

Inflation Forecast Targeting: An Alternative Approach to Estimating the Inflation-Output Variability Tradeoff

James S. Fackler* and W. Douglas McMillin†

We suggest a new way of computing the inflation-output variability tradeoff under inflation forecast targeting. Our approach is based on dynamic, stochastic simulations of the average inflation rate over a two-year horizon using the moving average representation of a vector autoregressive (VAR) model. Using real-time data over two samples, we estimate the inflation-output variability tradeoff for the United States and show that it has shifted favorably over time. We analyze the policy interventions required to achieve target inflation in each sample and compare these interventions over time.

JEL Classification: E52, E58

1. Introduction

Over the past decade and a half, most major central banks around the world have adopted monetary policy frameworks that include either explicit or implicit inflation targets. The type of inflation targeting implemented by central banks is best characterized as flexible inflation targeting, where central banks are simultaneously concerned about the variability of output around its natural level as well as the variation of inflation around target. Furthermore, Svensson (1997) has argued that inflation targeting should be implemented as inflation forecast targeting because of lags in the effect of monetary policy on inflation and because of imperfect control of inflation by the central bank. Finally, Bernanke (2004), in comparing and contrasting use of “simple feedback policies” (instrument rules) and “forecast-based policies,” concludes that a forecast-based approach to policy implementation “has become increasingly dominant in the monetary policymaking of leading central banks.... [T]he Fed relies primarily on the forecast-based approach for making policy.”¹

Following Taylor’s (1994) suggestion for estimating the inflation-output variability tradeoff for a variety of policy rules, recent research includes estimation of the tradeoff between inflation variability and output variability. This research investigates policy rules in the presence of a loss function for the central bank in which there is often a substantial weight on the variation of

* Department of Economics, University of Kentucky, Lexington, KY 40506-0034, USA; E-mail JamesFackler@gmail.com.

† Department of Economics, Louisiana State University, Baton Rouge, LA 70803-6306, USA; E-mail eodoug@lsu.edu; corresponding author.

The authors thank two anonymous referees, Scott Frame, Ellis Tallman, Bill Lastrapes, Stephen Cecchetti, and John Rogers for helpful comments and suggestions.

Received July 2009; accepted November 2010.

¹ Svensson (2003) provides a detailed critique of instrument rules in an inflation-targeting setting.

inflation around target relative to the weight on output variability. In this literature, a common approach is to vary the weight on the inflation target, derive the associated instrument rule, simulate a structural model within which the rule is embedded, and then compute inflation and output variability for alternative values of the relative weight on the inflation target. For example, Rudebusch and Svensson (1999) estimate empirically the inflation-output variability tradeoff for a variety of rules.² For each rule, some of which can be categorized as inflation forecast targeting rules, the tradeoff is derived by varying the relative weight on output.

In contrast to the instrument-rule approach, we present a novel way of estimating the inflation-output variability tradeoff (hereafter “variability tradeoff”) that uses forecast-based counterfactual simulations. Our motivation is the real-world policy making process described by Svensson (2003) in which policy makers do not write down an explicit loss function but instead evaluate alternative paths for the policy instrument and then pick a path that, in Svensson’s terminology, “looks best.” Given that central banks don’t announce a specific loss function and the weights therein, and following an analogy presented in Svensson (1997), we consider inflation target bands of varying width as proxies for changes in the relative weight on inflation versus output stabilization in a loss function. In this approach, wider target bands are analogous to a relatively smaller weight on the inflation target. Following McCallum’s (1988) suggestion that “rules” should be evaluated in a variety of models, we illustrate our technique using a vector autoregressive (VAR) model³ and compute the policy interest rate needed to keep forecast average inflation within several prespecified target bands.⁴ For each target band, the policy innovations needed to keep average inflation within the band are computed.⁵ These innovations, along with typical shocks for the other variables, are then used in a dynamic, stochastic out-of-sample forecast of the VAR model to determine the time paths for the system variables. For each of the inflation target bands, we simulate the model over 1000 trials, allowing us to compute the variances of output and inflation. By specifying several bands, we generate the variability tradeoff, a menu of options from which the policy maker can choose. As with other procedures for empirical policy evaluations, the Lucas critique is a concern, and we consider our results in light of this critique.

We illustrate our procedure in two experiments. In order to mimic the policy process, we construct a real-time data set for each experiment and employ Blinder’s (1997) description of the policy process as our template for the setup of each experiment.

The first experiment begins in 1983:10, using a model estimated over the pre-Great Moderation period 1962:1–1983:9 in which there was substantial variation in both inflation and output growth. Though not part of the implicit inflation-targeting period denoted by Goodfriend (2005) as emerging from the Federal Reserve (Fed) under Alan Greenspan, an

² The loss function employed maintains a weight of unity on the inflation objective, allows varying weight on the output goal, and also includes weight for interest rate smoothing.

³ Although we illustrate the technique with a VAR model, the technique can be applied to structural models as well.

⁴ In practice, central banks appear to be willing to tolerate some variability in inflation resulting from either noise or temporary factors affecting the price indexes, so that at least implicitly the goal can be interpreted as an average inflation target. An example of a central bank that employs average inflation targeting is the Reserve Bank of Australia. Reserve Bank Governor Fraser (1993, p. 2) argued that “if the rate of inflation in underlying terms could be held to an average of 2–3% over a period of years, that would be a good outcome.” He reiterated this position a year later, arguing that “In our judgment, underlying inflation of 2 to 3 per cent is a reasonable goal for monetary policy. These figures, incidentally, are not intended to define a (narrow) range; rather, they are indicative of where we would like to see the average rate over a run of years” Fraser (1994, p. 21).

⁵ The inflation target bandwidths are specified as the goal inflation rate, plus or minus prespecified tolerance bands that include a degenerate bandwidth of zero in which the goal rate is to be met precisely.

experiment beginning in 1983:10 is a useful benchmark for the second experiment described below. Specifically, we view the period beginning in the early 1980s, which followed a period of accelerating inflation that culminated with double-digit inflation as the decade began, as one with a relatively high weight on inflation control in the underlying policy maker preference function. Furthermore, as emphasized by the unusual Saturday evening Federal Open Market Committee (FOMC) meeting in October 1979, substantial concern existed with regard to inflation expectations. In broad terms, then, the objectives at the beginning of the first experiment were much the same as in inflation-targeting regimes.

The second experiment computes the variability tradeoff policy makers would have faced in 2001:1 using a model estimated over 1980:1–2000:12, a period characterized by a monetary policy focus on reducing the inflation rate, maintaining it at a low level, and reducing variability in inflation and output. The variability tradeoff in this period shifts inward compared with the 1962:1–1983:9 tradeoff.

The models used in these experiments were estimated over periods with varying intensity of concern for inflation and, especially in the first experiment, substantial variability in both inflation and output. The Lucas critique would thus seem applicable. We address this issue within the context of the discussion by Sims and Zha (2006b) and provide some empirical evidence about the relevance of the critique to our experiment using in part the Leeper and Zha (2003) approach to evaluating modest policy interventions.

In section 2, we present the theoretical underpinnings of our experiments. In section 3, we present the intuition behind our counterfactual analysis, with technical detail relegated to an Appendix that is available on request. We also discuss antecedents in the literature to our use of counterfactual simulations to evaluate alternative policies. In section 4, we specify the empirical model, discuss data, and examine the plausibility of the estimated model through a presentation of the impulse responses to a policy shock. We include in section 5 the main statistical results, culminating in presentation of the variability tradeoffs for each experiment. In section 6, we provide evidence on whether the Lucas critique is a concern for our experiments by computing the “modesty statistic” proposed by Leeper and Zha (2003). Our results suggest that the variability tradeoff constructed for the first experiment can be usefully compared with the tradeoff in the second experiment. Finally, in section 7, we summarize the results and discuss possible explanations for them.

2. Theoretical Background

Our empirical analysis is based on two premises. First, in keeping with the dual mandate of the Fed, we assume the Fed takes output stabilization into account, at least in the short-run, in monetary policy decisions even if the primary goal of the Fed is price stability. In the context of our experiments, the narrower (wider) the inflation bandwidth, the less (more) concern is implied for output stabilization. In addition, the shorter (longer) the horizon for moving the average inflation rate to target, the less (more) concern for output stability. Second, we assume Blinder’s (1997) idealized description of policy making is appropriate as a template of the policy process.

Attaining inflation objectives has long been a goal of monetary policy, and it can be argued that the Fed was an implicit inflation targeter from the early 1980s until at least the onset of the current financial crisis. At the beginning of the 1980s, Fed policy was clearly

focused on disinflationary strategies. More recently, as inflation-targeting policy frameworks have become popular around the world, Goodfriend (2005, p. 321) has argued that “the Greenspan Fed adopted, gradually and implicitly, an approach to monetary policy characterized as inflation targeting.” Goodfriend cites congressional testimony by Alan Greenspan that the inflation objective would be achieved when “the expected rate of change of the general level of prices ceases to be a factor in individual and business decision making.”⁶ Thus, that the United States used a targeting framework, at least implicitly, during the time periods for our experiments is a reasonable assumption.

With regard to the policy process, Blinder (1997, p. 9) describes a two-step procedure for policy makers:

First, you must plan an entire hypothetical path for your policy instrument, from now until the end of the planning horizon, even though you know you will activate only the first step of the plan. It is simply illogical to make your current decision in splendid isolation from what you expect to do in subsequent periods. Second, when next period actually comes, you must appraise the new information that has arrived and make an entirely new multiperiod plan. If the surprises were trivial, that is, if the stochastic errors were approximately zero, step one of your new plan will mimic the hypothetical step two of your old plan. But if significant new information has arrived, the new plan will differ notably from the old one. Third, you must repeat this reappraisal process each and every period.

We interpret this as follows. When the FOMC meets, it should evaluate, *inter alia*, the Greenbook forecasts and the various policy options contained in the Bluebook. This is the “first step of the plan” by which “an entire hypothetical path for the policy instrument, from now until the end of the planning horizon” is considered. Note that the forecasts and assessments of the policy alternatives in the planning horizon extend well beyond the next FOMC meeting.⁷ It is this “first step” of the Blinder two-step procedure we are empirically modeling. The second step of the Blinder procedure would come at the next FOMC meeting, when “new information” has arrived and “an entirely new multiperiod plan” is implemented.

Svensson (1997, 1999) has developed a simple, analytical model of inflation targeting that provides a useful technical background for the first step of Blinder’s policy description under the dual mandate. Specifically, he presents a model of inflation targeting implemented as inflation forecast targeting. We consider the key points in the Svensson model to fix ideas prior to our empirical implementation. While in Svensson’s model the lag between the policy variable and inflation is two periods, in our empirical work we allow the policy variable to influence inflation over the 24-month period we adopt as Blinder’s planning horizon.

⁶ In detail, Goodfriend argues that the Greenspan Fed (i) approached its inflation goal gradually, trying to avoid disruptions to output, consistent with flexible inflation targeting; (ii) would not have *deliberately* (Goodfriend’s emphasis) undertaken policies that would raise inflation above 2% after achieving personal consumption expenditures (PCE) inflation in that range in the mid-1990s; (iii) would have been unlikely to *deliberately* (his emphasis) aim at inflation below 1% given the costs associated with deflation and problems associated with the zero nominal interest rate bound; and (iv) when actual and expected inflation were well contained, “aggressively” pursued countercyclical policy.

⁷ Interestingly, by the time of our second experiment in 2001, the Bluebook presentation of policy alternatives presents long-range policy paths for the federal funds rate and the implications for broad macroeconomic aggregates through 2005, an “entire hypothetical path for the policy instrument ... even though you know you will activate only the first step of the plan.” Thus, at least part of the strategy suggested by Blinder in 1997, which we model below, seemed to be part of the operating procedures at the time of our second experiment. In 1983, during our first experiment, the longer-range horizon was shorter, however, by about a year. There was also a heavier focus on monetary aggregates though with qualitative discussion of the interest rate paths consistent with the alternative forecasts.

Svensson’s model specifies a setting where the policy interest rate affects inflation with a two period lag. Accordingly, the interest rate is set at a level today consistent with *forecasts* of inflation two periods later. The two period lag arises by assuming that the current interest rate setting affects the output gap with a one period lag, which in turn affects inflation with a subsequent one period lag. Recall from the introduction that Bernanke (2004) characterized forecast-based approaches as “dominant” in policy making.

In Svensson’s setting, the aggregate supply/Phillips curve relationship is given by

$$\pi_{t+1} = \pi_t + \alpha_y y_t + \varepsilon_{t+1}, \tag{1}$$

while the aggregate demand/IS is given by

$$y_{t+1} = \beta_y y_t + \beta_z z_t - \beta_r (i_t - \pi_{t+1|t} - \bar{r}) + \eta_{t+1}, \tag{2}$$

where π_{t+1} is the inflation rate in period $t + 1$, y_t is the output gap, ε_{t+1} is a random shock to aggregate supply, i_t is the nominal interest rate, $\pi_{t+1|t}$ is the expected inflation rate conditioned on information at time t , \bar{r} is the equilibrium real interest rate, η_{t+1} is a random shock to aggregate demand, and z_t is a vector of exogenous variables given by $z_{t+1} = \gamma_z z_t + \theta_{t+1}$, with γ_z being a conformable vector and θ_t a random vector with the same dimensionality as z_t . The period loss function, reflecting the dual mandate, is

$$-\frac{1}{2} [\lambda y_t^2 + (\pi_t - \pi^*)^2], \tag{3}$$

where λ is the relative weight on output stability and π^* is the target inflation rate. Svensson (1997) shows that the first-order condition for a minimum is

$$\pi_{t+2|t} - \pi^* = -\frac{\alpha_y c(\lambda)}{1 - c(\lambda)} y_{t+1|t}, \tag{4}$$

where $c(\lambda)$ is a function of model parameters with the properties that $0 \leq c(\lambda) < 1$, $\partial c/\partial \lambda > 0$, $c(0) = 0$, and $c(\infty) \equiv \lim_{\lambda \rightarrow \infty} c(\lambda) = 1$. The form of Equation 4, with the two-period-ahead inflation forecast and the one-period-ahead output gap forecast, is the result of the lag structure in the model; setting the interest rate today influences only future values of these variables.

Consider the meaning of Equation 4. As a result of minimizing the specified loss function subject to the existing tradeoff of inflation for output in Equation 1, Equation 4 is a consequence of the equality between the marginal rate of substitution in the policy maker preferences with the marginal rate of technical substitution. In practice, Svensson (2003, p. 451) notes that “the loss function is not specified in this detail.” Rather, given available information and any judgment to be applied, various forecasts for alternative paths for the policy interest rate are generated as in the Fed’s Bluebook, and the policy making committee then picks the “combination of forecasts that ‘looks best’, in the sense of achieving the best compromise between stabilizing the inflation gap and stabilizing the output gap, that is, implicitly minimizing [the loss function]” (Svensson, 2003, p. 451).⁸ Our empirical work below will exploit this interpretation of the policy process rather than explicitly specifying a loss function.

⁸ Svensson (2009) argues that Equation 4 implies $\pi_{t+\tau,t} - \pi^* = c(\lambda)(\pi_{t+\tau-1,t} - \pi^*) = c(\lambda)^\tau (\pi_t - \pi^*)$, where $\pi_{t+\tau,t}$ is the $t + \tau$ inflation forecast that “looks good” in period t . In Svensson’s analysis, the policy is engineered so inflation asymptotically approaches its target; driving inflation to the target within a specified policy horizon may not be efficient. In our empirical analysis, we adopt relatively long (48- and 36-month) transition periods to the target in the spirit of the asymptotic approach.

Given period t information, Equation 2 implies that the one-period-ahead output gap forecast is

$$y_{t+1} = \beta_y y_t + \beta_z z_t - \beta_r (i_t - \pi_{t+1|t} - \bar{r}), \tag{5}$$

and from Equation 1 the two-period-ahead inflation forecast is

$$\pi_{t+2|t} = \pi_{t+1|t} + \alpha_y y_{t+1|t}. \tag{6}$$

Subtracting π^* from both sides of Equation 6 and using Equation 4 to substitute for $\alpha_y y_{t+1|t}$ gives $\pi_{t+2|t} - \pi^* = c(\lambda)(\pi_{t+1|t} - \pi^*)$. Substituting the right-hand side of Equation 6 for $\pi_{t+2|t}$ and then the right-hand side of Equation 5 for $y_{t+1|t}$ in the resulting expression and then solving for i_t yields

$$i_t = \bar{r} + \left(1 + \frac{1 - c(\lambda)}{\alpha_y \beta_r}\right) (\pi_{t+1|t} - \pi^*) + \pi^* + \frac{\beta_y}{\beta_r} y_t + \frac{\beta_z}{\beta_r} z_t.$$

Finally, taking the expectation of Equation 1 conditional on time t information and substituting for $\pi_{t+1|t}$ gives

$$i_t = \bar{i} + \left(1 + \frac{1 - c(\lambda)}{\alpha_y \beta_r}\right) (\pi_t - \pi^*) + \left[\alpha_y \left(1 + \frac{1 - c(\lambda)}{\alpha_y \beta_r}\right) + \frac{\beta_y}{\beta_r}\right] y_t + \frac{\beta_z}{\beta_r} z_t, \tag{7}$$

where $\bar{i} = \bar{r} + \pi^*$. Equation 7 is the setting for the interest rate that satisfies the first-order condition, Equation 4, and implies that the optimal interest rate setting depends not only on the output gap and the deviation of inflation from target, but also on other model variables included in z . Note also that Equation 7 is not a modified Taylor rule, since (i) z represents the remaining variables in the system and (ii) this interest rate setting is implied by an optimization process (as reflected in the coefficients on the inflation and output gaps). The Taylor rule is generally simpler (i.e., omits the variables in z or includes only a small subset) and has coefficients on the inflation and output gap terms that are policy decision variables.

As noted, Svensson (2003) argued that central bankers generally don't explicitly reveal a loss function, a step we avoid here as well. Nonetheless, two comments about extending models such as that above to include non-degenerate inflation bands are in order.

First, in a typical loss function, such as Equation 3, λ represents the relative weight on output. Svensson (1997, p. 1135) states that a "wide [inflation] band could then potentially indicate that the central bank has a relatively high λ and hence a significant output stabilization goal. A narrow [inflation] band could indicate a commitment to a low or even zero λ ." Alternatively, we could normalize directly on inflation variability, where Λ is the relative weight, with a loss function such as $-(1/2)[y_t^2 + \Lambda(\pi_t - \pi^*)^2]$. Our bandwidth parameter, τ , is negatively related to Λ ; the narrower our bandwidth, the less tolerant the policy maker is to inflation variability, and implicitly the higher the loss associated with inflation variability.⁹ As in Equation 4, where the setting of λ determines the slope of the tradeoff between inflation and output deviations from desired levels, as we set τ to various values to trace the variability tradeoff, we are implicitly varying the value of Λ .

⁹ For the limiting case of τ approaching 0, Λ would become arbitrarily large. In this case, the policy maker would be a strict inflation targeter, abandoning the dual mandate.

Second, Flood and Isard (1989) present a model quite similar to that of Svensson detailed above with a positive weight on inflation variability, such as Λ in the above paragraph. The policy maker pursues one type of policy when the target variable (the exchange rate in their case) is within a band but another policy when this variable is shocked outside the band. They derive optimal policies with a loss function for two practical cases: discretion and partially state contingent. They demonstrate the existence of a mixed strategy (a probability-weighted average of the two), which is time consistent under certain conditions. That is, an optimal policy exists that is different outside that band than inside and still consistent with optimization. Thus, theoretical models underlying our empirical application consistent with optimization exist.

3. Methodology

Estimation of the Variability Tradeoff

We next present the basic intuition behind the methodology used to estimate the variability tradeoff for our inflation target. Consistent with Blinder, we plan a nominal interest rate path over the planning horizon; consistent with Svensson, we compute the policy interest rate needed to attain a forecast of the average inflation rate over a 24-month period on or within a prespecified band. Technical details on the computation of the interest rate setting (analogous to Eqn. 7) needed to attain the inflation target are contained in an Appendix that is available on request.

We specify the inflation goal in terms of a 24-month average inflation forecast.^{10,11} That is, policy makers are forward-looking, planning interest rate policy to maintain what we will call the forecast average inflation rate (FAIR) to be consistent with the inflation target on or within a specified inflation band. Aiming for inflation on or within a band is consistent with current practice by some inflation-targeting central banks and reveals a willingness to accept deviations from the midpoint of the target band. That is, if the forecast for inflation over the next 24 months is consistent with the inflation target, that is, on or within the band, no policy intervention is undertaken. If the FAIR is outside the band, then an intervention is used to return this measure of inflation to the band.¹²

Since a forecast of the inflation process is needed, we need to specify the source of these forecasts in our analysis. We begin with a structural model

$$y_t = A_0 y_t + A_1 y_{t-1} + \dots + A_p y_{t-p} + \varepsilon_t. \quad (8)$$

¹⁰ Svensson (2003) notes that both the Bank of England and the Sveriges Riksbank have used a two-year inflation forecast horizon for conducting monetary policy.

¹¹ Note that there is no necessary reason the “planning horizon” need be equal to the time frame over which the average inflation rate is targeted by the policy authority. We set both equal to 24 months to roughly mimic reality. Specifically, the Greenbook forecasts are now for a period of at least several years. And, as in the previous footnote, some inflation-targeting central banks employ two-year forecasts for inflation as the inflation objective. There is no reason, for example, that the forecasts evaluated at a policy meeting could not extend to, say, three years while announcing and maintaining a goal for a two-year average inflation rate over this planning horizon.

¹² While we do not do so in our simulations, it would be straightforward in practice to allow for judgment in the forecast by including an adjustment for factors that are outside the model but deemed by policy makers to be important for the immediate policy exercise.

In Equation 8, y_t is an $(N \times 1)$ vector of variables, including the inflation rate and the federal funds interest rate.¹³ The elements of the A_i matrices represent the structural coefficients, and the elements of ε_t are structural shocks. We assume that $E(\varepsilon\varepsilon') = \Omega$ is diagonal. The reduced form of Equation 8 is $\Pi(L)y_t = e_t$, where $\Pi(L) = I - \Pi_1L^1 - \dots - \Pi_pL^p$. Reduced-form coefficient matrices are given by $\Pi_i = (I - A_0)^{-1}A_i$, and reduced-form shocks are given by $e_t = (I - A_0)^{-1}\varepsilon_t$. The moving average matrix is defined as $C(L) = [\Pi(L)]^{-1}$, with $C_0 = I$. D_s is equal to $C_s(I - A_0)^{-1}$. The moving average representation of Equation 8, expressed in terms of the structural shocks, is

$$y_t = \sum_{s=0}^{\infty} D_s \varepsilon_{t-s}. \tag{9}$$

Fundamental to our analysis is the historical decomposition, which in its basic form is found by advancing Equation 9 by n periods and then decomposing the resulting expression into two terms:

$$y_{t+n} = \sum_{s=0}^{n-1} D_s \varepsilon_{t+n-s} + \sum_{s=n}^{\infty} D_s \varepsilon_{t+n-s}. \tag{10}$$

The second term on the right-hand side of Equation 10 is the dynamic forecast or base projection of y_{t+n} conditional on information at time t . The first term on the right-hand side shows the influence on y_{t+n} of the shocks to the variables in the system between periods $t + 1$ and $t + n$. Even though the expected values of these shocks are zero, policy makers know that the realizations of these shocks over any particular period are likely to be nonzero, which provides the motivation for the stochastic part of our simulation. These shocks, drawn from the estimated residuals, represent the source of variability around the base projection. Given a set of shocks to the system, we obtain monthly inflation rates from the relevant equation in system (Eqn. 10), which are then averaged to obtain the FAIR.

If the FAIR deviates from target, a policy action is called for. The goal for policy is to remain on or within the range $(\pi^* \pm \tau)$ where τ is half the bandwidth (including the case of a zero bandwidth), with policy aiming at $(\pi^* - \tau)$ when inflation is below the band and aiming at $(\pi^* + \tau)$ when inflation is above the band. For a non-degenerate bandwidth, our policy experiments return FAIR to the edge of the band rather than the midpoint for four reasons. First, although we don't model the loss function explicitly, our presumption about the dual mandate means that a more aggressive policy action—needed to return the FAIR to the midpoint of the band rather than the edge—induces additional variability in output, raising the overall loss.¹⁴ Second, if there is multiplicative uncertainty about the economy, in the sense of Brainard (1967), then the policy authority may not necessarily aim at the midpoint of the range. That is, if there is not certainty equivalence, then aiming at the midpoint is no longer necessarily

¹³ For simplicity, we do not explicitly consider expectational variables in our analysis. We do note, however, that under some conditions models with expectations of variables can be solved for a VAR of the type estimated here; see Fernandez-Villaverde et al. (2007) and Ireland (2004) for examples.

¹⁴ This is essentially the point of opportunistic policy described in Orphanides et al. (1997). For additional discussion, see also result 12 in Clarida, Gali, and Gertler (1999).

optimal.¹⁵ Third, if policy makers want to minimize their impact on financial markets, returning to the edge of the inflation band requires a smaller interest rate innovation and thus helps minimize interest rate movements. That is, we undertake the smallest policy action needed to attain the inflation objective over the two-year horizon. Of course, the tradeoff is that our smaller financial market interventions may be more frequent than relatively aggressive actions aimed at returning to the midpoint of the band. Fourth, as Blinder noted, during his Federal Reserve tenure there was a lack of consensus on the desirable inflation rate among the governors as well as a tendency to wait for opportunistic shocks to lower the inflation rate, in practice restraining inflation that is “too high” but not necessarily moving it aggressively toward a firm number such as 2%.

We emphasize that in our analysis the policy objective is the FAIR over the 24-month period rather than either the current or any particular future monthly inflation rate. Current inflation is the result of past decisions by both policy makers and private agents in the economy and is presumably not directly affected by current policy. Reported inflation, or a forecast of a particular monthly inflation rate, may deviate from the inflation target without necessarily calling for a policy action as long as the FAIR suggests that the longer-run objective will be satisfied. However, if the FAIR deviates from target or lies outside the band, a policy action is called for. In our application, we will use the federal funds rate as the policy instrument to control the inflation rate, and a policy action in a particular month is defined as an intervention in the funds rate equation in that month.¹⁶ Owing to interaction with other system variables via system dynamics, a policy action in a particular month will affect inflation over the remainder of the horizon. That is, even if the funds rate has a relatively small contemporaneous effect on inflation, marginal changes in this rate can still have substantial effects on long-run inflation.

Diagram 1 gives a schematic presentation of our model of the first stage of the Blinder process in an inflation forecast targeting regime that desires to keep inflation on or within a target band. Period 1 on the horizontal axis is the first period of the planning horizon, which in practice would be the date of an FOMC meeting. The vertical axis is the inflation rate.

Two alternative inflation bands are sketched on Diagram 1. The band centered at 2%, with $\pm\tau\%$ bands around this midpoint, represents our presumed long-run goal for policy. Of course, it is possible that at the time an inflation target is adopted, actual inflation will be above this long-run target range; after all, inflation targets are adopted to try to control an inflation problem. In Diagram 1, the negatively sloped lines (for simplicity) represent a linear transitional inflation band in which inflation is to be gradually lowered, so that after some (possibly publicly announced) period of time, inflation will be brought into the long-run range. In Diagram 1, a policy of gradualism over 48 months is depicted. This period could be shorter or longer depending on policy maker preferences.

¹⁵ Despite the fact that the “Brainard conservatism principle” (as Blinder dubs it) does not apply in all cases, he notes that as a central banker, he viewed the principle as “extremely wise.” The wisdom of the basic Brainard result is included as well in Clarida, Gali, and Gertler (1999, p. 1688), result 11. Specifically, they argue that “parameter uncertainty may reduce the response of the policy instrument to disturbances in the economy.” That is, the reduction in the response may lead to aiming for the edge of the band rather than the midpoint. Further, Barlevy (2009) has shown that when robust control techniques are applied in circumstances similar to those emphasized by Brainard, robust control implies an even more “conservative” policy response than does Brainard.

¹⁶ As is shown in the technical Appendix, the desired change in the funds rate is implemented by replacing the residual term in the funds rate expression with an appropriately sized shock that brings the funds rate to the level consistent with the inflation band.




Diagram 1. Schematic of Inflation Forecast Targeting Regime

Consider period 1 in Diagram 1. The policy maker has real-time data through period 0 and is interested in policy simulations conditional on data through period 0, a procedure roughly similar to that undertaken at a given FOMC meeting. Assume that the inflation rate is about 10%. The policy maker makes a forecast of inflation over the next 24 months and computes the average of the monthly inflation forecasts. The period 1–24 is “underlined” beneath the horizontal axis. If the FAIR is within the transitional inflation band, such as point “x,” then no policy intervention is undertaken. Alternatively, if the FAIR is at a point like “w,” a policy intervention is needed that will bring the current forecast to the upper edge of the band.¹⁷ During period 0, in planning for period 2 a new simulation is made conditional on what was done in period 1, intervention or not. If an intervention was needed in period 1, the magnitude of this intervention must be incorporated into the analysis to properly forecast inflation over the next 24-month period in which average inflation is to be on or inside the band. That is, the

¹⁷ If the forecast inflation rate is below the band, then in a transitional period such as that characterized by the negatively sloped inflation band in Diagram 1, the policy maker may choose to conduct policy consistent with opportunistic disinflation. This would entail lowering the transitional range of the inflation target so that the period over which the gradualist policy need be conducted is shortened.

period 0 forecast for inflation over periods 2–25 is computed, as indicated by the second “underline” beneath Diagram 1. If the forecast is within the band, similar to “x” in period 1, then no intervention is needed; if it is above the band, similar to “w” in period 1, an appropriate intervention is conducted to bring the 24-month average inflation forecast within the band. If, at a given policy meeting in period 0, Blinder’s “planning horizon” is two years [“even though you know you will only activate the first step of the plan” Blinder (1997, p. 9)] then inflation forecasts are conducted through the end of this horizon, so a period 0 forecast is needed for month 24, which extends over months 25–48, the last underlined period in Diagram 1.

Our choice of 48 months as a gradual transition to the long-run inflation target is based on both theoretical considerations and observation of central bank practices. Though not suggesting a specific length of the transition period, Svensson (1997) argues theoretically that a positive weight on the output gap in the loss function implies that optimal disinflationary policy will be one of gradualism, as in our downwardly sloped transition bands in Diagram 1. Given the “dual mandate,” U.S. policy makers should then approach inflation targets gradually. In practice, according to Bernanke and Mishkin (1997), central bankers behave as suggested by Svensson. Bernanke and Mishkin note (p. 99): “Initial announcements of inflation targeting generally allow for a gradual transition from the current level of inflation to a desired steady state, usually the level deemed consistent with price stability.” Furthermore, Bernanke and Mishkin (p. 101) later note that after the 1979 oil shock, the German Bundesbank “announced the ‘unavoidable’ inflation rate to be 4 percent, then moved its target gradually down to 2 percent over a six-year period.” In the United States, Goodfriend (2005, p. 8) indicated that an “inflation scare” in 1987 due to the infusion of liquidity after the October 1987 stock market crash took the Greenspan Fed “about five years to overcome.” Our choice of 48 months as the transition period is a bit shorter than, but not at great odds with, these descriptions of the behavior of inflation-targeting central banks. However, we also consider a shorter transition of 36 months to the inflation target; this shorter transition is consistent with a higher relative weight on achieving the inflation target as opposed to output stabilization.

We now present the setup of our experiments, each of which produces a point on the variability tradeoff the central bank would have faced at a particular time. Each experiment corresponds to our model of the “first step” in the Blinder policy process.

For each of our two time periods, we estimate a VAR model using real-time data that end with the period before the start of the simulation. We compute the base projection at the end of the estimation period, mimicking the real-time forecasting process just prior to an FOMC meeting. This forecast of y_{t+1} through y_{t+24} is represented by the second right-hand-side term in Equation 10 and is estimated from the lagged historical residuals from the VAR. Since the base projection is based on historical residuals, it does not change across the trials of a given experiment.

For each time period, we conduct four experiments for both the 48-month and 36-month transition periods. Each experiment has 1000 trials with a prespecified bandwidth. For each trial, we draw (with replacement) a vector of residuals of length 48 from the estimated residuals for each equation in the system.¹⁸

Each experiment starts with the actual inflation rate in the period prior to the experiment and gradually lowers the inflation target to 2% over a 48- (36)-month period (as in Diagram 1), subject to the bandwidth. The alternative bandwidths are 0%, 1%, 2%, and ∞ where a

¹⁸ As noted above, we need 48 residuals since during the last month of the two-year planning horizon, policy makers want to know the FAIR for the subsequent 24 months.

bandwidth of 0 means that the focus is on the target itself and a bandwidth of ∞ represents a stochastic simulation using the draw from the historical residuals without any policy intervention. So, in Diagram 1, given the negatively sloped transition path, the dashed lines could assume one of the indicated bandwidth values. The initial 24 terms of the drawn vector of shocks are used to compute the first term on the right-hand side of Equation 10. When combined with the base projection, Equation 10 gives the path the economy, as represented by the system of equations, would follow under this trial. Combining the base projection and the initial 24 elements of the vector of draws from the residuals gives the policy maker a forecast of inflation for each of the next 24 months for this particular trial. The policy maker averages these 24 individual monthly inflation rates to compute the FAIR. If this rate is inside the band or equal to target in the case of a 0% band, no policy intervention is needed. On the other hand, if this inflation rate is outside the band, a preemptive policy action of sufficient magnitude to return the 24-month average forecast inflation rate to the closest edge of the band is calculated.¹⁹ That is, if a policy intervention is needed to attain the band, the drawn residual from the interest rate equation is replaced with one that is computed to assure that the FAIR attains the policy objective.²⁰ This policy action, computed as detailed in the Appendix, produces an interest rate analogous to Equation 7 above. Furthermore, this policy intervention is carried along for the remainder of the trial, affecting all the system variables in later time periods. Also, as is shown in the Appendix, the computation by which the needed intervention is done is conditional on the shocks from the random draw for later time periods; the current policy action needed to attain the inflation target is “identified” using the remaining residuals from the draw. Finally, note that since the residuals in Equation 10 are structural, the residual drawn for the policy equation can be replaced with the needed policy action without implications for the random shocks to the other equations, since it was assumed there is no contemporaneous correlation among the structural shocks.^{21,22}

For each month in the planning period, a policy action is either needed or not. Either way, the dynamic path of the economy is computed. After passing through the planning horizon, intervening as needed, we obtain at the end a path for the system of variables in which policy is

¹⁹ We note that attempting to use the funds rate to control the inflation rate at very short horizons may lead to instrument instability. The intuition behind this statement is that the interest rate is not an important component of measured prices and that the contemporaneous effect of a change in the interest rate on aggregate demand is very small. Consequently, an interest rate change would have a relatively small near-term impact on the inflation process, requiring large interest rate movements to affect short-term inflation. With a longer-term inflation objective, say one of several years as we employ here, a current interest rate change has lagged effects on the inflation rate, consistent with system dynamics. This point is recognized by central bankers, who generally implement policy via interest rate innovations that are allowed to work their way through the dynamics of the economy.

²⁰ As we will discuss further below, this procedure is the same as that adopted by Leeper and Zha (2003).

²¹ An alternative approach for obtaining a desired average inflation rate would be to employ a “constant interest rate” approach, which would take the base projection and adjust it by imposing a constant interest rate over the 24-month horizon that brings about the desired average inflation rate. This approach thus implicitly imposes an entire path for the shocks to the interest rate equation. In our analysis, we identify the current policy shock needed to attain the objective (given the rest of the draw) while the constant interest rate approach implicitly identifies a vector of shocks, current and for the remainder of the horizon, needed to maintain a constant interest rate and simultaneously attain the inflation objective. The constant interest rate approach thus imposes more policy action than needed to attain the policy objective. It imposes interest rate smoothness, while our approach allows the path of rates to be determined by the response of the policy maker to the forces that may drive the FAIR outside the band.

²² It is possible to model correlations among the structural shocks, as in Bernanke and Mihov (1998). If such modeling included contemporaneous correlation between the policy innovation and other variables, then other structural shocks would be affected when a policy shock needed to attain the FAIR is imposed. We do not model such contemporaneous correlations here.

used to attain the inflation target. By construction, this counterfactual path over the planning horizon is consistent with the inflation objective of the policy authorities. With 1000 trials for each bandwidth, we can compute the variance of each element of the vector of variables.

Over the 1000 trials in a given experiment we determine the variances of inflation (around the trial mean) and the output gap (also around the trial mean). At a point in time and for a given inflation band, we use these variances of inflation and output to represent a point on the variability tradeoff. Using the four alternative bandwidths then allows us to plot the available inflation-output variability frontier. The policy maker can then select the value for τ that “looks best” in the characterization of Svensson (2003, 2009). By using different time periods, we can also discover how the point estimate of the frontier has changed over time.

To summarize, in the first step of Blinder’s process, the policy maker takes into account the results of the entire 24-month planning horizon since (in his words) “[i]t is simply illogical to make your current decision in splendid isolation from what you expect to do in subsequent periods” (Blinder 1997, p. 9). This longer-range planning process, analogous to the evaluation of the Green and Bluebooks at the FOMC meeting, may be aimed at issues such as whether there is instrument instability for a given policy that might not show up at shorter horizons, whether there are undesirable characteristics of the implied interest rate interventions (are they too frequent? do they impart too much variability into the financial market?), and so on. At the same time, the policy maker also knows that it is likely to “activate only the first [portion] of the plan.”

In the second step of Blinder’s procedure, “new information that has arrived” is assessed, and an “entirely new multiperiod plan” is made. In real time, this would include additional economic data arriving subsequent to an FOMC meeting, and the new plan would incorporate information contained in the shocks to these data. If we were to proceed to this second step in our analysis, we would collect a new real-time data set and then at the time of the next FOMC meeting re-do the experiments described above.

Comparison with Previous Studies

As referenced earlier, the use of counterfactual analysis to evaluate policy alternatives has several precedents in the literature. Fackler and Rogers (1995) were the first to suggest the general approach used here, though their analysis was in a simpler setting and was conducted in-sample in contrast to the out-of-sample analysis in the current article. Specifically, in the Fackler-Rogers analysis, a policy shock in a period was selected to attain a target exchange rate for that period rather than aiming at a 24-month average for the target variable. Christiano (1998) used a similar in-sample approach, computing the shocks to the interest rate equation needed in a given period to keep the money supply on a constant growth rate path during the Great Depression. He reported that the resulting path for the system variables other than the money supply “oscillated so wildly” that his presented results used an *ad hoc* method of combining a weighted average of the counterfactual shocks with the historical shocks rather than pure counterfactual residuals as is done here.

Leeper and Zha (2003) evaluated the effects of changes in the target funds rate beginning in 1990 and in 1994–1995. Using a different empirical model that employed a richer specification of the reserves market than Leeper and Zha, Fackler and McMillin (2002) compared with no-change policies the effects of specified changes in the target federal funds rate on the time paths of output, the price level, and other model variables in 1995, and again in

1998. In contrast to these two studies, which computed “own” shocks to the funds rate required to achieve an arbitrary target funds rate, the current article computes the shocks to the funds rate required to generate the time path of the funds rate that achieves an inflation target. Finally, we note that Leeper and Zha contributed importantly to the policy evaluation literature by constructing a “modesty statistic” that allows one to test whether the Lucas critique is a concern for policy analysis, and we employ this statistic to evaluate our counterfactual experiments.

Sims and Zha (2006b) considered the effects of changes in the monetary policy reaction function in a VAR model on the economy’s response to shocks to non-policy variables. They first computed the effect of the policy and non-policy shocks on the model variables assuming a monetary policy reaction function that featured lagged response of the policy variable (the funds rate or alternatively the three-month T-bill rate) to all model variables and contemporaneous response to only a subset of the model variables. They then separately computed the response of the economy to non-policy shocks assuming the central bank keeps the policy rate constant. Our analysis differs in two regards: (i) as noted earlier, we compute the policy shocks required to generate the (non-constant) path of the funds rate that achieves the inflation target and (ii) since shocks hit all sectors of the economy simultaneously, we compute the funds rate shocks in the presence of shocks to the entire system rather than consider shocks to each non-policy variable as a separate case.

4. Empirical Model

The variables in the VAR model we estimate include those in the typical New Keynesian model: the output gap, the inflation rate, and the federal funds rate. Additionally, we include the rate of change in a commodity price index for two reasons. First, we add commodity prices following earlier literature that addresses the well-known “price puzzle” often found in VAR models. Second, since commodity price volatility is often used to represent supply shocks, as a first (and likely crude) approximation, we use this variable to help control for changes in output and inflation volatility emanating from sources outside the policy process. In order to establish the usefulness of the model for monetary policy evaluation, the macroeconomic effects of monetary policy are estimated by computing impulse response functions (IRFs) for shocks to the federal funds rate.

The model is estimated using monthly real-time data over two time periods: 1962:1–1983:9 and 1980:1–2000:12. Our first set of counterfactual inflation-targeting experiments begins in 1983:10, a year after the end of reserve targeting that characterized the October 1979–October 1982 period, thus allowing for adjustment to the new operating procedure to be basically completed before initiating the experiments. The second counterfactual begins in 2001:1. This starting point was chosen for two reasons. One is that there was considerable uncertainty about the macroeconomic effects of the decline in stock prices that began in 2000. The second is that it allows a year’s transition from the temporary effects of the year-2000 (Y2K) preparations of the Federal Reserve and the subsequent volatility in the growth rate of the monetary base.²³ In estimating the VAR, 12 lags of all variables are employed.

²³ There was a big spike in total reserves in the system at the end of 1999, and the growth rate of the monetary base jumped sharply. Reserves quickly returned to the pre-Y2K level, and monetary base growth fell sharply over 2000 and even became negative toward the end of 2000, although it had begun to rise by the end of the year.

The transformations of the variables in the model follow the transformations of the variables in the typical New Keynesian model.²⁴ Measuring the output gap at a monthly frequency is problematic, and we considered three alternative measures. The first measure was constructed by subtracting the quadratic trend of log real GDP from the log of real GDP at a quarterly frequency and then interpolating to monthly values.²⁵ The second measure was the log of real GDP minus the log of Hodrick-Prescott filtered real GDP, again interpolated from quarterly to monthly frequency. This filter is commonly used despite the potential problems in its use noted by Cogley and Nason (1995), among others. The third alternative was the monthly total unemployment rate. Results for the model with the first measure are reported in the text, and results for the models using the other two real macroeconomic activity measures, which are quite similar to the first model, are in the Appendix. Since central banks tend to focus on longer-term inflation, the inflation rate is measured by the year-over-year rate of the change in the personal consumption expenditure deflator, a key series in the Fed's evaluation of inflation.²⁶ This avoids filtering out longer-run inflation information as would occur, for example, if we had used the annualized monthly rate of change in the price level. The federal funds rate is the monthly average of the daily rate. The rate of change in commodity prices is calculated as the annual difference of the log of this series. A description of the real-time data and sources of the data are provided in the data Appendix, which is available on request.

Monetary policy shocks are identified using a Choleski decomposition. The ordering is as follows: rate of change in commodity prices, output gap, inflation rate, and then federal funds rate. Placing the funds rate last is based on a suggestion by Bernanke and Blinder (1992) and allows a contemporaneous response by the Fed to movements in the other three variables while simultaneously imposing a lagged effect of monetary policy on these variables.²⁷

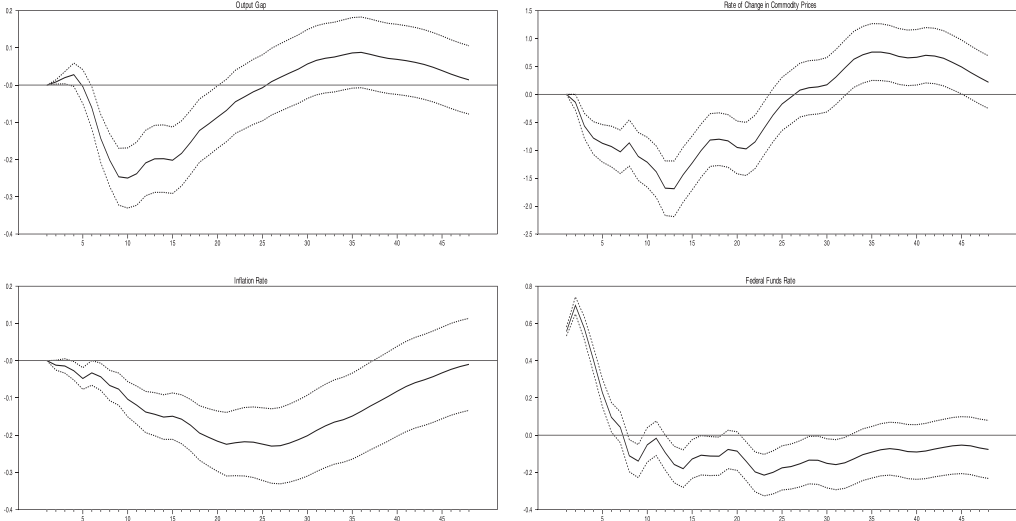
²⁴ Given the considerable debate about the power of unit root and cointegration tests, we chose to employ the transformations of the variables used in the typical New Keynesian model.

²⁵ The interpolation used the random walk option of the `distrib.src` procedure from WinRATS 6.02b. A commonly used alternative measure of the output gap is the log of real GDP minus the log of potential output constructed by the Congressional Budget Office (CBO). However, since the CBO measure of potential output is not available on a real-time basis, we used the real GDP output gap measures described in the text. Although series such as industrial production and personal income are available monthly, we concentrated on real GDP and interpolated the output gap measures to monthly frequency since the focus of policy is on broad economic activity and real GDP is better measure of broad economic activity than are industrial production or personal income.

²⁶ Although policy makers focus on the core personal consumption expenditure deflator, the core series was not used in the model because it was not possible to construct a real-time version of this series for use in the experiments.

²⁷ Imposing a lag in the effect of monetary policy on inflation is not controversial. However, one concern about this ordering is that it does not allow monetary policy to have a contemporaneous effect on the commodity price index, which is comprised of auction-market type variables that may well respond within the period to monetary policy shocks. Other concerns include (i) the assumption that the central bank responds contemporaneously to current period movements in output and the price level, whereas data (even preliminary) on current period values of these variables is available only with a lag, and (ii) the constraint that output isn't allowed to respond contemporaneously to a shock to monetary policy. Because of these concerns, we estimated a Bernanke (1986)-type structural VAR, which differed from the Choleski described in the text by allowing a contemporaneous effect of monetary policy on commodity prices, by allowing a concurrent effect of monetary policy on output, and by imposing no contemporaneous response of the federal funds rate to output and inflation shocks. The federal funds rate was, however, allowed to respond contemporaneously to commodity price shocks. The point estimates of the IRFs for this structural VAR for a shock to the federal funds rate were plotted along with the confidence intervals for the Choleski decomposition. The point estimates for a monetary policy shock for all variables for both samples were within the Choleski confidence intervals except for a few very minor departures in the very short-run for output. Based on these results, we used the Choleski decomposition in all experiments.

1962:1-1983:9



1980:1-2000:12

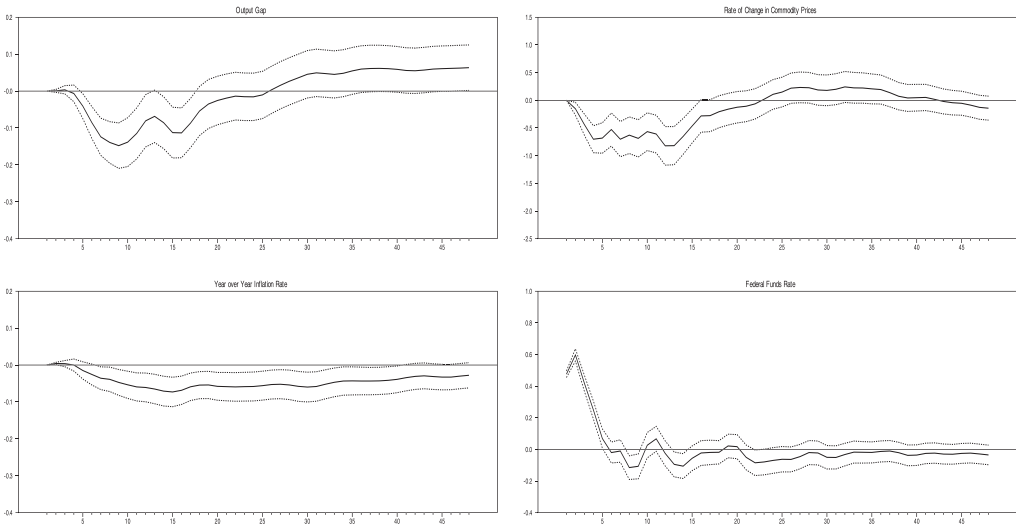


Figure 1. Impulse Response Functions

The IRFs for a shock to the federal funds rate for the model with the quadratic trend output gap for both estimation periods are presented in Figure 1. In each panel, the solid line is the point estimate and the dotted lines are one standard deviation confidence intervals computed using Monte Carlo simulations employing 10,000 draws. The general pattern of results is similar for each sample period, but the timing and magnitude of effects differs across samples. The magnitude of the one standard deviation federal funds rate shock is comparable across the two samples: 0.56 for 1962:1–1983:9 and 0.48 for 1980:1–2000:12. A positive shock

to the federal funds rate persists briefly, but the confidence interval for the funds rate spans zero within five to six months, which we interpret as a return to the initial value. There is a transitory negative effect on the rate of change in commodity prices, and the effect is stronger and more persistent for the 1962:1–1983:9 sample. The output gap becomes negative after several months but returns to its initial value over time. The magnitude of the effect is greater for the 1962:1–1983:9 sample than the second sample, but the time required for output to return to its trend and stay there is comparable for both samples. There is a transitory negative effect on the rate of change in the personal consumption expenditure price index, but while the magnitude of the effect differs between the two periods, the time required before the effect becomes significant and the time that lapses until the rate of inflation returns to its initial value is about the same across samples.

Since the VAR models are used to assess the quantitative implications of inflation targeting, it is important that the VARs produce paths of the model variables for shocks to monetary policy that are consistent with macroeconomic models in which monetary policy shocks can temporarily affect real variables. This appears to be the case for the VAR models used in this article.

5. Results

In this section, we present a variety of results from the inflation-targeting experiments and discuss their economic interpretations. The discussion focuses on what policy makers in real time would have seen had they employed our methodology. Specifically, we investigate the nature of the available tradeoffs between inflation and output variability and how these tradeoffs have changed in the two periods we consider. As detailed earlier, in each experiment we assume that a policy of gradualism to reduce inflation is employed, first for a 48-month transition period and then for a 36-month period. Since the standard deviations of output, inflation, and the interest rate for each target band in both experiments was only slightly greater for the 36-month period than the 48-month period and the inferences were the same, we present only the 48-month results. The 36-month results are in the Appendix. The benchmark policy is for the midpoint of each inflation band to approach 2% over a 48-month period with bandwidths varying between 0% and one that is arbitrarily large.

The first experiment begins in 1983:10. The top part of Figure 2 shows the actual inflation rate through 1983:9 and the base projection of the inflation rate along with the $\pm 1\%$ and $\pm 2\%$ bands moving to the long-run inflation target. The inflation rate as measured by the personal consumption expenditures deflator at the outset of this experiment was approximately 3.8%. Note that while the actual inflation rate was relatively low, the base projection suggested that inflation would quickly move outside the $\pm 1\%$ bands and approach the upper 2% band. Thus, for policy officials using real-time data in late 1983, the need for restrictive monetary policy looked highly likely. Such a policy would likely raise the specter of another recession following on the two at the outset of the decade, making empirical estimation of the variability tradeoff an important consideration.




Figure 2. Actual Inflation, Base Projections, and Target Bands

The second experiment begins in 2001:1. Even though inflation was reasonably well contained at approximately 2.5% when our second experiment begins and the base projection in the bottom part of Figure 2 puts inflation within the 2% inflation bands, uncertainty about the macroeconomic effects of the decline in stock prices that began in 2000 suggests it is worth considering the implications of inflation targeting in 2001.

Table 1. Frequency of Policy Interventions

	0% Band	1% Band	2% Band	∞ Band
A. 1983 Experiment				
Average interventions per 24-month trial	24.0	13.1	9.6	0
From above band	14.8	12.5	9.5	NA
From below band	9.2	0.6	0.1	NA
Average maximum consecutive interventions	24.0	12.0	9.2	0
Percentage of trials with any intervention	100.0	93.4	73.0	0
B. 2001 Experiment				
Average interventions per 24-month trial	24.0	8.2	1.7	0
From above band	17.0	8.1	1.7	NA
From below band	7.0	0.1	0.0	NA
Average maximum consecutive interventions	24.0	7.7	1.6	0
Percentage of trials with any intervention	100.0	61.2	18.8	0

Summary statistics and basic results for the four bandwidths for both periods are presented in Tables 1 and 2.²⁸ As detailed earlier, for each period and each bandwidth, these results are from 1000 trials starting with draws from the estimated residuals. Note that while the FAIR relative to the inflation band is used as the criterion of whether to intervene in a particular month, in order to be comparable to inflation data as commonly reported, the inflation statistics from our experiments reported in Tables 1 and 2 are for the underlying inflation rates for each particular month rather than the FAIR. The results in these tables are based on our 24-month characterization of Blinder's "planning horizon." That is, the information in these tables is analogous to the FOMC Greenbook forecasts on the longer-run implications of the current or proposed policy path.

Table 1 includes basic results. We note three main points. First, for both experiments, the percentage of trials with any intervention and the average number of interventions per trial both fall as the width of the bands increases. This result is as expected; the wider the inflation band the more likely the FAIR falls within the band and the less likely an intervention is needed to maintain average inflation inside the band. The number of interventions per trial starts at a maximum of 24 months (the planning horizon) when the bandwidth is zero and is zero when the band is arbitrarily wide (in which case it is not necessary to intervene). The percentage of trials with any intervention for the 1% and 2% bands is substantially smaller in the 2001 experiment than in the 1983 experiment, and for each of these bands the average number of interventions is substantially lower in 2001. As bandwidth increases, the decline in both the percentage of trials with any intervention and the average number of interventions is faster in the 2001 experiment than the 1983 experiment.

Second, for both samples, the average number of interventions from above the band is substantially greater than from below the band. Given the inflationary pressures suggested by

²⁸ We have excluded from the statistics in Tables 1 and 2 those trials in which a negative (nominal) interest rate would occur. Generally, the results that include trials in which negative interest rates occur are nearly identical to those reported below. Note that negative nominal rates do show up in real-world data on occasion. For example, Cecchetti (1988) discusses negative nominal interest rates on some Treasury securities in the 1930s and, more recently, Fleming and Garbade (2004) discuss repurchase agreements with negative interest rates. Casual analysis of our trials in which negative interest rates occur suggests that they were about the same order of magnitude as those that appeared in Cecchetti (1988) and Fleming and Garbade (2004).

Table 2. Detailed Analysis of Shocks

	1% Band		2% Band	
	Random Residuals	Policy Residuals	Random Residuals	Policy Residuals
A. 1983 Experiment				
5th percentile	-0.007	-0.006	-0.007	-0.005
95th percentile	0.009	0.020	0.009	0.015
Mean policy residual		0.0045		0.0036
B. 2001 Experiment				
5th percentile	-0.006	-0.004	-0.006	-0.004
95th percentile	0.008	0.013	0.008	0.013
Mean policy residual		0.0043		0.0037

the base projections in the two experiments, it is not surprising that for the $\pm 1\%$ and $\pm 2\%$ bands, the number of interventions needed to restrain inflation (positive policy shocks to the interest rate equation) outnumber the interventions needed to stimulate inflation in order to maintain inflation within the bands.

Third, for both samples, the average maximum number of consecutive interventions per trial falls as the width of the band increases, and the average maximum number of consecutive interventions for the 1% and 2% bands is much smaller for the 2001 experiment than the 1983 experiment. Consecutive interventions stem from our imposition of a mechanical “commitment” to the inflation target objective. In our experiments, there is no option for the policy maker to deviate from this objective when computing the intervention.²⁹ Once the inflation rate breaches the edge of the inflation band, depending on the inflation inertia several policy shocks may be needed to return average, long-run inflation to the specified level. Also note that even when there is a pattern of several consecutive interventions, given an initial intervention, it is less clear that there will necessarily be inertia in the interest rate itself, since (i) the interventions are partly a function of the random draws for all the variables, which can entail consecutive interventions but not necessarily of the same sign, and (ii) there is an endogenous component to the funds rate equation over and above the intervention term.³⁰

Table 2 provides a more detailed look at the policy interventions. Recall that each trial is initiated with a random draw from the estimated residuals, replacing the interest rate shock with a computed policy residual whenever the FAIR is outside the bounds of the inflation target band. The estimated residuals are, of course, zero mean since they are the result of ordinary least squares (OLS) regressions. The computed policy residuals need not be zero mean, and in our setting where the general pattern is for inflation to breach the upper bound of the target range (see Figure 2) the average policy intervention is a positive interest rate shock to restrain the economy and lower the average inflation rate. In Table 2, we take the entire set of residuals with which we end up from the trials for the 1% and 2% bands and separate them into those that were drawn from the historical residuals and the policy residuals we computed in

²⁹ See Equations A2 and A3 in the Appendix.

³⁰ Because each trial in each experiment that eventually builds up to the variability frontier always selects the policy innovation needed to attain the inflation goal, there is a risk of instrument instability. While our approach does not rule out instrument instability, it is also possible that inertia in the interest rate will result. It turns out that there is substantial smoothing of the interest rate in our counterfactual simulations, even though we have not imposed any features that would explicitly limit the magnitude of interest rate movements.

Table 3. Standard Deviations of Key Variables

	0% Band	1% Band	2% Band	∞ Band
A. 1983 Experiment				
Output	0.01984	0.01428	0.01047	0.00819
Inflation	0.00605	0.00883	0.01178	0.01492
Interest rate	0.04519	0.02636	0.01586	0.01167
B. 2001 Experiment				
Output	0.01692	0.01117	0.00967	0.00963
Inflation	0.00435	0.00614	0.00723	0.00733
Interest rate	0.04022	0.01489	0.01076	0.01039

order to keep average inflation at the acceptable level.³¹ For the 5th and 95th percentile shocks (used to avoid any outliers) we find that the 95th percentile policy shock is about one and a half to two times the size of the comparable estimated residual. Furthermore, note that the average policy shock is well within the range of the shocks of the estimated residuals. Policy innovations outside the historical normal range occur, but not on average.

Table 3 shows the fundamental results: the standard deviations of the key variables for each bandwidth of each experiment. These standard deviations, also plotted in Figures 3 and 4, show the basic results of the article: the estimated tradeoffs over time between inflation and output variability and inflation and interest rate variability. We note, for both periods, that as we move from the stochastic simulation of the historical policy in the estimated VAR (the infinity band, where no interventions are called for) to the 2% target band, inflation variability falls and output variability rises since more tightly controlling inflation implies the policy maker must accept the higher output volatility as a consequence. The absolute value of the change in output variability is less than the change in inflation variability—their ratio is 0.7 in both experiments. As we move from the 2% band to a 1% band, the absolute value of the ratio of the change in output variability to the change in inflation variability is 1.3 for the 1983 experiment and 1.4 for the 2001 experiment. Finally, as we move from a 1% band to achieving the target without variability (the 0% band), the absolute value of the ratio of the change in output variability to the change in inflation variability is two for the 1983 experiment and three for the 2001 experiment. The point on the tradeoff that would be chosen by the central bank clearly depends on the central bank's preferences regarding output and inflation variability—the point that, in the earlier words of Svensson, “looks best.”³² Our technique provides a data-driven methodology for estimating the objective tradeoff faced by the central bank.³³

Figure 3 plots the tradeoff data from Table 3 and shows a roughly parallel downward and inward shift of the inflation-output variability tradeoff over time. For example, in the 2001 experiment, inflation variability equal to 0.006 (1% band) is associated with output variability

³¹ Note that for the 0% band, all the innovations are (almost certainly) policy innovations, and for the arbitrarily wide band, there are no policy innovations, so the 1% and 2% bands discussed in Table 2 are the only relevant comparisons.

³² Implicitly, movement along the tradeoff reflects alternative policy maker preferences. In Svensson's setup as summarized in Equation 4, the tradeoff between inflation and output is a function of λ , Svensson's relative weight on output variability. An analogous point holds here where we have instead normalized the implicit loss function on inflation variability as expressed in Λ in section 2 above.

³³ Results analogous to those in Table 3 for the case of a 36-month transition horizon are included in Appendix Table A2. The implications are virtually the same as those reported in the text for the 24-month case.




Figure 3. Inflation-Output Standard Deviations over Time: 2% Target

of 0.011, whereas for the 1983 experiment output variability implicitly associated with the same inflation variability is approximately 0.019 (0% band), about 70% higher than for the 2001 experiment.

In Figure 4, similar results are found for the inflation-interest rate variability tradeoffs. Smaller inflation-targeting bands are associated with greater interest rate variability since more frequent and larger interest rate adjustments are required for more precise inflation control. As in Figure 3, the inflation-interest rate variability tradeoff also shifts inward in a roughly parallel

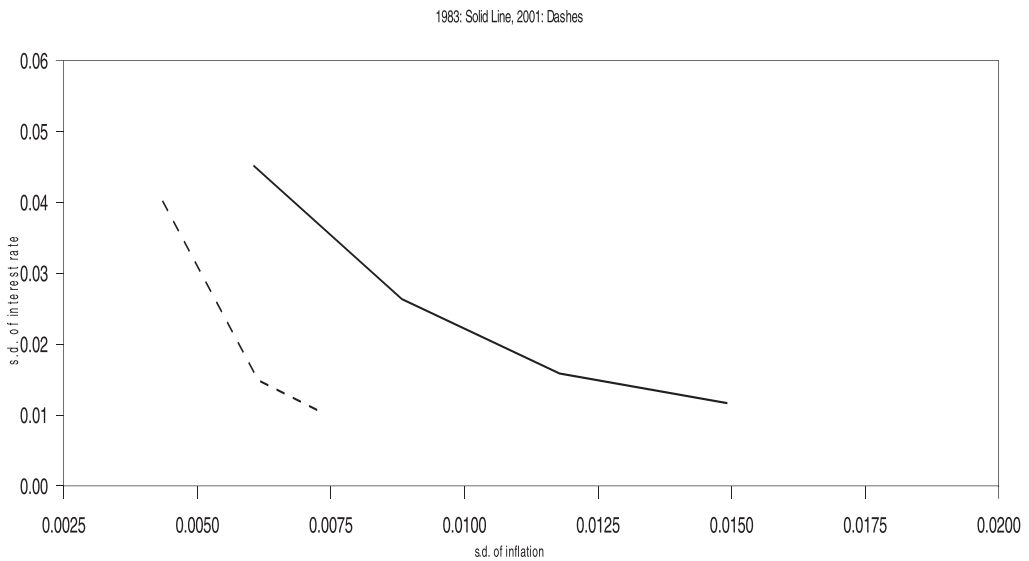


Figure 4. Inflation-Interest Rate Standard Deviations over Time: 2% Target

fashion; for given inflation variability, we observe much lower interest rate variability in the second period.

Given the differences in methodology, data, and estimation periods, only rough comparisons between our analysis and other work can be made. A particularly difficult methodological issue is the translation of weights in a loss function into explicit values for the width of inflation target bands; absence of an exact correspondence between the width of target bands and specific weights on output stabilization make direct comparison with research based on loss function weights difficult. However, we will discuss our variability tradeoff with another in the literature that, as closely as we can find, approximates our presentation.

Rudebusch and Svensson (1999) (RS) start with a loss function that has inflation deviations from target, the output gap, and the change in the interest rate as arguments. They minimize this loss function subject to a model of the economy that includes a Phillips curve/aggregate supply specification and an IS/aggregate demand equation. Policy focuses on the relevant feedback function for the nominal interest rate in the IS curve. Among the 11 simulated policy rules (some with several variants) that are compared with the optimal setting derived from the formal minimization are two that focus on inflation forecast targeting, their “flexible inflation forecast targeting (FIFT)” and “flexible inflation forecast targeting with smoothing (FIFTS)” policy rules. Their FIFT rules assume no interest smoothing, while the FIFTS rules allow various degrees of such smoothing. Since our model does not preclude such smoothing and our interest rate paths show that interest rate variations are relatively small (i.e., the data seem to select a path for the policy interest rate that exhibits patterns consistent with partial adjustment), we will briefly compare our results with the RS FIFTS model.

Since we have targeted a two-year inflation forecast, and given that RS use quarterly data, our comparison focuses on their eight-quarter forecast rule (FIFTS(8)). As noted in Table 3 and as plotted in Figure 3, when the standard deviation of inflation in our experiments is 0.006 (for the 1% bandwidth in the 2001 experiment and for the 0% bandwidth for the 1983 experiment) output variability was 0.011 (2001) and 0.019 (1983). From the RS results in their tables 5.3–5.7 (using various sets of weights on loss function arguments), for about the same inflation variability,³⁴ the RS output variability³⁵ varies from 0.0187 (0.0066 inflation variability) to 0.0267 (0.0050 inflation variability). We also note that as a rough point estimate, when we examine their figure 5.2, which plots the variability tradeoff for FIFTS(8) as well as other rules, we see that inflation variability of 0.006 (2.4 in the figure) is associated with output variability of about 0.002. Thus, our results seem reasonably close to theirs for this particular comparison, given the differences in methodologies and differences in sample periods for estimation. However, as noted earlier, conclusions from this comparison should be tempered by the fact that there is no straightforward way to directly translate loss function weights in their study to the bandwidths in ours.

Finally, an important question is “What is the source of the inward shift in the inflation-output variability tradeoff between the two periods that we find?” Although we cannot definitively identify the reason(s) for the shift, we note that much of the sample used in the estimation of the VAR for the later period spans the period of the “Great Moderation” in

³⁴ RS compute the average inflation rate as $400 \times [\ln p_t - \ln p_{t-1}]$. We translate their inflation standard deviation into one approximately equal to ours by dividing their standard deviations by 400. Across their experiments, their inflation standard deviation ranges from 0.0050 to 0.0066.

³⁵ RS use the output gap in percentage points, while we use decimal form. Thus, dividing their output gap standard deviations by 100 produces roughly comparable results.

macroeconomic volatility. Explanations for the Great Moderation include good luck in terms of reduction in the frequency and magnitude of economic shocks and better monetary policy, and evidence has been presented in support of both explanations. The consensus from studies such as Stock and Watson (2002), Ahmed, Levin, and Wilson (2004), and Sims and Zha (2006a) is that “good luck” in the form of reduced variability of non-monetary policy shocks is primarily responsible for the increased stability of the real economy since the mid-1980s. However, Stock and Watson (2002) attribute from 10% to 25% of the reduced variability in the real economy to improved monetary policy, and Ahmed, Levin, and Wilson (2004) find that, although “good policy” doesn’t seem to explain much of the reduced variability in real output, it is important in understanding the reduction in inflation variability.

Regarding the “good policy” explanation, we note that much of the recent literature on monetary policy suggests that inflation targeting allows central banks to gradually gain credibility. Clarida, Gali, and Gertler (1999) suggest that credible policy “enables the central bank to stabilize the economy with relatively modest movements in the short rate” (pp. 1689–90), and Carlstrom and Fuerst (2005) present simulations showing that central bank credibility allows the Fed to achieve given objectives with smaller policy interventions than in the case where credibility is lacking. Cecchetti, Flores-Laguna, and Krause (2006), based on a cross-section of both developed and developing countries, find that, for most countries, monetary policy was more efficient in the 1990s than in the 1980s and that better monetary policy led to a significant improvement in macroeconomic performance.³⁶

Unfortunately, our technique does not allow us to distinguish between the “good luck” and “good policy” explanations, but we note the fact that the tradeoff can shift substantially over time suggests that central banks cannot regard the tradeoff as fixed over time.

6. The Counterfactual Experiments and the Lucas Critique

An important issue in any counterfactual experiment is whether the Lucas critique would invalidate the results. With respect to the first experiment, we noted earlier that the 1962:1–1983:9 period over which the model was estimated was a period in which there was a great deal of variation in output and inflation, especially during the immediate post-October 1979 months at the end of the estimation period for this experiment. Romer and Romer (2002) characterized the estimation period as one in which the norm was substantial fluctuations in monetary policy variables, and Sims and Zha (2006a) found the 1979–1983 period to be distinct. The period covered by the estimation sample for our second experiment, ending in December 2000, appears to be more settled in terms of policy, with interest rate targeting effectively governing most of the period, culminating with substantial attention to the Taylor rule and its role in the Great Moderation. As an initial attempt to examine the importance of Lucas-critique problems, following Dufour (1980, 1982), we re-estimated the model over 1962:1–1983:9 and 1980:1–2000:12, adding a 0–1 dummy variable for each month in which instability is suspected. For the first sample, we added a separate dummy variable for each month of the post-October 1979 period included in the sample, 1979:10 through 1983:9, the period of nonborrowed reserve

³⁶ In their study, Cecchetti, Flores-Laguna, and Krause (2006) estimate inflation variability–output variability tradeoffs derived from a small structural model and an explicit loss function.

Table 4. Modesty Statistics

	0% Band (y, p, i)	1% Band (y, p, i)	2% Band (y, p, i)	∞ Band (y, p, i)
1983 Experiment	0.92, 0.53, 0.71	0.81, 0.48, 0.68	0.70, 0.49, 0.65	0.75, 0.61, 0.52
2001 Experiment	1.54, 0.71, 0.76	1.25, 0.66, 0.65	1.06, 0.70, 0.53	1.12, 0.77, 0.38

targeting plus an approximate one-year transition period from the end of nonborrowed reserve targeting. For the second sample, we included separate monthly dummies for 1980:1 (the first observation in the estimation period) through 1983:12 (roughly the period of nonborrowed reserve targeting in the second sample and an approximate one-year transition period from the end of nonborrowed reserve targeting). Following Sims and Zha (2006b), we used the Akaike Information Criterion (AIC) and the Schwarz Information Criterion (SIC) to compare the models with and without the dummies and thereby to provide information about the stability of the model. The AIC indicated the model with dummies was preferred for both samples, hence indicating instability. However, the SIC suggested that the model without the dummies was preferred for both samples. Thus, evidence regarding the stability of the model over the estimation periods is mixed.

While stability tests can usefully shed light on potential Lucas concerns within the estimation period, in our real-time setting the policy maker is concerned in addition with whether a proposed policy action will trigger among agents in the economy the perception that a proposed policy would be interpreted as a regime shift. If so, then the variability tradeoffs we presented in Figures 3 and 4 would be of little use in evaluating the available tradeoffs to policy makers. To this end, Leeper and Zha (2003) have introduced a “modesty statistic” intended to evaluate whether a prospective policy initiative is likely to be viewed as a modest policy intervention.

The Leeper and Zha theoretical approach is a Markov-switching model, with each regime a linear model of the economy (a VAR in their case). Within a regime, the effect of a policy intervention is as described by the first term on the right-hand side of our Equation 10, the impact of the proposed policy relative to the base projection. Specifically, picking a policy sequence $\{\varepsilon_{t+1}, \varepsilon_{t+2}, \dots, \varepsilon_{t+n}\}$, computing the expression $\sum_{s=0}^{n-1} D_s \varepsilon_{t+n-s}$, and then scaling by $\sqrt{\sum_{s=0}^{N-1} D_s^2}$ provides their “modesty statistic.”³⁷ Leeper and Zha (2003) note that the “modesty statistic” has a standard normal distribution, so when this statistic is less than two, the policy innovation embedded in the ε path over $t+1$ to $t+n$ does not induce agents to change their assessments about the policy regime in place.³⁸

We have computed the Leeper-Zha modesty statistic in the context of our counterfactual experiments, modifying their approach by using the randomly drawn disturbances to the other equations under which our policy interventions are computed rather than assuming that the shocks to the non-policy equations are all zero. Under this condition, in the trials underlying our basic results in Figures 3 and 4, the largest computed modesty statistics (in absolute value) for each experiment are presented in Table 4, where in each cell the relevant statistics for the

³⁷ Consistent with our approach, Leeper and Zha use the ε shock to the policy equation as the policy innovation and assume as we do that “although the policy advisor chooses [the ε -innovation], private agents treat it as random” (Leeper and Zha, p. 1678).

³⁸ Of course, alternative policy regimes can be “close” to each other, so that distinguishing between these regimes may be difficult. Thus, a modesty statistic of less than two is necessary but not sufficient to claim that no important Lucas effects are present.

impact of the policy sequence for output (y), the inflation rate (p), and the interest rate (i) are given, respectively. All the computed statistics are well below 2.0, so that there is no evidence that our proposed inflation bands and the policies needed to achieve them would have triggered Lucas critique–type concerns among agents in the economy. This implies, importantly, added confidence that the tradeoffs presented in Figures 3 and 4 can be reasonably interpreted as valid tradeoffs available to policy makers.

7. Summary and Conclusion

Our focus in this article is twofold: (i) illustration of how a VAR model can be used to implement and evaluate inflation forecast targeting and (ii) the derivation in real time of the output-inflation variability tradeoff available to the central bank under inflation forecast targeting and estimation of how this tradeoff has changed over time. Tolerance bands of varying widths around transitional inflation targets constructed to achieve 2% inflation are considered.

Our inflation forecast targeting approach is based on dynamic, stochastic simulations of the average inflation rate over a two-year horizon using the moving average representation of the VAR model. Deviations of the forecast average inflation rate from target generate interventions in the form of changes in the federal funds rate designed to gradually push the forecast inflation rate back to target, and we compute the required adjustments to the federal funds rate.

The technique is illustrated through two counterfactual experiments using real-time data. The first experiment begins in 1983:10 and is based on a VAR estimated over 1962:1–1983:9, whereas the second experiment begins in 2001:1 using a model estimated over 1980:1–2000:12. In terms of technical results, we find (i) fewer interventions are needed as the width of the bands increases, and fewer interventions are needed in the 2001 experiment than in the 1983 experiment; (ii) more interventions are needed to reduce the inflation rate than are needed to raise the inflation rate, reflecting the inflationary pressures during the periods, which in turn makes targeting exercises important; (iii) fewer consecutive interventions are required as the width of the target band increases, and fewer consecutive interventions are needed for the 2001 experiment than the 1983 experiment; and (iv) a given inflation variability is associated with lower output and interest rate variability in the 2001 experiment than in the 1983 experiment; that is, the tradeoffs between inflation variability and output variability and inflation variability and interest rate variability shifted favorably between our time periods. As noted earlier, the source of the shift might stem from “good luck” or “good policy,” but a critical lesson is that significant shifts in the tradeoffs have occurred over time and hence that the objective tradeoffs faced by the central bank should not be assumed to be fixed.

References

- Ahmed, Shaghil, Andrew Levin, and Beth Anne Wilson. 2004. Recent U.S. macroeconomic stability: Good policies, good practices, or good luck? *Review of Economics and Statistics* 86:824–32.
- Barlevy, Gadi. 2009. Policymaking under uncertainty: Gradualism and robustness. Federal Reserve Bank of Chicago. *Economic Perspectives* 33:38–55.

- Bernanke, Ben S. 1986. Alternative explanations of the money-income correlation. *Carnegie-Rochester Conference Series on Public Policy* 25:49–100.
- Bernanke, Ben S. 2004. The logic of monetary policy. Remarks before the National Economists Club, Washington, D.C., December 2.
- Bernanke, Ben S., and Alan S. Blinder. 1992. The federal funds rate and the channels of monetary transmission. *American Economic Review* 82:901–21.
- Bernanke, Ben S., and Ilian Mihov. 1998. Measuring monetary policy. *Quarterly Journal of Economics* 113:869–902.
- Bernanke, Ben S., and Frederic S. Mishkin. 1997. Inflation targeting: A new framework for monetary policy? *Journal of Economic Perspectives* 11:99–116.
- Blinder, Alan S. 1997. What central bankers could learn from academics—and vice versa. *Journal of Economic Perspectives* 11:3–19.
- Brainard, William C. 1967. Uncertainty and the effectiveness of policy. *American Economic Review* 57:411–25.
- Carlstrom, Charles T., and Timothy S. Fuerst. 2005. Considerable period of time: The case of signalling future policy. Federal Reserve Bank of Cleveland. *Economic Commentary*. November.
- Cecchetti, Stephen G. 1988. The case of the negative nominal interest rates: New estimates of the term structure of interest rates during the Great Depression. *Journal of Political Economy* 96:1111–41.
- Cecchetti, Stephen G., Alfonso Flores-Laguna, and Stefan Krause. 2006. Has monetary policy become more efficient? A cross-country analysis. *Economic Journal* 116:408–33.
- Christiano, Lawrence J. 1998. Discussion of Christopher A. Sims' "The role of interest rate policy in the generation and propagation of business cycles: What has changed since the 1930s?" In *Beyond shocks: What causes business cycles*, edited by Jeffrey C. Fuhrer and Scott Schuh. Boston: Federal Reserve Bank of Boston, pp. 161–70.
- Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. The science of monetary policy: A new Keynesian perspective. *Journal of Economic Literature* 37:1661–1707.
- Cogley, Timothy, and James M. Nason. 1995. Output dynamics in real business cycle models. *American Economic Review* 85:492–511.
- Dufour, Jean-Marie. 1980. Dummy variables and predictive tests for structural change. *Economics Letters* 6:241–47.
- Dufour, Jean-Marie. 1982. Generalized Chow tests for structural change: A coordinate free approach. *International Economics Review* 23:565–75.
- Fackler, James S., and John H. Rogers. 1995. Output, inflation and stabilization: A counterfactual analysis. *Journal of International Money and Finance* 14:619–40.
- Fackler, James S., and W. Douglas McMillin. 2002. Evaluating monetary policy options. *Southern Economic Journal* 68:794–810.
- Fernandez-Villaverde, Jesus, Juan F. Rubio-Ramirez, Thomas J. Sargent, and Mark W. Watson. 2007. ABCs (and Ds) of understanding VARs. *American Economic Review* 97:1021–26.
- Fleming, Michael J., and Kenneth D. Garbade. 2004. Repurchase agreements with negative interest rates. Federal Reserve Bank of New York. *Current Issues in Economics and Finance* 10:1–7.
- Flood, Robert P., and Peter Isard. 1989. Monetary policy strategies. *IMF Staff Papers* 36:612–32.
- Fraser, B. W. 1993. Some aspects of monetary policy. Reserve Bank of Australia. *Bulletin* April:1–7.
- Fraser, B. W. 1994. The art of monetary policy. Reserve Bank of Australia. *Bulletin* October:17–25.
- Goodfriend, Marvin. 2005. Inflation targeting in the United States? In *The inflation-targeting debate*, edited by Ben S. Bernanke and Michael Woodford. Chicago: University of Chicago Press, pp. 311–37.
- Ireland, Peter. 2004. A method for taking models to the data. *Journal of Economic Dynamics and Control* 28:1205–26.
- Leeper, Eric M., and Tao Zha. 2003. Modest policy interventions. *Journal of Monetary Economics* 50:1673–1700.
- McCallum, Bennett T. 1988. Robustness properties of a rule for monetary policy. *Carnegie-Rochester Conference Series on Public Policy* 29:173–204.
- Orphanides, Athanasios, David H. Small, Volker Wieland, and David W. Wilcox. 1997. A quantitative exploration of the opportunistic approach to disinflation. Federal Reserve Board of Governors Finance and Economics Discussion Series Paper 1997-36.
- Romer, Christina D., and David H. Romer. 2002. The evolution of economic understanding and postwar stabilization policy. In *Rethinking stabilization policy*. Kansas City: Federal Reserve Bank of Kansas City, pp. 11–78.
- Rudebusch, Glenn D., and Lars E. O. Svensson. 1999. Policy rules for inflation targeting. In *Monetary policy rules*, edited by John B. Taylor. Chicago: University of Chicago Press, pp. 127–54.
- Sims, Christopher A., and Tao Zha. 2006a. Were there regime switches in U.S. monetary policy? *American Economic Review* 96:54–81.
- Sims, Christopher A., and Tao Zha. 2006b. Vintage article: Does monetary policy generate recessions? *Macroeconomic Dynamics* 10:231–72.
- Stock, James H., and Mark W. Watson. 2002. Has the business cycle changed and why? *NBER Macroeconomics Annual* 17:159–218.

- Svensson, Lars E. O. 1997. Inflation forecast targeting: Implementing and monitoring inflation targets. *European Economic Review* 41:1111–46.
- Svensson, Lars E. O. 1999. Inflation targeting as a monetary policy rule. *Journal of Monetary Economics* 43:607–54.
- Svensson, Lars E. O. 2003. What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules. *Journal of Economic Literature* 41:426–77.
- Svensson, Lars E. O. 2009. Transparency under flexible inflation targeting: Experiences and challenges. *Sveriges Riksbank Economic Review* 1:5–44.
- Taylor, John B. 1994. The inflation/output variability tradeoff revisited. In *Goals, guidelines, and constraints facing monetary policymakers*. Proceedings of a conference held in June 1994. Boston: Federal Reserve Bank of Boston, pp. 21–38.