

Optimal Monetary Policy in a Model with Habit Formation

Jeffrey C. Fuhrer*

Federal Reserve Bank of Boston

November 2000

Abstract

A number of recent papers have explored monetary policy options, including inflation targeting and inflation forecast targeting (notably Svensson (1999a, 1999b, 2000)) and price level targeting (Wolman 2000, Batini and Yates 1999, Blinder 1999). Most papers explore “optimal” monetary policy in the context of a single model. However, a number of conclusions made in the literature depend strongly on the model specification used. In addition, most papers have used the efficient policy frontier concept to define optimal monetary policy. This paper investigates the behavior of a variety of small structural macro models under a variety of targeting rules. The paper examines both minimum variance policy frontiers and utility-maximizing policy. In the latter case, an explicit model of consumer behavior with inflation-induced tax distortions is explored. The paper examines the improvement in utility from an optimal price-level target and re-examines the improvement in utility in moving from a positive to a zero target inflation rate.

(JEL D12, E52, E43)

*Preliminary and Incomplete
Not for Quotation*

* Vice President and Economist, Research Department, Federal Reserve Bank of Boston, Boston, MA, 02106. e-mail: Jeff.Fuhrer@bos.frb.org. I thank Richard Kopcke, Scott Schuh, and Geoffrey Tootell for their helpful comments, and Ann Ferris for excellent research assistance. Thanks to Dan Feenberg for providing tax rate data from TAXSIM. The views expressed in this paper are the author’s and do not necessarily reflect those of the Federal Reserve Bank of Boston or the Federal Reserve System.

“Nothing breeds failure like success.”

The success of monetary policy in controlling inflation in the developed economies over the past 20 years has led many to wonder how to retain that success, or even to improve upon it. Certainly it is reasonable to be concerned about the extent to which favorable monetary policy has arisen because of the idiosyncratic traits of particular central bankers. Going forward, it would be preferable to design monetary institutions so that successful practice would be incorporated in central bank charters, rather than embodied in individuals.

As a result, a large literature has emerged to propose optimal monetary policy strategies. The work of Svensson (1999a, 1999b, 2000) and Bernanke, Laubach, Mishkin, and Posen (1999) has focused debate on the primacy of inflation targeting as an organizing principle. For many years, the idea that central banks might directly target the price level was dismissed. The reasoning was roughly as follows: Because stabilizing the price level requires moving the inflation rate above and below its target, and because doing so requires corresponding movements in output (according to most models), price level targeting would entail unacceptable variability in output.

But more recent work has called into question the generality of this proposition. Wolman (2000), for example, shows that when the nominal interest rate hits its zero lower bound, a price level target may help to lower *ex ante* real interest rates, because it will imply a future rise in expected inflation. The more general proposition that price-level targeting might provide a “free lunch,” in which inflation and output variability are reduced without consideration of the zero lower bound, has been advanced by Svensson (1999b). This result, however, is model-dependent, and as I will show below, does not hold up in data-consistent macro models.

Most of the literature relies on an *ad hoc* monetary policy loss function that is a weighted average of variances of inflation (about its target), output (about

potential), and the price level (about its deterministic trend).¹ While Woodford (1999) suggests that the weighted sum of variances of output and inflation provide a fairly good second-order approximation to actual welfare, this approximation is not likely to hold up in all specifications. In particular, the time-inseparable utility function explored in this paper will alter the appropriate second-order approximation.² In any event, this aversion to computing welfare directly from the posited utility function seems somewhat odd, given the long-standing emphasis on optimization-based macro models.

In this paper, I review some variance-based optimal policy results. I demonstrate the model dependence of some key results, particularly those relating to price level targeting. Next, I utilize a data-consistent optimizing model of consumption developed in Fuhrer (2000), augmented by inflation-induced tax distortions, to compute utility-maximizing monetary policy. I examine the robustness of my findings to a number of changes in specification. I then compare the results to the results from the variance-based loss function. Finally, I look at the utility gains to price-level targeting over inflation targeting, and to a zero inflation target versus a low, positive inflation target.

Section ?? describes the models for which I compute optimal variance trade-offs; section ?? presents the minimum weighted variance results; section ?? develops the habit-formation model with inflation-induced tax distortions; section ?? presents utility-maximizing monetary policy results for the baseline model and alternative specifications; and section ?? concludes.

¹One exception is Rotemberg and Woodford (1997), although they consider only real interest rate shocks. In addition, Woodford (1999) demonstrates an approximate equivalence between weighted variances and model-based social welfare for a particular specification.

²In the habit specification used below, for example, the variance of the growth rate of consumption may appear in the second-order approximation.

1 The Models

1.1 Fuhrer-Moore model

The model is described in detail in Fuhrer and Moore (1995b). Here I sketch its essential components. The price specification may be summarized in the definition of the price index, p_t , as a moving average of current and past nominal contract prices x_{t-i} :

$$p_t = \sum_{i=0}^3 f_i x_{t-i}$$

where the weights f_i are as described in Fuhrer and Moore. The real contract price index v_t is defined as the weighted average of current and past real contract prices, $x_{t-i} - p_{t-i}$

$$v_t = \sum_{i=0}^3 f_i (x_{t-i} - p_{t-i})$$

The weights follow the same distribution as that for the definition of the price index. Price-setters determine the current real contract price as a function of the real contract prices expected to be in effect for the duration of the contract, and as a function of expected demand conditions, \tilde{y}_{t+i} , over the duration of the contract

$$x_t - p_t = \sum_{i=0}^3 f_i E_t(v_{t+i} + \gamma \tilde{y}_{t+i})$$

The aggregate demand relation makes the output gap a function of its own lags and the *ex ante* real interest rate ρ_t

$$\tilde{y}_t = y_1 \tilde{y}_{t-1} + y_2 \tilde{y}_{t-2} + y_\rho (\rho_{t-1} - \bar{\rho}) + \epsilon_{yt}$$

where $\bar{\rho}$ is the equilibrium real interest rate, and ρ_t is defined according to the pure expectations hypothesis as a weighted average of future short-term real interest rates, $r_{t+i} - E_t \pi_{t+i+1}$

$$\rho_t = d \rho_{t+1} + (1 - d)(r_t - E_t \pi_{t+1})$$

where d determines the weights on future short-term real interest rates. The model is closed with a monetary policy rule with the short-term interest rate r_t as the policy instrument. Alternative rules are discussed below.

1.2 Rudebusch-Svensson model

I follow Rudebusch and Svensson (1998), using their mildly restricted VAR equations for “I–S” and Phillips curves. The I–S curve makes the output gap a function of two of its own lags and a weighted average of recent *ex post* short-term real interest rates. The Phillips curve is conventional in making inflation a function of its own lags (constrained to sum to one in this model) and one lag of the output gap.

$$\begin{aligned}\tilde{y}_t &= b_1\tilde{y}_{t-1} + b_2\tilde{y}_{t-2} - b_\rho\left(.25\sum_{i=0}^3(r_{t-i} - \pi_{t-i}) - \bar{\rho}\right) + \epsilon_{yt} \\ \pi_t &= \sum_{i=1}^4 c_i\pi_{t-i} + c_y\tilde{y}_{t-1} + \epsilon_{pt}\end{aligned}$$

The coefficients for the two equations are estimated from quarterly data as described in the data appendix. The model is closed with a policy rule in the short-term interest rate, as described below.

1.3 Habit Formation model

This model is fully described in Fuhrer (2000). The model includes a nondurables and services consumption sector that is characterized by an approximate log-linear consumption function derived from a utility function that exhibits habit formation. The utility function is

$$U = \sum_{i=0}^{\infty} \beta^i \frac{1}{(1-\sigma)} \left(\frac{C_i}{C_{i-1}^\gamma} \right)^{(1-\sigma)}$$

and the log-linearized consumption function is

$$c_t - y_t = E_t \sum_{j=1}^{\infty} \mu^j (\Delta y_{t+j} + a_1^* (p_{t+j+1} - p_{t+j}) + a_2^* (z_{t+j+1} - z_{t+j}) - \delta^* r_{t+j+1} + \kappa^*)$$

where p_t and z_t are as defined in Fuhrer (2000). The rest of GDP—durable goods consumption plus business investment, net exports, and government consumption and expenditures—is modeled as in the Fuhrer-Moore (1995b) model. The price specification is from Fuhrer and Moore (1995a), and the model is closed with an interest rate reaction function.

1.4 McCallum-Nelson model

This simple model, dubbed an “optimizing IS–LM model” for monetary policy analysis by McCallum and Nelson (1999), comprises a forward-looking I–S curve of the form

$$\tilde{y}_t = E_t \tilde{y}_{t+1} - \sigma (r_t - E_t \pi_{t+1} - \bar{r}) + \epsilon_{yt}$$

and a forward-looking Phillips curve

$$\pi_t = E_t \pi_{t+1} + \gamma \tilde{y}_t + \epsilon_{pt}$$

The properties of this model are discussed in Estrella and Fuhrer (1999). Because it is purely forward-looking, the model imparts very little persistence to output or inflation. As a result, it will imply markedly different variance trade-offs from the other models considered. It is included, however, because monetary policy regimes that target the deviation of the price level about a trend may not yield the inflation persistence that characterizes post-war inflation-targeting monetary regimes. The model is closed with an interest rate reaction function.

2 Minimum Variance Frontiers

The convention in most recent contributions to the optimal monetary policy literature has been to assume a loss function that sums the weighted variances of

key variables.³ I will discuss an alternative below, but here I adopt the minimum variance convention. A number of alternative loss functions and policy strategies are considered.

2.1 Alternative Loss Functions and Policy Strategies

Loss Functions

A fairly standard period loss function takes the weighted average of output gap, inflation, and interest rate variances:

$$\mathcal{L}_t = \mu_1 V(\pi_t) + \mu_2 V(\tilde{y}_t) + (1 - \mu_1 - \mu_2) V(r_t)$$

Then the loss function is the discounted sum of expected weighted variances. In what follows, I use the unconditional variance for the same set of variables—the limit of \mathcal{L}_t as $t \rightarrow \infty$. While the computations are simplified somewhat by this assumption, this simplification makes little or no difference in the results I discuss below.

A second loss function that I consider adds to the above loss function the unconditional variance of the log price level around its (possibly upward-sloping) trend. Of course, this loss function only makes sense if the policy strategy includes some response to price level deviations. If not, the deviation of the price around its trend follows a random walk, and its variance grows with t . Thus the loss function in this case is:

$$\mathcal{L}_t = \mu_1 V(\pi_t) + \mu_2 V(\tilde{t}_t) + \mu_3 V(r_t) + (1 - \mu_1 - \mu_2 - \mu_3) V(p_t - p_t^*)$$

where the $V()$ s may be interpreted as unconditional variances.

In the exercises described below, I assume that the loss function includes some weight on the price level deviation, unless the monetary authority does not respond

³In a recent paper, Woodford (1999) shows that such a loss function may be derived from a particular optimizing problem.

explicitly to the price level (as with pure inflation targeting or inflation forecast targeting).

Policy Strategies

For a given loss function, the monetary authority may pursue one of many proposed “policy strategies,” which here simply mean different manifestations of a linear instrument rule. For example, as Svensson has pointed out, an inflation forecast targeting rule can be an optimal rule even when the loss function incorporates concern about both inflation and output deviations. Alternatively, monetary policymakers might choose a backward-looking reaction function that responds to all state variables. The actual performance of these strategies, given a loss function, will depend on details of the assumed model specification and the characterization of shocks hitting the economy.

Each of the policy rules considered is a somewhat restricted version of the “optimal” rule that would arise. In each case, lags of output, inflation, and the price deviation would appear in the optimal rule. However, the presence of a lagged interest rate term implies (infinite) lags of each of these variables in the policy rule. But the lags appear in restricted form. Other authors have shown that the empirical distinction between “simple” rules like those below and complete state variable rules is small (see Tetlow and von zur Muehlen 1999, and Williams 1999, for example).

1. The Kitchen Sink

$$r_t = \rho r_{t-1} + (1 - \rho)(\pi + \alpha_\pi(\pi_t - \bar{\pi}) + \alpha_y \tilde{y}_t + \alpha_p(p_t - p_t^*) + \bar{\rho}) + \epsilon_{rt}$$

2. Inflation Targeting

$$r_t = \rho r_{t-1} + (1 - \rho)(\pi + \alpha_\pi(\pi_t - \bar{\pi}) + \bar{\rho}) + \epsilon_{rt}$$

3. Inflation Forecast Targeting

$$r_t = \rho r_{t-1} + (1 - \rho) \left[\alpha_\pi \mathcal{F} \sum_{i=0}^3 (\pi_{t+i} - \bar{\pi}) + \alpha_\pi (1 - \mathcal{F}) \sum_{i=1}^4 (\pi_{t-i} - \bar{\pi}) + \pi_t + \bar{\rho} \right] + \epsilon_{rt}$$

This strategy nests simple backward-looking inflation targeting via the parameter \mathcal{F} . Policymakers choose the weight on the inflation forecast versus the past, as well as the weight on inflation generally. With $\mathcal{F} = 1$, monetary policy responds entirely to the forecast of inflation over the next year.

4. Price Level Targeting

$$r_t = \rho r_{t-1} + (1 - \rho)(\alpha_p(p_t - p_t^*) + \pi_t + \bar{\rho}) + \epsilon_{rt}$$

This rule makes the short-term rate a function of deviations of the price level from its long-term trend, p_t^* . The trend may have a positive slope, so that the trend grows over time at the target rate of inflation $\bar{\pi}$

$$p_t^* = \bar{\pi} + p_{t-1}^*$$

Until recently, conventional wisdom held that price level targeting would entail an undesirable amount of inflation and output variability. Inflation variability would increase because inflation would have to be brought below target to return the price level from above trend. Because most models link inflation and real output, an increase in inflation variability would entail a commensurate increase in output variability. However, Wolman (2000) shows that, for some specifications, price-level targeting can reduce inflation and output variability.

2.2 Results I: Price Level Targeting Is Not a Free Lunch (for most specifications)

Figure 1 displays output/inflation variance frontiers for the habit formation model with a “Kitchen Sink” policy strategy.⁴ Each frontier depicts the outcomes for a

⁴The PDF version of the paper, available on the FRB Boston Research web site, displays color versions of the figures.

different value of α_p . The black line depicts the frontier for $\alpha_p = 0$, and increasing values of α_p imply frontiers that shift up and to the right (the colored lines). As the figure indicates, increasing emphasis on the price level deviation from its (upward-sloping) trend entails increasing output and inflation variability. There is no free lunch.⁵ This result also holds for the Fuhrer-Moore model and the Rudebusch-Svensson model.

Figure 2 displays the corresponding output/*price level* variance frontiers for the same model. Here we see a feature common to these three models as well: Increased emphasis on the price level deviation (not surprisingly) reduces price level variability (about trend). For a given level of price level emphasis, increasing inflation emphasis eventually raises *both* price level and output variability. This is true for all values of α_p in the figure. If price-level deviations about trend are an important component of the monetary policy loss function, too much emphasis on inflation will backfire in these models.

The McCallum-Nelson model, with flexible inflation and output, produces a different result. First of all, the feasible combinations of inflation and output variance lie well inside those attainable for models with sticky inflation and output (the minimum inflation standard deviation of nearly zero is attained at an output standard deviation of 3.5; in the sticky inflation model, the minimum inflation standard deviation is above 2 and arises for output standard deviations above 5). Practically speaking, increasing emphasis on the price level deviation causes no deterioration in output and inflation variability, as Figure 3 shows. In fact, for extremely inflation-averse preferences (the upper left portion of the frontier), additional emphasis on the price level slightly *improves* output and inflation variability.⁶ In addition, the price level/output variability trade-off does not bend back-

⁵Note that this exercise abstracts from the possible benefits to price-level emphasis when the nominal interest rate is at the zero lower bound.

⁶Macleon and Pioro (2000) find similar results in the Bank of Canada's QPM models, by varying the expectations assumptions employed between backward-looking and fully forward-looking.

wards in this model, as it does in the others. I find this result in none of the other specifications.⁷

2.3 Results II: Inflation Targeting versus Inflation Forecast Targeting

In a series of papers, Svensson (1998, 1999, 2000) has raised some critical issues regarding the theoretical differences between inflation targeting and inflation forecast targeting. Figures 4 and 5 display the benefits that accrue to an inflation-forecast-targeting central bank, compared to one that targets observed current and past inflation. Recall that the inflation forecast targeting strategy described above allowed for endogenous estimation of the mix between backward and forward weights on inflation in the reaction function.

Interestingly, this exercise yields two insights about forecast targeting. First, for moderate values of preferences, the weight on expected future inflation is generally one, while that on past inflation is zero. Second, the gains (in reduced inflation/output variability) from pursuing an inflation forecast strategy are fairly small. The red dashed lines in Figures 4 and 5 show the feasible variance combinations under simple inflation targeting; the solid lines depict the variances attainable under inflation forecast targeting. Each pair of lines represents a different emphasis on the price level deviation, with higher emphasis resulting in an outward shift in the frontier—i.e., higher inflation and output variances. As the figures indicate, the gains from inflation forecast targeting are small. And in some sense, the indicated gains are optimistic: In this exercise, agents know the “true” model and use it to forecast inflation. If the model used to forecast inflation is misspecified, then these gains may be even smaller.

The reason for the similarity of inflation targeting and inflation forecasting

⁷Estrella and Fuhrer (2000) document the strongly counterfactual implications of this and other specifications that make inflation and output purely forward-looking.

results is straightforward. The rational expectations forecast of inflation in the model is a particular linear combination of the observable variables in the model. Because the two models in Figures 4 and 5 embody data-consistent inflation persistence, the actual inflation rate today will be highly correlated with the rational forecast of inflation over the next four quarters. As a result, targeting today's observed inflation, rather than forecasted inflation, does not yield significantly worse performance.⁸

2.4 Results III: Pure Price Level Targeting

While few are currently proposing such an extreme policy, I consider pure price level targeting. Figure 6 compares a policy that responds only to the price level deviation—which will, of course, control inflation—with a “kitchen sink” policy. The results for the habit formation model (echoed in the F-M and R-S models) are clear: Price level targeting (the solid black lines) raises both inflation and output variability relative to the “kitchen sink” policy (the dashed red lines).

3 Explicit Inflation Distortions and Utility-Maximizing Monetary Policy

Why do economic agents care about inflation? In the models examined above, the only reason for agents to care about inflation is that they know that the monetary authority cares. As a result, the monetary authority will move output above and below potential in order to keep inflation near its target. The resulting real disruptions might well figure into consumers' expected utility calculations, for example. But in this case, an optimal monetary policy from the consumer's point of view is one that smoothes consumption as much as possible. We don't know why you

⁸Of course, these results do not address other potentially important benefits to explicit inflation forecast targeting.

feel the need to disinflate, the consumer might say to the monetary authority, but if you do it, do it gradually.

This section explores a very simple model in which the level of inflation matters explicitly for real consumption decisions. I eschew the convention of putting real money balances in the utility function, or imposing a cash-in-advance constraint. Instead, I focus on the distortions that arise from imperfect indexation. I pursue this avenue because if all transactions in the economy—wage determination, asset exchange, taxation—were perfectly indexed, much if not all of the cost of inflation would disappear.

In particular, I approximate the distortions that arise through the interaction of inflation with the U.S. tax code. This source of inflation cost has been emphasized recently by Feldstein (1997). In the early part of my sample, this included both “bracket creep”—the failure of tax rate brackets to keep pace with inflation—and the taxation of nominal asset returns. In the latter part of the sample, the labor income tax brackets are fairly well indexed (although with a lag), but the taxation of nominal asset income remains.

Thus, the distortion I consider enters through the budget constraint, rather than through the utility function. The log-linearized consumption function that I derive includes a term that implies that the log consumption-income ratio will be lower, the higher is the discounted expected sum of future inflation rates. The parameters that link expected inflation to real consumption are derived from effective marginal tax rates and their relationship to inflation across time. However, for a suitable change in parameters, the consumption function may be thought of as representing any inflation distortion (in the level of inflation) that enters through the budget constraint of a forward-looking consumer.

3.1 Derivation of the Model

Following Campbell and Mankiw's (1989) derivation of the log-linear budget constraint in the standard life-cycle model, assume that aggregate nominal wealth may be represented as⁹

$$(1) \quad W_t = P_t + Y_t$$

The rate at which nominal returns on wealth are taxed is τ . The effective marginal tax rate varies linearly with the rate of inflation

$$\tau(\pi_t) = \tau_0 + \tau_1 \pi_t$$

so the after-tax return on wealth is

$$(2) \quad R_{t+1} = (1 - \tau(\pi_{t+1}))(P_{t+1} + Y_{t+1})/P_t$$

Combining equations ?? and ?? yields a nonlinear difference equation in wealth

$$(3) \quad W_{t+1} = \frac{R_{t+1}(W_t - Y_t)}{(1 - \tau(\pi_{t+1}))}$$

Dividing through by W_t and log-linearizing yields

$$w_{t+1} - w_t = r_{t+1} + \log(1 - \exp(y_t - w_t)) - \log(1 - \tau_0 - \tau_1 \pi_{t+1})$$

where lowercase letters denote logarithms. Employing Campbell and Mankiw's linearization for $1 - \exp(y_t - w_t)$, and a log-linearization for the tax terms around $\pi = 0$ yields

$$w_{t+1} - w_t \approx r_{t+1} + \kappa^* + (1 - 1/\rho)(y_t - w_t) + (\tau_1/(1 - \tau_0))\pi_{t+1}$$

We can use the identity $\Delta w_{t+1} \equiv \Delta y_{t+1} + (y_t - w_t) - (y_{t+1} - w_{t+1})$ to substitute out for Δw_{t+1} to obtain

$$\frac{\Delta y_{t+1} + (y_t - w_t) - (y_{t+1} - w_{t+1})}{1} = r_{t+1} + \kappa^* + (1 - 1/\rho)(y_t - w_t) + (\tau_1/(1 - \tau_0))\pi_{t+1}$$

⁹As in Campbell and Mankiw, W_t represents all wealth, including human capital. I have normalized population at 1 for convenience.

Solved forward and simplified, this yields

$$(4) \quad y_t - w_t = E_t \sum_{j=1}^{\infty} \rho^j (r_{t+j} - \Delta y_{t+j} + \frac{\tau_1}{(1 - \tau_0)} \pi_{t+j}) + \kappa^*$$

Substituting equation ?? into Campbell and Mankiw's equation 3.6, we obtain the budget constraint

$$(5) \quad c_t - y_t = E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j} - \Delta c_{t+j} - \frac{\tau_1}{(1 - \tau_0)} \pi_{t+j}) + \kappa^*$$

I combine this budget constraint, which incorporates the cost of inflation-induced tax distortions, with the expression for Δc_{t+j} from the habit formation specification in Fuhrer (2000) to obtain the log-linear consumption function

$$(6) \quad c_t - y_t = E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j} + a_1^* (p_{t+j+1} - p_{t+j}) + a_2^* (z_{t+j+1} - z_{t+j}) - \delta^* r_{t+j+1} - \frac{\tau_1}{(1 - \tau_0)} \pi_{t+j}) + \kappa^* + \epsilon_{ct}$$

where the a_i^* , δ^* , p , and z are as defined in Fuhrer (2000). The interpretation of the consumption function is as in Fuhrer (2000), with the addition of the tax feature. When inflation is expected to be high in the future, the tax distortion is larger, real disposable income is lower, and consumption is lower relative to current income. The amount of the distortion depends on the slope and intercept of the function that relates taxes to inflation. We provide some data-based estimates of these parameters in the next section.

An alternative approximation for imperfect indexation makes the tax rate a function of both the level and the change in inflation.

$$\tau_t = \tau_0 + \tau_1 \pi_t + \tau_2 (\pi_t - \pi_{t-1})$$

This specification leads to the approximate log-linearized consumption function

$$(7) \quad c_t - y_t = E_t \sum_{j=1}^{\infty} \rho^j (\Delta y_{t+j} + a_1^* (p_{t+j+1} - p_{t+j}) + a_2^* (z_{t+j+1} - z_{t+j}) - \delta^* r_{t+j+1} - \frac{\tau_1}{(1 - \tau_0)} \pi_{t+j}) - \frac{\tau_2}{(1 - \tau_0)} (\pi_{t+j+1} - \pi_{t+j}) + \kappa^*$$

so that now the consumption-income ratio drops when the expected *change* in inflation is positive, reflecting its distorting effect on the tax rate.

3.2 Estimating the Dependence of Tax Rates on Inflation

I use two methods to estimate the dependence of effective marginal tax rates on inflation. The first is to indirectly estimate the effect by estimating the habit formation model with tax distortions in equation ???. Maximum likelihood estimation of τ_0 and τ_1 yields a tax distortion parameter $\tau_1/(1 - \tau_0) = 0.45$. If we allow for the future change in inflation to enter as well, the estimated coefficients on the future levels and changes of inflation are 0.24 and 4.09, respectively, with standard errors of 0.11 and 0.63.

Alternatively, I use data on average marginal tax rates for wage, interest, dividend, and pension income from the NBER's TAXSIM program to estimate the effect of inflation on effective marginal tax rates. To correspond to the concept of income used in the consumption sector, I aggregate the tax rates using the shares of these types of income in total income in the NIPA accounts. Figure 7 shows the relationship between the average aggregate marginal tax rate and inflation from the mid-1960s to 1994, the last year for which tax data are available.

Of course, effective marginal tax rates will vary with important changes in the tax code. In estimating the parameters τ_0 and τ_1 , I control for the changes in the U.S. personal tax code shown in Table 1.

Using dummies to control for these breaks in the tax code, I regress effective marginal tax rates on inflation, the first difference of inflation, real output growth, and the output gap. I find that the change in inflation enters insignificantly. The estimated ratio of τ_1 to $1 - \tau_0$ is 0.69, with asymptotic standard error 0.14.

3.3 The Government Budget Constraint

In this simple model, government expenditures are simply a component of "other GDP" (GDP excluding consumption of nondurable goods and services). Other

Table 1
Significant changes in US tax code

Year	Change
1969	<p><i>Tax Reform Act of 1969</i></p> <p>Increased personal exemption amount and lowered max rate on earned income. Increased standard deduction, established new rate schedule for single taxpayers, and repealed investment tax credit.</p>
1978	<p><i>Revenue Act of 1978</i></p> <p>Reduced individual taxes by widening tax brackets and reducing number of tax rates. Increased personal exemption, standard deduction, and capital gains exclusion.</p>
1981	<p><i>Economic Recovery Tax Act of 1981</i></p> <p>Significantly reduced taxes by reducing individual tax rates and indexing of bracket structure (beginning in 1985). Reduced top tax rate on capital gains, introduced a deduction for two-earner married couples, and expanded IRAs.</p>
1986	<p><i>Tax Reform Act of 1986</i></p> <p>Complete reform of tax code. Made major cuts in individual tax rates and repealed capital gains exclusion. Repealed investment tax credit, increased personal exemption amount and standard deduction.</p>
1990	<p><i>Omnibus Budget Reconciliation Act of 1990</i></p> <p>Increased excise and social insurance taxes. Created new 31% tax rate and max capital gains rate at 28%. Temporarily phased out personal exemptions.</p>
1993	<p><i>Omnibus Budget Reconciliation Act of 1993</i></p> <p>Increased tax rates for high-income taxpayers, with new brackets at 36 and 39.6%. Limited itemized deductions and extended phase-out of personal exemptions. Increased excise and social insurance taxes.</p>

GDP is modeled as a stationary process around a log-linear segmented trend.

$$\tilde{y}_t^o = d_1 \tilde{y}_{t-1}^o + d_2 \tilde{y}_{t-2}^o - d_\rho (\rho_{t-1} - \bar{\rho}) + \epsilon_{yot}$$

The process for the tax rate will be mean-reverting to τ_0 as long as inflation is stationary. With all of GDP reverting to its long-run trend, tax revenues will grow at the same rate as government expenditures. As a result, the long-run balance of the government sector is guaranteed.

3.4 Utility Maximization

For a given set of policy parameters and sequence of shocks, the model with habit formation and tax distortions implies a particular path of consumption $c_i, i = 0 \dots \infty$. I compute utility from the original nonlinear utility function as

$$U = \sum_{i=0}^{\infty} \beta^i \frac{1}{(1-\sigma)} \left(\frac{c_i}{c_{i-1}^\gamma} \right)^{(1-\sigma)}$$

In the figures that follow, I present “utility surfaces.” These surfaces plot utility as defined above as a function of policy parameter settings. For each setting of the policy parameters, I compute a stochastic simulation of the model, shocking it with shocks drawn from a multivariate normal distribution whose covariance matrix equals the covariance matrix estimated from the data. Discounted utility is computed from the simulated values for consumption. In order to avoid any small-sample distortions, I average the results over 1,000 simulations of length 100, which should provide a very accurate estimate of the true distribution implied by the model and error covariance matrix.

4 Utility-Maximizing Monetary Policy: The Baseline Case

The model is specified as described above, with an error covariance matrix computed for the sample from 1980 to the present. Only inflation-level distortions enter the model. The distortion parameters τ_1, τ_0 are set at their estimates from the marginal tax rate regressions, so that $\tau_1/(1 - \tau_0) = 0.69$.

The utility surfaces for the baseline case, which uses the sticky inflation specification with a “smaller” covariance matrix (excluding the 1970s supply shocks) and levels effects on tax rates, is shown in Figure 8. The figure displays variation in α_π and α_y on the x and y axes, with utility on the vertical axis. The surfaces differ in the value of the “smoothing” parameter ρ and in the value of the price-level targeting parameter α_p .

As indicated in the graph, utility is maximized for high values of α_y and low values of α_π . Utility rises rapidly as the response to inflation, α_π , falls below 2. For these parameters, consumers strongly prefer output to be smoothed. Interest rate smoothing that is more aggressive than the sample estimate of about 0.8 does not improve utility.

Interestingly, the “global” optimum in this grid of parameters arises when the monetary authority places significant weight on the deviation of the price level from the targeted trend path for prices. The utility gain from targeting the price level is relatively small—compared, for example, to the gain from not responding too aggressively to inflation. In this baseline case, the utility level for the optimal policy with price level targeting is $-.4316$, while the utility attained setting α_p to zero is $-.4511$, a difference of about $.02$. This compares to utility levels for moderate versus aggressive inflation targeting of $-.56$ and -2.38 , respectively, a difference of 1.82 .

The rationale for this result is simple: Recall that the form that the imperfect

indexation distortion (call it D) takes in the consumption function is:

$$D = E_t \sum_{j=1}^{\infty} \rho^j \left[-\frac{\tau_1}{(1 - \tau_0)} \pi_{t+j} \right]$$

The consumption–income ratio falls with the discounted sum of expected future inflation. But this term is very close to the expected future price level at a long horizon:

$$p_{\infty} = p_t + E_t \sum_{j=1}^{\infty} \pi_{t+j}$$

Minimizing the distortion that arises from imperfect indexation is nearly equivalent to keeping the price level stable in this model. The closer is the expected (discounted) sum (or average) of future inflation rates to zero, the closer will be the expected price level to the current level.

How robust is this result? The following section explores whether it can be overturned by changing the strength of the inflation distortion, the size and composition of the shocks buffeting the economy, or the stickiness of inflation in the model.

4.1 Robustness to Alternative Specifications

I consider a number of alternatives, which may be categorized as (1) specification options, and (2) error covariance options.

Specification Options

- **Larger inflation distortions:** I quadruple the size of the tax distortion parameter, setting the ratio of the tax parameters to 3.0.
- **Flexible inflation** I substitute a simple flexible inflation equation

$$\pi_t = \beta E_t \pi_{t+1} + \gamma y_t + \epsilon_{pt}$$

as in McCallum and Nelson (1999).

- **Myopic Consumers** I raise the discount rate used to discount future income, real rates, and inflation.

Error Covariance Options

1. **Larger shocks** I use an error covariance matrix estimated from the sample 1966-1979. By incorporating the 1970s oil shocks, this significantly increases the variance of the “supply shocks” buffeting the economy.

The utility surfaces for these options are displayed in Figures 9–12. As the figures show, the basic result, that consumers care more about monetary policy responses to output than inflation, holds for all of these alternatives. In addition, in all cases except the flexible inflation model, price level targeting dominates policies that ignore the price level, although the utility gains are relatively small. Larger shocks and a bigger value for τ_1/τ_0 simply magnify the penalties for extreme emphasis on inflation.

Figure 10 displays the utility surfaces for the flexible inflation model. It shares some features with the sticky inflation cases: Utility falls rapidly as α_π increases, and consumers prefer a fairly vigorous response to output. But in contrast to the sticky inflation results, the optimal policy here entails *no* response to price level deviations. Inflation is so flexible that a policy that targets the cumulative inflation rate rate (the price level) does no better than one that targets the current inflation rate. Inflation is also flexible enough that minimal responses to inflation yield stable current and future inflation with minimal tax distortions and smooth consumption.

Figure 11 displays utility surfaces for shocks drawn from a covariance matrix that includes data from the 1960s and 1970s. The only difference between this figure and the baseline is that the range of utility is larger. Poorly conducted monetary policies imply more dramatic losses of utility when the shock variances are larger.

4.2 Myopic Consumers with Sticky Inflation

Of course, a substantial increase in the rate at which consumers discount the future might reduce the importance of price-level targeting. In this case, the correspondence between the inflation distortion and the price level is weakened. In Figure 12, the discount rate ρ is set to 0.5, so that consumers' effective horizon is only a few quarters. Now the optimal policy entails no weight on the price level. The utility surfaces still suggest, however, that too much single-minded emphasis on inflation yields a tremendous loss in utility.

4.3 The Optimal Inflation Target

While the behavioral equations that determine the evolution of consumption, income, inflation, and interest rates in the model are linear, the utility function by which we gauge alternative policies is not. Thus different target inflation rates will imply different utilities. A key question is how much utility is affected by a shift from a low inflation target to a zero inflation target.

Figure 13 displays representative utility surfaces for the baseline model with a 2% inflation target and the baseline model with a zero inflation target. As is evident from the figure, the utility difference is negligible. The maximum utility attainable with the 2% target is -0.4316 , versus -0.4307 for the zero-inflation target case. Compare this to the difference in utility between the optimal policy under zero-inflation targeting and the sub-optimal vigorous inflation targeting policy. Utility declines from -0.4307 to -2.3593 when the monetary authority responds vigorously to inflation.

4.4 Comparing Weighted Variance and Utility Measures of Optimality

Figure 14 shows the optimal policy frontiers for the baseline model with tax distortions, for a variety of values of α_p . The contours are quite similar to those shown in Figure 1 for the model without the tax distortion, and the basic message is the same: Price level targeting does not deliver a “free lunch” with regard to output and inflation variability.

Figure 15 depicts the same variance frontiers along with the variances implied by the utility-maximizing points on Figures 8–11.¹⁰ These points are labeled with red stars and are taken from the “Baseline,” “Large τ ” (more prominent inflation distortion), “Large Ω ” (drawing shocks from the volatile 1970s), and “Flexible inflation” simulations, respectively.

The Baseline case lies somewhat outside the variance frontiers, even for very high values of α_p . This suggests that, from an optimal variance perspective, the utility-maximizing point implies a strong preference for low output variability, and a high tolerance for inflation variability. The large τ point lies on the frontiers and is naturally consistent with a lower tolerance for inflation variability, but it still lies below the -1 slope point identified by Taylor as a point reflecting a moderate trade-off between inflation and output variability. The flexible inflation point shows that utility-maximizing policy implies considerably smaller output and inflation variance when inflation is flexible than when it is sticky. Finally, the “large Ω ” point shows how much further out in variance space the utility-maximizing policy is if the economy is buffeted by larger shocks.

¹⁰These are computed as in the unconditional variance plots, taking the diagonal covariance matrix Ω assumed in the corresponding figure, the values of τ_1 and τ_0 assumed, and the flexible or sticky inflation specifications described above.

4.5 A Note on Finite-Sample Results

The final figure displays utility surfaces for two draws of length 220 quarters—about the length of the postwar sample. The shocks drawn from the random number generator are approximately normal and have the two sample diagonal covariance matrices, with the “true” covariance matrix displayed at right:

Shock	Small Sample 1	Small Sample 2	“True”
ϵ_p	5.1e-06	5.1e-06	5.2e-06
ϵ_{yo}	6.0e-05	8.1e-05	7.3e-05
ϵ_c	1.8e-04	1.9e-04	1.9e-04

Interestingly, these very similar covariance matrices, both of which approximate the variance of the “true” process quite well, produce very different implications for the nature of optimal monetary policy. One suggests that optimal policy places no emphasis on the price level with very little interest rate smoothing; the other implies that moderate price level targeting is important, and so is a sample-consistent degree of interest rate smoothing. Both agree that less attention to inflation and more to output improves utility.

The reason for this disparity (which we avoid by averaging over the large number of draws in the simulations above) is that the precise timing of shocks to the model matters. Because consumption, inflation, and interest rates are persistent, the timing and confluence of larger shocks can strongly influence the subsequent time path of these variables for a period as long as a decade. Equivalently, a 50-year sample really has only five ten-year episodes that constitute responses to significant shocks to prices or output. Even though the variances of the sample’s shocks will equal the large-sample “true” variance, the time path of the variables can be significantly different. As a result, the optimal monetary policy for two economies with identical structure and essentially identical shock variances can vary quite significantly, depending on the exact timing of large shocks.

5 Conclusions

The principal conclusions of this paper may be summarized as follows:

1. The Svensson proposition that price-level targeting may deliver a “free lunch”—lower price and inflation variability with no increase in output variability—is model specific. It depends on a degree of forward-looking behavior that Estrella and I have argued elsewhere is not data-consistent.
2. When gauged from a utility metric, price-level targeting delivers a very small, but non-zero, improvement in utility. This improvement is likely to arise in any model in which the inflation distortion is manifested as a concern for the discounted future path of inflation. In the model explored here, the utility improvement is about two-tenths of 1 percent.
3. The benefits of price-level targeting arise from the form of the inflation distortion used in the paper. For forward-looking consumers with standard discount rates, concern for the discounted present value of tax-distorting effects of high inflation rates is very similar to concern for the price level. An exercise that substantially increases the discount rate verifies that, if consumers are only moderately forward-looking, the benefits of price level targeting in this specification disappear.
4. For most of the specifications studied here (in particular, those with persistent inflation), utility-maximizing monetary policies involve relatively high responses to output, and quite low responses to inflation. This result does not depend importantly on the magnitude of the shocks buffeting the economy, or on the size of the inflation distortion.
5. This model suggests that the welfare improvement from reducing the inflation target from a low positive number (2 to 3 percent) to 0 is negligible.

6. Fifty years may be a short sample from which to gauge what might be an optimal monetary policy. Small-sample simulations suggest that economies with identical structures that are buffeted by shocks with the same covariance structure will benefit most from quite different monetary policies.

Of course, reasonable people can and should disagree on matters of model specification. Perhaps the strongest result to take from this paper is this: Given the uncertainties about key aspects of macroeconomic structure (sources of inflation persistence, degree of forward-lookingness, sources of inflation distortions), it is hard to make very strong statements about the optimal monetary policy.

References

- Batini, Nicoletta and Yates, Anthony, “Hybrid Inflation and Price Level Targeting,” working paper, August 1999.
- Bernanke, Ben, Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen, *Inflation Targeting: Lessons from the International Experience*, Princeton: Princeton University Press, 1999.
- Blinder, Alan S., “Monetary Policy in a Low Inflation Environment: Comment,” Federal Reserve System conference, October 1999.
- Campbell, John Y. and Mankiw, N. Gregory, “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence.” O. Blanchard and S. Fischer, eds., *NBER Macroeconomics Annual*. Cambridge: MIT Press, 1989, pp. 185–216.
- Estrella, Arturo and Fuhrer, Jeffrey C., “Dynamic Inconsistencies: Counterfactual Implications of a Class of Rational Expectations Models,” forthcoming (2000), *American Economic Review*.
- Feldstein, Martin, “The Costs and Benefits of Going from Low Inflation to Price Stability,” in *Reducing Inflation: Motivation and Strategy*, C. Romer and D. Romer, eds., Chicago: University of Chicago Press, 1997, pp. 123–56.
- Fuhrer, Jeffrey C., “Habit Formation in Consumption and Its Implications for Monetary Policy Models,” *American Economic Review*, **90**, No. 3, June 2000, pp. 367–90.
- Fuhrer, Jeffrey C. and Moore, George R., “Inflation Persistence”, *Quarterly Journal of Economics*, **110**, February 1995a, pp. 127–59.

- Fuhrer, Jeffrey C. and Moore, George R., “Monetary Policy Trade-Offs and the Correlation Between Nominal Interest Rates and Real Output”, *American Economic Review*, **85**, March 1995b, pp. 219–39.
- Maclelean, Dinah and Hope Pioro, “Price Level Targeting—The Role of Credibility,” prepared for the Bank of Canada Seminar on *Price Stability and the Long-Run Target for Monetary Policy*, June 2000.
- McCallum, Bennet T. and Nelson, Edward, “An Optimizing IS-LM Specification for Monetary Policy and Business Cycle Analysis”, *Journal of Money, Credit and Banking*, **31**, No. 3, August 1999, pp. 296–316.
- Rotemberg, Julio and Woodford, Michael. “An Optimization-Based Framework for the Evaluation of Monetary Policy,” B. Bernanke and J. Rotemberg, eds., *NBER Macroeconomics Annual*, Cambridge: MIT Press, pp. 297–346, 1997.
- Rudebusch, Glenn and Svensson, Lars E., “Policy Rules for Inflation Targeting,” in John B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press and NBER, 1998.
- Svensson, Lars E., “Inflation Targeting: Some Extensions,” *Scandinavian Journal of Economics*, **101**, No. 3, 1999a, pp. 337–61.
- Svensson, Lars E., “Price Level Targeting vs. Inflation Targeting: A Free Lunch?”, *Journal of Money, Credit and Banking*, **31**, August 1999b, 277–95.
- Svensson, Lars E., “The Zero Bound in an Open-Economy: A Foolproof Way of Escaping from a Liquidity Trap,” NBER Working Paper No. 7957, IMES Discussion Paper No. 2000-E-23, revised September 2000.
- Tetlow, Robert, and von zur Muehlen, Peter, “Simplicity Versus Optimality: the choice of monetary policy rules when agents must learn,” Finance and Eco-


nomics Discussion Series, Federal Reserve Board of Governors, paper No. 99–10, January 1999.

Williams, John, “Simple Rules for Monetary Policy,” Finance and Economics Discussion Series, Federal Reserve Board of Governors, paper No. 99–12, February 1999.

Wolman, Alexander L., “Real Implications of the Zero Bound on Nominal Interest Rates,” working paper, Federal Reserve Bank of Richmond, January 2000.

Woodford, Michael. *Interest and Prices*, unpublished manuscript, chapter 6, “Inflation Stabilization and Welfare,” 1999, pp. 4–23.

Figure 1
Comparison of Inflation/Output Std. Dev. Frontiers, Multivariate Objective Function
Kitchen sink targeting
Habit Formation Model



Note: Color versions of figures may be viewed in the PDF version, available on the FRB Boston Research Website

Figure 2
Comparison of Price Level/Output Std. Dev. Frontiers, Multivariate Objective Function
Kitchen sink targeting
Habit Formation Model




Figure 3
Comparison of Price Level/Output Std. Dev. Frontiers, Multivariate Objective Function
Kitchen sink targeting
McCallum–Nelson Model




Figure 4
Comparison of Inflation/Output Std. Dev. Frontiers, Multivariate Objective Function for Inflation Forecast Targeting (solid black) vs. Pure inflation targeting (dashed red) Habit Formation Model




Figure 5
Comparison of Inflation/Output Std. Dev. Frontiers, Multivariate Objective Function for Inflation Forecast Targeting (solid black) vs. Pure inflation targeting (dashed red) Fuhrer-Moore Model




Figure 7
Effective marginal tax rate versus inflation, 1965–1995




Figure 8

Utility surfaces, Baseline Model ($\tau = .69$), Baseline Ω , All shocks




Figure 9

Utility surfaces, Level Inflation Effects, $\tau = 3.0$, Baseline Ω , All shocks




Figure 10

Utility surfaces, Flexible Inflation Model, Baseline Ω , All shocks




Figure 11

Utility surfaces, Baseline Model ($\tau = .69$), Large Ω , All shocks




Figure 12

Utility surfaces, High discount rate, Baseline τ , Baseline Ω , All shocks




Figure 13

Utility surfaces, Baseline model versus SS $\pi = 0$




Figure 14
Comparison of Inflation/Output Std. Dev. Frontiers, Multivariate Objective Function
Kitchen Sink targeting
Model with tax distortion




Figure 15
Comparison of Variance Frontiers and Utility-Maximizing Policies
Model with tax distortions




Figure 16
Comparison of small-sample utility surfaces

High discount rate, Baseline τ , Baseline Ω , All shocks

