A Quantitative Exploration of the Opportunistic Approach to Disinflation*

Athanasios Orphanides, David H. Small, Volker Wieland, and David W. Wilcox
Board of Governors of the Federal Reserve System

June 1997

Abstract

A number of observers have advocated recently that the Federal Reserve take an "opportunistic" approach to the conduct of monetary policy. A hallmark of this approach is that the central bank focuses on fighting inflation when inflation is high, but focuses on stabilizing output when inflation is low. The implied policy rule is nonlinear. This paper compares the behavior of inflation and output under opportunistic and conventional linear policies. Using stochastic simulations of a small-scale rational expectations model, we study the cost and time required to achieve a given disinflation, as well as the steady-state distributions of inflation and output under the various rules.

KEYWORDS: Inflation, monetary policy, interest rates, policy rules.

JEL Classification System: E52

*The opinions expressed in this paper are the authors', and are not necessarily shared by the members of the Board of Governors of the Federal Reserve System nor by the other members of its staff. We are grateful for excellent research assistance from Dwight Bibbs, Adam Reed, and Kendrew Witt, and for helpful comments from John Williams, Andrew Levin, John Taylor, and participants in the Stanford Macro Lunch Group, the Konstanz Seminar on Monetary Theory and Policy and the 1997 meeting of the American Economic Association. Correspondence: Division of Monetary Affairs, Board of Governors of the Federal Reserve System, Washington, D.C. 20551. Orphanides: (202) 452-2654, aorphanides@frb.gov. Small: (202) 452-2659, dsmall@frb.gov. Wieland: (202) 736-5620, vwieland@frb.gov. Wilcox: (202) 452-2441, dwilcox@frb.gov.

1 Introduction

During the past 15 years or so, inflation has declined substantially in the United States, from more than 11 percent in 1980 to roughly 3 percent in each of the last few years.[1] When inflation is high, as it was at the beginning of this period, the consensus in favor of deliberate anti-inflationary action on the part of the central bank is very strong, notwithstanding the substantial associated cost in terms of lost output and employment. But when inflation is low, as is currently the case, the consensus for deliberate anti-inflationary action is less overwhelming.

A conventional view regarding the conduct of monetary policy is that a central bank should balance the objective of achieving and maintaining low inflation against the objective of stabilizing real activity around its sustainable level, and that the tradeoff between the two objectives should be roughly linear. This approach to policy determination implies, for example, that if the current rate of inflation is 3 percent and the target rate is less than 3 percent, the central bank should attempt to shade the economy toward operating a bit below its potential, in order to put downward pressure on inflation.

Recently, an alternative viewpoint has emerged. Adherents of this viewpoint argue that when inflation is low but still above the long-run objective, the monetary authority should focus exclusively on output and employment stabilization, and not worry about deliberately inducing any reduction in inflation. When disinflation occurs (perhaps because a favorable supply shock has occurred or an unanticipated demand-induced recession has taken place), the central bank should consolidate the gain, and—if necessary—move to stabilize output at potential. This alternative strategy has come to be known as "the opportunistic approach to disinflation."

A previous paper (Orphanides and Wilcox [1996]) examined possible theoretical rationales for the opportunistic approach. The primary objective of this paper is to compare the behavior of key macroeconomic variables under conventional and opportunistic policies,

[1] CPI-U, all items, 1983 forward; CPI-U-X1, which "backcasts" the measurement of homeowners' costs on a rental equivalence basis, 1980-83.

using stochastic simulations of a small-scale rational expectations model which we built for this purpose. We also provide some information on the performance of rules that attempt to confine inflation within a target zone. In our macro model, monetary policy has real effects because workers and firms sign long-term contracts. The model is large enough to provide a realistic description of some of the most basic features of the U.S. economy, and yet small enough to make stochastic simulations relatively convenient.

The rest of this paper is organized as follows. In Section 2, we give a more complete description of the opportunistic approach to disinflation, and contrast it with a more conventional approach. Section 3 describes the macroeconometric model we use as a laboratory for our comparisons. Section 4 uses some simple simulations to illuminate the differences between the opportunistic and conventional rules. In each of these simulations, the model is subjected to only one type of shock—for example, a shock to the contract wage equation or to aggregate demand. Section 5 conducts stochastic simulations using shocks to all behavioral equations in the model. We use these simulations to estimate probability distributions over the time required under each of the various rules to achieve price stability. We also calculate the distributions of inflation and the output gap in the stochastic steady state—that is, once the transition to price stability has been completed. Section 6 concludes.

2 Conventional and opportunistic decision rules

Much of the recent literature on monetary policy rules assumes that the policymaker wishes to stabilize output around potential and inflation around a long-run target. In the long run, there is no conflict between these objectives provided the long-run Phillips curve is vertical. At any given moment, however, they are in tension with one another. Henderson and McKibbin (1993) and Taylor (1993b) propose to resolve this tension using a reaction function of the following form:

$$i^s = r^* + \pi + \gamma_1 y + \gamma_2(\pi - \pi^*)$$

where i^s is the short-term nominal interest rate (the monetary authority's policy instrument), r^* is the equilibrium short-term real interest rate (that is, the short-term real interest rate consistent with output being at potential), y is the output gap (measured as log[actual output/potential output]), π is the rate of inflation, π^* is the policymaker's long-run target for inflation, and γ_1 and γ_2 are positive parameters. We refer to policymakers with linear decision rules of this class as "conventional."

We operationalize this rule for use with quarterly data by assuming that the policymaker reacts to the lagged output gap (y_{t-1}) and the lagged four-quarter inflation rate ($\pi_{t-1} \equiv p_{t-1} - p_{t-5}$, where p denotes the log of the price level). With these modifications, we rewrite the conventional decision rule as:

$$i_t^s = r^* + \pi_{t-1} + \gamma_1 y_{t-1} + \gamma_2(\pi_{t-1} - \pi^*). \tag{1}$$

Following Taylor (1993b), we set both r^* and π^* at 2 percent.

Equation (1) reflects the conventional view that both the inflation gap and the output gap should always influence the choice of policy stance.[2] An opportunistic policymaker does not subscribe to this view. On the contrary, she focuses on different objectives depending on the level of inflation: When inflation is high, she focuses on fighting inflation, but when inflation is low, she focuses on stabilizing output. Orphanides and Wilcox (1996) proposed the following decision rule as a formal description of opportunistic behavior:[3]

$$i_t^s = r^* + \pi_{t-1} + \gamma_1 y_{t-1} + f(\pi_{t-1} - \tilde{\pi}_{t-1}) \tag{2}$$

$$\tilde{\pi}_t = (1 - \lambda)\pi^* + \lambda \pi_t^\dagger \tag{3}$$

[2] Put slightly differently, the derivative of the short-term interest rate with respect to both gaps should be non-zero, no matter what the current value of either gap.

[3] The key assumptions involved in deriving this form of the opportunistic decision rule (see Orphanides and Wilcox [1996]) are that inflation expectations are formed adaptively, the policymaker sets the short-term interest rate before observing any shocks from the current period, and that the policymaker controls the output gap perfectly.

$$f(\pi - \tilde{\pi}) = \begin{cases} \gamma_3(\pi - \tilde{\pi} - \delta) & \text{if } \pi - \tilde{\pi} > \delta \\ 0 & \text{if } \delta \geq \pi - \tilde{\pi} \geq -\delta \\ \gamma_3(\pi - \tilde{\pi} + \delta) & \text{if } \pi - \tilde{\pi} < -\delta \end{cases} \tag{4}$$

where $\tilde{\pi}$ is an intermediate target, which is calculated as a weighted average of the long-run target and the inherited rate of inflation ($\pi_t^\dagger$). The latter is simply taken to be a backward-looking moving average of actual inflation.

The first three terms on the right-hand sides of equations (1) and (2) are identical. The only difference between the two specifications is in the reaction to deviations of inflation from target; here, there are two key differences. First, while the conventional policymaker reacts to the gap between actual inflation and her long-run inflation target, the opportunistic policymaker reacts to the gap between actual inflation and an intermediate target. Second, while the conventional policymaker responds to the inflation gap in a linear manner, the opportunistic policymaker's reaction to the gap between actual inflation and the intermediate target is nonlinear. If the discrepancy between lagged inflation and the intermediate target does not exceed a certain tolerance level (denoted as δ), she sets the inflation penalty equal to 0. If the discrepancy exceeds δ, she adjusts the interest rate by γ_3 percentage points for each percentage point of excess discrepancy. We refer to the range in which the opportunistic policymaker sets the inflation penalty to zero as the "zone of opportunism."

In the simulation exercises we conduct below, we focus mainly on two versions of the conventional linear specification given in equation (1) and one version of the opportunistic rule. The first version of the linear specification we simulate is Taylor's (1993b) rule, which we obtain by setting γ_1 and γ_2 equal to 0.5. The second version of the linear specification is a rule that was proposed by Henderson and McKibbin (1993); we obtain this rule by setting γ_1 equal to 2 and γ_2 equal to 1. Thus, the Henderson-McKibbin rule is more aggressive than the Taylor rule in combating both output gaps and inflation gaps, but more so with respect to output gaps than with respect to inflation gaps.

We calibrate the opportunistic rule as follows: We set the coefficient on the lagged output gap in the opportunistic rule at 2, and the slope of the inflation penalty function outside the zone of opportunism at 2.5. The penalty on output gaps is the same as in Henderson-McKibbin, while the penalty on inflation outside the zone of opportunism is considerably stiffer than the marginal penalty on inflation in Henderson-McKibbin. Our interest in calibrating the opportunistic rule relative to *some* specific conventional rule was motivated by results presented in Orphanides and Wilcox (1996) to the effect that optimizing conventional and opportunistic policymakers, confronted with the same economic environment, will choose closely related parameter settings. In the simple model Orphanides and Wilcox analyzed, the opportunistic policymaker chooses the same coefficient on the output gap as does the conventional policymaker, and sets a stiffer marginal penalty on inflation outside the zone of opportunism, consistent with our choices here. Our decision to maximize comparability specifically with the Henderson-McKibbin rule was motivated by the results reported in Levin (1996) and Brayton et al. (1996), as well as in Henderson and McKibbin (1993), to the effect that rules with more aggressive response coefficients than those in Taylor's rule seem to deliver better macroeconomic performance in at least some models.[4]

Figure 1 compares the opportunist's penalty on inflation, $f(.)$, with the penalties imposed by the two conventional policymakers. We show the deviation of inflation from target on the horizontal axis, and the implied increment to the short-term nominal rate on the vertical axis. The Taylor penalty function is shown with the dotted line, the Henderson-McKibbin with the dashed line, and the opportunistic with the solid line. (In superimposing the opportunistic penalty on top of the conventional ones, we are implicitly assuming that inherited inflation—and thus the intermediate target—are both equal to the long-run target.) The figure highlights that the opportunist behaves quite differently than either

[4] Henderson-McKibbin implies a much stronger response to output deviations than does Taylor. In most macroeconometric models, real output exhibits a relatively high degree of inertia, so that a high partial interest-rate elasticity with respect to output deviations can be expected to generate greater output stability. Two other minor notes on numerical assignment of parameter values for the opportunistic rule: We set λ equal to $^1\!/_3$, and thus put relatively more weight on the long-run target in the determination of the intermediate target. Also, we set δ equal to 1 percentage point; as a result, the zone of opportunism, which is centered on the intermediate target, has an overall width of 2 percentage points.

conventionalist: Whereas the conventionalists lean against inflation at least to some extent whenever current inflation deviates from the long-run target, the opportunist in this specification makes essentially no effort to move inflation closer to the intermediate target when the inflation gap is already small.

Although we devote the bulk of our attention to the three policy rules described above, towards the end of the paper we also present some results for three rules that can be interpreted as variations on the opportunistic theme. The first variation recognizes the influence that uncertainty about the location of the aggregate demand curve might have on the behavior of the opportunist. Orphanides and Wilcox (1996) argued that an opportunistic policymaker who took this type of uncertainty seriously would behave a bit more like the conventionalist than is suggested by Figure 1. Although the penalty function would remain flatter inside the zone of opportunism than outside, it would not be literally horizontal. In other words, the opportunist would expect to make some progress toward the long-run inflation objective even if inflation were already close to the target. Analytical derivation of a decision rule that allows for this type of uncertainty is difficult or impossible in a model as complex the one we describe below, so we approximate what we presume to be the actual shape of such a decision rule using a function that has the slope of the Taylor rule (0.5) inside the zone of opportunism, and the slope of the baseline opportunistic specification outside the zone of opportunism.

The other two variations are motivated by the recent literature on the concept of inflation targeting (see, for example, Leiderman and Svensson (1995), Svensson (1996)), as well as the rhetoric and actions of several central banks around the world (e.g., Australia, Canada, Sweden, and the United Kingdom) which tend to emphasize containing inflation within a target *zone* rather than confining it to a point (see Table 1 in Bernanke and Mishkin (1997)). In a mechanical sense, these variations can be derived from the opportunistic rule by fixing the intermediate target at the long-run target. Again, we consider two possibilities for the case in which inflation is close to the target: one in which the response function is flat and the other in which the response function is moderately sloped. The flat version might be most appropriate for a "zoner" who could control output perfectly, while the sloped version might be more appropriate for a zoner who exercises imperfect control over output.

We close this section with a technical note related to the algorithm we use to simulate the model with each of the policy rules. If this algorithm were fully satisfactory, the conditional expectation of inflation $E_t\pi_{t+N}$ would depend on the variances of the various shocks in the model *when the policymaker is using the opportunistic rule.* (The issue we are addressing here is not relevant when the policymaker is using a linear rule.) If at least some of those variances are positive, the conditional expectation would converge to π^* as N goes to infinity because economic agents would expect the opportunistic policymaker to react asymmetrically to future shocks, leaning against the unfavorable ones and accommodating the favorable ones. Given forward-looking behavior in wage setting, such expectations would drive the expectation of inflation to the long-run target. Moreover, the speed of convergence would depend on the variance of the various shocks in the model (larger variances being associated with a more rapid convergence of the conditional expectation to the long-run target). On the other hand, if the variances of all shocks in the model were zero, the limit of the conditional expectation would depend on initial conditions, but would not, in general, equal π^*. For example, if inflation initially were above π^* but still within the zone of opportunism and the model were otherwise in equilibrium, inflation would simply remain at its initial level.

Unfortunately, the simulation algorithm we have been using does not produce exactly these results. In fact, this algorithm causes the conditional expectations of all the state variables in the model to converge to user-specified terminal conditions, regardless of the variances of the shocks in the model. This obviously creates an uncomfortable internal inconsistency if the variances of the shocks in the model are declared to be zero: the conditional expectation of inflation converges to the long-run target even when it should not. We resolve this tension by using the following approximation to equation (4):

$$\begin{aligned} f(\pi-\tilde{\pi}) &\approx \gamma_3 g(\pi-\tilde{\pi}) \\ &= \gamma_3[0.05(\pi-\tilde{\pi}) + 0.475(-\delta+\pi-\tilde{\pi}+((-\delta+\pi-\tilde{\pi})^2)^{0.51}) \\ &\qquad + 0.475(\ \delta+\pi-\tilde{\pi}-((\ \delta+\pi-\tilde{\pi})^2)^{0.51})] \end{aligned} \tag{5}$$

The function $g(.)$ is continuously differentiable, and has a tiny bit of positive slope even when inflation is close to the intermediate target. That slope rationalizes the convergence of the conditional expectation of inflation to the long-run target even when the variance of future shocks is set equal to zero. This is, in fact, the function that is plotted in Figure 1; as can be seen there, the slope of the function inside the zone of opportunism is very slight, even though non-zero.

To be clear, we should emphasize that the variance of shocks in principle has both a direct and an indirect effect on the results. The direct effect is that a greater variance of shocks gives the opportunistic policymaker greater scope for asymmetric behavior. The indirect effect is that all agents (private as well as public) should be taking this into account when they form their expectations. The simulation algorithm we are currently using captures the direct effect but not the indirect one.

There are other solution algorithms for nonlinear rational expectations models available that do not impose certainty equivalence, but these alternative algorithms would be prohibitively costly to use with the current model given the large number of state variables therein. In future work with a much smaller model, we intend to explore the quantitative importance of the indirect effects.

3 A small model of the U.S. economy with rational expectations

This section describes the model that we use as a laboratory for the comparison of the rules described above. The main features of this model are that it accounts for the forward-looking behavior of economic agents, and that it is small enough to be usable for stochastic simulation studies of alternative policy rules. Monetary policy has temporary real effects in this model because wage contracts are staggered in the manner of Taylor (1980). The model is essentially a one-country version of the multi-country model which is presented and used for policy analysis in Taylor (1993a). A version of the model has been used by Taylor and Williams (1993) in a study of forecasting with rational expectations models. We alter the Taylor-Williams specification of the wage-price sector by following Fuhrer and Moore (1995a, 1995b) in assuming that workers and firms set the real wage in the first period of each new contract with an eye toward the real wage agreed upon in the first period of other contracts signed in the recent past and expected to be signed in the near future.[5] As Fuhrer and Moore show, models specified in this manner exhibit a greater (and hence more realistic) degree of inflation persistence than do models in which workers and firms care about relative wages in nominal terms. We use the parameter values that were estimated by Taylor and Williams (1993) and Fuhrer and Moore (1995a), and we simulate the model using a very efficient algorithm that was recently implemented in TROLL based on work by Boucekkine (1995), Juillard (1994) and Laffargue (1990).

Expectations of endogenous variables are formed rationally in our simulations, given the maintained assumption of certainty equivalence. Consequently, the simulations are immune to the Lucas Critique. In particular, agents' expectations fully reflect the choice of monetary rule. Rotemberg and Woodford (1997) take a similar approach using a model that is more explicitly grounded than ours in specific optimization problems confronted by firms, households, and policymakers. We note, however, that their model may be of only limited usefulness for the type of analysis we wish to undertake here because its parameters are calibrated rather than estimated.

The model is a simple linear flow model of the economy. We group the various equations under three headings: interest rates, aggregate demand, and the wage-price sector.

3.1 Interest rates

Three equations determine the various interest rates in the model. A policy rule determines the short-term nominal rate; a term-structure equation determines the long-term nominal rate; and a version of Fisher's equation determines the long-term real rate. We consider three different specifications of the policy rule: (a) an opportunistic rule, (b) Henderson

[5] By contrast, Taylor (1980) assumed that workers and firms set the *nominal* wage in the first period of each new contract with an eye toward the *nominal* wage settlements of recently signed and soon-to-be signed contracts.

and McKibbin's rule, and (c) Taylor's rule.

(1a) *opportunistic policy rule*

$$i_t^s = .02 + \pi_{t-1} + 2y_{t-1} + 2.5g(\pi_{t-1} - \tilde{\pi}_{t-1}) + u_{1,t}$$

$$\tilde{\pi}_t = (1-\lambda)\pi^* + \lambda\pi_t^\dagger$$

$$\pi_t^\dagger = \frac{1}{8}\sum_{j=0}^{7} \pi_{t-j}$$

$$\begin{aligned} g(\pi - \tilde{\pi}) \quad = \quad [.05(\pi - \tilde{\pi}) \quad &+ \quad .475(-\delta + \pi - \tilde{\pi} + ((-\delta + \pi - \tilde{\pi})^2)^{.51}) \\ &+ \quad .475(\ \ \delta + \pi - \tilde{\pi} - ((\ \ \delta + \pi - \tilde{\pi})^2)^{.51})] \end{aligned}$$

(1b) *Henderson and McKibbin's rule*

$$i_t^s = .02 + \pi_{t-1} + 2y_{t-1} + 1(\pi_{t-1} - \pi^*) + u_{1,t}$$

(1c) *Taylor's rule*

$$i_t^s = .02 + \pi_{t-1} + .5y_{t-1} + .5(\pi_{t-1} - \pi^*) + u_{1,t}$$

(2) *long-term nominal rate*

$$i_t^l = .02 + \frac{1 - .784}{1 - .784^{17}} E_t \sum_{j=0}^{16} (.784)^j i_{t+j}^s + u_{2,t}$$

(3) *long-term real rate*

$$r_t^l = i_t^l - E_t\pi_{t+4}$$

3.2 Aggregate demand

Aggregate demand Y_t is the sum of consumption C_t, fixed investment FI_t, inventory investment II_t, total (federal, state, and local) government purchases G_t and net exports NEX_t. Potential output Y_t^T is defined to be equal to a log-linear trend growing at an annual rate of 2.45 percent. Following Taylor and Williams, we scale each demand component by the level of potential output, and denote the result with lower-case letters.

(4) *consumption*

$$c_t = .221 + .669c_{t-1} + .269y_t^p - .040r_t^l + u_{4,t}$$

(5) *permanent income*

$$y_t^p = \frac{1-.8}{1-(.8)^9} E_t \sum_{j=0}^{8} (.8)^j y_{t+j}$$

(6) *fixed investment*

$$fi_t = .013 + 1.141fi_{t-1} - .271fi_{t-2} + .273y_t - .266y_{t-1} - .009r_t^l + u_{6,t}$$

(7) *inventory investment*

$$ii_t = .003 + .575ii_{t-1} + .352y_t - .316y_{t-1} - .011r_t^l + u_{7,t}$$

(8) *net exports*

$$nex_t = .000 + .898nex_{t-1} - .037y_t + .102y_t^w - .011e_t + u_{8,t}$$

(9) *government spending*

$$g_t = .012 + .936g_{t-1} + u_{9,t}$$

(10) *output gap*

$$y_t = c_t + fi_t + ii_t + nex_t + g_t - 1$$

Normalized consumption is modeled as a function of its own lagged value, permanent income, and the expected long-term real interest rate. The lagged dependent variable can be rationalized as reflecting a cost to changing consumption. Permanent income is the annuity value of expected income in the current and next eight periods. The investment equations are (nearly) of the accelerator type. Net exports depend on the level of income at home and abroad (y_t^w), and on the real exchange rate (e_t). Government spending follows a simple autoregressive process, with a near-unit root. All these equations are estimated in Taylor and Williams (1993) using GMM. Following Taylor and Williams, we hold foreign output and the real exchange rate constant in our simulations.[6]

[6] In future work, we intend to investigate whether we can bring into the model a slightly more realistic description of fiscal policy: As Taylor (1995) shows, the federal deficit is strongly countercyclical. This key feature is absent from the specification given above.

3.3 Wages and Prices

The wage-price block consists of two equations that determine the real wage to be paid in the current quarter under newly-signed contracts, and a markup equation determining the price in the current period.

(11) *index of real contract prices*

$$v_t = .372(x_t - p_t) + .291(x_{t-1} - p_{t-1}) + .210(x_{t-2} - p_{t-2}) + .129(x_{t-3} - p_{t-3})$$

(12) *current real contract wage*

$$x_t - p_t = E_t(.372v_t + .291v_{t+1} + .210v_{t+2} + .129v_{t+3})$$

$$+.008E_t(.372y_t + .291y_{t+1} + .210y_{t+2} + .129y_{t+3}) + u_{12,t}$$

(13) *aggregate nominal price*

$$p_t = .372x_t + .291x_{t-1} + .210x_{t-2} + .129x_{t-3}$$

Equations (11) and (12) specify that the real wage under contracts signed in the current period is set in reference to a centered moving average of initial-period real wages established under contracts signed as many as three quarters earlier as well as contracts to be signed as many as three quarters ahead. The real wage also depends on expected excess-demand conditions. Once contracts are signed, they remain in force for up to four quarters. As a result, the aggregate price p_t is a weighted average of the nominal wages that were negotiated in the current and previous three quarters (and thus that remain in force), with the weights reflecting the proportion of contracts outstanding from each quarter. Full-information maximum likelihood estimates for the parameters in the wage-price block are taken from Fuhrer and Moore (1995a).

3.4 The steady state of the model

In the steady state of this model, output is at potential, and the sectoral allocation of GDP is constant. Because we hold rest-of-world output and the real exchange rate constant, these conditions define a unique steady-state value of the long-term real interest rate. Taylor and Williams adjusted the intercept in the consumption equation so that the steady-state long-term real interest rate would be 4 percent.[7] Given the adjustment to the intercept in the consumption equation and the resulting 4 percent equilibrium long-term real rate, the steady-state consumption share is 66.5 percent, the fixed investment share is 15.6 percent, the inventory investment share is 0.4 percent, the net exports share is -1 percent, and the government spending share is 18.5 percent. The dynamic properties of the model are not sensitive to the calibration of the steady state.

The steady-state value of inflation is determined exclusively by the policymaker's reaction function; the wage-price block does not impose any restriction on the steady-state inflation rate. These conditions guarantee that the steady-state inflation rate will be π^* under all three of the rules we simulate.

4 Opportunism versus conventional policy: some illustrative examples

This section uses some simple simulations to illustrate the behavior of the model under each of the three decision rules described above. In each of these simulations, we either shut down all shocks completely, or we subject the model to only one type of shock (for example, an adverse supply shock or a negative shock to aggregate demand). In the next section of the paper, we consider simulations involving repeated shocks to all equations in the model.

4.1 Disinflation in the absence of shocks

We begin by examining the behavior of the model when the initial inflation rate is 4 percent (i.e., 2 percentage points above the long-run target) and there are no shocks to the economy during the simulation.[8] Figure 2 summarizes the results. The top-left panel shows the

[7] In the unadjusted model, according to Taylor and Williams, the steady-state long-term real interest rate is 7 percent.

[8] Note that initially output is set at potential and all other endogenous variables at their steady-state values given an inflation rate of 4 percent.

trajectory of inflation over the simulation period for each of the three policy rules, while the top right panel shows inflation under the opportunistic rule relative to the zone of opportunism. The middle panel shows the evolution of the output gap, while the bottom panel shows the path of the short-term nominal interest rate. In each panel, the data pertaining to the Taylor rule are shown with a solid line; those for the Henderson-McKibbin rule are shown with a dashed line, and those for the opportunistic rule are shown with a dotted line.

With the inflation rate initially at 4 percent, the Taylor-type policymaker sets the short-term nominal interest rate at 7 percent (2 percent for the real rate, 4 percent for the inflation rate, and 1 percent for the excess of inflation over the long-term target). This level of the nominal interest rate is sufficient to drive output nearly $^3/_4$ percent below potential after three quarters. In turn, this slack is sufficient to put inflation on a downward trajectory. After about three years, inflation has reached 3 percent—halfway to the long-run target—and output has returned about two-thirds of the way to potential. This policymaker perseveres in the fight against inflation, and eventually drives inflation down to its target of 2 percent.

The Henderson-McKibbin policymaker puts the nominal short-term rate at the start of the simulation at 8 percent (2 percent for the real rate, 4 percent for the prevailing inflation rate, and 2 percent for the excess of inflation over the long-run target). The resulting recession is only a bit deeper and more longer-lived, however, because the Henderson-McKibbin policymaker cuts the nominal rate quite aggressively once an output gap opens up. Inflation comes down slightly more rapidly under the Henderson-McKibbin rule than under the Taylor rule.

The opportunistic policymaker puts the short-term nominal rate at 6-$^2/_3$ percent at the start of the simulation, only slightly below the rate set under Taylor's rule. This level of the short-term rate reflects the following calculations on the part of the opportunist: With the long-run target set at 2 percent, an inherited rate of 4 percent, and a λ of $^1/_3$, the opportunist sets the intermediate target at 2-$^2/_3$ percent. And with a δ of 1 percent, the upper boundary of the zone of opportunism is 3-$^2/_3$ percent. Given a 4 percent prevailing rate of inflation and a coefficient of 2.5 on inflation outside the zone of opportunism, the inflation penalty equals $5/6$. Thus, the overall nominal interest rate of 6-$5/6$ percent makes allowance of 2 percentage points for the real rate, 4 percentage points for inflation, and $5/6$ percentage point for the inflation penalty.

Despite the fact that the initial short-term rate is nearly as high under the opportunist as it is under the Taylor-style policymaker, the opportunist endures a much smaller recession because the agents in both the bond and labor markets look ahead and correctly anticipate that in the absence of further shocks the opportunistic policymaker will adopt an easier policy stance in the future. Inflation declines considerably less rapidly under opportunism than under either of the conventional policies; indeed, inflation under the opportunistic policy reaches the 3 percent level only after 7 years. On the other hand, the cumulative output loss under the opportunist during that period is much smaller.

Thus, given our parameterization of these rules, the opportunistic policymaker tolerates more inflation than does either Taylor's policymaker or Henderson and McKibbin's. In return, the opportunist enjoys an economy that operates closer to potential. No one policy dominates the other two for all possible preference orderings over inflation and output.

4.2 Disinflation with favorable supply shocks: a first scenario

In Figure 3, we again assume an initial inflation rate of 4 percent (upper left panel). In this simulation, however, progress toward lower inflation is facilitated by a pair of favorable supply shocks. Specifically, we assume that the shock to the contracting equation, $u_{12,t}$, takes on negative values in each of the first two quarters of the simulation.

Under the Taylor-style policy, inflation reaches 3 percent within a year. (Recall that, in the absence of shocks, the inflation rate reached 3 percent only after about three years.) Because the policymaker does not anticipate the favorable shocks, she induces a recession that is not much shallower than the one she provoked when there were no shocks. However, she eases up much more quickly than she did in the no-shock scenario, and by the end of the second year of the simulation, output is essentially back at potential. The results for the Henderson-McKibbin policymaker are remarkably similar to those for the Taylor policymaker in this scenario. Evidently, the greater aggressiveness in response to output gaps about offsets the greater sensitivity to inflation gaps, yielding about the same net policy stance.

Under the opportunistic policy, the pace of disinflation during the first year of the simulation is very nearly as rapid as it is under either of the conventional policies. Once the disinflationary impetus from the shocks has disappeared, the inflation rate under the opportunistic rule flattens out at about 2-$^1\!/_2$ percent. As was the case in the absence of shocks, the opportunistic policymaker suffers much smaller output losses than does either conventional policymaker.

4.3 Disinflation with favorable supply shocks: a second scenario

Figure 4 examines a scenario in which inflation starts out at the upper end of the zone of opportunism. Two years into the simulation, the economy is hit with a pair of unanticipated favorable supply shocks. Thereafter, there are no shocks to the economy. In these circumstances, both the Henderson-McKibbin and Taylor-style policymakers drive output somewhat below potential before the shocks hit (see the middle panel) in order to bring the inflation rate down toward the long-run target. When the supply shocks hit, inflation actually falls below the long-run target, so these policymakers engineer a small boom in order to bring inflation back up to the long-run target. Eventually, the inflation rate converges to 2 percent, as expected.

Under this scenario, the opportunistic policymaker fares very well indeed. During the first two years of the simulation, the opportunistic policymaker sets output only slightly below potential, and inflation creeps toward the long-run target at a snail's pace.[9] When the favorable supply shocks hit, they are (by design) just large enough to drop the inflation rate to the long-run target level of 2 percent. Throughout the simulation, the opportunistic policymaker holds output very close to potential, and yet achieves the desired disinflation.

[9]This creeping would not occur if the simulation algorithm we use took explicit account of the variance of future shocks, and if agents expected the variance of future shocks to be zero.

In this case, the strategy of waiting for favorable shocks pays off handsomely: inflation is driven down to the long-run objective while output is held almost exactly at potential.

4.4 Rising inflation due to an adverse supply shock

Figure 5 shows a scenario in which the three policymakers are not nearly as fortunate as they were in Figure 4. Prior to the beginning of this scenario, the economy was in equilibrium, with output at potential and inflation at the central bank's long-run target of 2 percent. In the first quarter of the simulation, however, a huge adverse inflation shock hits the economy—sufficient to drive inflation up to 14 percent within four quarters. In response, both the Taylor-style and the Henderson-McKibbin-style policymakers raise the short-term nominal interest rate to about 20 percent. The increase in interest rates drives output down relative to potential; at its deepest point, four quarters into the simulation, the gap between actual and potential under these two policies reaches 3-½ percent. Thereafter, with the initial shock no longer exerting any direct influence on the economy, interest rates begin to decline, inflation begins to subside, and output begins to recover, but the process is fairly protracted.

In this scenario, the opportunistic policymaker responds much more aggressively than does either the Taylor policymaker or the Henderson-McKibbin policymaker, because the adverse shock pushes inflation far above the upper boundary of the zone of opportunism. The nominal interest rate under the opportunistic rule peaks at 27 percent, and output dips more than 5 percent below potential. At first, inflation declines slightly more rapidly than under either the Taylor rule or the Henderson-McKibbin rule, but later on less so. The skimpiness of the inflation-reduction reward to the opportunistic policymaker for enduring a considerably deeper recession partly reflects the very limited role of the output gap in determining the real contract wage. Nonetheless, this simulation demonstrates that under the right circumstances, the opportunist will be tougher than either of the conventional policymakers.

4.5 A recession

We complete this set of illustrative scenarios by considering the impact of an unanticipated shortfall in aggregate demand. Specifically, we assume an initial position of 4 percent inflation, and we induce a recession by feeding a sequence of two negative shocks to the equation for fixed investment.[10] Figure 6 summarizes the results.

Like her other two colleagues, the Taylor-style policymaker fails to anticipate that the economy is about to plunge into recession. Accordingly, she sets the short-term interest rate in the first period as high as she did in the simulations summarized in Figures 2 and 3—high enough, indeed, to induce some disinflation even in the absence of any shocks. Once she recognizes the occurrence of the negative demand shocks, however, the Taylor-style policymaker cuts the short-term nominal interest rate quite rapidly. Even so, the ensuing recession is substantially deeper than it was in the no-shock scenario: After two quarters, output is more than 3 percent below potential.[11] Not surprisingly, inflation comes down more quickly than in the no-shock scenario; by two years into the simulation, inflation has declined to 2.9 percent, compared with 3.2 percent in the no-shock scenario. Once the negative demand shocks are no longer hitting the economy and the easier stance of monetary policy has begun to take effect, output recovers fairly rapidly, and by the end of the second year of the simulation, the recovery is essentially complete. Inflation nonetheless continues on a downward trajectory, mainly reflecting the momentum of the expectations mechanism.

Under the Henderson-McKibbin rule, the adverse impact of the recession on output is met with an aggressive reduction in the short-term nominal interest rate. As a result, the early phase of the recession is a bit less severe than under Taylor's rule. The relatively small differences between the output paths under the two conventional rules translate into minute discrepancies between inflation trajectories, especially in the early years of the simulation.

[10] Note that these illustrative shocks have a greater impact on inflation than the average demand shock based on the estimated model. There are two reasons for this, namely that the shocks are for two consecutive quarters (whereas our estimated shocks are uncorrelated), and that the shocks are to fixed investment – the type of demand shock with the greatest impact on inflation in the model.

[11] Relative to historical experience in the United States, this is a mild recession. The output gap at the depth of the 1982 recession was about 8 percent.

On the whole, the Henderson-McKibbin policymaker appears to have achieved a relatively favorable tradeoff in this simulation between output stabilization and inflation stabilization.

The opportunistic policymaker responds to the onset of the recession as vigorously in terms of the short-term nominal rate as does the Henderson-McKibbin policymaker. In terms of the short-term *real* rate, her response is markedly more vigorous, because inflation declines under the opportunistic rule at a more pedestrian pace. The output gap under the opportunistic rule is smaller than under either conventional rule, but at least to the naked eye the gain in output stabilization looks relatively meager compared to the loss in inflation reduction, especially compared to the experience of the Henderson-McKibbin policymaker.

5 Transition and stabilization in a stochastic economy

Under either conventional monetary policy rule, the model we have been working with is linear. Therefore, given an initial level of inflation (say 4 percent), it would be straightforward with either of these rules to calculate the expected time until inflation is within some neighborhood of the long-run target. In fact, the calculation could be done exactly as in Figure 2, by setting all future shocks equal to their expected value of zero and simulating the model. In contrast, the model under the opportunistic policy rule is not linear, so the same "certainty equivalence" approach to calculating the expected time to arrival at some particular level of inflation cannot be taken.

We answer this question by conducting stochastic simulations. In preparation for those simulations, we first computed the structural residuals of the model presented in Section 3 based on U.S. data from 1980 to 1993.[12] Then we calculated the covariance matrix of those structural residuals. Using this covariance matrix, we generated artificial normally-distributed shocks and conducted 1,000 stochastic simulations of the opportunistic policy rule.[13] After each simulation, we recorded the time until inflation was brought within some

[12]The process of calculating the structural residuals would be straightforward if the model in question were a purely backward-looking model. For a rational expectations model, however, structural residuals can be computed only by simulating the full model based on historical data. The structural shocks differ from the estimated residuals to the extent of agents' forecast errors.

[13]For the linear rules we conducted 10,000 simulations to check whether the mean of the stochastic

neighborhood of the long-run inflation target. Then, as shown in section 5.1, we calculated the distribution of these arrival times. We use the same general methodology in section 5.2 to determine the distributions of inflation and output in the stochastic steady state. In section 5.3 we show the relative efficiency of the alternative rules in reducing the volatility of inflation and output.

5.1 Time to disinflation

The upper panels of Figure 7 show the probability distributions of the time taken to disinflate from 4 percent to 3 percent under each of the three monetary policy rules we have been studying. The horizontal axis indicates the number of quarters taken to reach an inflation rate of 3 percent (left panel) or 2 percent (right panel). The probability densities shown are smoothed versions of the discrete distributions generated by the simulations.

The densities to reach 3 percent inflation have modes at about 10 quarters for all three models. The Henderson-McKibbin rule has the greatest concentration of probability mass around the 10-quarter mark, with the Taylor rule not far behind. As expected, the opportunist policymaker acts more leisurely to bring inflation down to 3 percent–not only is the mode number of quarters slightly higher, but the upper tail of the distribution for opportunism has more probability mass than in the case of the alternative rules. In the upper right panel, the distribution of the time taken to disinflate from 4 percent to 2 percent has a mode of about 15-20 quarters, with opportunism again having a fatter upper tail.

The lower right panel shows the average (or expected) path of inflation for each rule. As shown, the expected path of disinflation under opportunism is noticeably slower than under the other rules. To reach an inflation rate of 3 percent, the expected time is about five years, while for the other two rules it is a bit less than three years.[14] The lower right panel plots the expected cumulative sum of the output gaps for each rule. As shown, at each point in time the output loss is slightly smaller for the opportunistic rule than for the

simulations converged to the deterministic "zero shock" path. Convergence was achieved.

[14]Still, this exercise shows that the opportunistic strategy will succeed in driving inflation down within a reasonable time period.

conventional rules. The vertical dotted lines match those in the lower left panel, indicating when each rule achieves an inflation rate of three percent. At those times, the output loss (and therefore the sacrifice ratio) is about the same for all three rules.[15]

5.2 Stochastic steady-state distributions of output and inflation gaps

A separate issue of interest concerns the distributions of inflation and output in the stochastic steady state. Our earlier analysis suggests that the distribution of the output gap should have more mass around zero under opportunism, and that the inflation distribution should be more diffuse within the opportunistic range, though not necessarily more diffuse outside that range. Our objective in this set of simulations is to verify that analysis, and to quantify the tradeoffs available to the policymaker.

Figure 8 presents density distributions in which there are only supply shocks (left panels) and only demand shocks (right panels).[16] For each type of shock, the distribution of inflation is shown in the top panel and the distribution of the output gap in the lower panel. The supply shocks reveal the dominant feature of opportunism most clearly–opportunism lets inflation vary from its long-run target more than do the other rules. But also, the distribution of inflation under opportunism is not a normal distribution–even though the shocks are normally distributed. Near the long-run target, the inflation distribution is flatter than a normal distribution because, when inflation is in the opportunistic region, the opportunistic policymaker does not actively try to bring inflation back to target. Rather, she attempts to keep output at potential, as shown in the lower left panel by the concentration of probability at the output gap at zero. The conventional rules keep inflation more tightly concentrated about its target, but at the expense of a wider distribution for output.

As shown in the right panels of Figure 8, the distributions derived with demand shocks

[15] Of course this need not be the case in a model with rational expectations such as the one employed here. Rather in such models the sacrifice ratio depends on the policy rule specification as well as other structural parameters. Consequently, the convenient similarity observed here may be sensitive to the specific parameter values.

[16] The densities under opportunism are based on 1,000 stochastic simulations of the nonlinear model, while the unconditional standard deviations under the linear rules are computed analytically using the method developed Swamy and Tinsley (1980). Since the shocks are drawn from a normal distribution and the model with the conventional rules is linear, the unconditional distributions of inflation and output are also normal.

differ less across alternative rules than do the distributions of supply shocks. In particular, demand shocks do not have much effect on inflation under any rule. The more aggressive Henderson-McKibbin rule clearly dominates the Taylor rule with respect to output and inflation stabilization. The former rule is the more appropriate reference point for opportunism because it responds in the same way to output gaps. Opportunism has a somewhat wider distribution of inflation than the Henderson-McKibbin rule and only a slightly narrower one for output.

The reason why the outcomes with demand shocks do not look that much different across the policy rules while those with supply shocks do is examined in Figure 9. The upper left panel shows the response of the output gap over time to a temporary, stimulative demand shock, while the upper right panel is the response of inflation to the same shock.[17] The simulation exhibits two surprising properties. The first one is that even though output is pushed one percent above potential, there is very little inflationary impact – inflation goes up by little more than one-tenth of a percentage point under all rules. The second reason is that the conventional rules do not need to create much of a shortfall in output to return inflation to target. Instead, they return output to potential nearly as fast as does the opportunistic rule, with little need to overshoot. Part of the reason for the small inflationary impact of temporary increases in output is the weak response of the real contract wage to the output gap in the Fuhrer-Moore contract wage equation (12). More important however is the forward-looking nature of agents' expectations. In a purely backward-looking model, such as the model in Orphanides and Wilcox (1996) that is based on an accelerationist Phillips curve, an increase in inflation due to a temporary demand shock can only be offset by an output gap of the same size but opposite sign. This does not hold anymore in a model where expectations are formed in a forward-looking manner. In this case, wage-setters incorporate the policymaker's inflation target as well as long-run potential output in forming their expectations about future inflation and output. Consequently, they expect

[17]Fixed investment is shocked by 0.25 percent (of output), which amounts to about two standard deviations.

inflation to return to target and output to return to potential. These expectations accelerate the return of inflation to target after a purely temporary output deviation, and reduce the need for opening up an output gap in order to bring inflation back.

As shown in the lower panels, the case is much different for temporary supply shocks (i.e., shocks to the contract wage equation). They have a direct and persistent effect on inflation and leave the policymaker with the difficult problem of trading off inflation and output losses. The conventional policymakers drive output below potential. The opportunistic policymaker however keeps output at potential as long as inflation stays in the opportunistic zone.

In simulations with the mixture of supply and demand shocks reflecting those we estimate to have hit the U.S. economy over the past decade or so, the distributions of inflation and the output gap are as shown in Figure 10. The effect of supply shocks producing a more diffuse distribution of inflation under opportunism shows through in the upper panel. However, the "reward" to the opportunistic policymaker appears to be mild in that, as the lower panel shows, the distribution of the output gap about zero is only slightly more concentrated for her than it is for the Henderson-McKibbin policymaker. This is because most of the dispersion of the output gap under all rules occurs in response to demand shocks, and that dispersion is not much different under opportunism and the conventional policy rules (lower right panel of Figure 8).

5.3 Output and inflation variability tradeoffs available to opportunistic and conventional policymakers

Having provided a detailed analysis of the properties of three specifically parameterized policy rules, we systematically explore alternative parameterizations of these rules. To focus the discussion, we concentrate on a comparison of the variability of inflation and output in steady state. In the long run, the fundamental tradeoff implicit in the choice among alternative policy rules is with regard to the relative variability between inflation and output. Within a class of policy rules, this tradeoff can be illustrated by identifying the parameterizations of the rules which yield the best combinations of inflation and output variability—that is a frontier which identifies the smallest possible variability of inflation consistent with a given variability of output and vice versa.[18]

For the family of conventional policy rules

$$i_t^s = r^* + \pi_{t-1} + \gamma_1 y_{t-1} + \gamma_2(\pi_{t-1} - \pi^*),$$

which encompasses the Taylor and Henderson-McKibbin parameterizations, we construct a grid of the output and inflation response parameters, γ_1 and γ_2. We then compute the steady state distributions of inflation and output corresponding to each element of the grid. Throughout, we assume that the economy is hit with the demand and supply shocks from the distributions utilized in the simulations described in section 5.2.

The solid line in Figure 11 shows the resulting inflation-output variability frontier, using the standard deviation as a measure of variability. Points on this line identify the best combinations of variability of inflation and output that are feasible within this class of rules provided no restriction is imposed on the choice of response parameters, γ_1 and γ_2. For instance, it appears feasible in this model, to achieve a standard deviation of output as low as 0.6 percent consistent with a standard deviation of inflation of 1.2 percent. However, the policy response parameters γ_1 and γ_2 needed to achieve such outcomes are very large and imply policies with extremely volatile nominal interest rates.[19] Since, in practice, policy appears to indicate a preference towards low variability of interest rates, it is useful to construct variability frontiers restricted so that the implied interest rate variability not exceed certain bounds. Two such frontiers are shown in Figure 11, corresponding to standard deviations for the nominal interest rate of 2 and 3 percent respectively.[20]

Figure 11 also identifies the variability outcomes associated with alternative policy rules. As shown, the Taylor and the Henderson and McKibbin (HM) specifications fall just within

[18]The usefulness of comparing alternative policies in terms of such a frontier has been highlighted by Taylor (1994) and more recently by Fuhrer (1997) and Williams (1997).

[19]The implied standard deviations of interest rates are in excess of 10 percent.

[20]By comparison, the standard deviation of the federal funds rate since the end of the non-borrowed reserves targeting period in 1982 has been 2.3 percent.

the 2 and 3 percent restricted variability bounds respectively.[21] As can be seen, the restricted frontiers become quite flat to the right of the two rules. For instance, with respect to the Henderson and McKibbin specification, conventional rules tolerating considerably larger variability of inflation but restricted to achieve comparably low interest rate variability, provide only a marginal reduction in the variability of output. Such rules correspond to a smaller inflation response coefficient, γ_2, than the Henderson and McKibbin value of 1 but only a marginally higher output response coefficient, γ_1.

This observation becomes particularly relevant for comparing the Henderson and McKibbin specification with with our baseline specification of the opportunistic rule (Opp). As can be seen, measured in terms of standard deviations, the opportunistic rule yields about the same variability of output as the Henderson and McKibbin specification but greater variability of inflation. In our benchmark, the opportunistic rule has the same output response coefficient as the Henderson and McKibbin specification but assigns no penalty to inflation deviations from target within the opportunistic region. Over this region, therefore, the opportunistic rule resembles a conventional rule with a smaller inflation response coefficient which, as indicated above would be expected to yield about the same standard deviation of output as the Henderson and McKibbin specification but a larger standard deviation of inflation. Although the penalty is proportionally more severe than the one corresponding to the Henderson and McKibbin specification when inflation deviates substantially from the target, this effect is reflected mainly in the higher moments of the distribution of inflation (as negative excess kurtosis).

Next we turn to the three variants of the opportunistic rule briefly described in section 2. In all three cases, we modify the policy reaction function by considering alternative responses to inflation but leaving the response to output unchanged (and equal to the response in the Henderson and McKibbin specification). The first variant, recognizes that an opportunistic policymaker who does not control aggregate demand perfectly (as is indeed the case in

[21]Indeed, the implied interest rate variabilities for the two rules are 1.9 and 2.8 percent, and both rules would be on the corresponding restricted frontiers.

practice), will be sensitive to deviations of inflation from the intermediate target even when inflation is inside the opportunistic region. The extent of this sensitivity is determined, in part, by the degree of uncertainty faced by the policymaker and is difficult to quantify in practice. Qualitatively, however, the effect is a policy reaction function with some slope inside the opportunistic region but greater slope outside the region. The rule corresponding to the point denoted as *Opp B* in Figure 11 assumes an inflation response similar to the one corresponding to the Taylor specification of 0.5 inside the opportunistic region but the same response as our baseline specification of the opportunistic policy outside the region. The result is to steepen, on net, the response of policy to inflation which as shown in the Figure reduces the standard deviation of inflation while leaving the standard deviation of output essentially unchanged.

The other two variants of policies we examine employ the same inflation penalty functions corresponding to the two variants of opportunistic policies but with a fixed inaction zone. Such policies, essentially describe policy makers who actively pursue inflation stabilization when inflation falls outside a target zone but allow for flexibility towards short run output stabilization when inflation is within the target zone.[22] With our parameterization, in these two variants we effectively model a policy rule with an inflation target zone two percentage points wide. Such a policy differs importantly from an opportunistic policy when inflation is above the target zone. In that case, the inflation target zone policymaker will appear more similar to a conventional policymaker as she will deliberately disinflate regardless of the recent history of inflation, whereas an opportunistic policy maker would tend to wait for favorable disinflation shocks. Fundamentally, however, inflation zone targeting exhibits an asymmetric response to shocks when inflation is close to its long run target which is inconsistent with conventional policy rules.[23]

[22] Bernanke and Mishkin (1997) identify the desire for this flexibility as an important reason why most central banks associated with "inflation targeting" specify a *range* rather than a *point* target.

[23] That is to say, a policymaker with an inflation zone target of one to three percent, may appear to behave like a conventional policymaker with a target of three percent as long as inflation exceeds three percent. It would be incorrect and obviously quite misleading to employ a conventional policy reaction function with a target of three to describe the steady state properties associated with such an inflation zone targeting regime.

The effect of shifting from an opportunistic inflation penalty to an inflation zone targeting one represents, on net, a more vigorous response of policy to inflation deviations from the long run target. In Figure 11 the points labeled *Zone* and *Zone B* indicate the standard deviations of the two policies, corresponding respectively, to a flat and a sloped penalty inside the inflation zone. As shown this reduces the standard deviation of inflation but also raises slightly the standard deviation of output.

6 Conclusion

Broadly speaking, the results of this paper can be summarized as follows: First, we find that disinflation under an opportunistic rule proceeds more slowly, on average, than under a conventional linear rule. At the same time, however, output losses along the way toward meeting the inflation target are commensurately smaller. On balance, the sacrifice ratio is about the same for the opportunist as it is for the conventional policymaker.

Second, once the steady state has been achieved, the relative experience of the opportunist compared to the conventionalist depends greatly on the distribution of shocks. In the wake of supply shocks, the distribution of the output gap is substantially more concentrated if the opportunist is at the monetary helm than if the conventionalist is in charge; on the other hand, the distribution of inflation around the long-run target is more diffuse. In the wake of demand shocks, the opportunist gains little in terms of output stabilization, given our model of the economy, while suffering some increase in inflation dispersion.

When we simulate the model with both demand and supply shocks, we find that the resulting steady-state inflation distribution is somewhat more diffuse under the opportunistic rule than it is under the conventional rule, while the distribution of the output gap is about the same. This result makes sense because our model estimates imply that the bulk of the variation in output during the 1980s and 1990s mainly derived from demand shocks, while the bulk of the variation in inflation reflected supply shocks. As a result, the combined distribution for inflation to a close approximation can be inferred from the supply-shocks-only simulations, while the combined distribution for the output gap can be inferred from the demand-shocks-only simulations.

Third, we examine the efficiency of inflation targeting compared to both the conventional and opportunistic policy rules in stabilizing output and inflation. We find that, for the parameterizations we explore, the inflation "zoner" falls about midway between the conventionalist and the opportunist: He achieves about the same degree of output stabilization as does either alternative, and fares somewhat better on the inflation front than the opportunist, though somewhat worse than the conventionalist.

Finally, we emphasize that our results shed no light on the issue of whether an opportunistic policy is as credible as a conventional policy: In every simulation we conduct here, private-sector agents are assumed to know both the specification and the parameterization of the monetary policy rule, and are assumed to believe that it will remain in force in perpetuity. A separate question, worthy of future study, is whether an opportunistic policymaker might suffer from credibility problems because—under certain circumstances—she is observed failing to take anti-inflationary actions even when inflation remains above her long-run target.

References

Anderson, Gary S. (1987) "A Procedure for Differentiating Perfect-Foresight-Model Reduced-Form Coefficients" *Journal of Economic Dynamics and Control* 11(.) 465-81.

_______ and George R. Moore (1985) "A Linear Algebraic Procedure For Solving Linear Perfect Foresight Models" *Economics Letters* 17(.) 247-52.

Bernanke, Ben S., and Frederic S. Mishkin (1997) "Inflation Targeting: A New Framework for Monetary Policy?" *Journal of Economic Perspectives* 11(2) 97-116.

Boucekkine, Raouf (1995) "An Alternative Methodology for Solving Nonlinear Forward-Looking Models," *Journal of Economic Dynamics and Control* 19(4) 771-734.

Brayton, Flint, Andrew Levin, Ralph W. Tryon, and John C. Williams (1996) "The Evolution of Macro Models at the Federal Reserve Board," presented at the Carnegie-Rochester Conference on Public Policy, November 1996.

Fair, Ray, and John B. Taylor (1983) "Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models" *Econometrica* 51(.) 1169-85.

Fuhrer, Jeffrey C. and George R. Moore (1995a) "Inflation Persistence" *Quarterly Journal of Economics* 110(1) (February) 127-59.

_______ and _______ (1995b) "Monetary Policy Trade-offs and the Correlation between Nominal Interest Rates and Real Output" *American Economic Review* 85(1) 219-39.

Fuhrer, Jeffrey C. (1997), "Inflation/Output Variance Trade-Offs and Optimal Monetary Policy," *Journal of Money, Credit, and Banking*, 29(2) 214 - 234.

Henderson, Dale and Warwick J. McKibbin (1993) "A Comparison of Some Basic Monetary Policy Regimes for Open Economies: Implications of Different Degrees of Instrument Adjustment and Wage Persistence", Carnegie-Rochester Conference Series on Public Policy (39) 221-318.

Juillard, Michel (1994) "DYNARE - A Program for the Resolution of Non-linear Models with Forward-looking Variables," Release 1.1, mimeo, CEPREMAP.

Laffargue, Jean-Pierre (1990) "Résolution d'un modèle macroéconomique avec anticipations rationnelles," *Annales d'Economie et de Statistique* 17, 97-119.

Leiderman, Leonardo, and Lars E. O. Svensson, eds., *Inflation Targets.* London: Centre for Economic Policy Research, 1995.

Levin, Andrew (1996) "A Comparison of Alternative Monetary Policy Rules in the Federal Reserve Board's Multi-Country Model," BIS Conference papers Vol. 2, pp. 340-366.

Orphanides, Athanasios and David W. Wilcox (1996) "The Opportunistic Approach to Disinflation," Finance and Economic Discussion Series, 96-24, Board of Governors of the Federal Reserve System, May.

Svensson, Lars E.O. (1996) "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," NBER Working Paper Series No. 5797.

Swamy, P.A.V.B. and Peter A. Tinsley (1980), "Linear Prediction and Estimation Methods for Regression Models with Stationary Stochastic Coefficients," *Journal of Econometrics*, (12), 103-142.

Rotemberg, Julio and Michael Woodford (1997), "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," mimeo, May 1997.

Taylor, John B. (1979) "Estimation and Control of a Macroeconomic Model with Rational Expectations", *Econometrica*, 47(5), 1267-86.

_______ (1980)"Aggregate Dynamics and Staggered Contracts" *Journal of Political Economy* 88(1) 1-23.

_______ (1993a) *Macroeconomic Policy in the World Economy: From Econometric Design to Practical Operation* New York: W.W. Norton.

_______ (1993b) "Discretion versus Policy Rules in Practice" *Carnegie-Rochester Conference Series on Public Policy* 39 195-214.

_______ (1995) "Monetary Policy Implications of Greater Fiscal Discipline," in *Budget Deficits and Debt: Issues and Options*, Federal Reserve Bank of Kansas City, 151-70.

_______ and John C. Williams (1993) "Forecasting with Rational Expectations Models" mimeo, Board of Governors of the Federal Reserve System.

Williams, John C. (1997) "Simple Rules for Monetary Policy" mimeo, Board of Governors of the Federal Reserve System.

Figure 1

The Interest Rate Penalty for Inflation

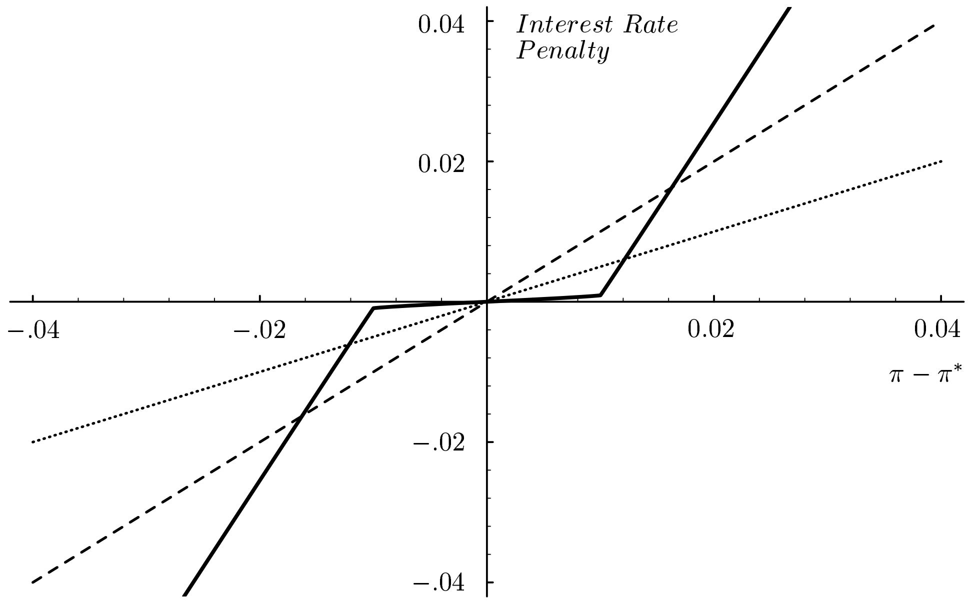

Note: The dotted and dashed straight lines depict the interest rate penalty levied by conventional policy rules with slopes 0.5 and 1, respectively. The piecewise linear solid line describes the penalty levied by the opportunistic rule. The flat segment of the latter schedule corresponds to what we refer to in the text as the "opportunistic region," in which the monetary authority aims to hold output at potential. The penalty slope outside this region equals 2.5.

Figure 2

Disinflation With No Shocks

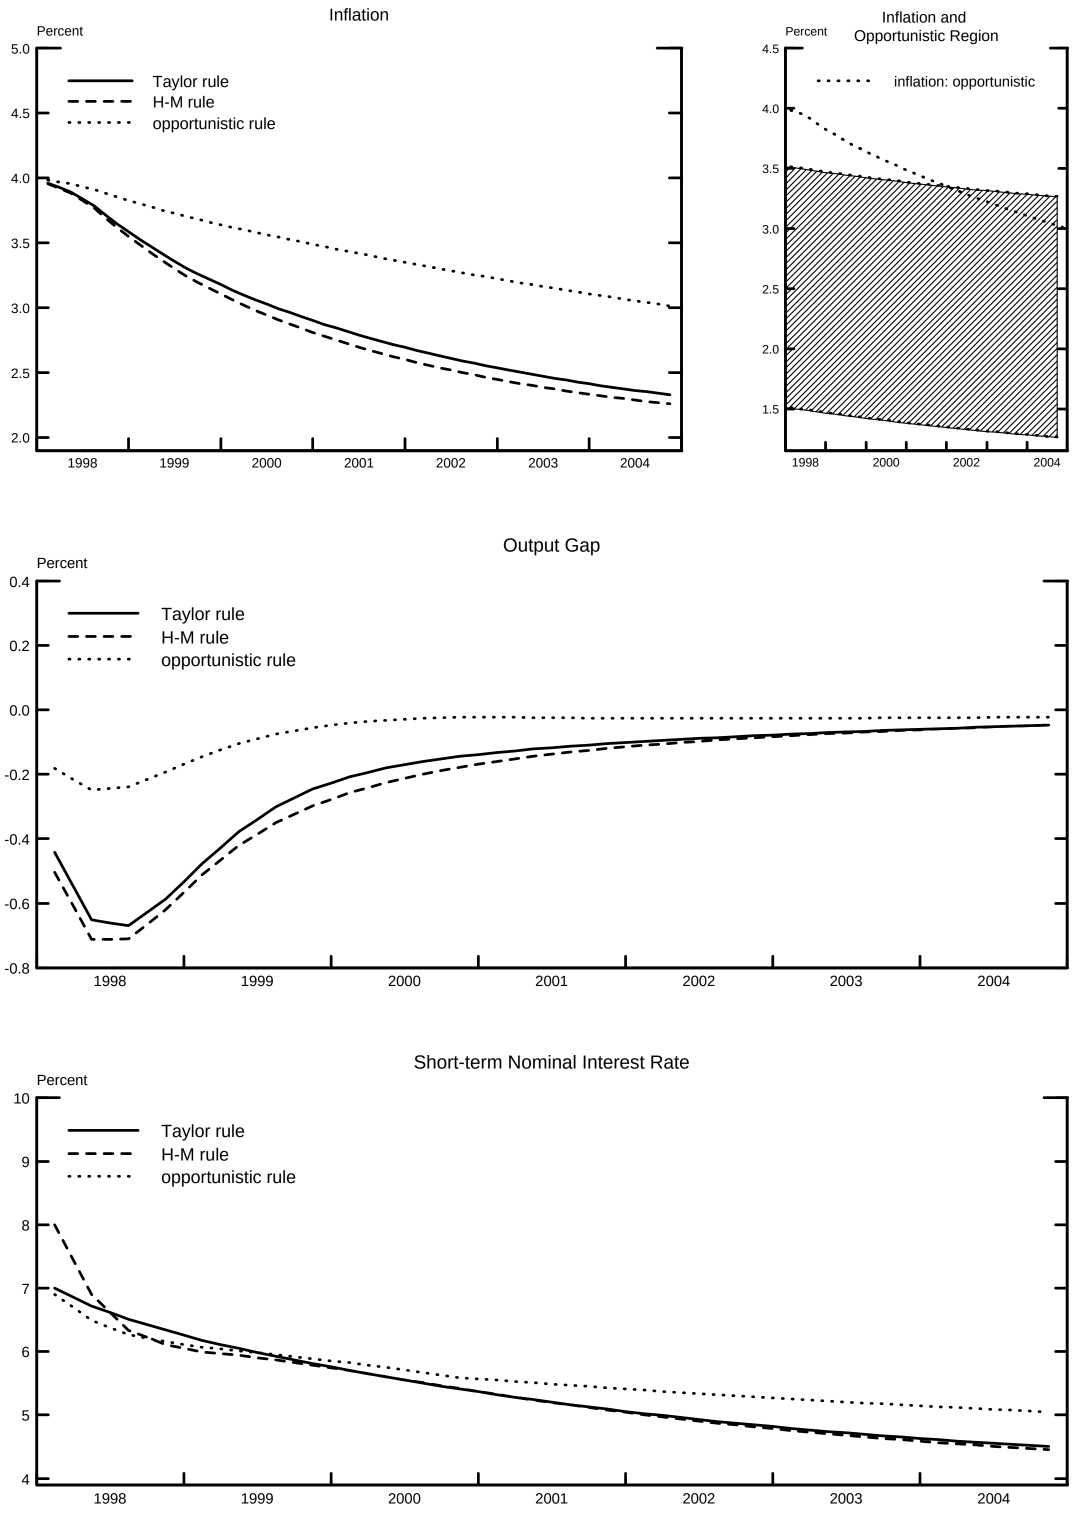

Figure 3

A Favorable Supply Shock: Scenario I

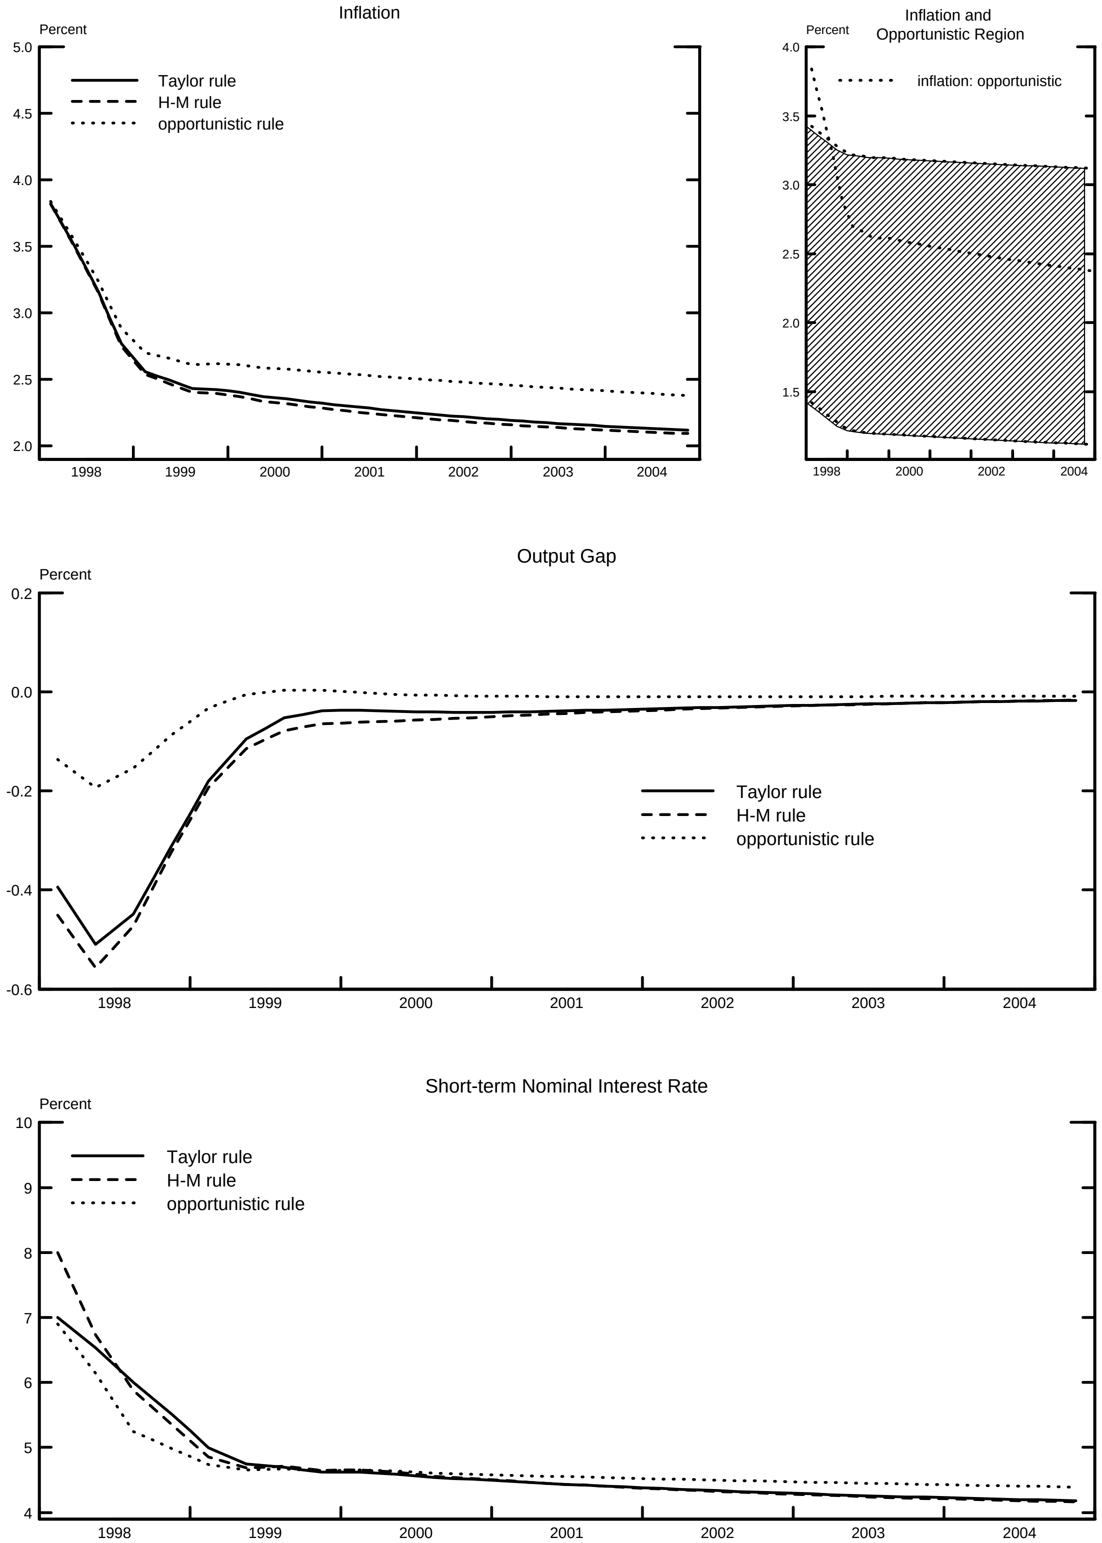

Figure 4

A Favorable Supply Shock: Scenario II

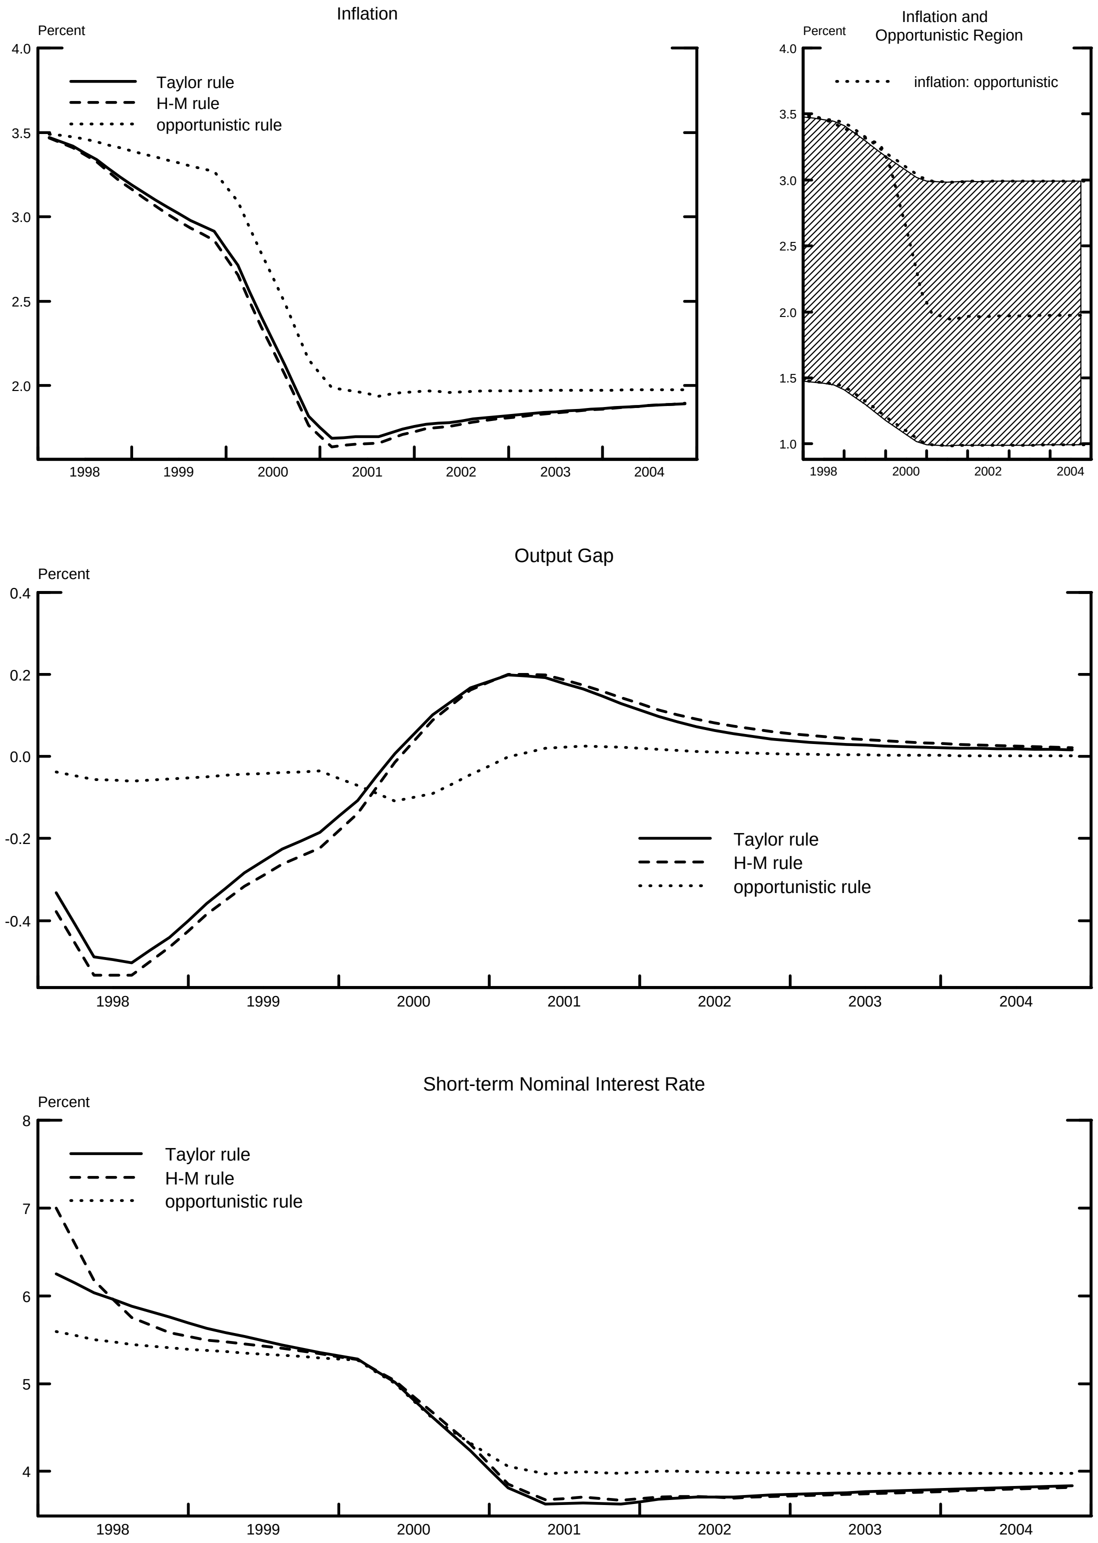

Figure 5

An Unfavorable Supply Shock

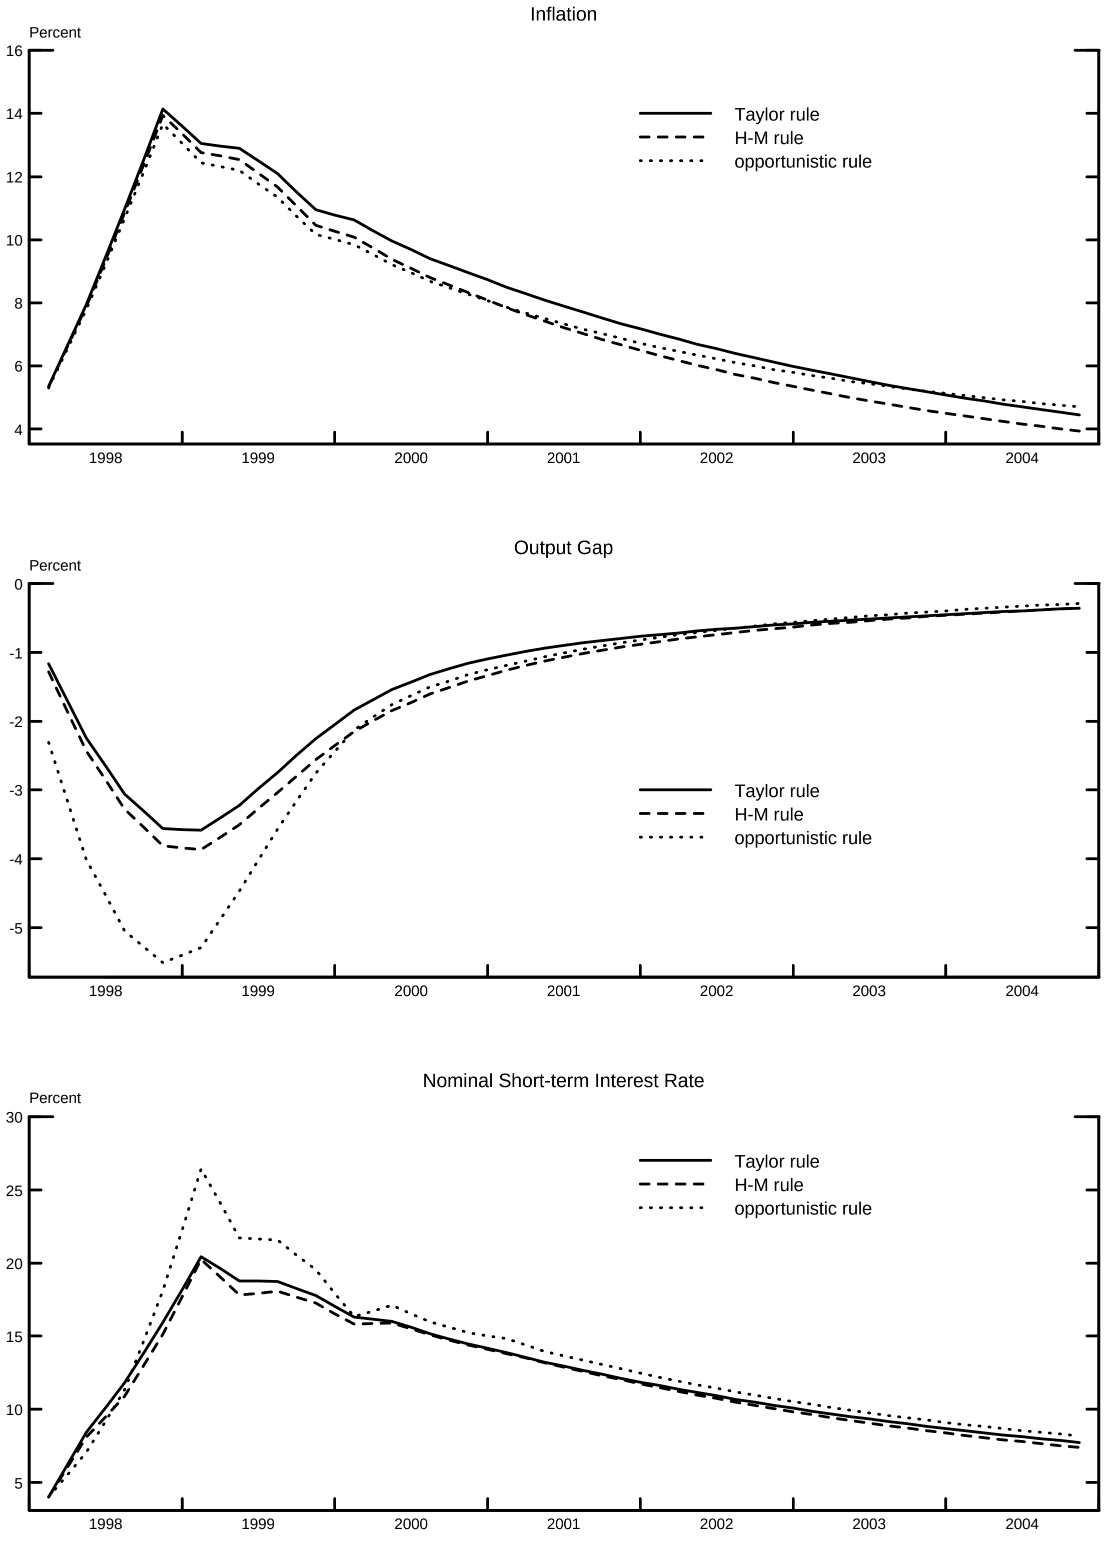

Figure 6

A Negative Demand Shock

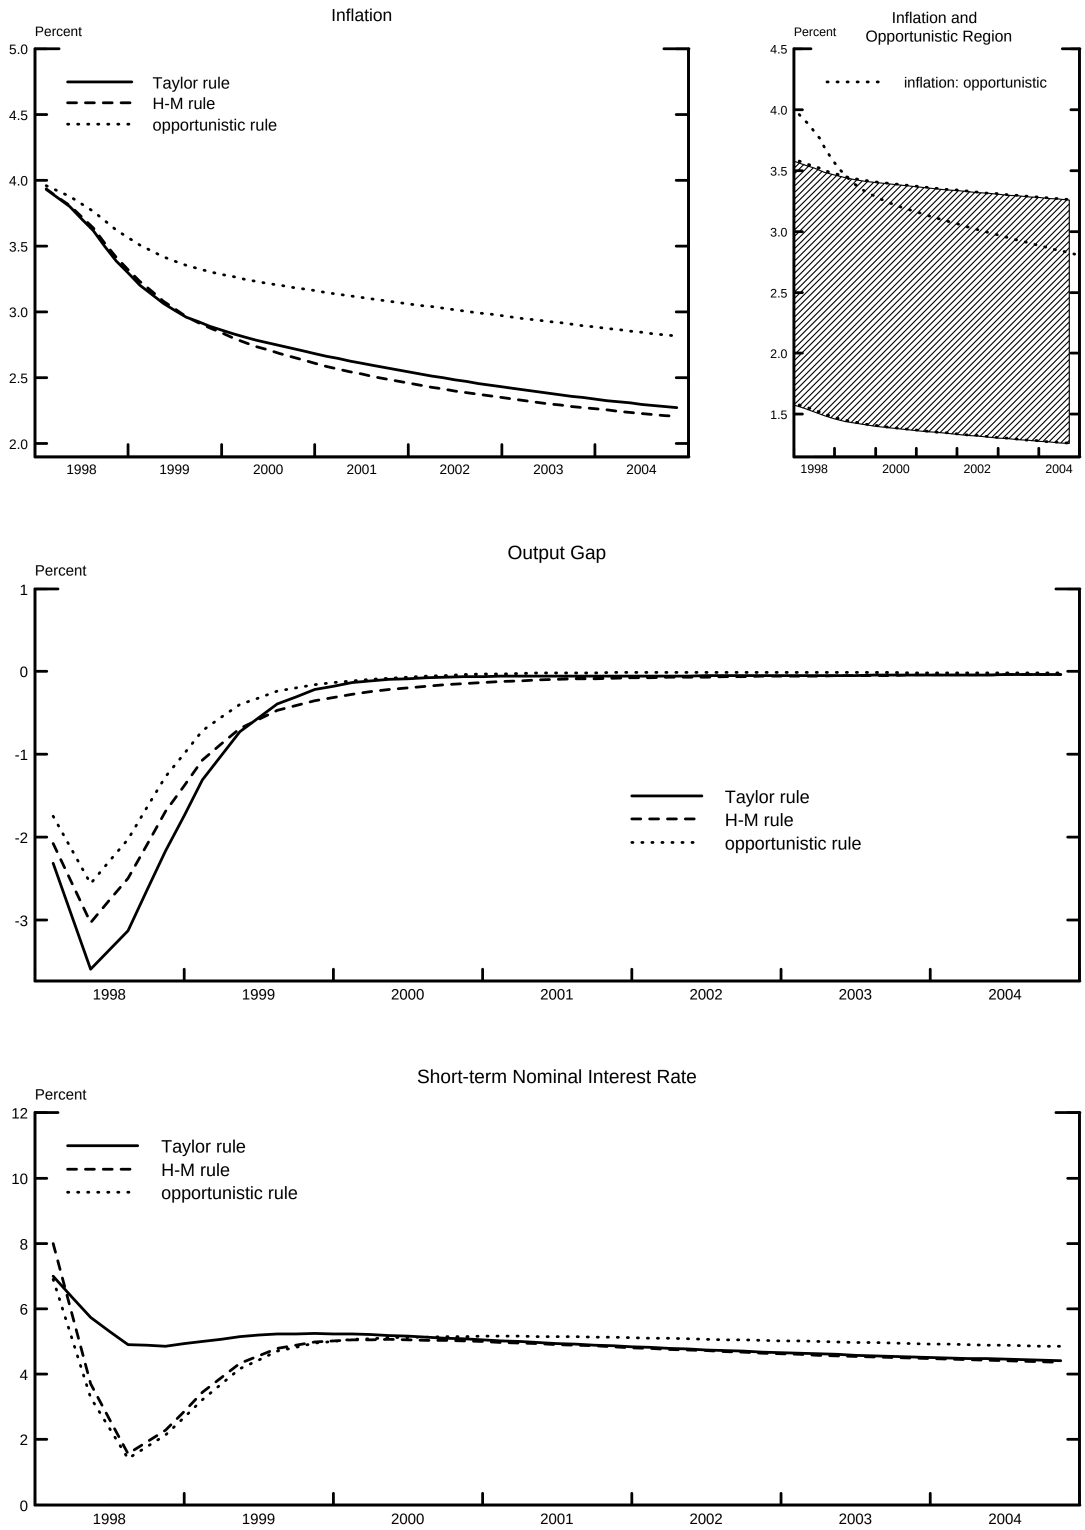

Figure 7

Distributions of Various Variables During a Disinflation

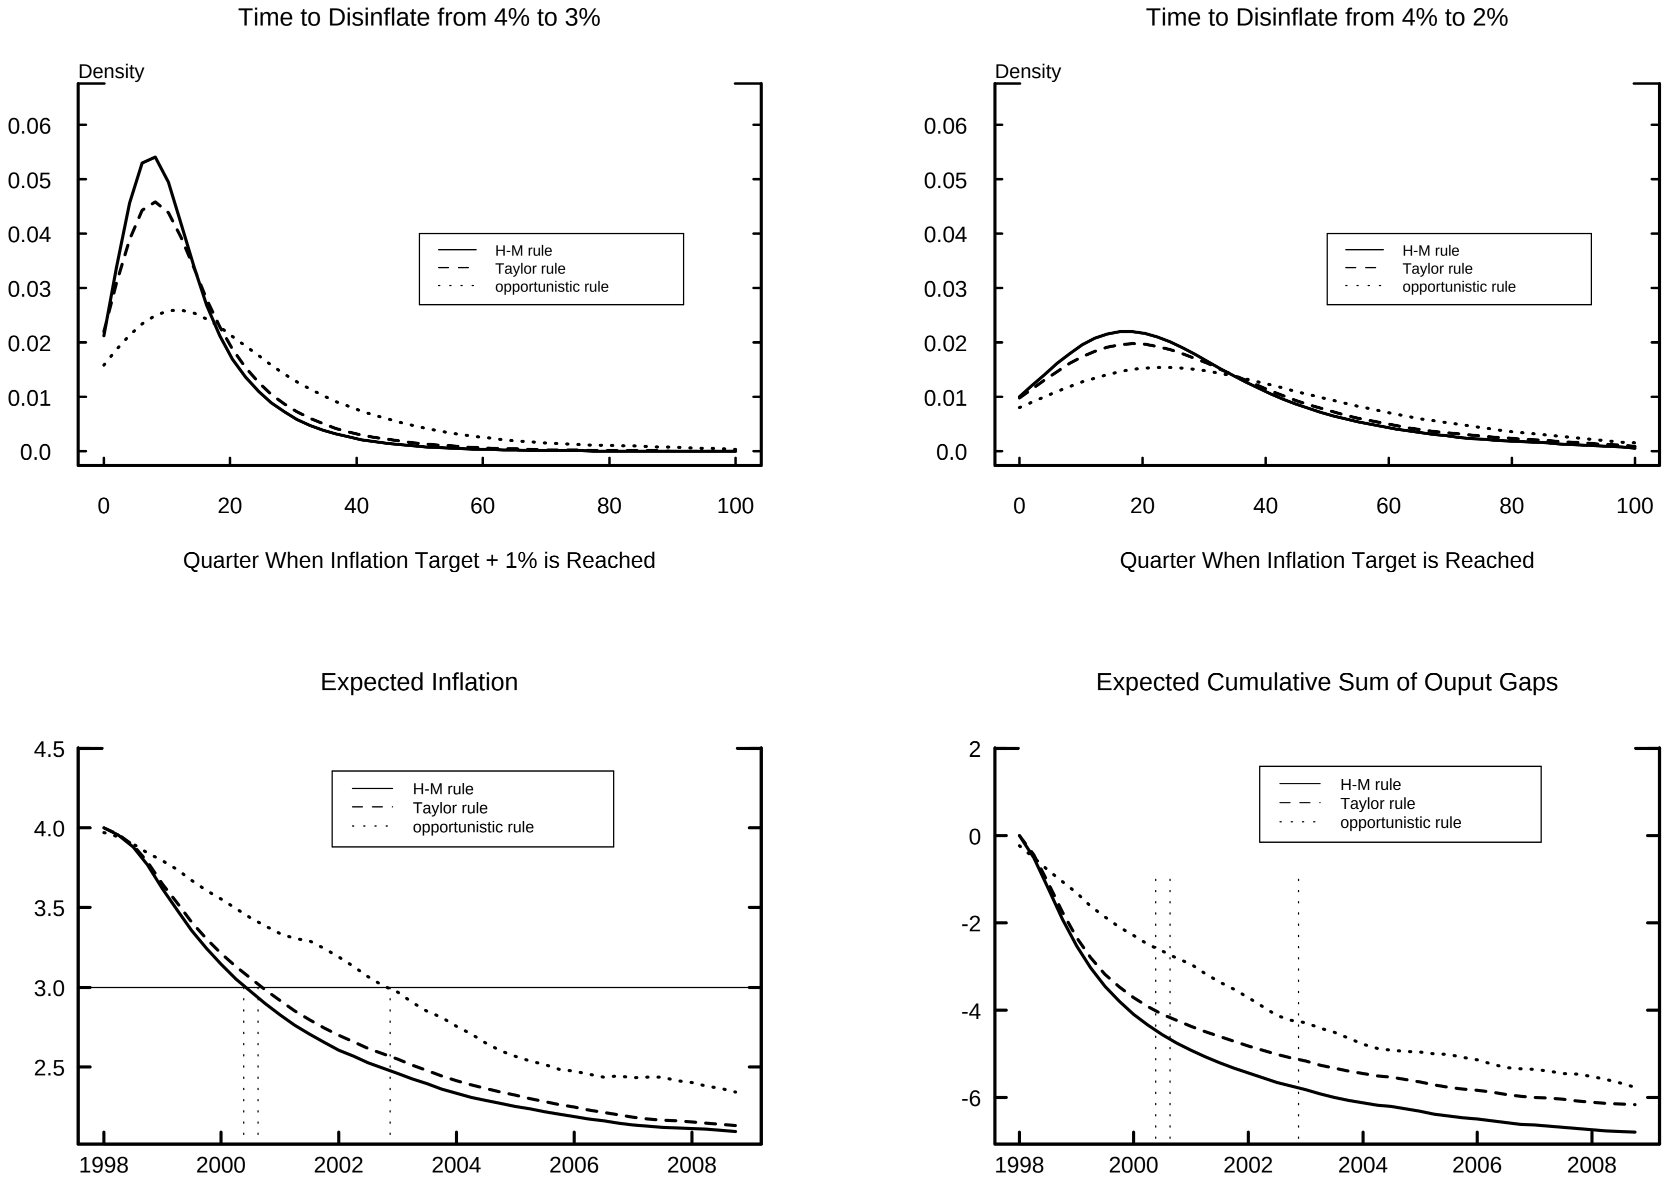

Figure 8

Steady-State Distributions: Supply and Demand Shocks Separately

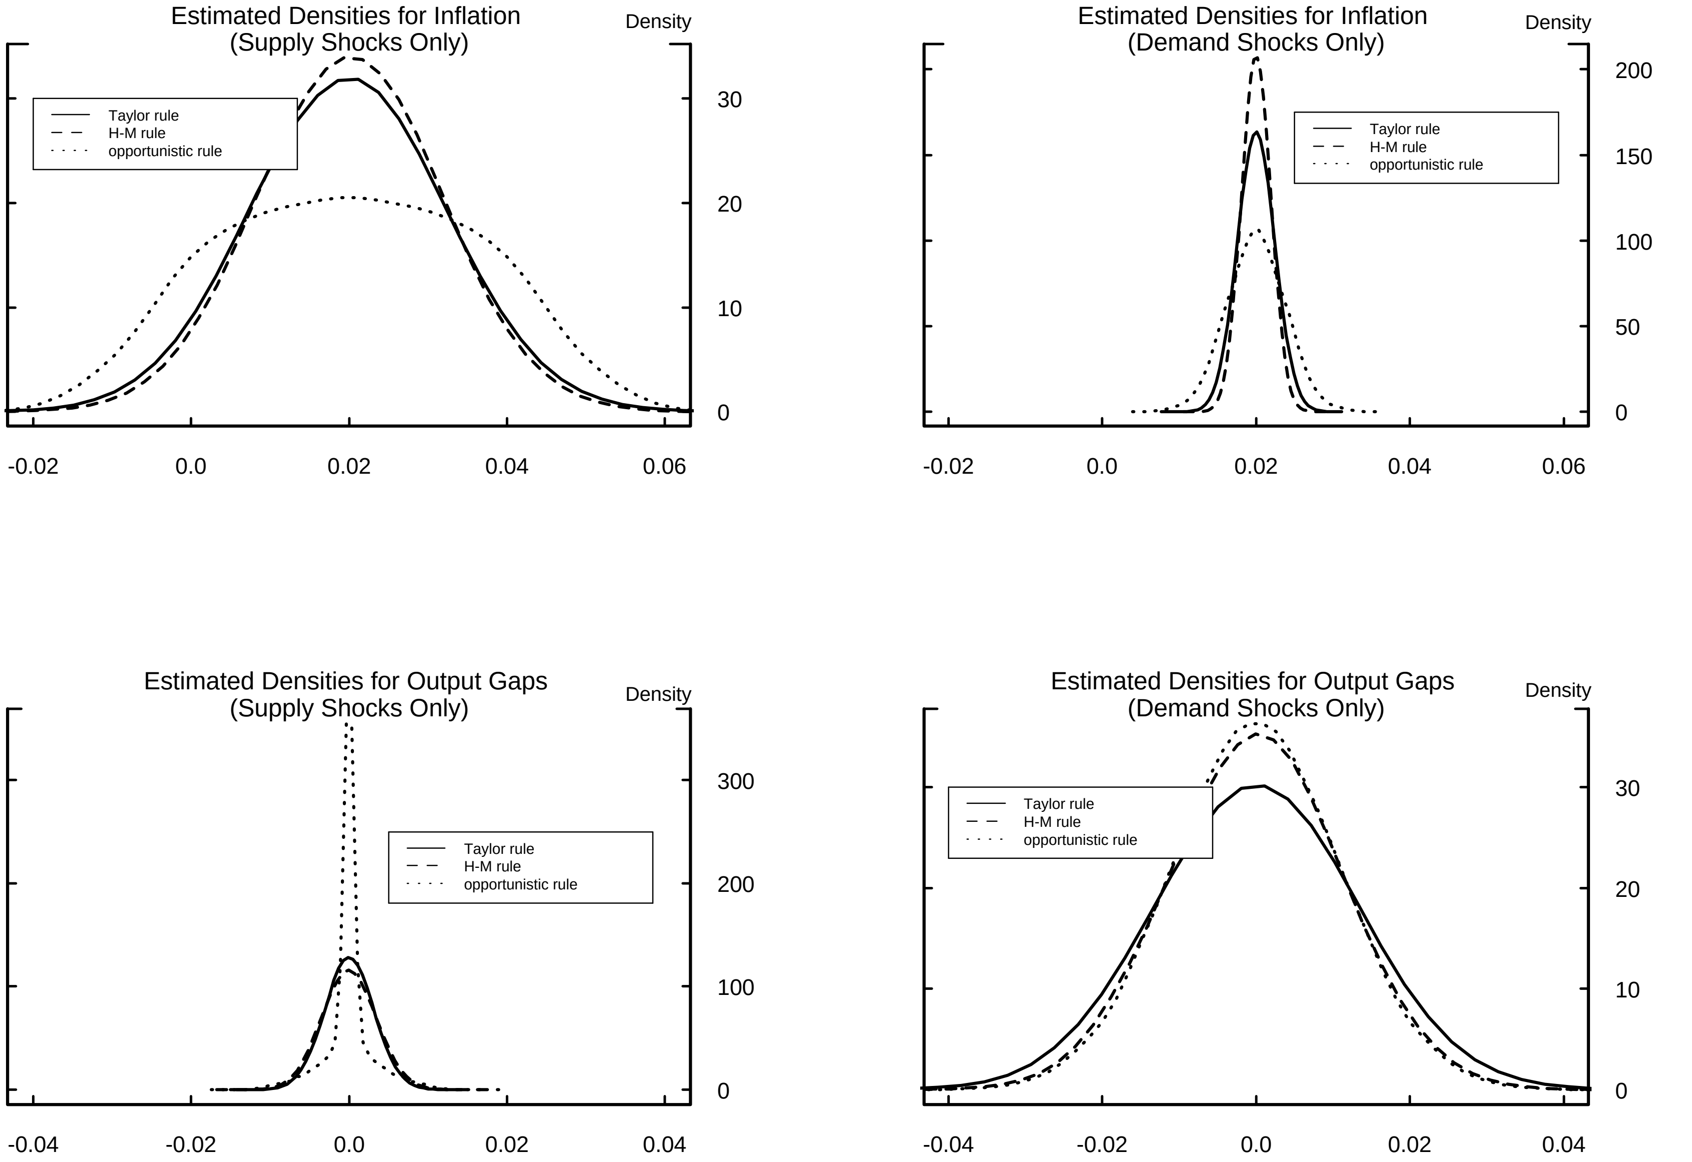

Figure 9

Consequences of Temporary Demand and Supply Shocks for Output and Inflation

Temporary Positive Demand Shock

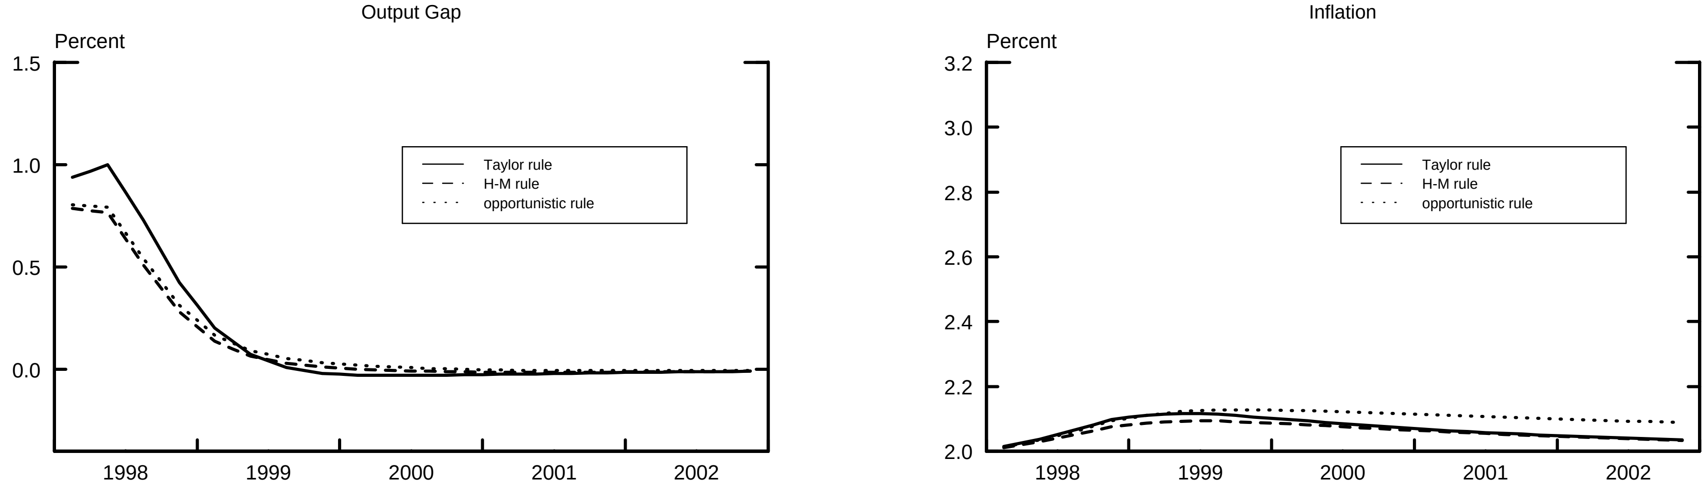

Temporary Adverse Supply Shock

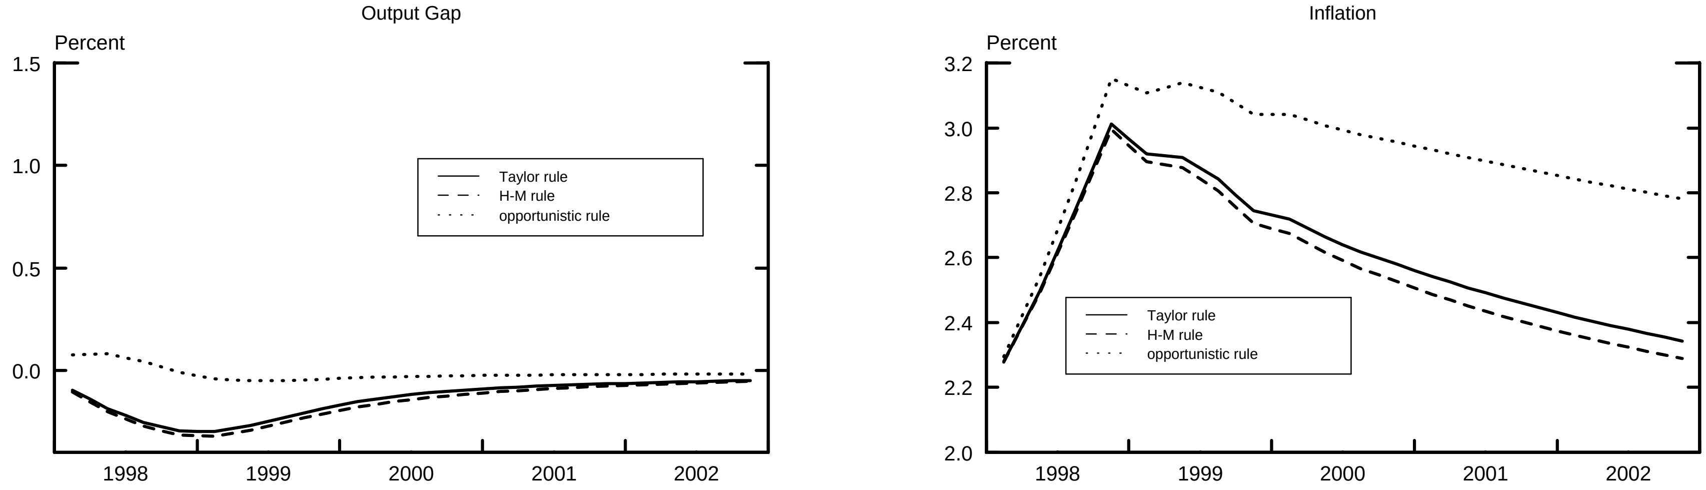

Figure 10

Steady-State Distributions: All Shocks

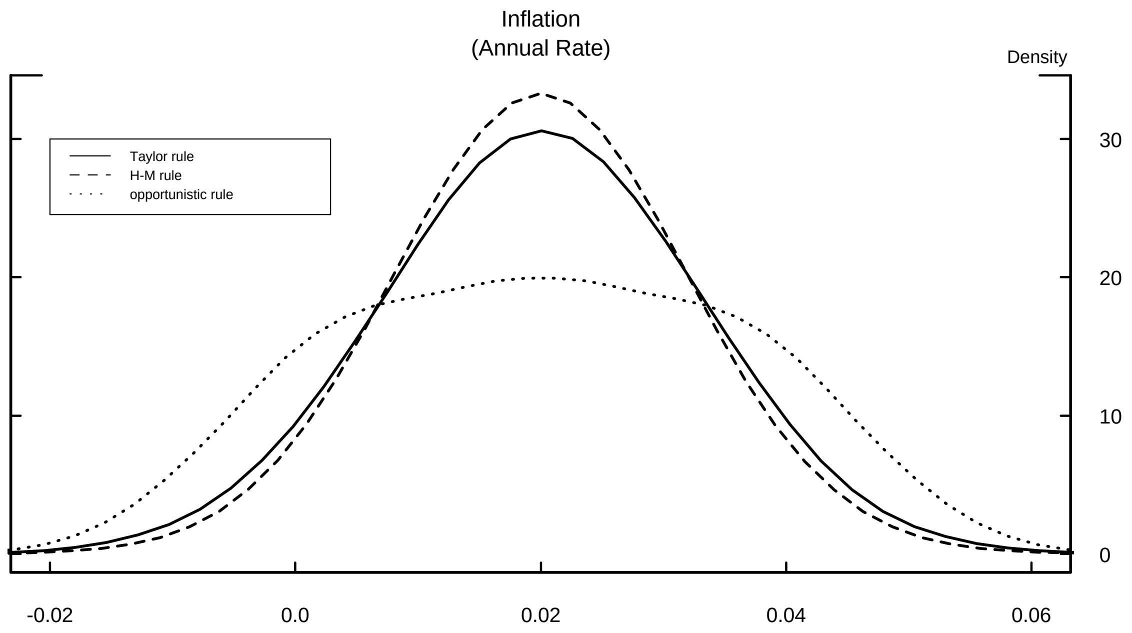

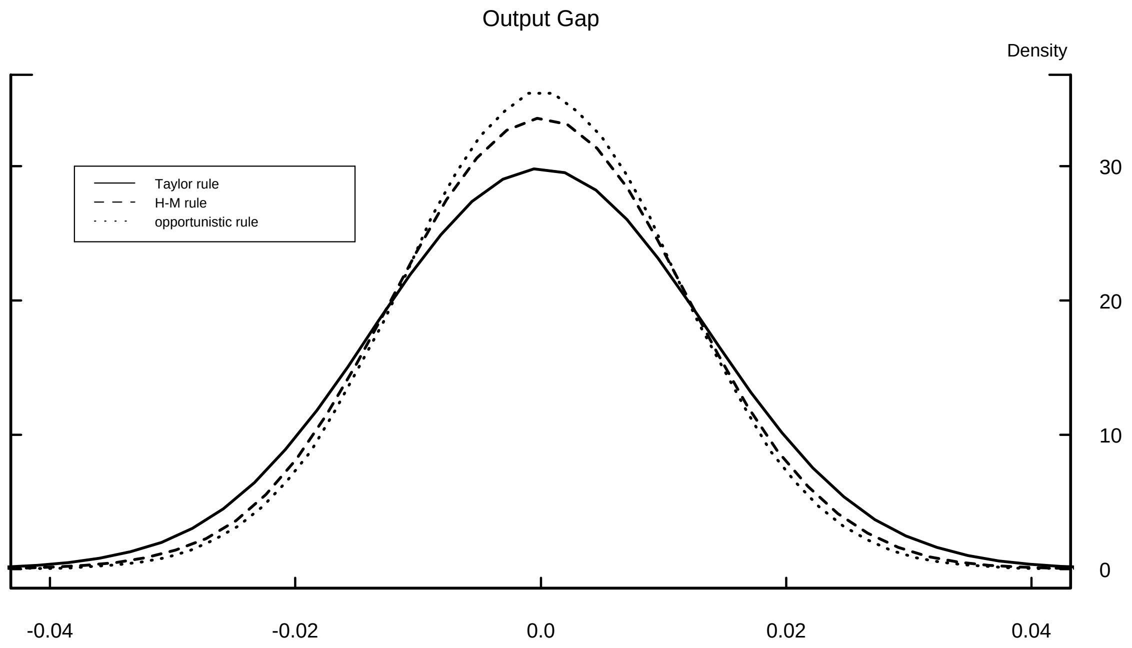

Figure 11

Variability of Inflation and Output in Steady State

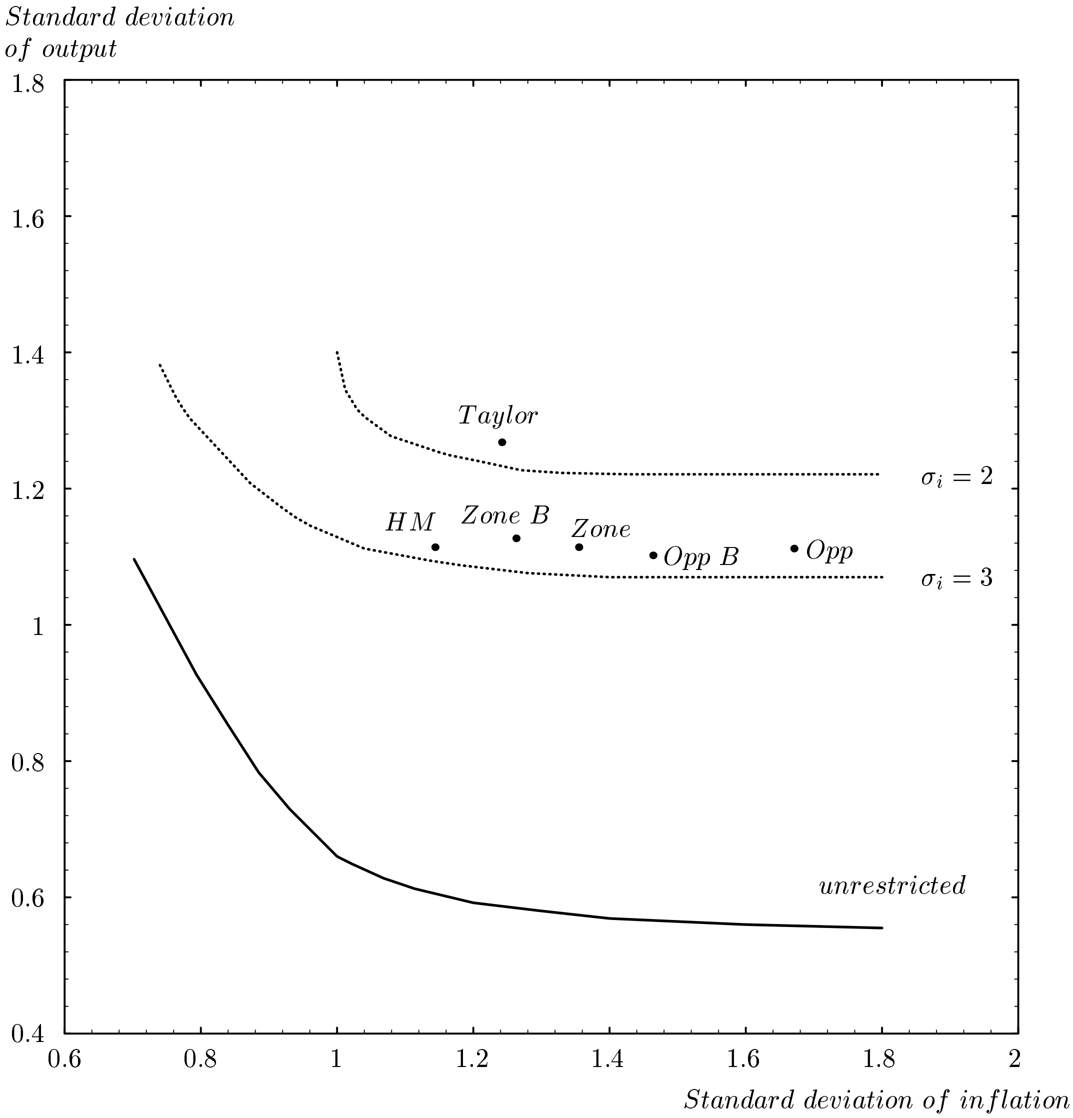

Note: The graph plots the standard deviation of inflation and the output gap corresponding to the alternative policy rules. The three lines indicate frontiers of outcomes for these standard deviations obtained by simulating the conventional rule with varying weights on the inflation and output. The solid line indicates the unconstrained frontier. The two dotted lines show frontiers constrained by policies for which the standard deviation of the federal funds rates is at most 2% and 3% respectively.