to determine the lowest value of π^* consistent with any given value of $\text{var}\{\hat{R}\}$) is of the same size as the resulting increase in L .²⁷ The constrained optimal equilibrium associated with this particular value of λ achieves the minimum value of $L + \pi^{*2}$ among all allocations consistent with the structural equations of our model. The variability of inflation, output, the funds rate, and long-run price level growth in this allocation are indicated by point I in figures 2.1, 2.2, and 2.3 above, and the row labeled I in table 2.1.

Observe that the optimal allocation does not involve complete stabilization of inflation or of the long-run price level. This is not because complete stabilization is impossible in principle, but because complete stabilization would require too great a degree of volatility of the funds rate, and consequently too high an average inflation rate.²⁸ Thus the concern expressed by Summers (1991)—that the desire to maintain a very low average rate of inflation conflicts with the desire to use interest rates as an instrument of stabilization, given the existence of a zero nominal interest rate floor—matters quantitatively in the context of our model. On the other hand, our results imply that it is possible, at least in principle, to stabilize both inflation and the funds rate—and thus both the average rate of inflation and the variance of inflation—to a greater extent than has been achieved by historical policy. This can be seen

27. As reported in Rotemberg and Woodford (1997), the value of λ is approximately 0.2249.

28. As explained in subsection 2.1.2 above, complete stabilization of the path of the price level would require that $\hat{R}_t = \hat{p}_{t-1}$ each period. Given our estimated shock processes, this would imply a standard deviation of funds rate variations of 27 percentage points—10 times the funds rate volatility associated with historical policy. (See table 2 in Rotemberg and Woodford 1997.) Using eq. (39) to determine the minimum required value for π^* , we conclude that the average inflation rate would have to equal 58 percent per year.

from the relative locations of points H and I in figure 2.2. Similarly, table 2.1 shows that both L and π^* are lower with the optimal policy than with historical policy.

2.3.2 Implementing the Optimal Allocation

While the optimal allocation is consistent with equation (49), it is important to stress that equation (49) does not represent a viable policy proposal, even if the Fed could directly observe the structural disturbances and infer the history of the shocks e_{it} . Such a way of setting interest rates would, instead, result in price level indeterminacy because the path of the funds rate would be exogenously specified, with no feedback from the evolution of prices or real activity.²⁹ Thus the construction of a feedback rule for the funds rate that implements the optimal allocation—that is not only consistent with it but also renders it the unique stationary equilibrium consistent with the proposed policy rule—remains a nontrivial problem. Furthermore, it is of considerable interest to ask how policy should make use of the information revealed by the evolution of inflation and output, as in the various variants of the Taylor rule discussed in section 2.2. Thus we are especially interested in finding a rule of the form

$$(51) \quad C(L)\hat{r}_t = A(L)\hat{\pi}_t + B(L)\hat{Y}_t,$$

where $A(L)$, $B(L)$, and $C(L)$ are finite-order lag polynomials, that implements the optimal allocation.

To do this, we first consider whether any rule of this form is consistent with the stochastic processes for interest rates, inflation, and output that characterize this allocation. Substituting expression (47) into (20) and (36), we can write $\hat{\pi}_t$ and \hat{Y}_t as moving averages of $e_{2,t}$ and $e_{3,t}$. These moving average representations for the optimal evolution of inflation and output can be written compactly as

$$(52) \quad \hat{Z}_t \equiv \begin{bmatrix} \hat{\pi}_t \\ \hat{Y}_t \end{bmatrix} = \Phi_Z(L)e_{t-1},$$

where

$$e_t \equiv \begin{bmatrix} e_{2,t} \\ e_{3,t} \end{bmatrix}.$$

Similarly writing equation (49) as $\hat{R}_t = \Phi_R(L)e_{t-1}$, it would then seem natural to attempt to obtain a representation of the form

$$(53) \quad \hat{R}_t = \theta(L)\hat{Z}_t$$

29. The result that equilibrium is indeterminate in this case can be observed from the fact that the point $a = 0, b = 0$ is in the zone of indeterminacy in figs. 2.4 through 2.8, or similarly that the point $a = 0, c = 0$ is in the zone of indeterminacy in figs. 2.9 through 2.12.

by writing $\theta(L) \equiv \Phi_R(L)\Phi_Z^{-1}(L)$. Unfortunately, $\Phi_Z(L)$ does not prove to be invertible, since the polynomial $|\Phi_Z(z)|$ has a root inside the unit circle.³⁰ This root is $1/c$, where c is approximately 1.3267. We can, however, write $\Phi_Z(L) = (1 - cL)D(L)$, where $D(L)$ is invertible so that $(1 - cL)e_{t-1}$ is equal to $D(L)^{-1}\hat{Z}_t$. This means that

$$(54) \quad (1 - cL)\hat{R}_t = (1 - cL)\Phi_R(L)e_{t-1} = \tilde{\theta}(L)\hat{Z}_t,$$

where $\tilde{\theta}(L) \equiv \Phi_R(L)D(L)^{-1}$.

This gives us a relation of the form (51) between the funds rate, its own past values, and current and past values of inflation and output. The two elements ($i = 1, 2$) of the matrix lag polynomial $\tilde{\theta}(L)$ can be written as

$$\tilde{\theta}_i(L) \equiv \sum_{j=0}^{\infty} \tilde{\theta}_{ij}L^j,$$

where the coefficients $\tilde{\theta}_{ij}$ are square-summable, so that long lags j contribute only a small amount to the overall variation in the right-hand side of equation (54). However, the coefficients $\tilde{\theta}_{ij}$ die out for large j only relatively slowly; they evolve asymptotically according to the difference equation

$$\tilde{\theta}_{ij} = -\tilde{a}\tilde{\theta}_{i,j-1} - \tilde{b}\tilde{\theta}_{i,j-2},$$

where the coefficients \tilde{a} and \tilde{b} are approximately equal to 1.0404 and 0.9643, respectively.³¹ These values imply that the characteristic equation $z^2 + \tilde{a}z + \tilde{b} = 0$ has a pair of complex roots with modulus approximately equal to 0.9820. Thus the coefficients $\tilde{\theta}_{ij}$ decline in magnitude only at an average rate of less than 2 percent per quarter; a very long distributed lag is required for an accurate approximation to the exactly optimal rule of the form (54). The length of the distributed lag that is needed can be reduced significantly by further quasi differencing of \hat{R}_t , yielding a rule of the form (51), where $A(L) \equiv a_0 + a_1L + a_2L^2 \dots \equiv (1 + \tilde{a}L + \tilde{b}L^2)\tilde{\theta}_1(L)$, $B(L) \equiv b_0 + b_1L + b_2L^2 \dots \equiv (1 + \tilde{a}L + \tilde{b}L^2)\tilde{\theta}_2(L)$, and $C(L) \equiv 1 - c_1L - c_2L^2 \dots \equiv (1 + \tilde{a}L + \tilde{b}L^2)(1 - cL)$. The coefficients in the matrix lag polynomials $A(L)$ and $B(L)$ then become negligible much sooner.

Ignoring the constant, the policy rule that we derive in this fashion can then be written as

30. We demonstrate this numerically by truncating the infinite lag polynomial $\Phi_Z(L)$ at a finite number of lags and solving for the roots of $|\Phi_Z(z)|$. In our numerical work, we use the terms for $j = 0$ through 130 in eq. (47). We stop at lag 130 because both $\Phi_Z(1)$ and our estimate of the root of $|\Phi_Z(z)|$, which lies inside the unit circle, are little affected by the addition of further terms.

31. We determined this by inspection of the coefficients $\tilde{\theta}_{ij}$, which can be computed recursively. The coefficients that we compute obey the stated recursion, up to four decimal places of accuracy, for both $i = 1$ and 2, for all values of j between 61 and 92. After this, the recursion breaks down, presumably because a small numerical error in our estimate of c introduces a nontrivial error into our computation of $\tilde{\theta}_{ij}$ for larger values of j .

$$\begin{aligned}
\hat{r}_t = & .29\hat{r}_{t-1} + .42\hat{r}_{t-2} + 1.28\hat{r}_{t-3} \\
& + .22\hat{Y}_{t-1} - .25\hat{Y}_{t-2} + \dots \\
(55) \quad & + .16\hat{\pi}_t + 1.00\hat{\pi}_{t-1} + 2.45\hat{\pi}_{t-2} - 1.45\hat{\pi}_{t-3} + .74\pi_{t-4} \\
& - .08\hat{\pi}_{t-5} + .25\hat{\pi}_{t-6} + .33\hat{\pi}_{t-7} + .23\hat{\pi}_{t-8} + .25\hat{\pi}_{t-9} \\
& + .19\hat{\pi}_{t-10} + .17\hat{\pi}_{t-11} + .13\hat{\pi}_{t-12} + .09\hat{\pi}_{t-13} + .06\hat{\pi}_{t-14} + \dots,
\end{aligned}$$

where the omitted terms in \hat{Y}_{t-j} are all of size 0.01 or smaller (to two decimal places) and the omitted terms in $\hat{\pi}_{t-j}$ are all of size 0.03 or smaller (to two decimal places).³² Supposing that the monetary policy rule is given by expression (55), we find that our model has a unique stationary rational expectations equilibrium. Furthermore, this unique equilibrium involves responses of output, inflation, and interest rates to the real shocks that closely approximate the optimal responses derived in the previous subsection. Thus we conclude that rule (55) does belong to the admissible class of interest rate feedback rules resulting in a determinate equilibrium; that it represents a good approximation to the optimal rule within the general class of rules of the form (51); and that the optimal rule within this class implements the optimal allocation as defined above. Rule (55) is accordingly the optimal policy rule, labeled I in table 2.1 and in figures 2.1, 2.2, and 2.3.

Several features of this optimal interest rate feedback rule are worth noting. First, the coefficient on $\hat{\pi}_t$ is a small positive number, which means that the optimal rule calls for some immediate tightening in response to an observation of inflation above the target level. However, most of the tightening prescribed by the rule in response to an inflation rate above target occurs later. This subsequent tightening is reflected both in the series of positive coefficients on lagged inflation deviations $\hat{\pi}_{t-j}$ in rule (55), and in the series of positive coefficients on lagged deviations of the funds rate itself. Even putting aside the consequences of the lagged funds rate terms, the $\hat{\pi}_{t-j}$ terms in (55) prescribe a much larger response to lagged inflation than to current inflation; for example, these terms place an average weight of 0.7 on the rate of inflation over quarters 1 through 4 prior to the quarter in which the funds rate is being set, or four times the weight that is placed on inflation in the current quarter.

Second, the coefficient on \hat{Y}_t is exactly zero. This means that the optimal response to an innovation in \hat{C}_t that increases output relative to what it would have been forecast to be a quarter earlier is to keep the interest rate at t unchanged. This does not mean that there is no optimal response to observed

32. For $j = 50$ and above, the terms in $\hat{\pi}_{t-j}$ are all .0001 or smaller, while the same is true of the \hat{Y}_{t-j} terms for $j = 19$ and above. Whether the small nonzero values that we still obtain for large j indicate that further quasi differencing is needed in order to obtain a rule of the form (51), with finite lag polynomials, or are simply due to numerical error, we have not been able to determine.

variations in output relative to trend, but that the optimal response is a delayed one. Moreover, the interest rate ought to respond more to the *growth rate* of output a quarter earlier than to the *level* of output.³³

Finally, the lag polynomial $C(L)$ has a root inside the unit circle, equal to the reciprocal of $c = 1.33$. Thus, just as in our optimal generalized Taylor rules, the optimal rule calls not simply for interest rate smoothing but for an explosively growing response of the funds rate to deviations of inflation from target. These explosions are avoided only if subsequent deviations with the opposite sign eventually counteract the effects of an initial deviation. If the inflation rate were permanently above its target, interest rates would grow asymptotically as $(1.33)^t$, just as if we chose $c = 1.33$ in the case of the family of simple rules (41). This explosive behavior is of course exactly what we concluded was desirable in our previous discussion of simple rules, and indeed the value $c = 1.33$ is not too different from the most desirable value of c in the case of simple rules.

One way of comparing the implications of rule (55) with those of other candidate interest rate rules is to plot its implications for the cumulative response of the funds rate to a sustained deviation of either inflation from target or output from its trend level. This particular way of describing the various feedback rules has the advantage of being independent of the degree of quasi differentiation that may have been used in the way that the rule is stated; for example, it treats expressions (54) and (55) as equivalent. The prescribed cumulative responses of the funds rate to sustained 1 percent deviations in the two variables are displayed in the two panels of figure 2.19. Each panel compares the prescribed response of the funds rate under four different rules: our estimate of historical U.S. policy over the period 1979–95, the rule proposed by Taylor (1993) as a rough description of recent U.S. policy, the optimal rule E_0 within the family (41), and our unrestricted optimal rule (55). We see in the top panel that, after two quarters of inflation being above target, the first two rules (which are quite similar to each other in this respect) involve much smaller responses of the interest rate to the inflation deviation than the latter two. Our unrestricted optimal rule is actually less aggressive than the optimal rule of the form (41) over the horizon displayed in this panel. Indeed, the initial reaction to inflation is actually smaller in this unrestricted optimum than it is in the case of the simple Taylor rule. This serves to highlight once again the fact that our model recommends postponing the reaction of interest rates while simultaneously increasing the absolute magnitude of these delayed reactions.

The bottom panel shows the responses to a sustained output deviation. Here the Taylor rule and our estimate of actual policy involve much stronger reactions of the interest rate over the first three quarters than are implied by either

33. Interestingly, our estimated historical policy rule for the United States, reported in Rotemberg and Woodford (1997), also implies more response to the growth rate than to the detrended level of real GDP; but the historical rule can be more accurately described as making the funds rate respond to $\hat{Y}_t - \hat{Y}_{t-2}$ rather than to $\hat{Y}_{t-1} - \hat{Y}_{t-2}$.




Fig. 2.19 Cumulative interest rate response under alternative policies

of the optimal rules. For the unrestricted optimal rule, the reaction remains more muted for the entire six-quarter horizon displayed here. This indicates an important difference between actual policy, at least as either Taylor or we have characterized it, and optimal policy according to our model: our model suggests that interest rate responses to output above trend should be much weaker, at least in the first few quarters, than they actually are. On the other hand, this does not mean that optimal policy would not involve interest rates eventually

being raised. For the optimal policy in the class (41), interest rates are actually higher after five quarters of high output than they would be under actual policy or the simple Taylor rule. If one extends the plot a few more quarters, this is also true of the unrestricted optimal policy, and both optimal rules (unlike the two characterizations of actual policy) imply that the funds rate eventually explodes.

A final feature of the optimal rule that is worth pointing out is its implication for long-run price level stability. We observe that the optimal rule has the form of an inflation-targeting rule, rather than a price-level-targeting rule, and indeed it does not imply trend stationarity of the price level. On the other hand, it *does* imply a tendency for unexpected increases in the price level to be subsequently offset by (forecastable) price level declines. This is indicated by a coefficient β^∞ that is negative, indicating that unexpected price level increases are eventually more than completely offset by subsequent price level declines, as in the case of the optimal simple rules E_0 and E_1 . As a result, optimal policy involves a significant degree of stabilization of the rate of change of long-run price level forecasts—the standard deviation of Δp^∞ is reduced by a factor of four, relative to our estimate of historical policy.

Finally, it is worth asking to what extent our analysis implies that a simple rule such as E_0 or E_1 can be improved upon by using additional information. We have already observed that, according to our structural model, the history of inflation and output variations alone, if observed with sufficient accuracy and timeliness, provide all of the information needed to implement the optimal equilibrium. Thus a sufficiently flexible rule of the form (51) suffices. But as a practical matter, it is probably even more interesting to observe that our results imply that the unrestricted optimal rule is not too different from, and not too much better than, the optimal rule within a simple family such as (41). Ninety-nine percent of the nearly 15-fold reduction in the size of the dead-weight loss $L + \pi^{*2}$ that is achievable by going from actual policy to optimal policy can be obtained by adopting the simple rule E_0 . Furthermore, if E_0 is not considered operational due to its reliance on measures of the current quarter's inflation and output, the simple rule E_1 , which requires only the previous quarter's data, results in performance that is nearly as good. Thus our analysis supports the view that simple policy rules, variations on the sort of rule proposed by Taylor (1993), have highly desirable properties both from the point of view of stabilizing inflation, interest rates, and the long-run price level, as well as from the point of view of economic welfare.

2.4 Conclusions

Our results offer a number of conclusions of importance for the design of a monetary policy rule. All of our conclusions are subject, of course, to the caveat that the seriousness with which they should be taken depends on one's confidence in the extent to which the specification of our structural model is not grossly incorrect.

Probably our most important conclusion is that a simple interest rate feedback rule of the kind proposed by Taylor (1993) can achieve outcomes nearly as good as are achievable in principle by *any* policy, assuming that the commitment of the monetary authority to the rule can be made sufficiently credible. At least in the context of the simple structural model that we consider, an interest rate feedback rule that uses only information about the recent behavior of inflation and output does quite well (and only the response to the recent level of inflation matters much for this). Furthermore, performance under the best rule of this kind is not significantly reduced if *lagged* inflation data are used. Thus lags in the availability of accurate measurements of inflation are not necessarily a serious problem for the implementation of such a rule.

It is worth noting in particular that a “backward-looking” rule, in which interest rates respond to measures of inflation that has already occurred, rather than to forecasts (of one sort or another) of future inflation (as in the rules considered by Rudebusch and Svensson, in chap. 5 of this volume, and Batini and Haldane, in chap. 4), do quite well. We show that, at least in our simple model, the theoretically optimal policy has a backward-looking representation, given by expression (55). Perhaps more to the point, even very simple backward-looking rules, such as rules E_0 and E_1 in table 2.1, are quite good approximations to optimal policy.

It is interesting to note that we obtain this result despite using a structural model that implies that monetary policy has no effects on inflation until the following quarter (and the largest effect only after two quarters), and no effects on real activity until after two quarters. Lags in the effects of a monetary policy change do *not* imply that an effective policy must be “forward looking.” The crucial insight is that there is no need for *policy* to be forward looking as long as the *private sector* is. A commitment to raise interest rates later, after inflation increases, is sufficient to cause an immediate contraction of aggregate demand in response to a shock that is expected to give rise to inflationary pressures. This channel should be effective as long as aggregate demand depends on expected *future* interest rates (or, equivalently, on *long* rates) and not simply on current short rates; as long as the monetary authority is understood to be committed to adhering to the contemplated policy rule in the future, and not only at the present time; and as long as private agents have model-consistent (or “rational”) expectations. Indeed, if, as our model implies, aggregate demand is affected *only* by expectations of future interest rates, and not by unexpected interest rate variations (either immediately or with a lag), then a credible commitment to systematically respond in the future is the only way in which monetary policy can be effective. But when one conceives policy in these terms, there is no need for that commitment about future action to involve a commitment to be forward looking at that future date.

Despite our general support for the type of policy rule proposed by Taylor, our analysis suggests that the best rules differ from the specific rule that he proposes in important respects. Probably the most important difference is our conclusion that short-term interest rates should depend not only on deviations

of inflation from target but also on their *own* past values—ideally with a coefficient even greater than one. A less radical-sounding version of our proposal would be to make the *change* in the funds rate, rather than the *level* of the funds rate, a function of deviations of inflation from its target value, as is also found to be desirable in the forward-looking models studied by Levin, Wieland, and Williams (chap. 6 of this volume). It is interesting to note that in forward-looking models of these kinds, such dependence, even with a coefficient greater than one on the lagged value, does not lead to instrument instability. This result contrasts sharply with the conclusion that one would obtain using a traditional, purely backward-looking macroeconomic model, such as the one considered by Rudebusch and Svensson (chap. 5 of this volume).

In our analysis, the desirability of such dependence on the lagged funds rate does not rest on any assumption that variability in the change in the funds rate from one period to the next is a bad thing in itself. Rather, it represents a way of allowing the central bank to commit itself to raise interest rates later, in response to an increase in inflation that is not offset by a subsequent (and sufficiently prompt) inflation decline, without having to have much of the *eventual* interest rate response occur immediately. Assuming that the private sector understands this commitment and is forward looking in its behavior, this allows the central bank to have a large effect on aggregate demand without having (in equilibrium) to move interest rates very much. This in turn is desirable if one wishes to maintain a low volatility of interest rates. We argue that a low volatility of the funds rate is in fact desirable as it allows a given degree of inflation stabilization to be consistent with a lower *average* rate of inflation, due to the zero floor for the nominal funds rate.

Our results here plainly depend on the assumption not only that the private sector is forward looking but that private agents fully understand and believe in the central bank's policy rule. One might wonder whether such an analysis gives a correct account of the consequences of adopting such a rule, especially in the short run, given that it would represent a significant departure from present policy (according to our estimates). Nonetheless, our analysis shows that the possibility of achieving a significant degree of stabilization without a great deal of interest rate volatility through this channel is an important advantage of a high degree of credibility for the central bank's commitment to a monetary policy rule. This helps to clarify why the design of arrangements under which such a rule could be credible could have significant benefits.

Another respect in which our conclusions differ from Taylor's proposal is that we find little gain from making interest rates depend on the current level of economic activity. We find that optimal rules within our simple families involve a small positive response to the level of detrended output, but it is much more modest than the sort of response suggested by Taylor, or indicated by our estimate of actual U.S. policy. The reason it is undesirable to respond to output deviations, in our model, is that deviations of output from trend have so little to do with deviations of output from *potential* (which, according to our esti-

mates, is quite volatile). It is possible that an alternative interpretation of the residuals in our aggregate supply equation, under which they would not all represent variations in the *efficient* level of output, would increase the role for responses to output variations in an optimal rule. Alternatively, it is possible that if we considered other real variables (such as employment) along with variations in detrended output, we would be able to construct a better proxy for deviations of output from potential (as proposed, e.g., by McCallum and Nelson in chap. 1 of this volume), to which it would be desirable for interest rates to respond.

Finally, our results shed light on the debate about the relative advantages of price level targeting and inflation targeting. We find that under a desirable policy, the central bank should consistently act to subsequently reverse any movements of inflation above its target level, rather than simply preventing further price increases without undoing the ones that have already occurred. Nonetheless, according to our analysis, there is no special significance to the goal of returning the price level to a deterministic target path. Our optimal policy rules actually imply that an unexpected increase in inflation should *decrease* the expected long-run price level. Such an outcome is obtained by a policy that involves no reference to a target price level path. It follows simply from the dependence of the funds rate on the lagged funds rate, mentioned above, which has the consequence that, in equilibrium, inflation increases must be followed by subsequent, and even greater, inflation declines in order to avoid causing the funds rate to grow explosively.

Appendix

Derivation of the Utility-Based Loss Function

Here we present further details of the derivation of equations (30), (31), (33), and (34), which describe our utility-based loss function $L + \pi^{*2}$. We begin with the derivation of equation (30) as a second-order Taylor series approximation to (29). Note that our objective function is of the form $W \equiv Ew_t$, where w_t is the average utility flow (integrating over the continuum of households) each period. This utility flow may be written as a function solely of the pattern of real activity $\{y_t(z)\}$ within a period, and the exogenous shocks:

$$(A1) \quad w_t = u(Y_t - G_t; \xi_t) - \int_0^1 v(y_t(z); \xi_t) dz.$$

We begin by considering a Taylor series expansion for each of the two terms in this expression, expanding around the levels of output $y_t(z) = \bar{Y}$ for each z and the values $G_t = \bar{G}$ and $\xi_t = 0$ for the exogenous shocks. Here \bar{Y} represents the level of output in an *optimal* steady state; it represents the constant equilibrium level of output in an equilibrium with no variation in the values of G_t and

ξ_t around their steady state values, a constant price level, and a tax rate $\tau = \tau^* \equiv -(\theta - 1)^{-1}$ that perfectly offsets the distortion resulting from firms' monopoly power. (As we shall see, our loss function takes an especially simple form in this case, and we wish to direct attention to the terms in it that survive even under these ideal circumstances. We leave for further work the analysis of how the welfare effects of monetary policy change when one considers possible interactions between monetary policy and distortions other than the one resulting from sluggish nominal price adjustment.) The steady state value \bar{G} is chosen to equal $E(G_t)$, and the shocks ξ_t are normalized so that $E(\xi_t) = 0$; thus the steady state values of the exogenous variables equal their unconditional means.

A second-order Taylor series expansion for the first term on the right-hand side of equation (A1) is given by

$$\begin{aligned}
 (A2) \quad u &= u(\bar{C}; 0) + u_c(C_t - \bar{C}) + u_\xi \xi_t \\
 &+ \frac{1}{2} u_{cc}(C_t - \bar{C})^2 + u_{c\xi}(C_t - \bar{C})\xi_t + \frac{1}{2} u_{\xi\xi} \xi_t^2 + \mathcal{O}(\|\xi\|^3) \\
 (A3) \quad &= u(\bar{C}; 0) + u_c \bar{Y}(\hat{Y}_t + \frac{1}{2} \hat{Y}_t^2 - \tilde{G}_t) + u_\xi \xi_t \\
 &+ \frac{1}{2} u_{cc} \bar{Y}^2 (\hat{Y}_t - \tilde{G}_t)^2 + u_{c\xi} \bar{Y} (\hat{Y}_t - \tilde{G}_t) \xi_t + \frac{1}{2} u_{\xi\xi} \xi_t^2 + \mathcal{O}(\|\xi\|^3) \\
 (A4) \quad &= u_c \bar{Y} \cdot \hat{Y}_t + \frac{1}{2} (u_c \bar{Y} + u_{cc} \bar{Y}^2) \hat{Y}_t^2 \\
 &- u_{cc} \hat{Y}_t^2 (\tilde{G}_t + s_c \bar{C}_t) \hat{Y}_t + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3) \\
 (A5) \quad &= u_c \bar{Y} \cdot \hat{Y}_t + \frac{1}{2} (u_c \bar{Y} + u_{cc} \bar{Y}^2) \hat{Y}_t^2 \\
 &- u_{cc} \hat{Y}_t^2 \cdot \hat{G}_t \hat{Y}_t + \text{t.i.p.} + \text{unf.} + \mathcal{O}(\|\xi\|^3).
 \end{aligned}$$

In (A2) we simply expand in terms of the index of aggregate consumption C_t , where $\bar{C} \equiv \bar{Y} - \bar{G}$, and each of the partial derivatives is evaluated at the steady state values $(\bar{C}; 0)$. Here the term $\mathcal{O}(\|\xi\|^3)$ indicates that we neglect terms that are of third or higher order in the deviations of the various variables from their steady state values. In the case of a monetary policy rule that implies $\pi^* = 0$ and a tax rate $\tau = \tau^*$, the variables will deviate from these values in an equilibrium only because of fluctuations in the shocks G_t and ξ_t around their steady state values. In this case, the omitted terms are all of third or higher order in the size of the exogenous shocks (and we use $\|\xi\|$ to indicate the a measure of the size of these shocks, where the size of fluctuations in G_t is intended to be included). More generally, the omitted terms also include terms that are of third or higher order in deviations of π^* from the value zero and of

τ from the value τ^* ; but we shall (for now) retain terms that are of first or second order in perturbations of those assumptions about long-run aspects of policy. In equation (A3) we rewrite the expressions in terms of $\hat{Y}_t \equiv \log(Y_t/\bar{Y})$ and $\hat{G}_t \equiv (G_t - \bar{G})/\bar{Y}$, using the Taylor expansion

$$Y_t = \bar{Y}(1 + \hat{Y}_t + \frac{1}{2}\hat{Y}_t^2) + \mathcal{O}(\|\xi\|^3).$$

In equation (A4) we suppress the terms that are independent of policy (because they involve only constants and exogenous disturbances), denoted “t.i.p.” as in the text, and make use of the definition $u_{c\xi}\xi_t = -u_{cc}\bar{C}_t$ to obtain a scalar representation of the disturbance to the marginal utility of consumption. Finally, in equation (A5) we recall the notation

$$\hat{G}_t \equiv \tilde{G}_t + s_c E_{t-2}\bar{C}_t = \tilde{G}_t + s_c \bar{C}_t + \text{unf.},$$

where “unf.” stands for an unforecastable term (i.e., a term x_t with the property that $E_{t-2}x_t = 0$). Unforecastable terms may be neglected because we are ultimately interested only in the unconditional expectation of each of the terms in equation (A5).

Similarly, a second-order Taylor series expansion of household z 's disutility of working is given by

$$\begin{aligned} v &= v(\bar{Y}; 0) + v_y(y_t(z) - \bar{Y}) + v_{\xi}\xi_t \\ &+ \frac{1}{2}v_{yy}(y_t(z) - \bar{Y})^2 + v_{y\xi}(y_t(z) - \bar{Y})\xi_t + \frac{1}{2}v_{\xi\xi}\xi_t^2 + \mathcal{O}(\|\xi\|^3) \\ \text{(A6)} \quad &= v_y\bar{Y} \cdot \hat{y}_t(z) + \frac{1}{2}(v_y\bar{Y} + v_{yy}\bar{Y}^2)\hat{y}_t(z)^2 - v_{yy}\bar{Y}^2 \cdot \hat{y}_t(z)\bar{Y} \\ &+ \text{t.i.p.} + \mathcal{O}(\|\xi\|^3), \end{aligned}$$

where now $\hat{y}_t(z) \equiv \log(y_t(z)/\bar{Y})$ and \bar{Y}_y , defined by the relation $v_{y\xi}\xi_t = -v_{yy}\bar{Y}_y$, provides a scalar measure of disturbances to the marginal disutility of supply. Integrating equation (A6) over z we obtain

$$\begin{aligned} \int_0^1 v(y_t(z); \xi_t) dz &= v_y\bar{Y} \cdot E_z\hat{y}_t(z) + \frac{1}{2}(v_y\bar{Y} + v_{yy}\bar{Y}^2) \cdot [E_z\hat{y}_t(z)]^2 \\ \text{(A7)} \quad &+ \frac{1}{2}(v_y\bar{Y} + v_{yy}\bar{Y}^2) \text{var}_z\hat{y}_t(z) - v_{yy}\bar{Y}^2 \cdot E_z\hat{y}_t(z) \\ &+ \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned}$$

Next we wish to express the terms in equation (A7) involving the population average $E_z\hat{y}_t(z)$ in terms of the Dixit-Stiglitz output aggregate \hat{Y}_t instead. To do so, we first compute a Taylor series expansion for the right-hand side of the aggregator equation (2), obtaining

$$\hat{Y}_t = E_z \hat{y}_t(z) + \frac{1}{2} \frac{\theta - 1}{\theta} \text{var}_z \hat{y}_t(z) + \mathcal{O}(\|\xi\|^3).$$

Solving this equation for $E_z \hat{y}_t(z)$ and substituting into equation (A7) yields

$$\begin{aligned} \int_0^1 v(y_t(z); \xi_t) dz &= v_y \bar{Y} \cdot \hat{Y}_t + \frac{1}{2} (v_{yy} \bar{Y} + v_{yy} \bar{Y}^2) \hat{Y}_t^2 \\ (A8) \quad &+ \frac{1}{2} (\theta^{-1} v_y \bar{Y} + v_{yy} \bar{Y}^2) \text{var}_z \hat{y}_t(z) - v_{yy} \bar{Y}^2 \cdot \hat{Y}_t \bar{Y} \\ &+ \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned}$$

Substituting equations (A5) and (A8) into (A1), we obtain

$$\begin{aligned} w_t &= u_c \bar{Y} \left[\hat{Y}_t + \frac{1}{2} (1 - \sigma) \hat{Y}_t^2 + \sigma \hat{G}_t \hat{Y}_t \right] \\ &- v_y \bar{Y} \left[\hat{Y}_t + \frac{1}{2} (1 + \omega) \hat{Y}_t^2 + \frac{1}{2} (\theta^{-1} + \omega) \text{var}_z \hat{y}_t(z) - \omega \bar{Y} \hat{Y}_t \right] \\ (A9) \quad &+ \text{t.i.p.} + \text{unf.} + \mathcal{O}(\|\xi\|^3) \\ &= -\frac{1}{2} u_c \bar{Y} [(\sigma + \omega) \hat{Y}_t^2 - 2(\sigma + \omega) \hat{Y}_t^s \hat{Y}_t + (\theta^{-1} + \omega) \text{var}_z \hat{y}_t(z)] \\ &+ \text{t.i.p.} + \text{unf.} + \mathcal{O}(\|\xi\|^3). \end{aligned}$$

Note that in deriving (A9) from the line above we have (at last) used the assumption that \bar{Y} is the efficient level of output, so that $u_c = v_y$, and the definition

$$\hat{Y}_t^s \equiv (\sigma + \omega)^{-1} (\sigma \hat{G}_t + \omega E_{t-1} \bar{Y}_t) = (\sigma + \omega)^{-1} [\sigma \hat{G}_t + \omega \hat{Y}_t] + \text{unf.}$$

Then taking the unconditional expectation of equation (A9), we obtain

$$\begin{aligned} (A10) \quad W &= -\frac{1}{2} u_c \bar{Y} [(\sigma + \omega) \text{var}\{\hat{Y}_t - \hat{Y}_t^s\} + (\sigma + \omega) [E(\hat{Y}_t)]^2 \\ &+ (\theta^{-1} + \omega) E \text{var}_z \hat{y}_t(z)] + \text{t.i.p.} + \mathcal{O}(\|\xi\|^3). \end{aligned}$$

As promised, we have obtained a welfare measure that allows us to compute all second-order or lower terms in W using only a first-order (log-linear) approximation to the equilibrium solution for the pattern of activity $\{y_t(z)\}$, since no terms of order $\mathcal{O}(\|\xi\|^2)$ in the solution for $y_t(z)$ have any effect on terms of order lower than $\mathcal{O}(\|\xi\|^3)$ in equation (A10).³⁴ If we furthermore assume that the

34. This result depends on our having linearized around the efficient \bar{Y} , since otherwise our expression for W would contain a term that is linear in $E(\hat{Y}_t)$. However, even without this choice, we could have obtained the same result by assuming that tax policy adjusts in response to any change in the monetary policy rule in order to preserve a particular value for $E(\hat{Y}_t)$, where this quantity then becomes one of the terms independent of the monetary policy rule. In fact, we assume that taxes respond to keep output fixed in the work reported here in any event.

tax rate τ (or some other aspect of “long-run” policy) is adjusted so as to guarantee that $E(\hat{Y}_t) = 0$ (i.e., $\log Y_t$ equals $\log \bar{Y}$ on average) *regardless* of the monetary policy rule, then the $[E(\hat{Y}_t)]^2$ term in equation (A10) can also be suppressed, as this term is also independent of the monetary policy rule. Our decision to assume this results from a belief, as discussed in the text, that monetary policy is not an appropriate instrument with which to seek to affect the long-run average level of economic activity, given the existence of other instruments with which policymakers may more directly seek to offset the distortion resulting from suppliers’ market power. Finally, noting that \hat{Y}_t equals $E_{t-2}\hat{Y}_t$ plus a forecast error term that is both unforecastable and independent of monetary policy, one can show that

$$\text{var}\{\hat{Y}_t - \hat{Y}_t^s\} = \text{var}\{E_{t-2}(\hat{Y}_t - \hat{Y}_t^s)\} + \text{i.i.p.}$$

Substitution of this into equation (A10), along with the stipulation that $E(\hat{Y}_t) = 0$ regardless of monetary policy, then yields equation (30).

Some might prefer instead an alternative that would assume a tax rate τ that remained invariant under alternative monetary policy rules. In this case, we would not be able to drop the $[E(\hat{Y}_t)]^2$ term in equation (A10). However, our model implies that

$$E(\hat{Y}_t) = \frac{1 - \beta}{\kappa} E(\hat{X}_t) + \mathcal{O}(\|\xi\|^2) = \frac{1 - \beta}{\kappa} E(\pi_t) + \mathcal{O}(\|\xi\|^2),$$

where the first equality follows from taking the unconditional expectation of all terms in equation (21), and the second from taking the unconditional expectation of all terms in (20). (Note that in the log-linear approximations to the model equations reported in the paper, we routinely suppress terms of order $\mathcal{O}(\|\xi\|^2)$.) Thus the only difference in the alternative case would be the presence of an additional negative term in π^{*2} in equation (31). This would have no effect on the definition of the loss from incomplete stabilization L in equation (33), but it would mean that in (32) we would have $L + \mu\pi^{*2}$ instead of $L + \pi^{*2}$, for a certain $\mu > 1$, as our overall deadweight loss measure. This would imply that the optimal point on the $L - \pi^*$ frontier (discussed and graphed in Rotemberg and Woodford 1997) would be slightly different from the one that we assume here, involving a slightly smaller, though still slightly positive, value of π^* . Such a change makes no qualitative difference, however, in the conclusions announced here about the nature of optimal policy.

Next we turn to the derivation of equation (31) from (30). As noted in the text, we need to show that the dispersion of levels of production across differentiated goods is a function of the degree of variability of the aggregate price level. We begin by noting that output dispersion follows from price dispersion, since equation (3) implies that

$$(A11) \quad E \text{var}_z\{\log y_t(z)\} = \theta^2 E \text{var}_z\{\log p_t(z)\}.$$

To relate the cross-sectional variance of prices to the variability over time of the price index P_t , we begin by recalling that in any period t , a fraction α of suppliers charge the same price as at $t - 1$ (and the distribution of their prices is the same as the distribution of period $t - 1$ prices); a fraction $(1 - \alpha)\gamma$ charge a common new price p_t^1 chosen at $t - 1$, and a fraction $(1 - \alpha)(1 - \gamma)$ charge a common new price p_t^2 chosen at $t - 2$. Then, introducing the notation $\bar{p}_t \equiv E_z \log p_t(z)$, we obtain

$$\begin{aligned}
 \text{var}_z\{\log p_t(z)\} &= \text{var}_z\{\log p_t(z) - \bar{p}_{t-1}\} = E_z\{[\log p_t(z) - \bar{p}_{t-1}]^2\} + (\Delta\bar{p}_t)^2 \\
 &= \alpha E_z\{[\log p_{t-1}(z) - \bar{p}_{t-1}]^2\} + (1 - \alpha)\gamma(\log p_t^1 - \bar{p}_{t-1})^2 \\
 \text{(A12)} \quad &+ (1 - \alpha)(1 - \gamma)(\log p_t^2 - \bar{p}_{t-1})^2 + (\Delta\bar{p}_t)^2 \\
 &= \alpha \text{var}_z\{\log p_{t-1}(z)\} + (1 - \alpha)\gamma(\log p_t^1 - \bar{p}_{t-1})^2 \\
 &+ (1 - \alpha)(1 - \gamma)(\log p_t^2 - \bar{p}_{t-1})^2 + (\Delta\bar{p}_t)^2.
 \end{aligned}$$

Taking the unconditional expectation of both sides of (A12) then yields

$$\begin{aligned}
 \text{(A13)} \quad E \text{var}_z\{\log p_t(z)\} &= \gamma E[(\log p_t^1 - \bar{p}_{t-1})^2] + (1 - \gamma)E[(\log p_t^2 - \bar{p}_{t-1})^2] \\
 &- (1 - \alpha)^{-1}E[(\Delta\bar{p}_t)^2].
 \end{aligned}$$

Similar reasoning as is used in deriving equation (A12) also yields

$$\begin{aligned}
 \bar{p}_t - \bar{p}_{t-1} &= E_z[\log p_t(z) - \bar{p}_{t-1}] \\
 \text{(A14)} \quad &= \alpha E_z[\log p_{t-1}(z) - \bar{p}_{t-1}] + (1 - \alpha)\gamma(\log p_t^1 - \bar{p}_{t-1}) \\
 &+ (1 - \alpha)(1 - \gamma)(\log p_t^2 - \bar{p}_{t-1}) \\
 &= (1 - \alpha)\gamma(\log p_t^1 - \bar{p}_{t-1}) + (1 - \alpha)(1 - \gamma)(\log p_t^2 - \bar{p}_{t-1}).
 \end{aligned}$$

Taking the expectation of (A14) conditional on date $t - 2$ information, one obtains

$$\text{(A15)} \quad E_{t-2}(\bar{p}_t - \bar{p}_{t-1}) = (1 - \alpha)(\log p_t^2 - \bar{p}_{t-1}) + \mathcal{O}(\|\xi\|^2),$$

using the facts that $\log p_t^1 = E_{t-2} \log p_t^1 + \mathcal{O}(\|\xi\|^2)$ and that all date $t - 1$ prices are known at $t - 2$. This combined with (A14) implies that

$$\begin{aligned}
 \text{(A16)} \quad (\bar{p}_t - \bar{p}_{t-1}) - (1 - \gamma)E_{t-2}(\bar{p}_t - \bar{p}_{t-1}) &= (1 - \alpha)\gamma(\log p_t^1 - \bar{p}_{t-1}) \\
 &+ \mathcal{O}(\|\xi\|^2).
 \end{aligned}$$

Furthermore, given that we are expanding around a steady state with zero inflation, the right-hand sides of both equations (A15) and (A16) consist solely of terms of order $\mathcal{O}(\|\xi\|)$. Thus by squaring (A15) and taking the unconditional expectation, we obtain

$$\begin{aligned} E[(\log p_t^i - \bar{p}_{t-1})^2] &= (1 - \alpha)^{-2} E[(E_{t-2} \Delta \bar{p}_t)^2] + \mathcal{O}(\|\xi\|^3) \\ &= (1 - \alpha)^{-2} \text{var}\{E_{t-2} \Delta \bar{p}_t\} + (1 - \alpha)^{-2} (E \Delta \bar{p}_t)^2 + \mathcal{O}(\|\xi\|^3). \end{aligned}$$

A similar expression for $E[(\log p_t^i - \bar{p}_{t-1})^2]$ is implied by (A16). Substituting these expressions into (A13) then yields

$$\begin{aligned} (A17) \quad E \text{var}_z\{\log p_t(z)\} &= \frac{\alpha}{(1 - \alpha)^2} [\text{var}\{E_{t-2} \Delta \bar{p}_t\} + (E \Delta \bar{p}_t)^2] \\ &+ \frac{1 - \gamma(1 - \alpha)}{\gamma(1 - \alpha)^2} \text{var}\{\Delta \bar{p}_t - E_{t-2} \Delta \bar{p}_t\} + \mathcal{O}(\|\xi\|^3). \end{aligned}$$

Finally, the definition of the price index (4) implies that

$$\bar{p}_t = \log P_t + \mathcal{O}(\|\xi\|^2).$$

Making this substitution in equation (A17), we obtain

$$\begin{aligned} (A18) \quad E \text{var}_z\{\log p_t(z)\} &= \frac{\alpha}{(1 - \alpha)^2} [\text{var}\{E_{t-2} \pi_t\} + (E \pi_t)^2] \\ &+ \frac{1 - \gamma(1 - \alpha)}{\gamma(1 - \alpha)^2} \text{var}\{\pi_t - E_{t-2} \pi_t\} + \mathcal{O}(\|\xi\|^3). \end{aligned}$$

Substitution of equations (A11) and (A18) into (30), and the fact that (as a consequence of our definition of π^*) $E \pi_t = \pi^* + \mathcal{O}(\|\xi\|^2)$, then yields equation (31). This last expression can in turn obviously be written in the form (32), where L is defined by (33), and

$$\Omega \equiv \frac{1}{2} u_c \bar{Y} \theta (1 + \theta \omega) \frac{\alpha}{(1 - \alpha)^2}, \quad \Lambda \equiv \frac{(1 - \alpha) \kappa}{(1 - \alpha \beta) \theta}.$$

Finally, we can rewrite L so that it depends only on the stochastic process for the relative price variable \hat{X} . To do this, note first that equation (21) implies that

$$(A19) \quad E_{t-2}(\hat{Y}_t - \hat{Y}_t^s) = (1/\kappa) E_{t-2}(\hat{X}_t - \beta \hat{X}_{t+1}).$$

At the same time, equation (20) implies that

$$(A20) \quad \pi_t - E_{t-2} \pi_t = \psi(\hat{X}_t - E_{t-2} \hat{X}_t),$$

so that

$$(A21) \quad \text{var}\{\pi_t\} = \text{var}\{E_{t-2} \hat{X}_t\} + \psi^2 \text{var}\{\hat{X}_t - E_{t-2} \hat{X}_t\}.$$

Substituting the expressions in (A19), (A20), and (A21) into equation (33), we obtain (34).

References

- Bean, Charles R. 1983. Targeting nominal income: An appraisal. *Economic Journal* 93 (December): 806–19.
- Bénabou, Roland, and Jerzy D. Konieczny. 1994. On inflation and output with costly price changes: A simple unifying result. *American Economic Review* 84 (March): 290–97.
- Bernanke, Ben S., and Michael Woodford. 1997. Inflation forecasts and monetary policy. *Journal of Money, Credit and Banking* 24 (November): 653–84.
- Beveridge, Stephen, and Charles R. Nelson. 1981. A new approach to the decomposition of economic time series into permanent and transitory components with particular attention to measurement of the “business cycle.” *Journal of Monetary Economics* 7, no. 2 (March): 151–74.
- Calvo, Guillermo. 1983. Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12 (3): 383–98.
- Clarida, Richard, Jordi Galí, and Mark Gertler. 1998. Monetary policy rules in practice: Some international evidence. *European Economic Review* 42:1033–67.
- Dixit, Avinash K., and Joseph E. Stiglitz. 1977. Monopolistic competition and optimum product diversity. *American Economic Review* 67 (3): 297–308.
- Feldstein, Martin. 1997. Capital income taxes and the benefit of price stability. NBER Working Paper no. 6200. Cambridge, Mass.: National Bureau of Economic Research, September.
- Fuhrer, Jeffrey C., and Geoffrey Moore. 1995. Monetary policy trade-offs and the correlation between nominal interest rates and real output. *American Economic Review* 85 (March): 219–39.
- Gertler, Mark. 1996. Discussion of Friedman and Kuttner, “A price target for U.S. monetary policy?” *Brookings Papers on Economic Activity*, no. 1:126–34.
- Goodfriend, Marvin. 1991. Interest rate smoothing in the conduct of monetary policy. *Carnegie-Rochester Conference Series on Public Policy* 34:7–30.
- Goodhart, Charles A. E. 1997. Why do monetary authorities smooth interest rates? In *European monetary policy*, ed. S. Collignon. London: Pinter.
- Hall, Robert E., and N. Gregory Mankiw. 1994. Nominal income targeting. In *Monetary policy*, ed. N. Gregory Mankiw. Chicago: University of Chicago Press.
- Henderson, Dale, and Warren McKibbin. 1993. An assessment of some basic monetary policy regime pairs. In *Evaluating policy regimes: New research in macroeconomics*, ed. R. Bryant. Washington, D.C.: Brookings Institution.
- Kerr, William, and Robert G. King. 1996. Limits on interest rate rules in the IS model. *Federal Reserve Bank of Richmond Economic Quarterly* 82 (spring): 47–76.
- King, Robert G., and Mark W. Watson. 1996. Money, prices, interest rates and the business cycle. *Review of Economics and Statistics* 78 (February): 35–53.
- King, Robert G., and Alexander L. Wolman. 1996. Inflation targeting in a St. Louis model of the 21st century. *Federal Reserve Bank of St. Louis Review* 78 (3): 83–107.
- Koenig, Evan F. 1987. A dynamic optimizing alternative to traditional IS-LM analysis. Discussion Paper no. 87-07. Seattle: University of Washington, Department of Economics, May.
- Lucas, Robert E., Jr. 1976. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy* 1:19–46.
- McCallum, Bennett T. 1997. Comment on Rotemberg and Woodford, “An optimization-based econometric framework for the analysis of monetary policy.” In *NBER macroeconomics annual 1997*, ed. Ben S. Bernanke and Julio J. Rotemberg, 355–59. Cambridge, Mass.: MIT Press.

- McCallum, Bennett T., and Edward Nelson. Forthcoming. An optimizing *IS-LM* specification for monetary policy and business cycle analysis. *Journal of Money, Credit and Banking* 31.
- Rotemberg, Julio J., and Michael Woodford. 1995. Dynamic general equilibrium models with imperfectly competitive product markets. In *Frontiers of business cycle research*, ed. Thomas Cooley. Princeton, N.J.: Princeton University Press.
- . 1997. An optimization-based econometric framework for the evaluation of monetary policy. In *NBER macroeconomics annual 1997*, ed. Ben S. Bernanke and Julio J. Rotemberg, 297–346. Cambridge, Mass.: MIT Press.
- . 1998. An optimization-based econometric framework for the evaluation of monetary policy: Expanded version. NBER Technical Working Paper no. 233. Cambridge, Mass.: National Bureau of Economic Research, May.
- Rudebusch, Glenn D. 1995. Federal Reserve interest rate targeting, rational expectations and the term structure. *Journal of Monetary Economics* 35:245–74.
- Sack, Brian. 1998. Does the Fed act gradually? A VAR analysis. Finance and Economics Discussion Paper no. 98-17. Washington, D.C.: Board of Governors of the Federal Reserve System, April.
- Shiller, Robert. 1997. Why do people dislike inflation? In *Reducing inflation: Motivation and Strategy*, ed. Christina D. Romer and David H. Romer. Chicago: University of Chicago Press.
- Summers, Lawrence. 1991. How should long term monetary policy be determined? *Journal of Money, Credit and Banking* 23 (August): 625–31.
- Taylor, John B. 1979. Estimation and control of an econometric model with rational expectations. *Econometrica* 47:1267–86.
- . 1993. Discretion versus policy rules in practice. *Carnegie-Rochester Conference Series on Public Policy* 39:195–214.
- Wicksell, Knut. 1907. The influence of the rate of interest on prices. *Economic Journal* 17:213–20.
- Woodford, Michael. 1996. Control of the public debt: A requirement for price stability? NBER Working Paper no. 5684. Cambridge, Mass.: National Bureau of Economic Research, July.
- Yun, Tack. 1996. Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of Monetary Economics* 37 (2): 345–70.

Comment Martin Feldstein

The paper by Julio Rotemberg and Michael Woodford is a complex and rich virtuoso performance. It is worth careful reading not only because of the difficult technical problem that they solve but also because of the extension that they propose to the basic Taylor rule, making the optimal interest rate a function of the past interest rate with a coefficient greater than one. Although I have reservations about specific aspects of the paper, I think it is an important contribution to the analysis of model-based monetary rules that will serve as a base for further useful developments.

Martin Feldstein is the George F. Baker Professor of Economics at Harvard University and president of the National Bureau of Economic Research.

Collections of Rules

Before commenting on the Rotemberg and Woodford paper itself, I will discuss the broader issue of the appropriate role of formal monetary policy rules. I think that such rules should not be viewed as substitutes for judgment by the monetary authorities but rather as inputs into that judgmental process. A good rule is therefore one that provides a useful starting point for central bank deliberations.

More specifically, I believe that a central bank can benefit from having a *collection* of alternative good rules, that is, rules that have optimal properties in a variety of models. I envisage central bank officials or staff using such a collection of rules each time a federal funds rate decision must be made. Before the decision is taken, the staff would calculate what the optimal federal funds rate would be according to each of the several different rules.¹ If the different rules all point to the same decision, the central bank's choice should be relatively easy. If they do not, the officials and staff have to dig deeper into the reasons for the differences and the authorities have to consider more carefully the decision that seems best in the current circumstances.

I emphasize the idea of using a collection of rules in this way because of the uncertainty inherent in the monetary policy process. There are two types of uncertainty for which this approach can be helpful: model uncertainty and situation uncertainty.

Ben McCallum has emphasized *model uncertainty* in his important writing on the choice of a monetary policy rule. McCallum's research imagines trying alternative rules in a variety of models and picking the single rule that does well under a variety of model assumptions. That approach may be too optimistic. There may be no rule that does uniformly better in a broad class of plausible economic models. Moreover, the rule that is optimal for the true structure of the economy may not do very well at all over a wide range of other possible models of the economy. McCallum's procedure may nevertheless help monetary authorities to develop and use a collection of rules by excluding some rules completely and by suggesting that certain rules deserve more weight in the central bank's thinking than others because of their robustness to different plausible models of the economy.

The second type of uncertainty, *situation uncertainty*, is the uncertainty about the current state of the economy and about where the economy would be going with no change in the federal funds rate. The expected values of uncertain estimates and of uncertain forecasts can be used in the decision rules only in very special circumstances. A prudent decision maker would therefore consider the optimal policy under different assumptions about the unobserv-

1. In the spirit of the Rotemberg and Woodford paper, I will discuss rules in terms of the federal funds rate rather than in terms of an optimal monetary aggregate, but the logic of what I say can obviously be transposed to models that indicate optimal monetary aggregates.

able state of the economy and about its future path. This would be done with a collection of rules since a single rule may not reflect the sensitivity of the optimal policy to the situation uncertainty.

All of this sounds like a lot of information for central bank decision makers to absorb. But it basically comes down to a list of the optimal federal funds value implied by each model in the collection, cross-classified by each set of alternative forecast scenarios (where forecasts include the current as well as future values of key economic magnitudes).

The Lagged Interest Rate and the Weakness of Monetary Policy

This brings me back to the Rotemberg and Woodford paper. An important feature of their work is that it extends the traditional Taylor rule by making the current optimal interest rate depend on the lagged interest rate as well as on inflation and output. In their analysis, the optimal response of the current optimal interest rate to the lagged value of the interest rate may have a coefficient greater than one because doing so significantly reduces the resulting average rate of inflation.

Although I find the analysis that leads Rotemberg and Woodford to this conclusion quite interesting, in the end I am not persuaded. Let me therefore review the logic behind their rule specification and then explain why I am not convinced. The starting point of the Rotemberg-Woodford analysis is the fact that the nominal interest rate that the central bank sets must be greater than or equal to zero. If the inflation rate is high, the mean value of this nominal interest rate will be high and the standard deviation of the fluctuations in that nominal rate can also be large. But if the inflation rate is low, the mean nominal interest rate will also be relatively low and the standard deviation of the interest rate must be small so that the actual nominal interest rate that the central bank sets is never required to be less than zero.

The Rotemberg-Woodford analysis shows that a monetary policy rule with a lagged value of the interest rate with a coefficient greater than one permits a low rate of inflation. This occurs because (in their rational expectations analysis) the public understands the rule and knows that when the Fed raises the interest rate it will go on raising the rate until the effect of this autoregressive process is dominated by a decline in actual inflation. This is sufficient to cause the actual inflation rate to decline before the rate is raised very much. They note that this low-inflation outcome comes with a high price in terms of increased variance of output.

I believe that their analysis overstates the extent to which low inflation requires volatile output. I think that a richer class of rules would also show that the autoregressive interest rate with the coefficient greater than one is also not necessary for low inflation.

A key limitation of the Rotemberg-Woodford analysis that they impose in order to carry out the complex rational expectations calculations is that the policy rule is linear. Thus the response of the interest rate must be the same to

low and high rates of inflation. This means that if a large interest rate decrease is not possible (because of the nonnegativity constraint), there cannot be a large interest rate increase to dampen inflation. The substitute for large interest rate increases is the lagged interest rate with a coefficient greater than one because, as I noted above, individuals understand the autoregressive rule and respond to small interest rate increases in a way that damps economic activity.

In actual practice (although not in their model) it would of course be possible to have an asymmetric (i.e., nonlinear) rule. The magnitude of the interest rate increase when inflation is regarded as too high could be much greater than the magnitude of the decrease when inflation is low and output is regarded as too low. This would make it unnecessary to have the lagged interest rate to signal the expected future tightening of interest rates when inflation is deemed to be too high.

The nonlinear rule means that high interest rates can be used to damp strong demand. But what about situations of weak demand? Rotemberg and Woodford are certainly correct that when inflation is low, the real interest rate cannot be reduced much because the nominal interest rate cannot be negative. But that need not mean that monetary policy is ineffective in dealing with weak demand. A central bank that cannot reduce the real interest rate can still increase demand through open market sales of the domestic currency for foreign currencies. Reducing the currency's value through such unsterilized intervention increases exports and reduces imports. The closed economy character of the Rotemberg-Woodford model precludes that, but it is a feature of the real world that needs to be taken into account.

Similarly, a sustained weakness of demand may be a reason for a fiscal stimulus—as in Japan today—which can be effective if the monetary authority keeps the real interest rate and the exchange rate from rising in response to the fiscal expansion. This too can prevent output declines even though interest rates cannot move down.

For these reasons, I remain to be convinced of the desirability of the lagged interest rate with a greater-than-one coefficient that Rotemberg and Woodford find in their optimal policy rule. I hasten to add that a lagged interest rate in the optimal response function may be a good idea for other reasons—uncertainty about economic conditions, the need to appear consistent, and the potential adverse effects of interest rate volatility on financial institutions—but that is a separate matter.

In a richer analysis, it may also be desirable to have asymmetric policy responses and unsterilized foreign exchange intervention. But those are matters that require further analysis.

Sticky Prices

The Rotemberg-Woodford analysis is both sparse and sophisticated. It is sparse in the sense that it boils down to two equations describing the behavior

of households and price-setting sellers. It is sophisticated in that it presents dynamic optimizing behavior of those agents with respect to a continuum of commodities in a rational expectations framework. But to make monetary policy effective in this framework and to create the observed stylized fact of a two-quarter lag in the impact of monetary policy, Rotemberg and Woodford introduce arbitrary lags into the behavior of both the households and the sellers.

I find this mixture of brilliant rational expectations optimizing behavior on the one hand and arbitrary lags on the other very disconcerting. Moreover, to the extent that the optimal policy rules are sensitive to the resulting lag pattern, as presumably they are since that is where the impact of monetary policy originates, how much weight should be put on rules that reflect arbitrarily imposed lag structures? What would happen if different assumptions were made about the lags in household and firm behavior?

Since the strength of the Rotemberg-Woodford model is its optimizing rational expectations framework, it would be interesting to drop the arbitrary lags and derive lags from the assumption that individuals and firms are following some kind of optimal Bayesian learning strategy. The resulting lags might produce a model that is too difficult for policy optimization of the type developed in section 2.3 of the paper, but such a change would give a more logically consistent basis for simulating alternative rules.

Since this paper already represents an enormous amount of complex work, it seems greedy of me to ask for more. The extension to Bayesian learning of the rational expectations model is certainly work for another paper.

Final Thoughts

But there is one exercise that could easily be done and that I think would be quite interesting. That would be to pick a few historical dates and calculate what the different rules described in this paper would mean for the optimal interest rate at those dates. At a minimum, it would be interesting to see how much the optimal interest rates differ across the different rules.

There is a final question that I would like to raise about the role of public confidence. I believe one of the reasons monetary policy has been effective in reducing inflation in the United States and in many other countries in recent years is that the Federal Reserve and other monetary authorities have stated the goal of reducing inflation and have, to a greater or lesser extent, specified numerical targets for future inflation.

To what extent would the effectiveness of monetary policy be enhanced in practice by using a rule or a collection of rules that emphasizes price stability? And to what extent is that effectiveness weakened by using an explicit monetary rule that attempts to optimize multiple criteria? Shifting from the existing informal emphasis on price stability to anything like a Taylor rule surely requires consideration of that issue.

Discussion Summary

In response to Martin Feldstein's point about the Lucas critique, *Michael Woodford* clarified that the assumption in the paper is not that agents do not act on the basis of knowledge of the monetary policy rule. The reason why the Lucas critique does not apply to the model is that optimal behavior is described by Euler equations that do not depend on the coefficients of the policy rule. These Euler equations are the first two equations in the model. This means that when changing the coefficients of the monetary policy rule, two of the three equations that characterize the equilibrium remain unchanged. *Glenn Rudebusch* noted that the Lucas critique still affects the results of the paper since the parameter estimates for the first two equations were obtained using a historical VAR reaction function. A different assumption about the historical reaction function would lead to different parameter estimates.

Lars Svensson asked whether the rules derived in the paper are only constrained optimal. *Woodford* responded that the rules are actually globally optimal assuming complete information about the underlying state variables, the real disturbances. The best rule in the family of rules that only respond to inflation and output observations is equivalent to the optimal rule. The intuition behind this result is that there are two real disturbances in the model and observing two endogenous variables gives enough information about these real disturbances.

Svensson also suggested an intuition for the response coefficient c on the lagged federal funds rate being above unity. Ideally, the central bank would like to affect long interest rates and thereby aggregate demand. The optimal way to do this is to have a large positive response coefficient on the lagged short rate. This will make future short rates react to the current short rate, which in turn will affect the long rate. Thus the rule works like a threat to do something in the future.

Donald Kohn remarked that the high coefficient on the lagged funds rate was interesting since central banks do not have a series of yield raises in mind when they reverse policy. They know they are too easy or too tight, but they have no idea by how much. Since they tend to think they are off by a fairly small amount initially, central banks often perceive bond markets as overreacting to their initial move. Looking at federal funds futures and forward rates, bond markets rarely call correctly the amount by which the Federal Reserve changes rates and almost always anticipate that the tightening or easing will persist beyond the point it actually stops and begins to be reversed.

Lawrence Christiano noted that a policymaker faced with a recommended rule might not follow the rule exactly but just a rule nearby. The graphs in the paper seem to suggest that if a rule with slightly different parameter values than the optimal rule of the model is picked, the result may be catastrophic. The problem arises in particular at the boundaries of indeterminacy and explosiveness regions. Also, the results may change dramatically when parameter

uncertainty is accounted for. *Woodford* responded that these features would have to be built into the model to provide answers to these questions.

Ben Friedman emphasized the importance of distinguishing between two kinds of rules: rules and rules of thumb. First, when evaluating *rules of thumb*—that is, some indicative guide used as input into the judgmental process—the whole class of arguments stemming from time inconsistency in support of the rule is unavailable. Second, if a model is solved on the presumption that the rule is a *rule*—that is, the dynamic equilibrium is solved based on the assumption that agents form expectations based on commitment—then the results cannot be used to justify a rule of thumb. *Ralph Bryant* noted that the evaluation of committed rules is still useful when thinking about rules of thumb. For example, the inference drawn about what variables to consider in a setting with committed rules can be carried over to rules of thumb. *Frederic Mishkin* remarked that the distinction is also important for the communication of policy advice to central banks. Central bankers often get upset when they hear the word “rules” because, given that the true model of the economy is not known, they want to have some discretion available to them.

Nicoletta Batini wondered how the rules in the paper compared to forward-looking rules. With the output predetermined at $t - 2$ and a coefficient on the lagged interest rate bigger than one, a feedback on current-dated variables may be de facto equivalent to a forward-looking rule. *Woodford* replied that in equilibrium the federal funds rate depends on expected future inflation. In this sense, the rule is similar to a forward-looking rule that depends directly on expected inflation.

Volker Wieland noted that, according to the paper, rules that are very effective in stabilizing inflation would require a fairly high degree of interest rate variability and thus frequently violate the zero-bound constraint on nominal interest rates—especially in a low-inflation environment. In recent work *Fuhrer* and *Madigan* (1997) incorporate the zero bound explicitly in a small macroeconomic model of the U.S. economy and use deterministic simulations to evaluate the effect of the zero bound in the event of a negative demand shock. They find only modest effects on output adjustment with a zero inflation target. However, *Orphanides* and *Wieland* (1998) find that in a stochastic framework with shocks similar to those in the 1980s and 1990s, the zero bound significantly increases the variability of output and inflation and introduces a long-run trade-off between inflation and output at inflation targets below 2 percent.

Bob Hall expressed three criticisms of the behavioral assumptions made by the paper. First, it is not clear why the public is devoted to a monetary unit. In fact, most intermediate product transactions are set in real terms, not in nominal terms. Second, firms change prices as frequently as they change quantities. Third, firms do not grant their customers call options. In wholesale, these options are not granted at all. In retail, while many call options are granted, these are with respect to limited quantities, and when firms run out, they reconsider their prices.

References

- Fuhrer, Jeffrey, and Brian Madigan. 1997. Monetary policy when interest rates are bounded at zero. *Review of Economics and Statistics* 79:573–85.
- Orphanides, Athanasios, and Volker Wieland. 1998. Price stability and monetary policy effectiveness when nominal interest rates are bounded at zero. Finance and Economics Discussion Paper no. 98-35. Washington, D.C.: Board of Governors of the Federal Reserve System.