

**For release on delivery
1:30 p.m. CST (2:30 p.m. EST)
January 16, 2002**

Rules and Discretion

Remarks by

Laurence H. Meyer

Member, Board of Governors of the Federal Reserve System

before the

Owen Graduate School of Management

Vanderbilt University

Nashville, Tennessee

January 16, 2002

An extensive literature addresses the question of whether it is preferable to implement monetary policy by a rule or by discretion. This question has traditionally been referred to as the issue of rules *versus* discretion.

In a strict interpretation of a rules-based regime, policymakers commit to how they will adjust their policy instrument in response to incoming data or to changes in the forecast. Once this rule is specified, their judgment no longer is relevant to the policy outcomes. In a discretionary regime, policymakers do not commit in advance to a specific course of action and instead apply their judgment, deciding on each occasion what policy is appropriate.

Even if we cannot imagine policymakers turning over the conduct of policy to a rule, research on rules might provide guidance to policymakers that could improve their judgmental adjustments to policy. I strongly believe this is the case. No one would argue, after all, that good policy is whimsical. On the contrary, good policy should be systematic. Good discretionary policy therefore should be, in some meaningful way, rule-like, though it might be impossible to write down in a simple or even complicated equation all the complex considerations that underpin the conduct of such a systematic monetary policy. In this spirit, I focus on rules *and* discretion, specifically how the study of rules can be helpful in informing discretionary policy decisions. This complementarity between rules and discretion follows a direction encouraged by John Taylor (1993) and is a central theme of his papers on rules.

I begin with a discussion of the ingredients of monetary policy: objectives, instruments, models, and strategy. Next I discuss the pros and cons of a commitment to a

rule and, in a discretionary regime, the advantages of a transparent and systematic policy. Then I introduce the Taylor rule as a simple example of a policy rule, one that has had considerable influence on how the public and indeed monetary policymakers think about monetary policy.

Although the Taylor rule has been a useful benchmark for policymakers, my experience during the last 5-1/2 years on the FOMC has been that considerations that are not explicit in the Taylor rule have played an important role in policy deliberations. In particular, forecasts clearly have played a powerful role in shaping the response of monetary policy in a way not reflected in the simple Taylor rule. Uncertainty about the estimate of the NAIRU has influenced the timing and aggressiveness of response of monetary policy to movements in the unemployment rate. In addition, changes in the equilibrium real interest rate and the potential implications of a very low nominal funds rate have been important considerations in policy deliberations. These concerns have motivated much of the research I discuss in the remaining sections of the paper. To my mind, this research has helped to refine the policy strategies embedded in rules, while maintaining the fundamental wisdom of the simple Taylor rule. This paper will highlight a few of the lessons that can be learned from this research.

I. Objectives, Instruments, Models and Strategy

Let's begin by identifying the objectives and instruments of monetary policy. In broad terms, the Congress has given the Federal Reserve a two-part or dual mandate--to promote "maximum employment" and to foster long-run price stability. The price stability objective has come to be interpreted as a low, stable rate of inflation. Maximum employment is usually interpreted as maximum "sustainable" employment, meaning the

maximum level of employment sustainable without upward pressure on inflation. This interpretation of the employment objective ensures it is compatible with the price stability objective in the long run.

A division of labor between monetary and fiscal policies has evolved over the last couple of decades, reflecting the constraints imposed by both economic theory and political realities. Monetary policy has been given full responsibility for achieving price stability. This reflects the widespread acceptance of the theoretical propositions that inflation in the long run is principally, if not exclusively, a monetary phenomenon and that price stability (as opposed to significant inflation or deflation) is the environment most conducive to achieving a high, sustainable rate of growth in output over time. With respect to aggregate demand management, fiscal policy still plays an important role through the operation of the automatic stabilizers (tax and to a more limited extent spending responses to changes in the level of activity), but discretionary fiscal policy has generally been less nimble than monetary policy and hence had not been used or even discussed much if at all until last fall. While generally relegated to the sidelines in the face of standard cyclical fluctuations, fiscal policy can be an important complement to monetary policy in the face of very large shocks or when interest rates are already low, limiting the scope for monetary stimulus.

Economic theory--specifically the neutrality of money--indicates the limits to any influence of monetary policy on the level and growth rate of real GDP in the long run, other than through achieving long-run price stability. At the same time, theoretical analysis and empirical evidence supports a role for fiscal policy--through control over the

structural budget deficit and through the structure of marginal tax rates and tax incentives--in influencing the level and the rate of growth of real living standards.

The instrument of monetary policy in the United States is the control over the level of reserves through open market operations. Almost all central banks use this or other instruments as a means of holding some very short-term interest rate as close as possible to a target level. In the case of the Fed, the target is set for the federal funds rate, an overnight rate on reserves lent from one bank to another. Because the Fed's ability to control the funds rate is very good, that rate, in effect, is the instrument of monetary policy. I will therefore focus only on rules that describe how the federal funds rate should be set.

The strategy of monetary policy involves thinking about how the federal funds rate should be adjusted in response to incoming data and/or evolving forecasts in order to promote the objectives of monetary policy. The best approach to making such adjustments will depend on the model that policymakers believe describes the determination of output and inflation. For example, if--as I will assume--that model indicates that the balance of potential aggregate supply and aggregate demand in the output and/or labor markets is a proximate source of upward or downward pressure on the inflation rate, then proxies for that balance will play a role in informing policymakers' views about appropriate adjustments in the policy rate. Another key property of many macro models is that, while there is no trade-off between inflation and output in the long run, there is a trade-off between output variability and inflation variability. The policy rule and guidelines that are suggested by the policy rule provide

direction to policymakers about how to balance their two objectives in light of this trade-off.

II. Commitment and Rules versus Discretion

A strict rules-based policy establishes an unequivocal commitment by policymakers to achieve their policy objectives, especially to meeting their inflation objective. Such a commitment, in turn, increases the transparency and accountability of monetary policy and thereby helps to pin down inflation expectations. In principle, the resulting credibility about policymakers' commitment to price stability could reduce the cost of disinflation, if inflation were to rise above the objective, and could reduce the spillover of supply shocks--that is autonomous shocks to the price level-- into broader price movements.

The gains from a strict rule-based policy arise in part from the ruling out problems of “time inconsistency.”¹ Time inconsistency refers to the incentive of policymakers to commit to one policy and then later to pursue another, different policy that is inconsistent with that commitment. Specifically, this view builds on the observation that, at least in certain models, monetary policymakers have an incentive to convince private agents of a commitment to price stability and then “cheat” on their commitment by driving the unemployment rate down, at the expense of higher inflation in the future. However, the public understands the temptation on the part of the policymaker to cheat at any given moment and so expects higher inflation than would be the case if policymakers were willing and able to commit themselves to follow a rule.

¹ The time-inconsistency approach is developed in Kydland and Prescott (1977) and Barro and Gordon (1983).

The time inconsistency issue seems germane to problems of a political business cycle. However, I have never found it to be a convincing description of the conduct of real-world monetary policymakers in countries where the central banks have a high degree of independence from the rest of government. The point of an independent central bank, after all, is to reduce the political pressures that could produce the time inconsistency result. Indeed, those that worry about time inconsistency may not fully appreciate the extent to which the Federal Open Market Committee might represent a form of the commitment technology they advocate. Absent such political influences, it is not evident that independent monetary policymakers would prefer an unemployment rate below the non-accelerating inflation rate of unemployment (NAIRU), as is often assumed in the time-inconsistency literature.

A more nuanced view of policy--and a more realistic one--is that monetary policymakers recognize that they play a repeated game with learning by private-sector decision-makers. Policymakers therefore pursue a strategy that uses that learning to make their policy more effective--for example, by providing information to guide expectations over time. In order to use learning by private decision-makers, policy needs to be transparent and systematic, what I have been calling rule-like. The combination of transparent and consistent policy responses allows the financial markets to better anticipate future policy actions. This, in turn, results in long-term rates moving more quickly and assuredly in response to changing economic conditions, in the expectation that monetary policy action will follow. The articulation of the logic implicit in a simple rule--the description of the normal course of policy--allows policymakers to communicate the broad rationale for their policy actions, and, at the same time,

regularizes policy responses to changing economic conditions. However, both policymakers and private decision-makers recognize that views of the state of economy evolve over time, as do more fundamental views of the correct model of the economy. Therefore, while monetary policy can follow a rule-like behavior, it can and should avoid the straightjacket of a quarter-to-quarter commitment to a strict rule. At the same time, while rule-like behavior allows policymakers to retain the flexibility to deviate from the rule, transparency would call for a special effort to explain the rationale for such deviations.

No one policy rule can anticipate the appropriate response to all possible circumstances before they arise. And, even if one could, the rule would be optimal only in the model for which it was designed. The research supporting the efficiency of commitment depends on some very strong assumptions, including the time-invariance of the model. For example, Westaway (1989) shows that if the economy undergoes an unforecastable persistent shift in some variable--for example, potential output--commitment to an ex ante optimal rule can be worse than discretion.² This seems particularly relevant to the economy in the second half of the 1990s when policymakers were adapting to new information about the underlying rate of productivity growth. A systematic monetary policy, informed by policy rules but flexible enough to adapt to structural changes and other real-world complexities, is therefore, in my view, the best direction for monetary policy.

² Even here, one could think of a rule that would incorporate how to update estimates of the effective NAIRU and potential output and hence respond to at least some kinds of structural change. But it would be difficult, indeed virtually impossible, to write a rule that incorporated all the possibilities for responding to structural changes, including changes in model specification.

III. The Taylor Rule

John Taylor (1993) introduced a simple rule that has deservedly received a lot of attention, both from policymakers and from others. An important aspect of Taylor's perspective is that he didn't advocate that policymakers commit to the rule, but rather use it to inform their discretionary decisions. The Taylor rule specifies how policymakers should set the level of the nominal federal funds rate, i . The key feature of the rule is the systematic response of the real federal funds rate to deviations of output and inflation from their respective targets.

The policymaker following this rule is aiming to stabilize inflation around its target and output around its so-called potential level. The objective for output is not a policy choice; rather it is dictated by the structure of the economy. It is sometimes referred to as potential output or the full employment level of output. Implicit in this concept is that an excess of output relative to its potential level pushes inflation higher; and of course, a deficiency of output relative to its potential level pushes inflation lower. Potential output is therefore the maximum level of output sustainable without upward pressure on inflation.

Lurking beneath the surface, but not very far below, is thus a version of the Phillips Curve, because the level of potential output, by Okun's Law, is closely connected to the concept of the NAIRU.³ Writing the level of real output as Y and potential output as Y^* , the deviation of output from its objective is usually expressed in terms of the percentage output gap, or percentage deviation of Y from Y^* . That is, $y = [(Y - Y^*)/Y^*]$ *

³ The connection between the output and unemployment gaps is often expressed in an equation referred to as Okun's Law: $y = a \text{UGAP}$, with a typically estimated at about 0.5. Empirically, versions of Okun's Law that allow for lags in the relationship between output and unemployment gaps seem to fit better.

100. In terms of the unemployment rate, the deviation of unemployment from its objective is $U - U^*$, where U^* is the NAIRU, or the level of the unemployment rate associated with full employment.

To define a corresponding gap for inflation, we have to identify an inflation objective. I will discuss how this should be set operationally below. For a given inflation target, p^* , we can then write the inflation gap as $p - p^*$, where p is the actual rate of inflation.

The Taylor rule assumes a constant equilibrium real federal funds rate, in effect a constant “neutral” level of that rate, appropriate when both price stability and full employment have been achieved. We will refer to this as r^* . When output or inflation deviate from their respective targets, policymakers vary the nominal funds rate (their instrument) to move the real funds rate relative to r^* .

The last ingredient for the Taylor rule is the numerical values of the response parameters that determine how much the real funds rate should be adjusted for a given deviation of output and inflation from their respective targets. Taylor assumed both parameters were equal to 0.5.

We can then write the Taylor rule as

$$(1) i = r^* + p + 0.5 y + 0.5 (p - p^*)$$

or, using the definition of the real interest rate, $r = i - p$,

$$(1') r = r^* + 0.5 y + 0.5 (p - p^*)$$

It is important to distinguish two possible uses for such a rule. First, the rule could have a *normative* application; that is, it could be used to prescribe systematic policy reactions to incoming data that research suggests have good stabilizing properties over

time across a number of different models. Second, the rule could be *descriptive*; that is, it might provide a reasonably close description of the way policymakers have made policy over some given period of time.

Taylor suggested this rule based on its normative qualities but then discovered it also described quite well the way policy had been made in the United States over the previous few years, specifically 1987-1992. Others have found that the Taylor rule also described the way policy had been made in a number of other countries.⁴

Taylor suggests two ways in which such a rule might be used at the Federal Reserve. First, policymakers could be given a memo prior to each FOMC meeting showing the funds rate path consistent with the Taylor rule and the staff's Greenbook forecast. Indeed, members of the FOMC do receive, in advance of each meeting, a one-page chart showing the funds rate prescribed by the simple Taylor rule, as well as the rate indicated by an estimated version of that rule.⁵

Taylor also suggested that policymakers could improve the outcome of their policy decisions by following guidelines that reflect the insights derived from the policy rule. This, it seems to me, potentially captures a lot of the value of monitoring the

⁴ See Clarida, Gertler and Gali (1998).

⁵ This latter exercise ignores deviations in rule-based policy from the policy path assumed in the Greenbook forecast. Such deviations are probably not very important for a few quarters, but they could generate significant inconsistencies further out. I believe that the FOMC would also benefit, therefore, from seeing an alternative simulation in which the FRB-US model was used to adjust the Greenbook forecast to make it consistent with the policy rule. In addition to providing the committee with a forecast consistent with the policy rule, the resulting forecast would also be useful to committee members when they prepare their semi-annual forecasts for the Monetary Policy report to the Congress that are supposed to reflect "appropriate" monetary policy.

prescriptions from specific policy rules. There are two important lessons that central bankers can learn directly from the simple Taylor rule.⁶

(1) Vary the real interest rate in response to deviations of inflation from its target.

A key insight from the Taylor rule for the nominal interest rate is that the coefficient on inflation must exceed one. Combining the two terms on inflation in equation (1), the coefficient in the simple Taylor rule is 1.5. If the coefficient were less than one in the equation for the nominal rate, an increase in inflation would be accompanied by a decline in the real interest rate, stimulating aggregate demand and reinforcing the upward pressure on inflation. This would be an unstable system. A coefficient greater than one ensures that higher inflation will be followed by a rise in the real interest rate that will, in turn, restrain aggregate demand and thereby contribute to restraining the rise in inflation. If there were only one lesson to teach central bankers about adjusting the policy rate, this would be it.

(2) Vary real and nominal interest rates in response to changes in resource utilization rates.

This guideline is relevant whether or not policymakers set an objective for the level of output. Resource utilization rates--such as the output gap or the unemployment gap--are an important factor in the dynamics of inflation. As long as policymakers have an inflation objective, they should respond to resource utilization rates as a leading indicator of inflation pressure. Of course, if policymakers are concerned directly with stabilizing output as well as inflation, then there is an added reason to respond to deviations of output from this objective. In this case, the response to resource utilization

⁶ These are two of the guidelines for policy that I discussed in an earlier paper, "The Strategy of Monetary Policy."

rates will be more aggressive than if utilization rates are only a leading indicator of inflation. Following this guideline would ensure that monetary policy leaned against the cyclical winds, with the real interest rate rising during periods of above-trend growth--even when the unemployment rate is still high--and with real interest rates falling when growth was below trend--even when the unemployment rate is still low.

IV. Methodology underlying Research on Policy Rules

Research on policy rules considerably predates Taylor's work, but Taylor's contribution renewed and invigorated this research. This research compares the performance under optimal and simple policy rules and under alternative parameter values in simple rules and investigates whether and under what circumstances performance can be improved by refinements to the specification suggested by Taylor. In this section I describe the basic methodology underlying some research on policy rules.

The first step in such research is to specify a loss function that describes the costs to society associated with deviations of output and inflation from their respective targets for each period of time. Such a loss function, for a given quarter, is usually specified as a weighted sum of squared deviations of output and inflation from their respective targets.

$$(2) L = y^2 + w (p - p^*)^2$$

where w is the relative cost associated with inflation compared to output deviations.

Squaring the deviations makes the cost associated with deviations depend only on their absolute magnitude and not their sign. In addition, it disproportionately penalizes larger deviations. In the dynamic setting in which monetary policy is typically analyzed, the total loss is a discounted sum of these weighted deviations over time.

The second step is to specify a model or a number of models of the economy.

Much recent research uses some version of a simple two-equation dynamic macro model.

The first equation is typically a dynamic specification of the standard IS curve, such as

$$(3) y_t = a + b y_{t-1} - c r_{t-1} + e_y$$

where e_y is the random error term and b and c are positive.⁷ The second equation is typically some variation of the expectations-augmented Phillips Curve in which inflation depends on both the output gap (a measure of the balance between aggregate demand and potential supply) and expected inflation, such as

$$(4) p = d y + [f p_{t-1} + (1-f) E_t(p_{t+1})] + e_p$$

where e_p is the random error term, $E_t(p_{t+1})$ is expected inflation next period, based on information available today, and d and f are positive. The bracketed term for expected inflation can be interpreted as a combination of sophisticated forecasts of inflation ($E_t(p_{t+1})$) and rule-of-thumb forecasts based on past inflation (p_{t-1}). This particular specification collapses to a backward-looking version of the Phillips Curve when $f=1$ and a rational expectations version when $f=0$.

In principle, the model can be combined with the loss function to derive an optimal rule. The optimal rule is the one that minimizes the loss function, given the model specification and parameters. Derivation of the optimal rule is useful because it identifies the variables to which policymakers should respond, given the model, and the factors that should guide how aggressive the response should be.

Because we do not have direct information on the value of w , the research often focuses on establishing *efficient* rules. Efficient rules are those that minimize a weighted

⁷Some prefer the forward-looking or new-Keynesian specification of the dynamic IS curve: $y_t = E_t y_{t-1} - ? r_t + e_y$.

sum of output variance and inflation variance for some choice of weights. Efficiency frontiers can be derived that show the minimum levels of the output variance that can be achieved for a given level of inflation variance (and vice versa). The optimal rule can then be determined by the point on that frontier that minimizes the policymakers' loss function, which will depend on their preference parameter w .

Because fully optimal rules can be quite complicated, involving current and lagged values of all the variables in the model, much of the research has focused on the performance of simple rules that involve only a small number of variables.

A simple approach to assessing the performance of simple rules is stochastic simulation. First, the model of the economy is completed by the addition of a simple rule, such as the Taylor rule. Some of the research uses simple models of the form described above. But a fair amount of research has been done at the Board using large-scale models, including the Board staff's FRB-US model. Second, simulations are run for a large number of random draws of the values of the two error terms, e_y and e_p , subject to information about the distribution of these errors over some historical period. Finally, the performance of alternative simple rules is evaluated using the loss function.⁸

Most of my discussion of policy rules will focus on research on simple rules. The focus on simple rules is justified by the finding that simple, well-designed rules are robust across a range of empirical macro models and come close to the performance of optimal rules for many models.⁹ Optimal rules, on the other hand, can be very sensitive

⁸ In much of the recent research on rules, the variances of output and inflation are calculated from the reduced form of the model. This approach is much less computationally intensive than running stochastic simulations. However, if the underlying model is non-linear, then stochastic simulations are run to derive the results.

⁹ See for example, Levin, Wieland and Williams (2000).

to model specification. These results suggest that the major gains to policymakers come from monitoring and interpreting the results for simple rules and that the incremental gain from more complicated optimal rules is relatively small.

VI. Operational specification and parameterization of the simple Taylor Rule

I will first focus on issues related to the empirical implementation of the simple Taylor rule and research on the optimal values of the parameters in that specification. In the following section I will discuss the research on alternative specifications that retain the spirit of the Taylor rule but try to either reconcile the practice of monetary policy with the evidence on optimal policy responses with a simple Taylor rule or otherwise improve upon the stabilizing properties of that rule.

To apply the Taylor rule empirically, we have to determine the specific measures for inflation and for the resource utilization rate, identify the target levels for the inflation and utilization rates respectively, and define and estimate the value of the equilibrium real interest rate.

The Measure of Inflation

There are several measures of inflation that have been used in simple policy rules. Taylor himself uses the aggregate price level underlying the definition of real GDP to compute his measure of inflation. It is a comprehensive measure of the average price level of goods produced in the United States and, as such, has a broader coverage of goods than the basket of goods and services the typical household consumes. Most inflation-targeting countries use a measure of consumer price inflation, comparable to our CPI, or a measure of core consumer price inflation that nets out the direct influence of volatile components whose movements are not closely related to the overall balance of

aggregate demand and supply. It is not clear theoretically whether a producer or consumer price measure is preferable. The preference for the CPI in inflation-targeting countries may be due to its visibility, facilitating the communication of policy objectives to the public. In any case, work at the Board suggests that there is not much difference in economic performance associated with the choice among producer or consumer price measures.¹⁰

There are, however, a couple of situations for which the choice of inflation rate in the policy rule does matter. A change in the exchange rate, for example, would immediately affect a consumption price measure, but not the production price measure of inflation. Whether or not the inflation rate directly captures the effect of a change in the price of imported goods may, in turn, affect whether or not the policy rule prescribes an immediate response of monetary policy to an exchange rate shock. The role of nominal exchange rate fluctuations in affecting monetary policy decisions is likely to be more important in smaller open economies, like Canada and New Zealand, than in the United States.

An oil price shock is a second example of a situation where the choice of the inflation measure in the policy rule will affect the policy response. One reason is that oil has a larger weight in consumption than in production in the United States. As a result, the aggressiveness of the immediate policy response prescribed by a policy rule would be greater for a consumption- than production-based measure of inflation. A second reason is that selecting a core measure of consumer prices would eliminate the direct effect of

¹⁰ In the United States, there is also a question as to whether to use the CPI or the PCE measure of inflation for either core or overall consumer prices. These two measures often provide somewhat different signals, related to their individual source data and differences in weighting schemes. For example, over the last 12 months, inflation measured by the core CPI is 2.8 percent, compared to just 1.6 percent for the core PCE.

swings in oil prices altogether and hence eliminate any direct response of monetary policy to the temporary inflationary impetus of a one-time rise in the price of oil. The choice of a core measure of consumer prices for the inflation rate in the policy rule is consistent with a third guideline for monetary policy that I discussed in my earlier paper on monetary policy strategy: Monetary policy should look through the direct effect and respond only to the secondary effects of supply shocks.

By expressing the Taylor rule in terms of core measures of consumer price inflation, policymakers are encouraged to “look through” temporary increases in energy or food prices, responding only to the extent this inflation shock spills over into broader price developments, as captured in the core measure of inflation. Note, however, that this consideration will be less important in forward-looking versions of the rule I discuss below, because forecasts also “look through” the temporary effects of supply shocks on inflation.

A final operational detail is whether the inflation rate should be measured in the quarter corresponding to the nominal funds rate or as an average over some longer period. The inflation measure will, in effect, determine the measure of the real interest rate, so the question in part is whether to use an ex post measure of the real interest rate or whether to use a measure that might better capture the expected inflation rate. An average over the past year that washes out some of the high frequency noise in quarterly inflation rates has been identified by a number of investigators as a pragmatic choice. In support of this proposition, Levin, Wieland and Williams (2000) find that using an average inflation rate over the previous four quarters in the policy rule, rather than the inflation rate in the current or previous quarter, can improve performance. Using an

average inflation rate reduces the risk of responding to noise in higher-frequency inflation data and, as a result, unnecessarily destabilizing output in an effort to control inflation.

The inflation target

Many central banks today operate with an explicit, numerical inflation objective. Identifying the inflation target--for both the policymakers and the public--is of course relatively easy in inflation-targeting countries. However, even in these cases, the target may be expressed as a range instead of a point. The simple solution in the case of the range is to use the mid-point when performing Taylor rule calculations.¹¹

The Federal Reserve does not have an explicit numerical inflation objective, which to my mind is unfortunate. But that was a topic for an earlier paper.¹² To implement a Taylor rule for the United States empirically, it is therefore necessary to make some assumption about the value of the FOMC's implicit inflation objective or to estimate the FOMC's inflation objective as part of an estimation of the parameters of the Taylor rule. The public, for example, could infer an implicit inflation target from past actions or recent statements of the FOMC. That leaves open the interesting question of how policymakers would implement a Taylor rule when they have not established an explicit inflation target. I have asked that question a number of times myself! In any case, Taylor assumed a 2 percent inflation target, for the GDP measure. That would be consistent, based on recent experience, with about a 2 percent inflation rate for the PCE measure and a 2-1/2 percent rate for the CPI. I have indicated a preference for a 2

¹¹ If the target is a range, the policy response might depend on whether there is some preference for being at the midpoint and, if so, what the implications are of the boundaries of the range. Orphanides and Wieland (2000) note that ranges may imply a nonlinear policy response, with a stronger response to inflation when it is closer to or outside the zone and milder when inflation is close to the mid-point.

¹² See Meyer (2001).

percent objective for inflation based on the CPI, corresponding to 1-1/2 percent for the PCE and GDP--based measures of inflation.

Measuring the utilization rate

One of the Fed's objectives is, as noted earlier, maximum sustainable employment. This is usually interpreted as an objective of smoothing output relative to potential output or the unemployment rate relative to its "full employment" level, the NAIRU. The consensus model that underlies most research on policy rules also assumes that the balance between potential supply and demand in the labor and/or output markets is a proximate source of movements in the inflation rate. So, even if there were not an output or employment stabilization objective, monetary policy makers would respond to the output or unemployment gaps in order to meet its inflation objective.

Unfortunately, we do not have direct measures of potential output or the NAIRU. In operational specifications of the model, we therefore estimate their values using some macroeconomic procedure, potentially subject to considerable error, and then use the estimated values to form measures of the utilization rates--the output gap and unemployment gap. The resulting measures of the output gap and unemployment gap are examples of proxies for the unobserved utilization rates.

Some models of inflation dynamics give primary weight to excess demand in the labor market and hence support the choice of the unemployment gap. Others focus directly on excess demand for goods and hence support a choice of an output gap. Because there is a well-known regularity linking with output and unemployment gaps,

summarized in Okun's Law, there is, in practice, relatively little difference between the two measures.¹³

A further complication, and one I believe has been very relevant during the period I have been on the FOMC, is that there can be transitory as well as more persistent changes in the level of potential output or the NAIRU. Research at the Board suggests that a persistent increase in the rate of structural productivity growth--as occurred in the second half of the 1990s--may have resulted in a temporary decline in the NAIRU. This arises because of the different speeds at which workers and firms perceive the changed prospects for the growth of real wages. Although the decline in the NAIRU is temporary, it can be quite persistent. As a result, it is useful to define a short-run or effective NAIRU that takes this temporary effect into account. Thus a permanent increase in productivity growth leads to a temporary but persistent decline in the SR-NAIRU.¹⁴ The short-run NAIRU is the variable that should enter the Taylor rule, because it pins down the inflation dynamics over the period relevant to policy and also identifies the unemployment rate consistent with price stability in the near term. It should be noted that this model of inflation dynamics--based on a possible divergence of the short-run from the long-run NAIRU--is not universally accepted, even among those who find the Philips

¹³ Still, the output measure in principle takes into account utilization rate of capital as well as of labor and, because the unemployment rate adjusts with a lag to changes in output, there is also some timing difference between the two measures of resource utilization. There are, in addition, some proxies for the balance between aggregate demand and potential supply of goods that do not directly build on the unemployment gap. Taylor, for example, used a measure of the output gap that did not rely on an estimate of the NAIRU. His measure of Y^* was just a measure of the longer-term trend in Y . Another independent measure of a utilization rate is the capacity utilization rate, though this measure applies only to goods production, not overall GDP.

¹⁴ See Braun (1984) for the original development of this idea and Meyer (2000) for a simple model incorporating this effect.

Curve a useful analytical tool. But, if this is the correct model--and I believe it is--then the utilization rate in the Taylor rule should be the gap between the unemployment rate and the short-run NAIRU or a measure of the output gap consistent with latter measure of the unemployment gap.

An important strand of recent research on policy rules considers the implications of uncertainty about the measurement of the utilization rate. I will return to this topic.

Defining and measuring the equilibrium real interest rate

The last of the variables in the Taylor rule is the equilibrium real interest rate. To determine an operational counterpart, we need to first define what is meant by an *equilibrium* real rate and then design a methodology for estimating it, since it is not directly observable.

One possible definition of the equilibrium real rate is the value consistent with balance between potential aggregate supply and demand over a period that allows time for the economy to move back to equilibrium after some disturbance. The timeframe should be at least as long as the average business cycle--at least five years. If the equilibrium rate is a constant, the real equilibrium rate can be measured by an average over some appropriately long period. Taylor used 2.0 as his estimate of the equilibrium real rate in his rule, close to the average from 1960 to the mid-1990s based on the GDP measure of inflation.

However, there are good theoretical reasons for expecting the equilibrium real interest rate to change over time. Incorporating a time-varying measure of the equilibrium real interest rate in the Taylor rule is one of the refinements I will discuss below.

Parameterization of the simple Taylor rule

There are two parameters in the Taylor rule: the aggressiveness of the response of the federal funds rate to changes in the output and inflation gap respectively. Taylor set the values of both these parameters to 0.5. One natural focus of the research on policy rules is on whether choosing different magnitudes for these two parameters could improve the performance of the economy.

In general, the studies of efficient policy rules find that policymakers should respond to inflation about as aggressively as Taylor assumed, but should respond to output gaps more aggressively. Studies of efficient rules based on Taylor's specification suggest that raising the coefficient on the output gap to about one reduces the variability of both output and inflation.¹⁵ The assessment of the size of this parameter remains an issue of considerable controversy and has been a focus of the recent research on rules that I discuss below.

An empirical specification of the simple Taylor rule

Let me sum up my preferences for the operational specification of the simple Taylor rule. I prefer to use the unemployment gap because it highlights a controversial issue, the estimate for the NAIRU. I would measure the unemployment gap as the difference between the unemployment rate and the short-run NAIRU. While the short-run NAIRU, in my view, fell well below the long-run NAIRU in the late 1990s, the recent slowdown in structural productivity growth has, in my view, pushed the short-run NAIRU back to close to its long-run value that I estimate to be about 5-1/4 percent. For

¹⁵ See, for example, Ball (1999). Henderson and McKibbin (1993) also favor a higher coefficient on the output gap. They set that parameter to 2.0 and, in addition, set the coefficient on inflation to 2.0.

the core CPI as the inflation measure, my target would be 2 percent. For a constant equilibrium real rate, I would use about 2.5 percent.¹⁶

For the response parameters I would use 0.5 on the inflation gap and 2.0 on the unemployment gap (the equivalent to 1.0 on the output gap). But I will discuss below the case for attenuating the response to the output or unemployment gap during periods of elevated uncertainty about their measurement.

VII. Issues Concerning Policy Rules

Some major questions about the specification of rules have prompted a good deal of research over recent years. For example, why don't policymakers behave in a manner consistent with what research tells us are optimal rules? I explore this question below in the section on interest-rate smoothing. Another set of questions is motivated by the challenges faced by policymakers in the second half of the 1990s and during period of economic weakness in 2001 and the related research attempts to reconcile policy rules with monetary policy in practice.

Interest Rate Smoothing

Rudebusch (2001), Sack (2000) and Ball (1999) find that historical changes in the federal funds rate are smaller, more inertial, and exhibit fewer reversals than would be expected under the optimal policy rules derived from many models.

¹⁶ When measuring average CPI inflation rates over such a long period, it is important to use the "research series using current methods"--a series that employs the current methodology consistently over the full historical period. The published data for the CPI, in contrast, are not adjusted backward when there are methodological changes in the way the index is computed. The average value of the real federal funds rate over 1961 – 2001 is 2.7% for all three of the inflation measures I have discussed--the chain price measure for GDP, the core PCE, and the current-methods core CPI. Over recent years, however, inflation measured by the core PCE has been about ½ percentage point lower than the inflation rate for the core CPI. Over the last year, the difference is wider than a full percentage point. Using an equilibrium real rate of 2.4% for a core CPI-based rule and 3.0% for a core PCE-based rule would lead to the same policy prescription on average over the past few years. In addition, the average of the two real rates would be 2.7%, equal to the average real rate for both measures over the longer period.

If the question is a matter of fitting the data, the simple Taylor rule can be revised so that it is consistent with the observed gradual adjustment of interest rates to changes in output and inflation. One way to do so would be the following partial-adjustment equation for the nominal interest rate:

$$(5) \ i_t - i_{t-1} = g (i_t^N - i_{t-1})$$

In a given period, policymakers adjust the federal funds rate to move it a fraction of the way toward its “notional value” i_t^N , which is in turn determined, for given values of y and p , by equation (1). Thus the partial adjustment specification yields a gradual adjustment of the funds rate to changes in y and p . Empirical researchers often employ the partial adjustment framework and typically find the lagged interest rate term to be overwhelmingly significant. In such estimated versions, the values of r^* and p^* are typically jointly estimated as part of the constant term of the equation, rather than imposed, as in the case of the simple Taylor rule.

Another specification that allows for gradual adjustment of the nominal interest rate is a generalized version of the Taylor rule that adds a term involving the lagged nominal funds rate to equation (1).

$$(6) \ i_t = r^* + p_t + a y_t + b (p_t - p^*) + h i_{t-1} + (1 - h) (r^* + p_t)$$

This specification has the advantage that the Taylor rule is a special case where $a = b = \frac{1}{2}$ and $h = 0$ and the first-difference specification studied by Levin, Wieland and Williams (2000) is a special case where $h = 1$.

The harder question is why do monetary policymakers appear to engage in interest rate smoothing. The answer may be that optimal rules leave out some important aspect of the policymaking process. The literature suggests three possibilities:

policymakers attach a cost to interest rate variability; the more gradual response yields better policy outcomes when policymakers are uncertain about the model parameters or the model itself; and the shocks to which policymakers respond are serially correlated.

In the first case, interest rate variability would be an additional term in the loss function and policymakers would aim to smooth interest rate movements as they sought to achieve the other objectives of monetary policy, perhaps because they were concerned that excessively volatile interest rates could threaten financial stability. In my experience, a simple aversion to large movements in interest rates has not played a prominent role in policy discussions or decisions.

In the second case, the appropriate response of policymakers to uncertainty will depend on its source. In terms of the simple consensus model introduced earlier, one source is referred to as *additive* uncertainty, captured in the two stochastic error terms in the IS and Phillips Curve equations, e_y and e_p . There could also be *parameter* uncertainty, uncertainty about the parameters of the model, the b , c , d , and f terms. There could be, in addition, *model* uncertainty, uncertainty about the structure of the model itself. Finally, there could be *measurement* uncertainty, uncertainty about the measurement of the variables to which monetary policy responds, y and p in the model.

If there is additive uncertainty alone--and the model is linear and preferences are quadratic, as assumed above--policymakers should assume the errors take on their mean or expected values (zero in this case) and then make policy as if they were perfectly certain about the relationships in the model. This is called the certainty equivalence principle. Sack (2000) finds that the optimal rule under certainty equivalence yields more aggressive interest rate movements than observed.

Brainard (1967) has demonstrated that, under some conditions, parameter uncertainty encourages policymakers to respond more cautiously. However, some empirical research has found that parameter uncertainty has little effect on the optimal rule.¹⁷ This research generally involves small models in which simple rules tend to be optimal. Other work using models with a larger number of parameters has found that parameter uncertainty does reduce the aggressiveness of policy responses to a greater extent.¹⁸ Sack (2000) finds that parameter uncertainty results in a policy response much closer to what has historically been followed, though it is still somewhat more aggressive than the optimal rule.¹⁹

Rules with partial adjustment may also be more robust to uncertainty about the structure of the model, at least for models with rational expectations, as shown by LWW (1999). They find that a small degree of partial adjustment can dramatically lower the volatility of the short-term interest rate with negligible effects on output and inflation variability in such models. The reason may be that partial adjustment makes future movements in interest rates more predictable and forward-looking behavior brings those effects forward. In effect, more predictable movements in short-term rates allow expectations of future short-term rates to be embedded today in long-term interest rates and asset prices. However, Rudebusch and Svensson (1999) found that the optimal

¹⁷ See Rudebusch (2001) and Estrella and Mishkin (1999).

¹⁸ See Sack (2000), Martin and Salomon (1999) and Soderstrom (1999a).

¹⁹ However, the finding that uncertainty calls for a less aggressive policy response is not a theoretical necessity and some other researchers have presented cases where parameter uncertainty may lead to a more aggressive policy response. See, for example, Tetlow and von zur Muehlen (2000) and Soderstrom (1999b). At the risk of oversimplification, this research suggests that policy outcomes may be improved by a more gradual policy response when the uncertainty is isolated in scope and well understood, but that outcomes might be improved by a more aggressive policy when uncertainty is ubiquitous.

partial-adjustment coefficient is near zero in an adaptive-expectations model, raising some questions about the robustness of the LWW results across a range of models.

Lastly, Rudebusch (2000) and others have argued that policymakers might appear to adjust interest rates gradually, even when they do not do so, because the estimated policy rule likely omits important variables that are serially correlated. This corresponds with my experience that the progressive movements in the federal funds rate during tightening or easing cycles have generally reflected progressive changes in the incoming data or the forecast. However, English, Nelson and Sack (2001) allow for both serially correlated errors and partial adjustment in an estimated policy rule and find that partial adjustment remains an important characteristic of the estimated policy rule even when serially correlated errors are included.

Overall, rules with partial adjustment yield policy rules closer to estimated rules based on actual policy decisions and, for the reasons suggested above, seem to be a constructive refinement of the simple Taylor rule.

Noisy Information

In light of my experience on the FOMC, the aspect of the Taylor rule that has been most controversial and most challenging is how to implement the prescription that policymakers should adjust the federal funds rate in response to movements in the output gap or the unemployment gap. Even among those who accept the underlying model that gives rise to this conclusion--some version of the Phillips curve as a model of inflation dynamics--the degree of policy response to changes in the output gap is a major source of controversy, given the considerable uncertainty about the estimate of the NAIRU and the related uncertainty about the size of the output gap.

I focus here on two key implications of measurement uncertainty or noisy information: the importance of using real-time data to estimate policy rules and the possible attenuation of the response to utilization rates.

Real time data and historical analysis of policymaking: Orphanides (2001a, b) constructed a real-time database to allow estimation of policy rules over historical periods. He demonstrates that rules estimated from *ex post* revised data can yield very different interpretations than rules estimates from the real time data available when the policy actions were taken. In particular, his results contradict the finding by Taylor (1999) and Clarida, Gali and Gertler (2000) that practice of monetary policy improved in the 1980s and 1990s, compared to the 1960s and 1970s. In the earlier period, they find that coefficient on inflation in the Taylor rule is less than one, implying that monetary policy actions destabilized rather than stabilized inflation. In the later period, they find that the coefficient on inflation is greater than one. Orphanides, in contrast, finds that, using real-time data, the coefficient is above one in both periods. Thus the assessment of policy in the earlier period appears distorted by the use of data that was not available to policymakers at the time. The key variable in question is the measure of the output gap.

Orphanides (2001b, 2002) concludes that the poor performance of monetary policy in the 1970s was not due to the failure of monetary policy to respond aggressively enough to inflation, but rather to an overly aggressive response to output gaps or unemployment gaps measured with considerable imprecision. His research supports concerns voiced earlier by Friedman (1968) and Meltzer (1987) that overly active monetary policy could turn out to be destabilizing.

I have some reservations about the characterization of a policy rule with a positive weight in the output gap as an “activist” rule. Of course, such a characterization seems natural. But let's think of monetary policy as being implemented by setting the growth rate of reserves. After all, this is the ultimate instrument of monetary policy. In this case a passive policy would be a constant rate of growth of reserves. The constant rate would of course be set to be consistent with the long-run price stability objective, but it would be independent of movements in output or inflation. Consider the implications of a positive aggregate demand shock under such a policy. Holding reserve growth constant, an autonomous increase in aggregate demand would increase nominal and real interest rates. To mimic this non-activist policy under an interest rate rule, there would therefore have to be a positive coefficient on output. A coefficient of zero would imply that monetary policy responds to a positive demand shock by increasing the rate of reserve growth to prevent a rise in interest rates, precisely the policy response that Meltzer and Friedman warned against. That is, a zero coefficient would imply that monetary policy should respond perversely to demand shocks with open market operations that reinforced the shock. This is not to deny that monetary policy could be overly aggressive and therefore destabilize output. But that becomes a more subtle question of how large the coefficient on output should be.

In addition, I am cautious about reaching conclusions relevant to policy today based on the experience in the 1970s. At that time, conventional empirical macro models did not yet embody the expectations-augmented Phillips Curve. As a result, there was only a rudimentary appreciation of how to define potential output or the NAIRU as we understand it today and no serious effort to update these estimates on the basis of

incoming data. Today, in my view, we both have a clearer understanding of the concepts of potential output and the NAIRU, and hence of inflation dynamics, and we have evolved better procedures for making real-time adjustments in our estimates of potential output and the NAIRU in response to incoming data.

Attenuation: Indeed, the research focused on measurement uncertainty suggests that policymakers should update their estimates of utilization rates, using all available information. This updating is often referred to as “filtering” the data--that is, using the available data to improve our estimate of some unobserved variable. If this “filtering” of the data is very good, then measurement uncertainty will have little effect on the optimal response output or unemployment gaps. In this case, we are back to close to the certainty equivalent result. If the filtering is very poor on the other hand, as it clearly was in the 1970s, then policy is likely to be improved, as Orphanides has suggested, by attenuating the response to output or unemployment gaps, that is, by downweighting the response to changes in output or unemployment gaps.

This result that policymakers should attenuate their response to utilization rates in the presence of elevated uncertainty about their measurement is found in two strands of the literature. Empirical work by Orphanides (1998) and Smets (1999), using simple policy rules, found that performance of these rules was improved by attenuating the response to the output gap in proportion to the degree of noise in that data. They also found that the response to inflation should be attenuated as well. Orphanides (1998) concludes that when uncertainty is especially high, a Taylor rule that omits the gap completely performs well.

Orphanides, et al (1999) compare the optimal response to the output gap under varying degrees of confidence in its measurement. They find that, in the absence of measurement uncertainty, policy should respond to output more aggressively than assumed in the simple Taylor parameterization. If, on the other hand, revisions in the output gap are assumed to be about the same as on average over the last 30 years, they find that the policy response should be significantly attenuated, but still larger than the 0.5 value of the parameter in the simple Taylor rule. Finally, if revisions are assumed to be at the upper end of historical experience, they find that policy would be improved by eliminating any response of the funds rate to movements in measured output gap.

A second strand of the literature on noisy information uses signal extraction models.²⁰ In Swanson (2000), inflation dynamics are assumed to depend on “excess demand,” an unobserved variable. Empirical models must therefore use imperfect proxies for excess demand, such as estimated output and unemployment gaps and inflation. Signal extraction models formally describe how policymakers estimate such an unobserved variable by extracting the signal about the theoretical but unobserved variable from noisy indicator variables. Swanson found that policy can be improved by attenuating the response to the proxy in proportion to the degree of lack of confidence in the accuracy with which it measures excess demand. He also found that policy is improved by raising the aggressiveness of the response to inflation to compensate for the increased uncertainty about utilization indicators.

²⁰ See Swanson (2000) and Svensson and Woodford (2000).

Swanson (2000) and Meyer, Swanson, and Wieland (2001) find that a nonlinear policy rule may outperform a linear rule in the presence of uncertainty about the NAIRU or potential output that is not normally distributed. When the level of the unemployment rate is close to the best estimate of the NAIRU, policymakers should significantly attenuate their response to the unemployment gap, reflecting the large amount of uncertainty they face about the magnitude and even sign of the gap. However, when the unemployment rate moves further away from their best estimate, policymaker's confidence increases that there is excess demand or supply and, as a result, policymakers should incrementally increase the aggressiveness of their response.

My conclusion from this research is that a rule that totally eliminated any response to the output gap would seriously underperform rules that included responses of the order that I have described above. On the other hand, there may still be a case for attenuating (but except in extreme cases not eliminating) the response to changes in output or unemployment gaps in periods of elevated uncertainty about their measurement.

The research on noisy information is very relevant, in my experience, to the challenges faced by the FOMC in the second half of the 1990s and beyond. This uncertainty, in my view, became an important consideration for monetary policy in 1999, once it had become clear that the Asian financial crisis and global financial instability were not slowing growth in the United States. The continued low level of the unemployment rate might have suggested a much quicker reversal of the 75 basis point easing that had been implemented in the fall of 1998 when financial market conditions deteriorated and the forecast of real economic activity suggested a risk of a significant slowdown. One could interpret the FOMC's reluctance to reverse the easing quickly as

reflecting an attenuation of the response to output gaps, measured on the basis of the prevailing estimate of the NAIRU that had become increasingly suspect.

Outcome-based versus forecast-based rules

The simple Taylor rule is often referred to as an outcome-based rule. That is, policy responds only to the incoming data or to realized outcomes of the economy. Outcome-based rules use a very small subset of the available information. Based on my experience on the FOMC, the simple outcome-based Taylor rule misses many of the factors that policymakers consider in their deliberations leading to policy actions. Much of the discussion at FOMC meetings is focused, after all, on the forecast of output and inflation. An alternative approach to policy rules more consistent with FOMC deliberations would therefore have policymakers respond directly to forecasts. Such forecast-based rules are information encompassing to the extent that the forecasts use all available information.

In forecast-based rules, the utilization and inflation rates are replaced by the forecast of the future values of these variables. It is easy to understand why there might be gains from responding to forecasts, in addition to their information-encompassing property. Policy affects output and inflation with a lag. Therefore, if policy actions responded directly to forecasts, this might speed the response of policy to shocks and thereby reduce output and inflation variability. Furthermore, a good forecast will cause policymakers to “look through” transitory fluctuations in inflation and resource utilization, and only respond to more persistent fluctuations.

Interestingly, much of the research with forecast-based policy rules concludes there is at best only a marginal gain from using forecasts. For example, Levin, Wieland and Williams (2001) find that optimized forecast-based policy rules perform only

marginally better than optimized output-based rules. They find that the optimal forecast horizon is relatively short, one year at most, and that rules using longer horizons are less robust with respect to model uncertainty.²¹

A possible explanation follows from the response of longer-term interest rates to expectations of future policy. To the extent that longer-term interest rates do accurately reflect future policy actions, the bond market will speed the response of the economy to actual interest rate moves implemented under an outcome-based rule, essentially replicating the results that otherwise would have been achieved if the policymakers themselves had implemented policy based on the forecast.

My experience suggests that during periods when both the data and forecasts are changing slowly, policy can generally be described adequately as responding to the incoming data, as in the outcome-based rule, although it could just as easily be described as responding to forecasts. As suggested by the empirical results, there is little difference between the two approaches much of the time and hence little incremental benefit to responding to forecasts as opposed to outcomes. However, in my experience, sharp movements in policy--as during the fall of 1998 and during early 2001--often reflect a response by policymakers to a discrete and sharp change in the forecast, one that is not apparent from simply looking at current aggregate measures for output and inflation. In such cases, policymakers appear to switch from simple outcome-based rules to more complicated forecast-based rules. Having said that, outcome-based policy rules with a coefficient of one on the output gap do a pretty good job of tracking the pace of easing last year.

²¹ Orphanides (2001a), on the other hand, finds that a forecast-based policy rule based on real-time data seems to describe actual policy better than comparable Taylor rule specifications.

Time varying equilibrium real interest rate

While the simple Taylor rule assumes a constant real equilibrium funds rate, research at the Board suggests that there has been significant time variation in the equilibrium interest rate over the past forty years. Laubach and Williams (2001) find that the equilibrium real rate was at a peak in the 1960s and again in the late 1990s.

Theoretical models and empirical evidence suggest that movements in the rate of structural productivity growth and in the structural budget deficit are important sources of the variation in the equilibrium real rate over time. The 1960s and the second half of the 1990s were the periods of peak underlying productivity growth over the last forty years.

Laubach and Williams find that incorporating a time-varying estimate of the equilibrium real federal funds rate into a policy rule significantly reduces the variability of output and to a lesser extent inflation.

This is another issue that became important in monetary policy deliberations in the second half of the 1990s. Once it became clear that the acceleration in productivity was playing a very important role, the question arose as to its implications for the equilibrium real interest rate. A number of committee members discussed this in terms of the relation between the “natural” rate and the policy rate. If the productivity acceleration raised the natural rate, holding the policy rate constant would result in an increase in monetary stimulus, measured by the gap between the natural and policy rates. A policy rule with a time-varying equilibrium real rate would encourage policy makers to adjust the policy rate to keep in line with the equilibrium real rate, except as justified by movements in the output and inflation gaps.

However, the technology for measuring a time-varying real interest rate was evolving at the same time and there was therefore considerable uncertainty about how large an increase in real rates was called for. Continued work in estimating a time-varying equilibrium rate would be an important contribution to policy rules and would in turn provide useful information to policymakers, at least during periods of significant changes to the underlying pace of productivity growth or when there are other developments that might affect the equilibrium real rate.

Implications of the zero nominal bound

As the nominal funds rate fell to very low levels at the end of 2001, it raised an issue that had been discussed in principle at a conference sponsored by the Federal Reserve in 2000 and that has confronted the Bank of Japan in recent years--the possibility that the effectiveness of monetary policy could be diminished in a low inflation environment.²² The typical linear specification of the Taylor rule assumes that the FOMC can move the nominal funds rate as appropriate to achieve its objectives. But the nominal funds rate cannot decline below zero. The real federal funds rate, equal to the nominal funds rate less inflation, can therefore only be as negative as the rate of inflation. As a result, if the inflation rate is zero, the zero nominal bound prevents the real rate from becoming negative and this constraint, at times, could limit the ability of policymakers to adjust policy aggressively enough to stabilize output. The situation would be even more serious if the economy were to slip into deflation; then the central bank would be prevented from pushing the real short-term rate below some positive number, set by the expected rate of deflation.

²² See Monetary Policy in a Low Inflation Environment (2000).

Reifschneider and Williams (2000) use simulations with the Board's FRB-US model to provide evidence on the degree of potential deterioration in the economy's cyclical performance inflicted by the zero nominal bound. They run stochastic simulations with alternative inflation targets in the model's policy rule. These simulations subject the economy to a range of disturbances that are typical to shocks that have been observed over the past forty years. For an inflation target of 4 percent, the zero bound is reached less than one percent of the time and the average duration of a spell of zero interest rates is about two quarters. As the inflation target falls toward zero, both the frequency and duration of zero interest rate episodes increase. The relationship is highly nonlinear, so there is not much affect as the target inflation rate falls to 2 percent; but when the target falls below 2 percent, such episodes become progressively more common and more prolonged. At a zero inflation target, the funds rate is at the zero bound 14 percent of the time and the average duration of these spells is one and a half years. And there is a clear tendency for the cyclical performance of the economy to deteriorate as the inflation target falls below two percent: the frequency of mild recessions declines and the likelihood of severe contractions correspondingly increases.

The first lesson to policymakers from this research is that the inflation objective should be positive, to provide some cushion for real interest rates to be negative if necessary to stabilize the economy. This is perhaps one reason why the "price stability" objective should be implemented as a "low inflation" objective, what I refer to as price stability plus a cushion. In practice, perhaps in part for this reason, central banks with

explicit inflation targets typically have inflation objectives with a mid-point of about 2 percent.²³

As the nominal funds rate declined to and then below 2 percent late in 2001, some argued that the FOMC should “hold its ammunition,” meaning that it should not implement cuts that it otherwise would have made in order to preserve the opportunity to respond to any further unexpected adverse shocks. The literature on policy rules, however, provided a different message. Reifschneider and Williams (2000) find that rules with more aggressive response parameters work better when the nominal bound is a potential constraint, in part because they quickly move the nominal funds rate toward zero when the economy weakens and inflation is low. Indeed, it is because aggressive response parameters increase the likelihood that nominal funds rate will be driven to zero that they do a better job of damping fluctuations in output and inflation and heading off potential deflations. An implication of this research is that the potential deterioration in cyclical performance as a result of the zero nominal bound can be compensated for by an asymmetric policy response, specifically by a more aggressive response to both declines in output and inflation when nominal rates are already very low.

A second approach suggested by Reifschneider and Williams involves a more fundamental respecification of the Taylor rule. The principle underlying the respecification is that monetary policymakers have to find a way to lower real long-term rates, even when the nominal interest rate is at the zero bound. They may be able to do

²³ A second reason for a positive inflation target is an upward bias in measured inflation rates. The Advisory Commission to Study the CPI (1996)--generally referred to as the Boskin Commission after its chairman--estimated that the bias in the CPI was about 1.1 percentage points. A recent update of the estimate of this bias by staff at the Board--Lebow and Rudd (2001)--following a series of methodological revisions to the index by the Bureau of Labor Statistics, put the central tendency for the bias at 0.6 percentage points.

so to the extent they can make a commitment today about their future policy actions, thereby altering expectation in the market today about the future path of short-term interest rates. Specifically, if they can commit to maintaining a lower path of the nominal funds rate in the future, even when the zero nominal bound is no longer a constraint, they may be able to lower real long-term interest rates today. Reifschneider and Williams offer a specific respecification of the policy rule that might achieve this outcome. The respecified rule hold down the funds rate in the future in proportion to how much it was prevented by the zero nominal bound from lowering the nominal funds rate during the period when the zero nominal bound was a constraint.

A second respecification suggested by Reifschneider and Williams uses a longer period for defining the average inflation rate in the rule. If the average inflation rate is constructed over a three-year period, policymakers will attempt to offset subperiods when inflation is below the target with inflation above the target over the subsequent subperiod. This means that if inflation falls below the target, the policymakers commit to a more stimulative policy for a while in the future and hence to a longer period of low nominal funds rate. This again could reduce long-term real interest rates today.²⁴

While both respecifications improve the performance of the economy during periods subject to the zero nominal bound, they raise a question about the credibility of the commitment implied by the rule. In particular, the effectiveness of such a commitment hinges directly on the ability of the central bank's promise of future actions (perhaps several years into the future) to influence the public's expectations today. In

²⁴ Wolman (1998) demonstrates that a price level target has the same property. That is, a price level target works well in a period of deflation because it involves a promise to reflate in the future to get the price level back to its desired level.

such a case, transparency may offer an important benefit. In particular, if workers, firms, and investors can be convinced through public statements that an unusual situation calls for unusual action, the central bank's ability to affect expectations about its future policy--when the promised future policy is different from its normal conduct--may be enhanced.

The role of other variables in the policy rule

Should policymakers respond only to actual or prospective changes in output and inflation, or would policy be improved if policy also responded directly to other variables? Two candidates often singled out for additional attention are the exchange rate and equity prices.

Monetary policy and exchange rates:

Ball (1999) and Batini and Haldane (1999) find that in small open economies policy outcomes may be improved by adjusting the short-term policy interest rate in response to changes in the nominal exchange rate, thereby reducing the impact of these movements on output and inflation. Canada for example used a monetary conditions index--a weighted average of the policy interest rate and the exchange rate--as its monetary policy instrument for some time. This implied that any change in the exchange rate would be automatically offset by appropriate movements in the policy rate, for unchanged values of output and inflation. However, the relevance of this consideration to the United States appears minimal. Moreover, the implications of a change in the exchange rate for movements in output and inflation depend on the reasons for the change. The difficulty of making this clear to markets prompted Canada to stop using monetary conditions index as its instrument.

Monetary policy and equity prices:

One of the most controversial issues in the conduct of monetary policy during the second half of the 1990s was whether or not monetary policy should directly respond to movements in equity prices. The conventional wisdom, especially among monetary policymakers, is that policy should respond to actual or expected effects of changes in equity prices on inflation and output, but not to the movements in equity prices themselves. This is consistent with simple policy rules that exclude equity prices, but include either realized values of inflation and output or forecast values of these variables.

The question is whether monetary policy should also respond directly to asset price movements--and specifically to discrepancies between asset prices and their estimated fundamental value--over and above their effects on output and inflation. This question became relevant during the second half of the 1990s when some feared an equity bubble was developing that was distorting economic decisions in the near term and the ultimate correction of which could have significant adverse effects on real economic activity. In fact, there was a significant correction and, in retrospect at least, the surge in equity prices in the technology sector and subsequent correction looks like a classic emergence and then bursting of an asset price bubble that appears to have contributed to the volatility of output. Naturally, the question arises whether monetary policy could have encouraged better outcomes by responding more directly to developments in the equity markets.

Bernanke and Gertler (1999) studied this question using simulations with a model that incorporates a policy rule. They conclude that trying to stabilize asset prices is problematic, in large part because it is nearly impossible for policymakers to infer

whether a given movement in asset prices are consistent with fundamentals or the product of some unsustainable speculative factors. Therefore, responding directly to asset prices is more likely to introduce noise and detract from the pursuit of policymakers' fundamental objectives. They argue that a disciplined monetary policy, focused on an inflation target, will tend to tighten during inflationary asset booms and ease during deflationary asset price busts, reducing the prospects of larger and more disruptive asset price bubbles and corrections.

While Bernanke and Gertler confirm the conventional wisdom among monetary policymakers, their research methodology also illustrates the challenges of analyzing the implications of equity price movements. They take equity prices movements as exogenous, including an assumed correction in equity prices. It is, after all, difficult to model equity price movements that are not driven by fundamentals.

Cecchetti et al (2000) reach the opposite conclusion--that central bankers should directly respond to asset prices. Unfortunately, this result is predicated on the assumption that the central bank can identify an asset bubble and knows in advance when the bubble would burst.

Cecchetti et al point out that the difficulty of estimating fundamental value for equities is not different in principle from the difficulty of estimating output or unemployment gaps. This suggests that the same issues that we discussed in the section on noisy information may apply here. First, to the extent policymakers consider responding to deviations of equity prices from fundamental value, they should do so based on updated estimates of fundamental value, using all available information to implement the updating. Second, if policy did respond to such deviations, there may be a

case for downweighting that response, including using a nonlinear response suggested in Meyer, Swanson and Wieland (2001), in periods of heightened uncertainty about fundamental value. Still, in my view, there is no convincing evidence to date that a direct response by monetary policy to equity prices, in any shape of form, would improve economic performance, measured in terms of output and inflation variability. But as I noted, it is very difficult to develop models that can definitively answer this question, and this topic remains a fertile area for future research.

There is, I believe, in addition, an important connection between the issue of responding to equity prices and that of a time-varying real interest rate. In particular, an acceleration in productivity--perhaps the dominant factor driving economic performance in the second half of the 1990s--might result in both a fundamentally-based rise in equity prices and, in addition, some speculative overshooting. A productivity acceleration also would be expected to raise the equilibrium real interest rate. To the extent that monetary policy allows the real funds rate to rise with any increase in the equilibrium real interest rate, it should provide some restraint to an open-ended speculative rise in equity prices.

VIII. Lessons from Research on Policy Rules

1. Simple policy rules provide useful guidance for monetary policymakers, are effective in communicating the rationale underlying conduct of monetary policy, and describe how policy has been conducted in several countries, including the United States, at least since the mid 1980s.
2. Simple rules work almost as well as optimal rules, are easier to communicate to the public, and are more robust to model uncertainty than optimal rules.

3. The most important feature of such rules is that policymakers should raise real interest rates in response to an increase in inflation.
4. The optimal response to utilization rates is more controversial. Except when uncertainty about the measurement of utilization rates is especially high, effective monetary policy involves some response to changes in utilization rates. Such response to utilization rates can reduce the volatility of both output and inflation.
5. Policymakers should continuously update their estimates of utilization rates, using all available data. However, even when such updating is implemented, uncertainty about utilization rates may justify some attenuation in the response to utilization rates, relative to the responses based on simple rules without measurement uncertainty.
6. When policymakers attenuate their response to utilization rates, they might improve policy outcomes by increasing the aggressiveness of their response to movements in inflation.
7. Gradually responding to movements in output and inflation gaps significantly smoothes interest rates without much loss in terms of output and inflation variability.
8. Research suggests that responding to forecasts as opposed to only incoming data may, on average, improve policy outcomes only marginally. However, I suspect that policy may be improved if policymakers are prepared to switch to a forecast-based approach following a sharp change in the forecast.
9. When the economy is confronted by an adverse shock at a time when interest rates are already low, policy should move more aggressively than normal.

References

- Advisory Commission to Study the CPI, "Final Report of the Advisory Commission to Study the Consumer Price Index," Washington, D.C.: U.S. Government Printing Office, December 1996.
- Ball, Lawrence, "Efficient Rules for Monetary Policy," *International Finance*, 1999, 63-83.
- Ball, Lawrence, "Policy Rules for Open Economies," in John B. Taylor, ed., *Monetary Policy Rules*, 1999, 127-144.
- Barro, Robert J, and David B. Gordon, "A Positive Theory of Monetary Policy in a Natural-Rate Model," *Journal of Political Economy*, August 1983, 589-610.
- Bernanke, Ben and Mark Gertler, "Monetary Policy and Asset Price Volatility," in *New Challenges for Monetary Policy*, Proceedings of a symposium sponsored by the Federal Reserve Bank of Kansas City, 1999, 77-128.
- Brainard, William C., "Uncertainty and the Effectiveness of Policy," *American Economic Review*, 1967, 411-425.
- Braun, Steven N., "Productivity and the NAIRU (and Other Phillips Curve Issues)," Working Paper Series, Federal Reserve Board, June 1984.
- Cecchetti, Stephen G., Hans Genberg, John Lipsky and Sushil Wadhvani, *Asset Prices and Central Bank Policy*, Geneva Reports on the World Economy 2, International Center for Monetary and Banking Studies, and Centre for Economic Policy Research. 2000.
- Clarida, Richard, Jordi Gali, and Mark Gertler, "The Science of Monetary Policy," *Journal of Economic Literature*, 1999, 1661-1707.
- Clarida, Richard, Jordi Gali and Mark Gertler, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, 2000, 147-180.
- Clarida, Richard, Jordi Gali and Mark Gertler, "Monetary Policy Rules in Practice: Some International Evidence," *European Economic Review*, 1998, 1003-1067.
- English, Bill, Bill Nelson, and Brian Sack, "Can Omitted Variables Explain the Significance of Lagged Interest Rates in Estimated Policy Rules?" Mimeo, Board of Governors of the Federal Reserve System.
- Estrella, A. and Mishkin, F., "Rethinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty," in John B. Taylor, ed., *Monetary Policy Rules*. NBER and Chicago Press, 1999.

Friedman, Milton, "The Role of Monetary Policy," *American Economic Review*, March 1968, 1-17.

Fuhrer, Jeffrey C. and Mark S. Sniederman, eds., *Monetary Policy in Low Inflation Environment*, *Journal of Money, Credit and Banking*, November 2000, Part 2.

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

Levin, Andrew, Volker Wieland, and John C. Williams, "Robustness of Simple Monetary Policy Rules under Model Uncertainty," in *Monetary Policy Rules*, ed. John B. Taylor, Chicago, University of Chicago Press, 1999, 263-299.

Levin, Andrew, Volker Wieland, and John C. Williams, "The Performance of Forecast-Based Monetary Policy Rules Under Model Uncertainty," Finance and Economics Discussion Series, 2001-29, Federal Reserve Board, August 2001.

Kydland, Finn E. and Edward C. Prescott, "Rules rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, June 1977, 473-491.

Lebow, David E. and Jeremy B. Rudd, "Measurement Error in the Consumer Price Index: Where Do We Stand?" Board of Governors of the Federal Reserve System, December 7, 2001.

Meltzer, Alan, "Limits of Short-Run Stabilization Policy," *Economic Inquiry*, 1987, 1-14.

Meyer, Laurence H., "The Economic Outlook and the Challenges Facing Monetary Policy," October 19, 2000.

Meyer, Laurence H., "The Strategy of Monetary Policy," 1998.

Meyer, Laurence H., Eric Swanson, and Volker Wieland, "NAIRU Uncertainty and Nonlinear Policy Rules," *American Economic Review*, May 2001.

Orphanides, Athanasios, "Monetary Policy and the Great Inflation," presented at the 2002 Meeting of the American Economic Association, January 2002.

Orphanides, Athanasios, "Monetary Policy Evaluation with Noisy Information", Finance and Economics Discussion Series 1998:50. Board of Governors of the Federal Reserve System, Washington, D.C., 1998.

Orphanides, Athanasios, "Monetary Policy Rules Based on Real-Time Data," *American Economic Review*, 2001.

Orphanides, Athanasios, "Monetary Policy Rules, Macroeconomic Stability and Inflation: The View from the Trenches," Finance and Economics Discussion Series, 2001-62, Board of Governors of the Federal Reserve System, December 2001.

Orphanides, Athanasios, Richard D. Porter, David Reifschneider, Robert Tetlow and Fredrico Finan, "Errors in the Measurement of the Output Gap and the Design of Monetary Policy," *Journal of Economics and Business*, 2000, 117-141.

Orphanides, Athanasios and Volker Wieland, "Inflation Zone Targeting," *European Economic Review*, 2000.

Reifschneider, David and John C. Williams, "Three Lessons for Monetary Policy in a Low-Inflation Era," *Journal of Money, Credit and Banking*, November 2000, Part 2, 936-996.

Rudebusch, Glenn, "Is the Fed Too Timid?" *Review of Economics and Statistics*, 2001, 203-217.

Rudebusch, Glenn and Lars Svensson, "Policy Rules for Inflation Targeting," in John B. Taylor, ed., *Monetary Policy Rules*. Chicago: NBER and Chicago Press 1999, 203-246.

Sack, Brian, "Does the Fed Act Gradually? A VAR Analysis," *Journal of Monetary Economics*, 2000, 229-256.

Sack, Brian and Volker Wieland, "Interest Rate Smoothing and Optimal Monetary Policy" *Journal of Economics and Business*, 2000, 205-228.

Sargent, T., Discussion of 'Policy Rules for Open Economies' by Lawrence Ball, in John B. Taylor, ed., *Monetary Policy Rules*. Chicago: NBER and Chicago Press, 1998.

Smets, Frank, "Output Gap Uncertainty: Does it Matter for the Taylor Rule?" in B. Hunt and A. Orr, eds., *Monetary Policy under Uncertainty*, Wellington, New Zealand: Reserve Bank of New Zealand, 1999.

Soderstrom, Ulf, "Should Central; Banks Be More Aggressive?" Central Bank of Sweden, manuscript, 1999a.

Soderstrom, Ulf, "Monetary Policy with Uncertain Parameters," Central Bank of Sweden manuscript, 1999b.

Svensson, L. and M. Woodford, "Indicator Variables for Monetary Policy," Mimeo, Princeton University. 2000.

Swanson, Eric, "On Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rules," Federal Reserve Board Finance and Economics Discussion Series 2000-32, 2000.

Taylor, John, "Discretion vs Policy Rules in Practice," *Carnegie Rochester Conference Series on Public Policy*, 1993, 195-214.

Taylor, John B., "Historical Analysis of Monetary Policy Rules," in John B. Taylor, ed., *Monetary Policy Rules*, 1998, 319-340.

Tetlow, Robert and Peter von zur Muehlen, "Robust Monetary Policy with Misperceived Models: Does Model Uncertainty Always Call for Attenuated Policy?" *Journal of Economic Dynamics and Control*, 2001, 911-949.

Westaway, Peter, "Does Time Inconsistency Really Matter?" in *Dynamic Modeling and Control of National Economies*, Edinburgh, Scotland: International Association of Automatic Control, 1989, 145-152.

Wolman, Alex, "Staggered Price Setting and the Zero Bound on Nominal Interest Rates," *Federal Reserve Bank of Richmond Economic Quarterly*, fall 1998, 1-24.