Monetary Policy When Interest Rates Are Bounded at Zero

by Jeffrey Fuhrer
and
Brian Madigan

No. 94-1 March 1994

Working Paper Series

Federal Reserve Bank of Boston

Monetary Policy When Interest Rates Are Bounded at Zero

by Jeffrey Fuhrer
and
Brian Madigan

March 1994
Working Paper No. 94-1

Federal Reserve Bank of Boston

Monetary Policy When Interest Rates Are Bounded at Zero

Jeffrey Fuhrer and Brian Madigan*

Revised
March 4, 1994

* Assistant Vice President, Research Department, Federal Reserve Bank of Boston, Boston, MA, 02106, and Associate Director, Division of Monetary Affairs, Board of Governors of the Federal Reserve System, Washington, D.C., 20551. The authors acknowledge George Moore's role in initiating this research and thank Thomas Sargent, Michael Woodford, and participants in the Stanford/San Francisco Fed conference on "Monetary Policy in a Low Inflation Regime" for helpful comments. The views expressed in this paper are the authors' and do not necessarily reflect those of the Federal Reserve Bank of Boston or the Board of Governors of the Federal Reserve System. Please do not quote without permission of the authors.

Abstract

This article assesses the importance of the zero lower bound on nominal interest rates for the conduct of monetary policy. The article employs a small, forward-looking model developed by Fuhrer and Moore. The model is simulated under several policy rules that involve either high- or low-inflation targets. We determine the extent to which the zero bound on nominal interest rates prevents real interest rates from falling in response to negative spending shocks, and thus cushioning aggregate output, when zero inflation results in low nominal rates.

In general, the results suggest that real long-term interest rates drop considerably in response to an adverse spending shock under a variety of policy rules and inflation rates. The extent of the decline in long real rates, and thus the ability of monetary policy to cushion such shocks, generally depends to only a modest extent on the level of inflation. For relatively small and short-lived spending shocks, as well as for permanent and large shocks, the path of output in the zero inflation case is only a little below that in the higher inflation. But for large shocks persisting a few quarters, differences in output paths across high- and low-inflation scenarios can be larger.

Without a doubt, these results are somewhat model-specific, and their real-world implications depend on how quickly a central bank can recognize shocks and how vigorously it can respond to them. Moreover, in situations when the zero bound on nominal interest rates does limit the ability of the central bank to stimulate the economy by reducing interest rates, other policy tools—such as fiscal policy—may still be effective. Nonetheless, this research suggests that the constraint on monetary policy posed by the zero bound is an issue that merits careful thought and perhaps further investigation in alternative model settings.

1 Introduction

This article assesses the importance of the zero lower bound on nominal interest rates for the conduct of monetary policy. In the context of arguing that the optimal rate of inflation is positive, Lawrence Summers [8] stated that a possible drawback of aggregate price stability is that the central bank would be constrained in its ability to offset adverse spending shocks because nominal interest rates cannot turn negative. Cushioning output appreciably in the face of a negative demand shock may require moving long-term real rates down significantly. If short-term nominal interest rates were already low before the shock because inflation were low, the central bank may not be able to reduce short-term real rates much. The argument assumes implicitly that the inability to lower short-term real rates significantly impedes the downward adjustment of long rates.

In this article, we assess this argument using a small forward-looking model.[1] The model was estimated by Jeffrey Fuhrer and George Moore [2]. It incorporates multi-period pricing contracts, a standard IS curve that depends on long-term real interest rates, and a forward-looking bond market in which real long-term rates are set consistent with market participants' expectations of future short-term real rates. This model and its characteristics are described in some detail in the next section.

We examine solution paths for the model under higher and lower rates of inflation and a variety of monetary policy reaction functions. We take the higher rate of inflation to be 4 percent and the low rate to be zero. (We have ignored biases in price indexes that may cause the desired measured rate of inflation to be positive rather than zero.) We assess differences in the high- and low-inflation scenarios by relative deviations of output from baseline.

[1]David Lebow examines this as well as other arguments relating to the zero bound on nominal interest rates [5].

identify in estimating the model. Evidence supporting (1) is provided in Fuhrer and Moore [2]. As for the shocks, the estimated model assumes a linear trend for potential output, so that permanent shocks to output are not identified. Thus we do not know if the permanent shocks entertained in this paper are consistent with the shocks underlying the estimated model. The temporary output shocks that we simulate, however, fall well within the estimated distribution of shocks to the output process.

We enforce the zero bound on nominal interest rates through two alternative techniques involving the monetary policy reaction function, rather than through a nonlinear money demand equation.[3] (In fact, the model includes no money demand equation or variable measuring the quantity of money.) The techniques are:

- The left-hand-side of the reaction function is specified in terms of log differences of the short-term nominal rate.
- The left-hand-side of the reaction function is specified instead in terms of levels of short-term nominal rates, but the response to nominal income is adjusted to keep nominal rates from becoming negative.

An advantage of the first technique is that such a policy rule can be specified fully in advance; no negative nominal rate is arithmetically possible. Thus, a reaction function with given weights on deviations of the targeted variable can be employed in high- and low-inflation scenarios. A possible disadvantage of this technique is that the numerical solution methodology we employed does not identify a stable solution for some values of the parameters. This problem in some cases may be a shortcoming of the numerical

[3]That is, a possible alternative procedure would involve a demand for central bank money that went to infinity as nominal short rates asymptotically approached zero combined with a reaction function specified in terms of money rather than in terms of short-term interest rates.

2 The Model

The simple structural model that we use comprises three sectors: an IS curve that relates output to the *ex ante* long-term real interest rate, a monetary policy reaction function that moves the short-term nominal interest rate in response to deviations of target variables from desired values, and a price contracting specification in which nominal price contracts are negotiated in real terms. The model has been estimated on postwar quarterly data for the 3-month Treasury bill, the deflator for nonfarm business output, and a measure of the output gap for nonfarm business output, defined as the residual from a regression of log per capita nonfarm output on a constant and a linear time trend. Maximum likelihood estimation yields significant estimates of all the structural parameters. The dynamics implied by the model, as summarized by the vector autocorrelation function, match the dynamics from an unrestricted vector autoregression very well. At the estimated parameter values, the model implies a sensible sacrifice ratio, about in line with the estimates in Gordon [4]. Overall, the model behaves similarly to a conventional macroeconometric model such as the MPS model, despite its forward-looking asset and price sectors. Fuhrer and Moore [2] present a more extensive discussion of the model and its properties.

2.1 The IS Curve

Let R_t be the yield to maturity on a coupon bond selling at par, and let M be the maturity of the bond at the end of quarter t. Then the duration of the bond is given by

$$D_t = \frac{1 - e^{-R_t M}}{R_t} \tag{1}$$

scale of long rate volatility, one should recognize that the volatility of long rates depends on the particular monetary policy rule, as well as the parameters of the model. Thus, the volatility of long rates observed in this paper's simulations may not correspond closely with that observed historically.

Given the definition of the expected long real rate, the real economy is represented as a simple IS curve that relates the output gap $\tilde{y}_t$ (the deviation of the log of output from the log of potential output) to its own lagged values and one lag of the long-term real interest rate, ρ_{t-1},

$$\tilde{y}_t = 0.017 + 1.254\tilde{y}_{t-1} - 0.415\tilde{y}_{t-2} - 0.798\rho_{t-1}\,, \tag{4}$$

where the parameters are taken from Fuhrer and Moore [2] and ρ_t is the rate on consols defined on the previous page.[6] In the steady state, $\tilde{y} = 0$, so the IS curve defines the equilibrium or "natural" real rate of interest, ρ^n, to be 2.1 percent (0.017/0.798). Note that the equilibrium real rate includes any real term premium built into the long rate.

One potential concern over using such a simple IS curve when inflation and nominal rates are near zero is that the linear representation will not capture an important nonlinear response of spending to interest rates when nominal rates are near zero. However, the period of estimation for the IS curve includes the 1975-79 period, during which short-term real rates varied from -6 percent to 0 percent, and the long-term real rate implied by the model dropped well below its equilibrium. The IS curve shows no sign of misbehaving during this period, suggesting that if the response exists, it is not of primary importance for total spending.

Another possible shortcoming of this IS curve is the omission of a real-

[6] These parameters are consistent, but inefficient, partial-information estimates. However, they differ insignificantly from the full information estimates presented in Fuhrer and Moore [3]. The full-information estimate of the interest elasticity parameter, for example, is −0.746, with a standard error of 0.25.

contracts currently in effect,[7]

$$v_t = \sum_{i=0}^{3} f_i (x_{t-i} - p_{t-i}) . \tag{7}$$

Agents set nominal contract prices so that the current real contract price equals the average real contract price index expected to prevail over the life of the contract, adjusted for excess demand conditions.

$$x_t - p_t = \sum_{i=0}^{3} f_i E_t (v_{t+i} + \gamma \tilde{y}_{t+i}) . \tag{8}$$

Substituting equation 7 into equation 8 yields the real version of Taylor's contracting equation, [8]

$$x_t - p_t = \sum_{i=1}^{3} \beta_i (x_{t-i} - p_{t-i}) + \sum_{i=1}^{3} \beta_i E_t (x_{t+i} - p_{t+i}) + \gamma^* \sum_{i=0}^{3} f_i E_t (\tilde{y}_{t+i}) . \tag{9}$$

In their contract price decisions, agents compare the current real contract price with an average of the real contract prices that were negotiated in the recent past and those that are expected to be negotiated in the near future; the weights in the average measure the extent to which the past and future contracts overlap the current one. When output is expected to be high, the current real contract price is high relative to the real contract prices on overlapping contracts.

The contracting specification is parameterized by s, the slope of the con-

[7] This is a convenient simplification from the theoretically preferable specification that defines the real contract price as the difference between the nominal contract price and the weighted average of price indexes that are expected to prevail over the life of the contract. The simplification yields an algebraically more straightforward model. The effects of the simplification on the empirical properties of the model are relatively small. See Fuhrer and Moore [2] for details on the alternative specification and associated empirical results.

[8] Compare equation 9 with equation 1 on page 4 of Taylor [9]. The coefficients in equation 9 are $\beta_i = \sum_j f_j f_{i+j} / (1 - \sum_j f_j^2)$ and $\gamma^* = \gamma / (1 - \sum_j f_j^2)$.

long-term real rate. The ultimate target is either nominal income or nominal income growth.

3 Simulations

3.1 Permanent Unanticipated Shocks

This section discusses simulations conducted under a permanent, unanticipated shock. The shock increases the output gap initially by 0.4 percent by reducing the natural rate of interest by 50 basis points. In the post-shock steady state, real long rates will be 1.6 percent (2.1 percent minus 50 basis points); short nominal rates will be 5.6 percent with 4 percent inflation and 1.6 percent with zero inflation.

3.1.1 Operating Instrument: Log-Difference Nominal Rates

Chart 1 illustrates a simulation using a log-difference reaction function and a nominal income target. The response of interest rate differences to deviations of the level of nominal income from target—λ—is set to a value of 60. This value is the maximum at which a simulation could be obtained. As shown by the solid line in the upper panel, the nominal short rate in the zero-inflation case adjusts down over a period of about a year by a total of nearly 1-1/4 percentage points. By contrast, the nominal rate in the high-inflation scenario (shown by the dashed line) falls about 3 percentage points. As would be expected with the log-difference reaction function, the *percent* reduction in nominal interest rates is similar in the two cases. The small difference reflects the slightly stronger nominal income in the high-inflation case and the feedback through the reaction function to the nominal rate.

The middle panel of Chart 1 shows that the long real rate overshoots in both cases—it initially falls by more than the 50 basis point decline in the natural real rate, as markets bring forward in time the lower short real rates in the future that will result from lower nominal short rates combined with only sluggishly declining inflation. In the zero inflation case, the long rate falls 56 basis points right away and then trends up gradually. In the high inflation scenario, the long rate initially drops a bit more—62 basis points. The long real rate then rises more steeply than in the low inflation case, reflecting the anticipated need for monetary policy to lean more heavily against the stronger upward burst of output shown in the lower panel. The drop in real rates is obviously similar across the two scenarios.

The drop in output in the first quarter is identical in the two cases—0.4 percent. The model incorporates a one-quarter lag in the response of demand to interest rates, so the drop in output in the first quarter represents solely the exogenous decline in demand, and hence is identical in the two scenarios. The slightly lower initial real rates of the high-inflation scenario cushion output in the second quarter, essentially preventing it from falling further as it does in the zero-inflation case; the difference, however, is slight—less than one-tenth of 1 percent of the level of output. Output subsequently recovers a little more steeply in the high-inflation case: The end of the recovery—defined as the point at which output "recovers" its pre-recession level—comes about a quarter earlier. The subsequent cycles are of greater amplitude in this case.

Chart 2 presents a simulation somewhat similar to that of Chart 1, except that policy reacts to nominal income *growth* rather than nominal income levels. Under nominal income *growth* targeting, the price level will be lower in the post-shock steady state than under its baseline rate, while the inflation rate will return to its baseline rate, which is equal to the targeted growth rate of nominal income. By contrast, under nominal income targeting, both the price level and the inflation rate ultimately return to baseline after a shock. The requirement that the price level return to baseline in the nominal income case induces additional cycles in the solution relative to the nominal income growth case, which only requires the inflation rate to return to its targeted level. These additional cycles are evident in a comparison of Charts 1 and 2.

With policy reacting only to income growth rather than income levels, the short rate is reduced by less and is brought up sooner to the vicinity of its new equilibrium. Real rates consequently drop a little less. In the low-inflation case, the real long rate drops immediately to, but not below, the new natural rate, while in the high-inflation case the long rate overshoots. The lower real rates of the high-inflation case bring the output gap to zero appreciably more quickly than in the zero-inflation scenario.

3.1.2 Operating Instrument: Nominal Short-Rate Levels

In section 3.1.1, the zero percent floor on nominal interest rates was enforced by considering reaction functions in log-difference form. In this section, the operating instrument is considered to be levels of short-term nominal rates, as in equation 11. The policy responsiveness coefficient, λ, is adjusted on a case-by-case basis in such a way as to allow the nominal short rate to fall to, but not below, zero.

Such a reaction function, i.e., one specified in terms of levels, seems most consistent with a view that the central bank can determine the true level of the natural rate of interest with a high degree of certainty. In such a situation, the central bank presumably would wish to move interest rates promptly to appropriate levels. By contrast, the previous section's log-difference reaction function, which embodies interest-rate smoothing, might better characterize policy as actually practiced, as monetary policymakers take into account uncertainty, any costs of interest rate variability, and perhaps an aversion to frequent reversals of course.

As shown by the dashed line in Chart 3, in the 4 percent inflation case, policy lowers the short rate to zero in the second period, after nominal income began to fall significantly below target in the first period. That is, nominal short rates fall 6.1 percentage points almost immediately. This responsiveness corresponds to a λ equal to 14. By contrast, the lower level of nominal rates in the low-inflation case permits a much less aggressive policy response: Short rates can fall only 2.1 percentage points, corresponding to a λ of 3.5.

The middle panel shows that the larger move in nominal rates in the 4 percent inflation case results in a sharper initial drop in the real long rate, by about 9 basis points for two periods. Long rates subsequently move up more strongly in the high-inflation situation. The larger drop in real rates in this case causes the recession (defined as the period during which the output gap is growing) to end after one quarter, whereas the recession in the low-inflation case lasts two quarters. The recovery similarly ends sooner in the 4 percent inflation scenario. Output overshoots and cycles a little in the high inflation case. Although output also overshoots slightly in the low inflation case, the approach to equilibrium is more gradual.

Chart 4 presents simulations for an interest rate levels operating target and a nominal income *growth* ultimate target. The high- and low-inflation cases use λ equal to 14 and 6.5, respectively. In both simulations, monetary policy drops nominal short rates sharply as an output gap opens and inflation falls below target, leading to a shortfall in nominal income growth. The easing is reversed quickly, however, as a drop in the real long rate prompts a rebound in real output that pushes nominal income *growth* roughly back to target. In both the high- and low-inflation cases, after a few quarters the nominal short rate actually gets pushed a bit above its new long-run equilibrium, and this overage is transferred to the real long rate. Very slowly, the real long rate drifts toward the new natural rate, bringing the output gap eventually to zero. Although output is a little higher in the second through sixth quarters in the high inflation case, the difference is small.

3.2 Temporary Unanticipated Shocks

This section generally considers a temporary 0.4 percent shock to aggregate demand; in most simulations, the shock occurs in the first period of the solution and last for one quarter. Section 3.2.2 also considers a longer-lasting temporary shock and a reaction function that includes forward-looking elements.

3.2.1 Operating Instrument: Log-Difference Nominal Short Rates

With a log difference reaction function and λ set equal to 30, short rates in the high inflation case decline about 4 percentage points over the span of a few quarters, while short rates fall 1-1/2 percentage points (to 65 basis points) in the low inflation case. (This simulation is shown in Chart 5.) Real long rates drop considerably farther initially with 4 percent inflation—35 basis points, as opposed to 20 basis points in the zero inflation case. The lower real rates permit a somewhat steeper recovery of output, but the pattern is not markedly different. The relatively modest variation in output across the two cases reflects the small difference in long rates measured in *percentage points*.

3.2.2 Operating Instrument: Nominal Short Rate Levels

Chart 6 shows results for temporary unanticipated shocks and an interest-rate-levels reaction function. Again, the high inflation case permits a substantially more aggressive response measured in terms of percentage point movement in nominal short rates. (The policy responsiveness parameter λ is equal to 10 and 1.65, respectively, in the two cases.) Consequently, the decline in real long rates is more than twice as steep. But because the percentage point difference in long rates is relatively small, the trajectory of

Chart 6
Unanticipated Temporary Shock
Nominal Income Targeting
Operating Target: Interest Rate Levels
(percent)

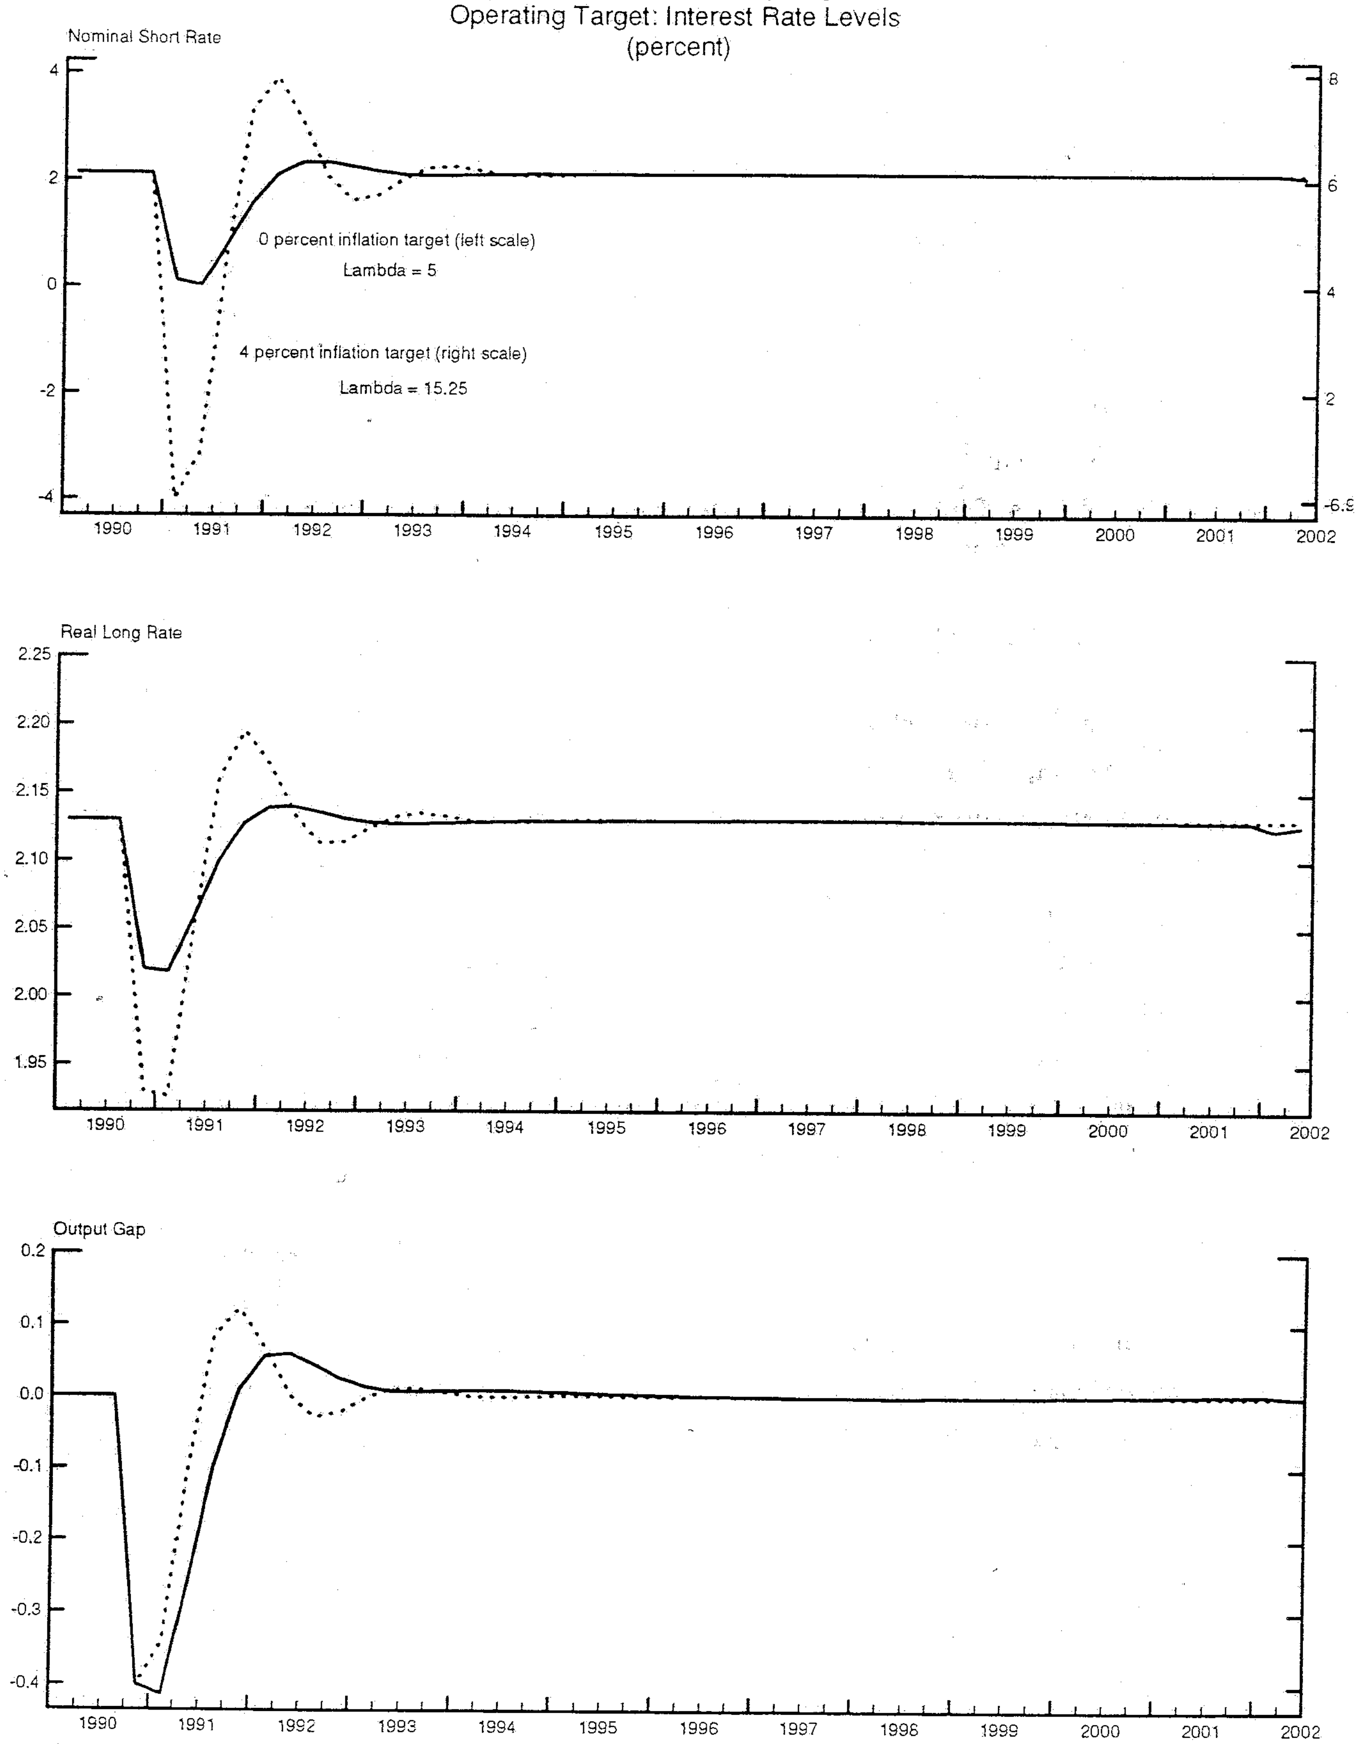

chart6.00.new,chart6.04.new

Chart 7
Initially Unanticipated Temporary (Four-quarter) Shock
Nominal Income Targeting
Operating Target: Interest Rate Levels
(percent)

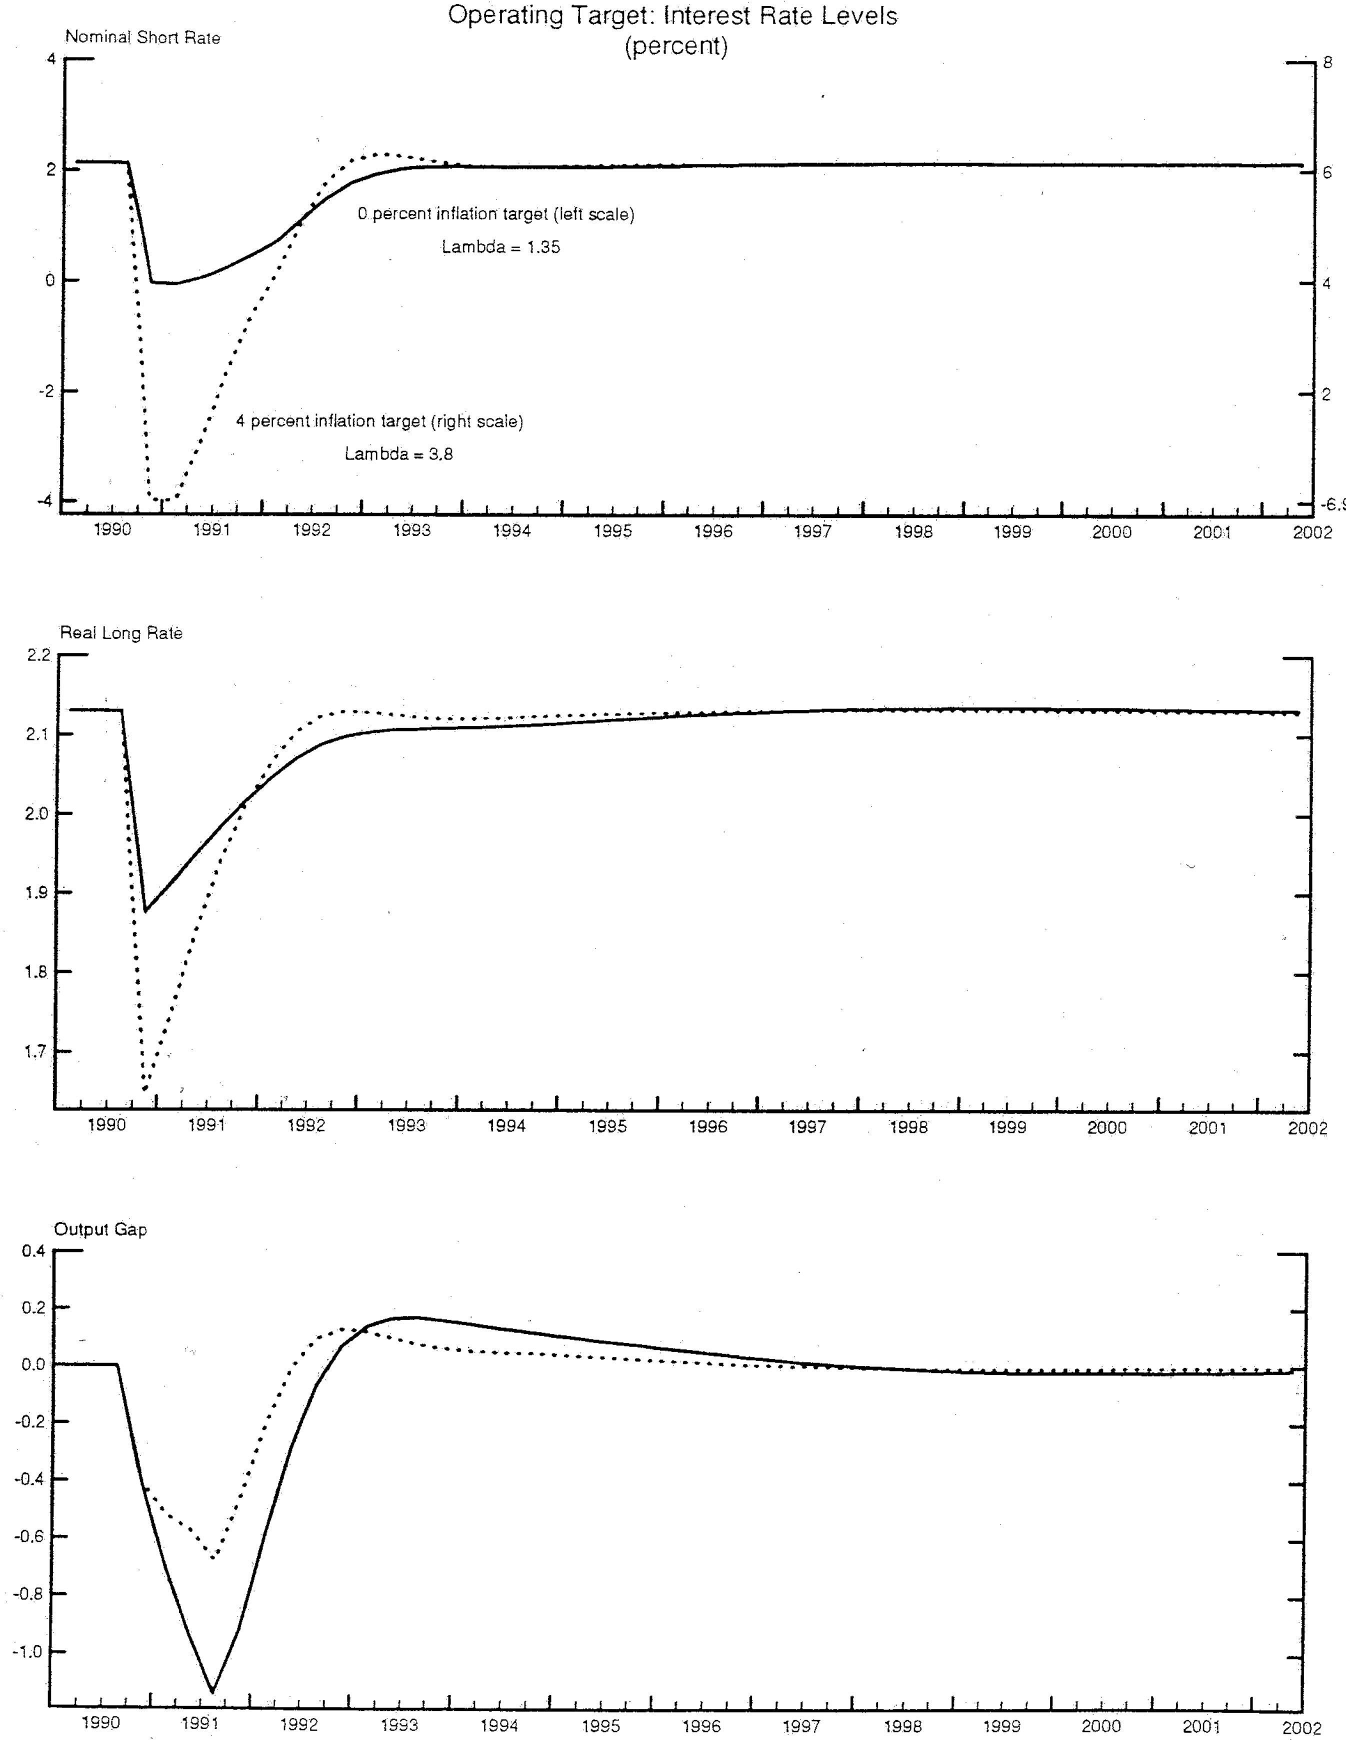

chart7.00.new,chart7.04.new

data.[9] In the 4-percent-inflation scenario, the output gap reaches a trough of 2.8 percent, while in the zero-percent-inflation case, output troughs at 3.6 percent. The 0.8 percentage point absolute difference in output gaps in the two scenarios probably would be regarded by many as economically meaningful.[10]

3.3 Temporary Anticipated Shocks

Chart 9 considers the case of temporary anticipated shocks with an operating instrument specified in terms of interest rate levels. The fact that demand will be depressed for one quarter by 0.4 percent becomes known four quarters in advance. In response, long-term real rates drop immediately by about 5 basis points. Output initially moves up as interest-sensitive spending responds to lower interest rates in advance of the spending shock. Monetary policy responds to the excess of nominal income over target by pushing up the short rate. After four quarters, output drops sharply as demand falls off temporarily. With a lag of one quarter, monetary policy eases. In the high-inflation case, the short rate is dropped extremely sharply, from nearly 10 percent to zero. In the zero-inflation case, the short rate falls from about 3.25 percent to zero. Anticipating the more aggressive easing of monetary policy in the high-inflation case, real long rates decline about 9 basis points further than under zero inflation. Consequently, the trough in output in the high inflation case is slightly above that in the low-inflation situation.

[9] Recession depths may be estimated either from a log detrended output series (with the trend broken in 1973) or from an output gap series implied by the unemployment rate and an inverted Okun's Law.

[10] Interestingly, the *relative* differences in the output gaps at the trough in Chart 8 are smaller than in Chart 7, while the *absolute* differences are larger. These differences illustrate perhaps a general point for reaction functions that bring nominal rates promptly to zero: As the size of the shock grows, the *relative* differences in output shrink.

4 Additional Considerations

The differences identified in the previous section between output under high- and low-inflation cases may depend in part on certain aspects of the model's specification and parameters. In this section, we consider the following modifications:

- The existence of a term premium in long-term interest rates, which implies a lower steady-state level of real short rates, may limit the ability of monetary policy to stimulate economic activity.
- If short real rates, in addition to long real rates, affect spending, the ability of long rates to jump down, cushioning the effects of an adverse spending shock, would be less relevant and the behavior of short rates would be more relevant.
- If bond markets are partly backward-looking, bond rates may be less apt to jump down when news becomes available about adverse spending shocks.

4.1 Term Premium in Interest Rates

As noted, the previous simulations assume that real long-term interest rates contain no term premium. If there is a term premium in long rates, however, the ability of monetary policy to ease in response to an adverse spending shock would be more constrained, because the equilibrium short real rate would be lower than the equilibrium long rate by the amount of the term premium. Whitesell [10] estimated that the equilibrium real rate on 3-month Treasury bills between 1978 and 1992 was between 1/2 and 1-5/8 percentage points below that on 10-year Treasury notes. In this section, we assume a constant term premium in long-term real rates of 1 percentage point. In contrast to the estimation results underlying the previous simulations, we

Chart 10
Unanticipated Permanent Shock
Nominal Income Targeting
Operating Target: Log Interest Rate Changes
(percent)
1 Percentage Point Premium in Real Long Rates
Lambda = 60

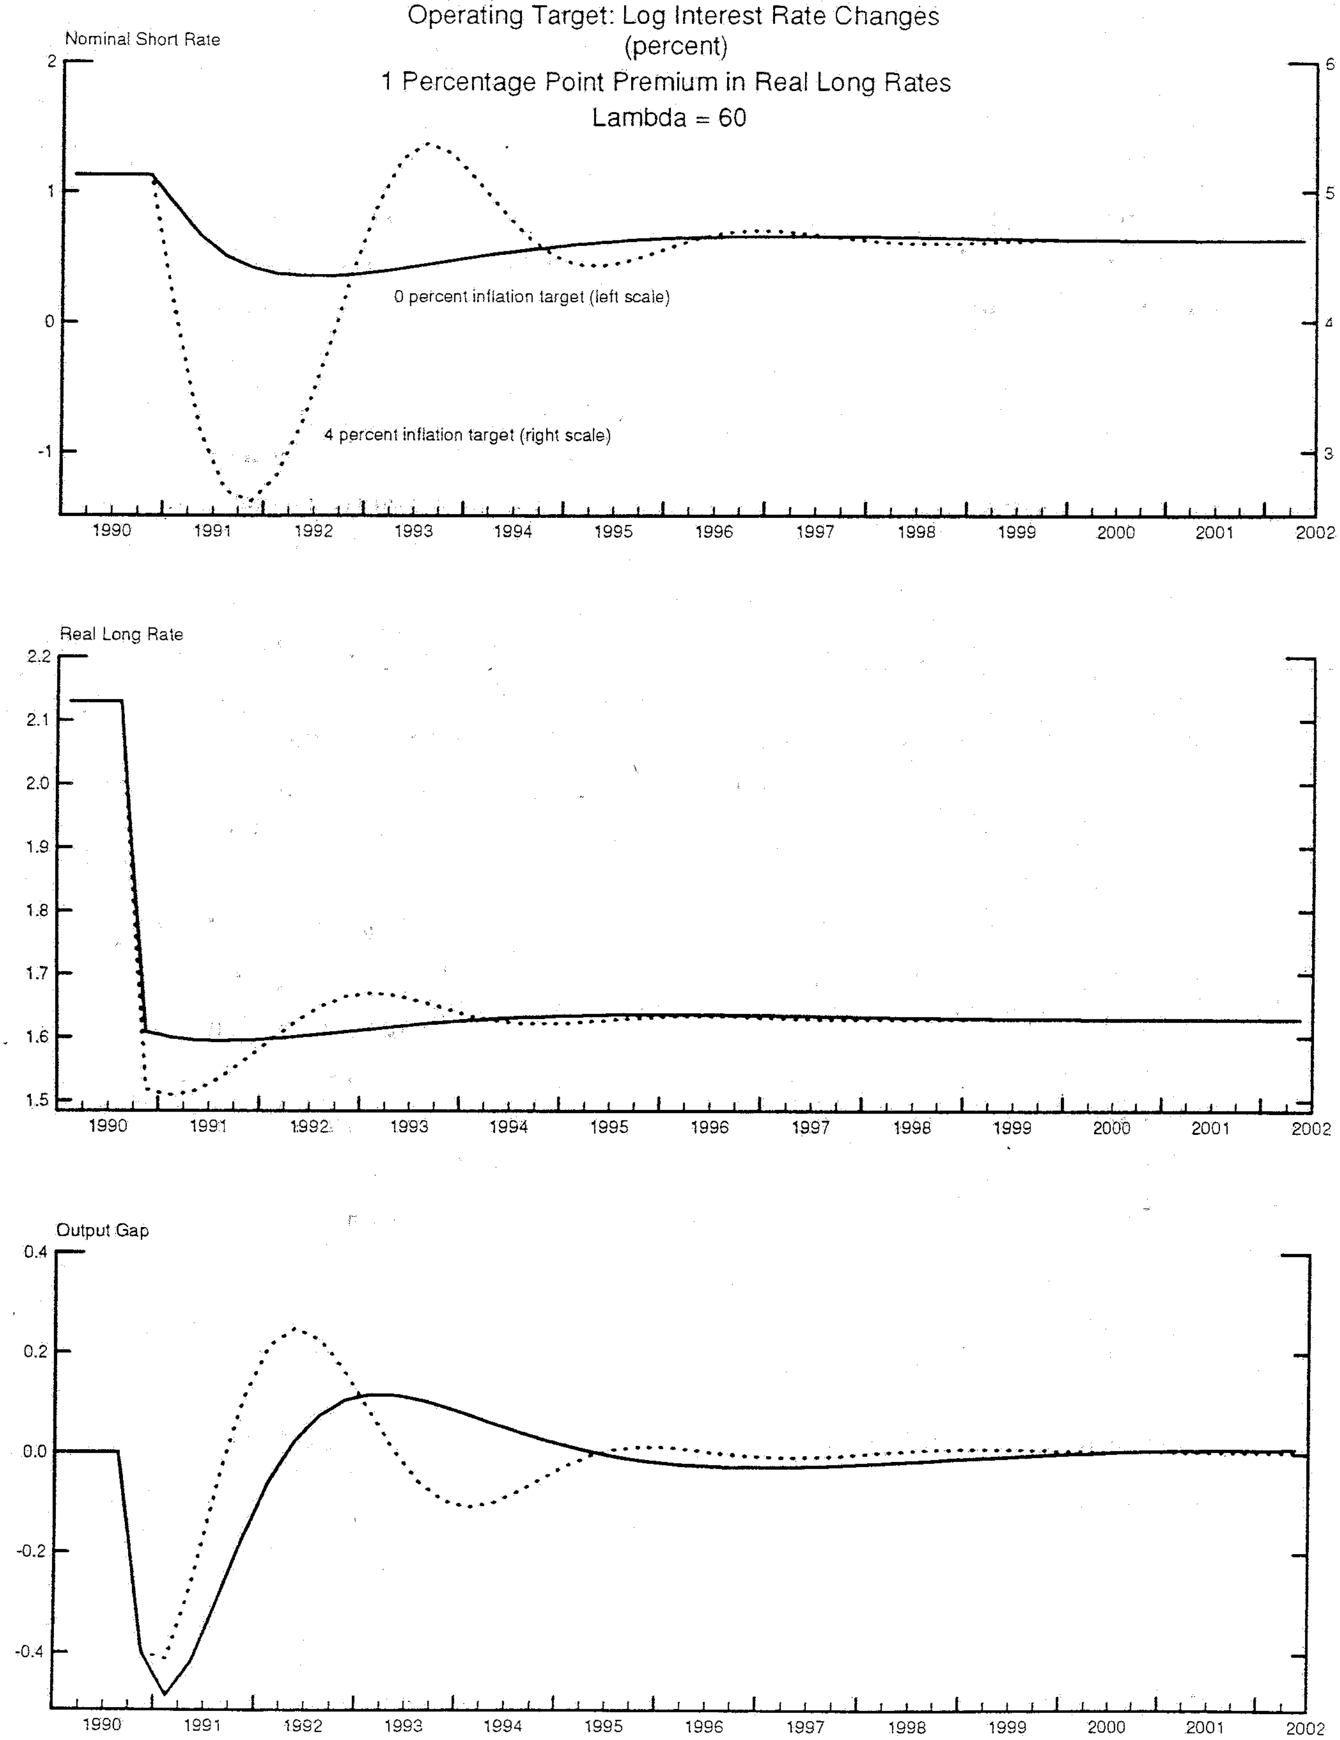

chart10.00,chart10.04

Chart 11 displays the results of this simulation. The model behaves about as it does in the simulations in Chart 1, although the movements are somewhat smaller in amplitude, reflecting the smaller size of the initial shock. Real long rates again jump down, overshooting their equilibrium values. The real two-year rate jumps down considerably and rebounds more vigorously in expectation of the movements in the short nominal rate, especially in the high-inflation case, reflecting the sharper contemporaneous and anticipated movements in nominal short rates. Despite the sharper movements of the shorter real rate, the path of the output gap under high inflation is stronger only to a modest degree in comparison with that of Chart 1. Evidently, the weight on short rates in the IS curve is not large enough to make an appreciable difference in output.

Overall, for our specification, the exclusion of shorter-duration real rates appears to be an unimportant omission. Including the short-duration real rate yields qualitatively similar behavior for moderately aggressive policy responses. However, a more disaggregated IS curve that separately modeled spending on consumer durables, business equipment, residential structures, and nonresidential structures might well find larger sensitivity to shorter-duration real rates than we have used.

4.3 Backward-Looking Bond Markets

The expected long-term real rate as defined in equation 3 is completely forward-looking, satisfying period-by-period arbitrage. As a result, the real rate jumps immediately in response to an unanticipated shock, as is clear in Charts 1 to 11. This feature of the model's long real rates may be a bit unrealistic; while holding period returns are unlikely to diverge over extended periods, they may not be equalized period by period. Thus we explore the robustness of our simulation results to a modified real rate specification that combines both forward-looking and backward-looking behavior. The "mixed" real rate is defined as

$$\rho_t^m = w\rho_t^b + (1-w)\rho_t^f \tag{13}$$

where ρ_t^f is the purely forward-looking real rate defined in equation 3, and ρ_t^b is a weighted average of past short real rates, with weights that sum to 1 and decline geometrically into the past

$$\rho_t^b = \delta\rho_{t-1}^b + (1-\delta)(i_{t-1} - \pi_t) \tag{14}$$

which expands to

$$\rho_t^b = (1-\delta)\sum_{j=0}^{\infty} \delta^j (i_{t-1-j} - \pi_{t-j}). \tag{15}$$

By varying the degree of backward-lookingness, w, in the real rate, and the rate at which the backward-looking weights, δ^j, decay into the past, we can get an idea of the sensitivity of our results to the real rate specification.

Interestingly, the stability of the model is sensitive to the exact combination of w, δ, and IS interest elasticity. For example, with $w = .5$, $\delta = .9$, and the interest elasticity in equation 4, the model does not have a stable, unique solution. The backward-looking long rate places too much weight on the recent past, and thus exerts a destabilizing force on output and inflation. Setting $\delta = .98$ (which implies weights decaying into the past at the same rate as the weights decay into the future for the forward-looking real rate), $w = .95$, $\lambda = 50$, and the interest elasticity to its benchmark value, the model is stable with a log-difference reaction function. Chart 12 displays the solution paths of the short nominal rate, the long real rate, and the output gap in response to an unanticipated permanent output gap shock.

As can be seen in the middle panel, the real long rate still jumps down considerably, despite the very large parameters on the backward-looking component of real rates. But it takes three to four quarters for the rate to reach its trough, unlike the case in Chart 1, where long rates reach their low point essentially immediately. As in Chart 1, real rates are lower in the high inflation scenario than in the low inflation case for the first four quarters or so, but the relative differences between the purely forward-looking case and the mixed backward/forward case are quite small. Consequently, the relative paths of output are similar, although the cycles are larger in Chart 12, given the sluggish cushioning effect of long rates.

Overall, incorporating backward-looking behavior in the bond markets does not alter the qualitative properties of the model simulations.

5 Conclusion

This article examined one argument that the optimal rate of inflation is positive, namely that the lower nominal rates of interest that would accompany zero inflation would limit the ability of monetary policy to ease in response to an adverse spending shock. To assess the argument, we utilized a small model of the U.S. economy that captures forward-looking behavior both in financial markets and in product markets and whose broad properties correspond with those of large-scale macroeconometric models. Our results indicate that the argument is correct, qualitatively speaking. Although long-term real rates in forward-looking bond markets do decline in response to news about adverse spending shocks, thus cushioning the reduction in output, the decline in real rates can be constrained by the inability of nominal rates to fall below zero.

We find that for relatively small and short-lived spending shocks, as well as for permanent and large shocks, the path of output in the zero inflation case is only modestly below that in the higher inflation case—on the order of a tenth or two of a percent; the recession and the recovery tend to be completed one quarter later with higher inflation rates. But for large shocks has provided an initial quantification and a starting point for future research on the issue of the relevance of the zero bound.

References

[1] Gary Anderson and George Moore. A linear algebraic procedure for solving linear perfect foresight models. *Economics Letters*, 17:247–252, 1985.

[2] Jeffrey Fuhrer and George Moore. Inflation persistence. FEDS Working paper 93-17, Board of Governors of the Federal Reserve System, May 1993.

[3] Jeffrey Fuhrer and George Moore. Monetary policy and the behavior of long-term real interest rates. FEDS Working Paper 93-16, Board of Governors of the Federal Reserve System, May 1993.

[4] Robert J. Gordon. Understanding Inflation in the 1980s. *Brookings Papers on Economic Activity*, pages 263–299, 1:1985.

[5] David Lebow. Monetary policy at near-zero interest rates. Manuscript, Board of Governors of the Federal Reserve System, April 1993.

[6] Brian Madigan. User's guide to resolver: A routine for solving nonlinear forward-looking models. Manuscript, Board of Governors of the Federal Reserve System, forthcoming.

[7] George Moore and Sandra Cannon. User's guide to aim: A routine for solving forward-looking models. Manuscript, Board of Governors of the Federal Reserve System.

[8] Lawrence Summers. How should long-term monetary policy be determined? *Journal of Money, Credit, and Banking*, 23(3):625–31, August 1991.

[9] John B. Taylor. Aggregate dynamics and staggererd contracts. *Journal of Political Economy*, 88:1–17, 1980.