

Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice*

Michael Woodford
Columbia University

Revised draft: December 20, 2007

Abstract

Forecast targeting is an innovation in central banking that represents an important step toward more rule-based policymaking, even if it is not an attempt to follow a policy rule of any of the types that have received primary attention in the theoretical literature on optimal monetary policy. This paper discusses the extent to which forecast targeting can be considered an example of a policy rule, and the conditions under which it would represent a desirable rule, with a view to suggesting improvements in the approaches currently used by forecast-targeting central banks. Particular attention is given to the intertemporal consistency of forecast-targeting procedures, the assumptions about future policy that should be used in constructing the forecasts used in such procedures, the horizon with which the target criterion should be concerned, the relevance of forecasts other than the inflation forecast, and the degree of robustness of a desirable target criterion for monetary policy to changing circumstances.

*Revision of a paper presented at the conference “John Taylor’s Contributions to Monetary Theory and Policy,” Federal Reserve Bank of Dallas, October 12-13, 2007. I would like to thank Olivier Blanchard, Ray Fair, Marc Giannoni, Rick Mishkin, Ed Nelson, Bruce Preston, Lars Svensson, and John Taylor for helpful discussions; Mehmet Passaogullari for research assistance; the NSF for research support through a grant to the NBER; and the Arthur Okun and Kumho Visiting Professorship, Yale University, for providing the time to write this paper.

For several decades, John Taylor has advocated an ambitious program of research on quantitative rules that could serve as guidelines for monetary policy. While the comparative study of historical performance under alternative policies has helped to shape Taylor's views (Taylor, 1999b), probably the most distinctive element of his approach to the problem has been his use of macroeconometric models for the normative analysis of alternative policy rules (Taylor, 1979, 1993b, 1999c). In addition to his important contributions to the technical methods for the development, estimation, and numerical analysis of such models, Taylor has constantly been concerned with the issue of the robustness of policy proposals to model uncertainty (Taylor, 1999a), and with the distillation of the results of the research literature into a form that could influence actual policy (Taylor, 1993a, 1998).

To what extent has the literature that Taylor has launched arrived at conclusions that are likely to help to improve the conduct of policy by central banks? In my view, the most important recent development with regard to the practical use of policy rules has been the development, at several central banks since the early 1990s, of methods of *forecast targeting*, both as a systematic approach to monetary policy deliberations and as a basis for communication with the public.¹ Forecast targeting relies heavily on the use of quantitative structural models of the effects of monetary policy, and in this respect at least the success of the approach relies heavily on the output of the research program promoted by Taylor. Moreover, forecast targeting implies a decision process focused on the achievement of specific quantitative objectives, and typically implies a greater degree of explicitness about both the goals of policy and the justification of particular policy decisions than is seen at other central banks; in these respects, it represents an important step toward the ideal of rule-based policymaking. I believe that the most likely way in which the research literature on monetary policy rules can help to improve actual policy is through contributing to the refinement of forecast-targeting procedures.

One cannot say, however, that forecast targeting, as developed at central banks

¹The banks that best illustrate this approach all have inflation targets, which play a prominent role in their targeting approaches; and Svensson (1997) refers to the type of policy regime with which I am concerned as “inflation-forecast targeting.” But I wish to discuss the optimal design of a forecast targeting regime without necessarily assuming that only the inflation forecast should play a role — indeed, I shall argue below that normative policy analysis indicates a role for other projections as well — and so I prefer to speak simply of “forecast targeting.” Earlier discussions of forecast targeting as a general approach include Svensson (2003, 2007) and Svensson and Woodford (2005).

in the early 1990s, represented an attempt to implement ideas from the theoretical literature on monetary policy rules. Nor has the extensive literature on monetary policy rules of the past two decades given a great deal of attention to the analysis of procedures of this kind. This paper seeks to examine the extent to which forecast targeting does represent a desirable policy rule, from the standpoint of what Taylor (2000) calls “the new normative macroeconomics,” and the extent to which theoretical analyses of optimal monetary policy can provide guidelines that should improve the operation of forecast targeting regimes. In seeking to find connections between the theoretical literature and a policy framework, developed on the basis of primarily practical considerations, that is currently used at a number of central banks, it represents an essay in what Taylor (1998) calls “translational economics.”

1 Forecast Targeting as a Policy Rule

It is first important to be clear about what I mean by forecast targeting. Lars Svensson (1997) introduced the term “inflation-forecast targeting” (IFT) to refer to a policy regime with more specific characteristics than the mere announcement of an inflation target, though the existence of a public target for some measure of inflation is an important feature of such regimes. First of all, IFT involves a commitment to a particular *decision procedure* for monetary policy: the central bank’s operating target for the policy instrument should be adjusted in the way that is judged necessary in order to ensure that the bank’s projections of the economy’s future evolution satisfy certain conditions, which I shall call the *target criterion*, at all times. The instrument of policy is typically an overnight interest rate in an interbank market, similar to the federal funds rate in the US, though this is not essential to the logic of forecast targeting, and neither is it a feature that distinguishes this approach from the current conduct of policy at most other central banks, such as the Fed. The target criterion should be a specific quantitative property of the projections, so that (in principle at least) there should be little debate about whether a given set of projections satisfy the target criterion, even if a great deal of judgment may be involved in producing the projections themselves. As an example of what is intended, the Bank of England has often described its decision procedure as checking that the projection for a particular variable (currently, CPI inflation) equals a particular value (the official inflation target, currently 2.0 percent) at a particular horizon (8 quarters

in the future).²

And secondly, forecast targeting involves a distinctive approach to *communication policy*, under which the central bank regularly publishes the quantitative projections on the basis of which policy has been judged to be on track, together with extensive discussion of the reasoning underlying these projections. This is a key feature of the *Inflation Reports* that are published three or four times per year by the leading practitioners of IFT, such as the Reserve Bank of New Zealand, the Bank of England, Sweden’s Riksbank, and the Norges Bank.³ Forecast-targeting central banks have led the way in increasing the transparency of monetary policy deliberations, most notably through these publications, and this is not fortuitous. For the decision procedure associated with forecast targeting both lends itself more easily to such communication (because it is highly structured) and is particularly dependent on transparency for its success (because the procedure would involve little discipline if the central bank did not have to discuss its projections with anyone outside its own walls).

To what extent can a regime of this kind be considered an example of a *policy rule*? One might think that it is not a rule at all — or at least, not a rule that is sufficiently well-specified to be subjected to the kind of quantitative analysis that is the hallmark of the Taylor research program — insofar as no precise recipe is given for the adjustment of a policy instrument, or even a precise specification of some “intermediate target” that can be influenced relatively directly by the central bank. And it is true that forecast targeting represents a different style of specification of a policy rule — what Svensson and Woodford (2005) call a “higher-level specification” of the policy rule — than such familiar examples of policy rules as Milton Friedman’s proposal of a constant target for money growth, or the “Taylor rule” (Taylor, 1993a).

²See, for example, Vickers (1998) and Goodhart (2001). The justifications given for policy decisions in the Bank of England’s *Inflation Report* more recently do not suggest quite so simple a target criterion; for example, there are frequent references to inflation projections beyond the 8-quarter horizon, as well as to the projection for output growth. However, the introduction to each *Inflation Report* still always includes a chart showing the projection for CPI inflation, just before the summary discussion of the most recent policy decision, and this chart always includes a horizontal line at the inflation target of 2.0 percent and a dashed vertical line at the horizon 8 quarters in the future, allowing easy visual inspection of the degree to which the simple target criterion is satisfied.

³The latter two central banks have recently changed the name of their publications to *Monetary Policy Report*, presumably in recognition of the fact that the inflation projection is not the sole focus of these publications.

Nonetheless, such a regime does serve many of the most important objectives that proponents hope to achieve through adoption of a policy rule.⁴ First, it does increase the systematic character of policy decisions. Even if actual forecast targeting regimes have not made policy decisions as simple as some proposals in the academic literature would, they represent a substantial movement in this direction relative to the procedures actually followed by other central banks. This has multiple advantages: in addition to potentially improving the reliability of policy decisions (both by structuring policy deliberations, and by allowing accumulated wisdom to be more efficiently transmitted to new members of the policy committee), it helps to reduce political interference in central-bank deliberations. (It is surely no accident that the Bank of England was granted independent authority to set interest rates only after five years of experience with its IFT regime.)

And secondly, a forecast targeting regime serves to make policy decisions more easily forecastable by the private sector. This increases the effectiveness of policy for two somewhat different reasons.⁵ On the one hand, the actual effects of monetary policy on spending decisions and ultimately on the rate of inflation occur largely as a result of the effect of central-bank actions and announcements on market expectations regarding the future *path* of short-term interest rates, rather than through direct effects of the current level of overnight rates itself. Hence achieving the effects of policy that are desired depends on private-sector expectations regarding agreeing with the intentions of the central bank. And on the other hand, the benefits of price stability very much depend on the degree to which economic actors remain *confident* of the stability of the monetary unit. Anchoring inflationary expectations is therefore important, and one of the most important arguments for commitment to a policy rule is the expectation that such a commitment should give people a better ground for expecting a particular future rate of inflation. But an IFT regime, if implemented in a credible way, is well suited to stabilize such expectations. In addition to making visible the central bank's commitment to a systematic procedure that is intended to ensure a stable rate of inflation over the medium run, the constant emphasis in public communications on the central bank's own forecast of future inflation can only help to

⁴Taylor (1998) provides a useful list of reasons for the growing consensus among monetary economists regarding the desirability of a policy rule.

⁵This issue is discussed in more detail in Woodford (2005).

ensure that the public understands this aspect of the policy's intended consequences.⁶

In thinking about whether a forecast-targeting procedure constrains central-bank behavior to a sufficient extent to achieve the benefits that one hopes to obtain from a policy rule, it is useful to recognize that the conduct of policy under such a regime can actually be described at three distinct levels. At the highest level of generality, a specification of the target criterion explains what projected outcome the central bank is seeking to achieve through its choice of its operating target for overnight interest rates, though it does not specify exactly what the interest-rate target should be. At a more specific level, one could instead describe the specific procedure that should be used to determine the appropriate interest-rate choice in each decision cycle, in order to satisfy the target criterion, but without specifying exactly which repurchase operations should be conducted on any given day. And finally, at the most specific level, one could describe the way in which the central bank decides each day on the appropriate quantity of cash to supply to the money markets through its repurchase operations, in order to achieve its operating target for the policy rate. At each successively lower level of the specification, one comes closer to saying precisely what the central bank ultimately must do. At each lower level, finer institutional details about the precise mechanism through which monetary policy affects the economy become relevant. And finally, at each lower level, it is appropriate for the central bank to be prepared to adjust course more frequently on the basis of more recent information. At the lowest level, a decision about the quantity of repurchases must be made daily (at most central banks), and occasionally even more frequently; at the intermediate level, the interest-rate operating target is ordinarily reconsidered only at intervals of several weeks; and at the highest level, the target criterion should remain fixed for years at a time, though here too reconsiderations will be appropriate from time to time in light of improved understanding or in response to structural change in the economy.

At which, or how many, of these levels must the nature of policy be spelled out, in

⁶A public commitment to a money growth rate target, by contrast, only directly helps to stabilize inflation expectations to the extent that members of the public understand the economic theory according to which a given rate of money growth should imply a particular rate of inflation, at least over long enough periods of time. Of course, any rule that succeeds in maintaining inflation at a low and stable rate over a period of time should eventually result in stable inflation expectations as a consequence of observed performance; but an IFT regime has this virtue to the same extent as other proposed rules.

order for the policy specification to count as a “rule”? I think there is wide agreement that there is no need for explicit description of the lowest-level specification of policy; the literature that compares the consequences of alternative policy rules generally takes it as given that any non-negative target for the policy rate can be implemented with a high degree of accuracy over time scales (a day or two) that are quite short compared to those that matter for the effects of interest rates on the basis of which the policy is to be judged, and that the details of the required open-market operations have little or no consequences for the objectives of policy.

It is less obvious that description of a policy solely at the highest of these three levels suffices, and the literature on the quantitative evaluation of policy rules has almost exclusively focused on rules specified as formulas to determine the value of a policy instrument that is under the relatively direct control of the central bank. Nonetheless, I think there are important advantages to considering rules that are specified by target criteria that need not involve *any* variable over which the central bank has direct control.

A first question is whether a mere specification of a target criterion suffices to *fully determine* outcomes under the policy, so that one can compare the outcomes associated with alternative policies. The answer is that it can, if one assumes that the target criterion will be satisfied at all times. One need not specify the exact actions of the central bank that result in its being satisfied, in order to ask how inflation, output and other variables would have to evolve in a rational expectations equilibrium of that kind. (In principle, the point is the same as when one proposes to analyze the consequences of a given interest-rate feedback rule, without specifying how the central bank will determine the size of repurchase operations that will be required each day in order to achieve overnight interest rates consistent with the rule — though admittedly, there is greater reason to question the degree of precision with which it is feasible to satisfy the target criterion, when it involves variables such as the overall rate of inflation.) One can use the target criterion itself as the “missing equation” that specifies monetary policy, allowing a solution (in the case of a suitably chosen target criterion) for a determinate rational expectations equilibrium (REE). Hence one can study the advantages of alternative target criteria, using the same methods as the literature initiated by Taylor has used to assess the advantages of alternative feedback rules.

Of course, I do not mean to claim that there should be a determinate REE as-

sociated with *any* target criterion whatsoever. For example, a criterion that only involves projected outcomes two or more years in the future is one that is unlikely to imply a determinate solution; there will be alternative paths by which the economy could reach a situation consistent with the criterion, and in such a case the target criterion fails to fully determine policy. In my view, it is important to adopt a target criterion that *does* fully determine (but not over-determine) a particular equilibrium. But this is a property that one can analyze given a specification of the target criterion alone; one need not specify the policy at a lower level in order to check this. And one should recall that there is also a question whether a given interest-rate feedback rule determines a unique equilibrium or not; one argument for the importance of choosing a rule that conforms to the “Taylor Principle” is that in many models, rules with weaker feedback from realized inflation to the interest-rate operating target have been found to result in indeterminacy of equilibrium (e.g., Woodford, 2003, chap. 4).

Nor do I wish to suggest that there are no important issues connected with the problem of implementation of target criteria. However, in the case of a target criterion that is found to satisfy the property just mentioned — that there exists a determinate REE associated with it — then there should exist a monetary policy that should satisfy the target criterion. (In fact, solution for the REE associated with the criterion should already indicate, among other things, the state-contingent evolution of the nominal interest rate that must obtain, at least under an REE, if the criterion is to be satisfied.) Thus it is possible to search for a desirable target criterion simply on the basis of a consideration of alternative policies specified at this level of generality, and to turn to the problem of implementation only once one has chosen a target criterion. (Some of the subtle issues that remain with regard to implementation are taken up in section 4 below.)

A second question is whether specification of a target criterion, rather than a reaction function, is a useful way of providing a *guideline for policymakers* in their deliberations. Of course, a monetary policy committee has to decide on the level of overnight interest rates, so the target criterion alone does not provide them sufficient information to discharge their duty. Nonetheless, a target criterion relating the paths of some of the variables that the policy committee wishes to stabilize seems the appropriate level of detail for a prescription that a policy committee can agree to use to structure its discussions, that can be explained to new members of the committee, and that can ensure some degree of continuity in policy over time. Special factors

are likely to be important at each meeting, in deciding upon the level of interest rates consistent with fulfillment of the target criterion; hence it is difficult to impose too much structure on this kind of deliberation, without the committee members feeling that their procedures are grossly inadequate to dealing with the complexity of the situation in which they find themselves. But the considerations involved in a judgment that a particular target criterion is sensible are less likely to constantly change.

Indeed, there are good theoretical reasons (discussed further in section 3) to expect that a desirable target criterion will depend on fewer details about the current economic environment than would a desirable specification of a reaction function. Giannoni and Woodford (2002, 2005) show how to construct robustly optimal target criteria, which implement an optimal response to shocks regardless of which types of shocks are more important, or of the degree of persistence, forecastability, and so on of the shocks that occur. The coefficients of an optimal reaction function will instead depend on the statistical properties of the shocks.⁷ Since each shock that occurs is somewhat different from any other, there will always be new information about the particular types of disturbances that have most recently occurred, making advance commitment to a particular reaction function inconvenient. The coefficients of the optimal target criterion may also change⁸ in the event of a shift in the central bank's estimate of structural parameters such as elasticities of supply or demand; but information of this kind is not likely to shift as dramatically so suddenly.

Of course, the Taylor rule was not proposed as a mechanical formula that would precisely determine the federal funds rate operating target at each point in time. Instead, Taylor (1993a) describes it as a “guideline” that indicates how policy should be conducted under normal conditions, but from which policymakers will frequently be justified in deviating, in response to special circumstances that arise. But an instrument rule subject to an open-ended escape clause of this kind deserves less to be called a policy *rule* than does a target criterion which is intended to be the focus of policy deliberations under virtually all circumstances, even if the considerations that should determine the policy rate consistent with the target criterion are not spelled out in advance. The target criterion approach provides a more consistent structure for policy deliberations; for example, the rule itself makes it clear when

⁷This is illustrated by Svensson and Woodford (2005) in the context of a simple example.

⁸Even this need not be so, as is illustrated by the discussion in section 3.

new circumstances justify a departure from standard rules of thumb for interest rate decisions: whenever the new developments cast doubt on one's normal expectations about the relation between the policy rate and the variables that enter the target criterion. The targeting approach should make the consequences of the central bank's decisions more predictable, since one should be able to count on the target variables satisfying the target criterion to a reasonable extent, even if the path of the policy rate that this involves will not always be highly predictable. And the targeting approach provides greater protection against political pressure on policy decisions, since one need not explain to the politicians why current circumstances should not provide yet another fine occasion for an exception to the usual rule of thumb; instead, the discussion can be kept on the plane of the relatively technical issue of which level of interest rates will lead to paths for inflation and real activity consistent with the target criterion.

Yet a third question is whether a target criterion represents a useful way of *explaining* the nature of a central bank's policy commitments to the public. Some might feel that a commitment to aim at satisfaction of a particular target criterion is less meaningful than a commitment to a particular instrument rule, or even than a commitment to a quantitative target for some "intermediate target" that can be fairly directly controlled by the central bank, on the ground that it is less specific about what the central bank will do, so that whether the central bank is actually complying is less directly verifiable. But there are two important counter-arguments to such a view. First, while the target criterion is less explicit about what the central bank will do with the instruments that it can most directly influence, the target criterion has *more* explicit implications for the evolution of "target variables" such as the inflation rate; and the main reason for wishing to establish a credible commitment to a policy rule is to anchor private-sector expectations regarding these variables. And second, while a "higher-level" description of the policy commitment might seem to reduce verifiability, this problem can be overcome, to an important extent, through a commitment to public *explanation* of how policy decisions have been determined by the target criterion — which is precisely the function served by the discussion of the bank's quantitative projections in a *Monetary Policy Report*. Moreover, accountability is increased, to the extent that the higher-level commitment represents one that can actually determine policy decisions more consistently, rather than being subject to so many "escape clauses" as will inevitably be required in the case of an explicit

instrument rule.

Another important argument for policy rules does not depend on any supposition that central bankers have any difficulty determining the action that would best serve their objectives on any given occasion, or that the private sector may fail to correctly understand the systematic character of policy. Instead, Kydland and Prescott (1977) argue that a process of sequential optimization, with no advance commitment regarding future policy actions, is inherently flawed, on the ground that a sequential optimizer will never have any reason to take into account the way in which his systematic (and hence predictable) response to current conditions has shaped prior expectations, as these expectations are already a historical fact by the time that the decision has to be made. Commitment to conduct policy in accordance with a rule, regardless of whether the required actions are those the policymaker would most prefer at the time that the actions are taken, can solve this problem. But to what extent are forecast-targeting procedures examples of systematic approaches to policy that avoid the pitfalls of discretionary policymaking identified by Kydland and Prescott?

Forecast targeting is a sequential decision procedure; rather than choosing a plan for policy over several years at one time and then sticking to it, new projections are computed and the targeting exercise is repeated afresh, several times per year. It is true that it requires a central bank to be clear about its objectives, and clear about what it expects the effects of its policy actions to be; but these are also features of “discretionary” policy in the sense of Kydland and Prescott. Does forecast targeting really represent anything other than a more “scientific” way of implementing discretionary policy — one that can therefore more closely approximate the theoretical model of sequential optimization proposed by Kydland and Prescott, but that does nothing to overcome the inherent flaws of sequential optimization that they identify?

The answer to this depends on exactly how forecast targeting is implemented. It might indeed correspond precisely to discretionary policy in the sense of Kydland and Prescott, and it could also correspond to something less coherent than that, and even less successful at achieving the bank’s stabilization goals. But if appropriately implemented, a forecast-targeting procedure can address the problem identified by Kydland and Prescott; in fact, under ideal circumstances, it can provide a convenient approach to implementation of the equilibrium that would result from a once-and-for-all commitment to an optimal state-contingent policy. But achieving or even approaching this ideal requires that one be careful about a number of details of what

is meant by forecast targeting.

2 The Problem of Intertemporal Consistency

An important potential advantage of forecast targeting, stressed above, is the possibility that the published projections can help to clarify what the private sector should expect, and thus prevent the bank's stabilization objectives from being thwarted by private actions based on mistaken forecasts. But the degree to which publication of central-bank projections can be expected to shape the expectations of private decisionmakers will depend on how credible these projections are as forecasts of the economy's likely evolution. Among the possible grounds for doubt is a tension inherent in the logic of the forecast-targeting procedure itself. Production of projections of the economy's evolution several years into the future requires that the central bank make assumptions about its conduct of policy not merely in the immediate future, but over the entire forecast horizon (and even beyond, in the case of a forward-looking model). But while the projections must specify policy far into the future each time they are produced, in each decision cycle policy is only *chosen* for a short period of time (say, for the coming month, after which there will be another decision).

This raises a question as to whether this decision procedure should be expected to actually produce the kind of future policy that is assumed in the projections. One might imagine, for example, a central bank wishing always to choose expansionary policy at the present moment, to keep employment high, while projecting that inflation will be reduced a year or two in the future, so that the expectation of disinflation will make it possible to have high employment with only moderate inflation. But if the procedure is one in which the disinflation is always promised two years farther in the future, private decisionmakers have no reason ever to expect any disinflation at all.

Thus one requirement for credibility of the central bank's projections is that the forecast-targeting procedure be *intertemporally consistent*. This means that the future policy that is assumed in the projections should coincide with the policy that the procedure itself can be expected to recommend, as long as those aspects of future conditions that are outside the control of the central bank turn out in the way that is currently anticipated. While this may seem an obvious requirement, a number of apparently sensible approaches to forecast-targeting fail to satisfy it.

2.1 Constant-Interest-Rate Projections

A popular approach in the early years of inflation-forecast targeting — used, for example, in the Inflation Reports of the Bank of England prior to August 2004 — was to construct projections conditional upon a constant interest rate over the forecast horizon (Vickers, 1998; Jansson and Vredin, 2003). The appropriate current interest-rate decision was then taken to be the interest rate that, if expected to be maintained over the forecast horizon, would lead to projections satisfying the target criterion (for example, 2 percent inflation 8 quarters in the future). This procedure had a number of advantages. First, a bank had only to consider variations in policy over a single dimension (alternative constant interest rates), with the consequence that a one-dimensional target criterion would suffice to identify the correct policy. Hence it was enough to specify what the inflation projection should be like some years in the future, without having to take a stand on the trickier question of how one might choose among alternative nearer-term transition paths. Second, contemplated changes in the current interest-rate decision would be predicted to have non-trivial consequences, given that any change was assumed (for purposes of the projection exercise) to be a permanent change. And finally, it was possible to construct projections without the bank's having to tip its hand as to the likely character of future policy.

But constant-interest-rate projections raise a number of conceptual problems (Goodhart, 2001; Leitemo, 2003; Honkapohja and Mitra, 2005; Woodford, 2005). The assumption that the nominal interest rate will remain fixed at some level, regardless of how inflation or other variables may evolve, is not a sensible one. Moreover, in forward-looking (rational-expectations) models of the kind that are now beginning to be used by central banks, the assumption of a constant nominal interest rate often implies an indeterminate price level, so that it becomes impossible to solve uniquely for an inflation forecast under any such interest-rate assumption.⁹ In models with backward-looking expectations, the model can be solved, but such policies often imply explosive inflation dynamics. Such difficulties appears to have been a frequent problem with the constant-interest rate projections of the Bank of England (Goodhart, 2001), which often showed the inflation rate passing through the target rate at the eight-quarter horizon, but not converging to it. Figure 1 provides an example. In such a case, it is not obvious why anyone should believe that policy is consistent

⁹See Woodford (2003, chap. 4) for examples of this problem.




Figure 1: The February 2004 CPI projection under the assumption of a constant 4.0 percent interest rate. Source: Bank of England, *Inflation Report*, August 2004.

with the inflation target, or expect that inflation expectations should be anchored as a result of a commitment to such a policy.

The most fundamental problem, however, is there will often be no reason to expect interest rates to remain constant over the policy horizon. Indeed, constant-interest rate projections themselves often imply that the people making the projections should not expect the interest rate to be maintained over the forecast horizon. Consider, for example, the inflation projection shown in Figure 1, a constant-interest rate projection on the basis of which the February 2004 Bank of England *Inflation Report* concluded that a 4 percent policy rate was appropriate at that time.¹⁰ The figure shows that under the assumption of a constant 4 percent policy rate, consumer price inflation was projected (under the most likely evolution, indicated by the darkest area) to pass

¹⁰In the February *Report*, only the projection up to the 8-quarter horizon was shown. The figure that has been extended to a horizon 12 quarters in the future is taken from the August 2004 *Inflation Report*, in which the Bank explained its reasons for abandoning the method of constant-interest-rate projections.

through the target rate of 2.0 percent at the eight-quarter horizon (indicated by the vertical dashed line), and then to continue rising in the following year. Thus, if the policy rate were to be held at 4 percent for a year, the Bank's expectation in February 2004 should have been that (under the most likely evolution, given what was known then) in February 2005 a similar exercise would forecast consumer price inflation to pass through 2.0 percent at the one-year horizon, and to exceed 2.0 percent during the second year of the projection. Hence, the Bank has essentially forecasted that in a year's time, under the most likely evolution, the policy committee would have reason to raise the policy rate. Thus the February 2004 projection itself could have been taken as evidence that the Bank should not have expected the policy rate to remain at 4 percent over the following eight quarters.

As these issues have come to be understood, a number of central banks that formerly relied upon constant-interest-rate projections (such as the Bank of England, since August 2004) have switched to an alternative approach. This is the construction of projections based on *market expectations* of the future path of short-term interest rates, as inferred from the term structure of interest rates and/or futures markets. In the case that the projections constructed under this assumption satisfy the target criterion, the correct current interest-rate decision is taken to be the one consistent with market expectations. The use of projections based on market expectations allows a central bank to avoid assuming a constant interest rate when there are clear reasons to expect rates to change soon, while still not expressing any view of its own about the likely future path of interest rates.

But the market expectations approach does not really solve the problem of internal consistency just raised.¹¹ One problem is that market expectations can at most supply a single candidate forward path for policy; it is not clear what decision one is supposed to make if that path does *not* lead to projections consistent with the target criterion. Thus the procedure is incompletely specified; and if it is only the projections based on market expectations that are published, even though the central bank has chosen to contradict those expectations, the published projections cannot be expected to shape private decisionmakers' forecasts of the economy's evolution.

Moreover, even if the forward path implied by market expectations does lead to projections that fulfill the target criterion, the exercise is not intertemporally

¹¹For further discussion of problems with this approach, see Woodford (2005) and Rosenberg (2007).

consistent if this path does not in fact correspond to the central bank’s own forecast of the likely future path of interest rates. Why should it count as a justification of a current interest-rate decision that this would be the first step along a path that *would* imply satisfaction of the target criterion, but that the central bank does not actually expect to be followed? And why should anyone who correctly understands the central bank’s procedures base their own forecasts on published projections constructed on such an assumption?

2.2 Choosing an Interest-Rate Path

In fact, there is no possibility of an intertemporally consistent forecast-targeting procedure that does not require the central bank to *model its own likely future conduct* as part of the projection exercise. Approaches like both of those just described — which introduce an artificial assumption about the path of interest rates in order to allow the central bank to avoid expressing any view about policy decisions that need not yet be made — necessarily result in inconsistencies. Instead, a consistent projection exercise must make assumptions that allow the evolution of the central bank’s policy instrument to be projected, along with the projections for inflation and other endogenous variables. In such a case, it would be possible, but somewhat awkward, for the central bank to remain silent about the implications of its assumptions for the forward path of interest rates; and so it is natural to include an interest-rate projection among the projections that are discussed in the *Monetary Policy Report*.¹² This has been done for the past decade now by the Reserve Bank of New Zealand, and is now done by the Norges Bank (since 2005) and the Riksbank (since 2007) as well. In the case of the latter two central banks, “fan charts” (similar to the one shown in Figure 1) are presented for the policy rate; this (among other things) makes it clear that the path is simply a forecast, rather than a definite intention that has already been formulated, let alone a promise.

But how should future policy be specified in such an exercise? It is sometimes

¹²Since one is talking about projections for the paths of endogenous variables, rather than announcing an intention, there is no reason why there need be a projection for only one interest rate, or even for the interest rate that is most emphasized to be the policy rate. Nonetheless, there are obvious advantages in giving primarily emphasis to only a small number of key variables; and it might seem disingenuous not to offer a view of the path of the policy rate, given that this is most directly under the bank’s own control.

suggested that the monetary policy committee should conceive of its task as the choice of a *path* for interest rates, rather than a single number for the current operating target, in each decision cycle. Discussions of the feasibility of such an approach have often stressed the potential difficulty of committee voting on a decision with so many dimensions.¹³ And when announcing its intention to begin publishing its own view of the path of the policy rate, the Riksbank (Rosenberg, 2007) indicated that it would publish “forecasts ... based on an interest-rate path chosen by the Executive Board.”¹⁴

However, the idea that one should simply ask the policy committee to decide which forward path for interest rates they prefer, presumably after asking their staff to produce projections for other variables conditional on each path that is considered, is problematic on several grounds that have nothing to do with the complexity of the decision or the need for a committee to agree among themselves. First of all, the specification of future policy by a simple path for a short-term nominal interest rate, independently of how endogenous variables may develop, is never a sensible choice, and is unlikely to lead to well-behaved results in a sensible model. (The problems mentioned above in connection with the assumption of a constant interest-rate path apply equally to *any* specification of an exogenous path; they do not result from the assumption that the interest rate does not vary with time, but from the assumption that it is independent of outcomes for inflation and other variables.) Moreover, the assumption of a specific path for interest rates, unaffected by future shocks, would seem to require one to publish a specific path for this variable, alongside the fan charts for variables such as inflation; but this would encourage the dangerous misunderstanding that the bank has already committed itself to follow a definite path

¹³See, for example, Goodhart (2005) for a skeptical view; Svensson (2007) responds by proposing a voting mechanism intended to overcome potential intransitivities in majority preferences over alternative paths.

¹⁴It is likely, of course, that this was only a loose way of speaking in a statement intended for a non-technical audience, and that the intention was to indicate that the Executive Board would have to endorse the assumptions about future policy involved in generating projections of an endogenous interest-rate path. The change in procedure does seem to have meant that the Executive Board is now required to approve the assumptions made in the projections in a way that was not previously true; this has made it necessary to allow for possible revisions in the projections following the meeting at which the policy decision is made (Sveriges Riksbank, *Monetary Policy Report* 2007/1, p. 21).

long in advance.

Even supposing that these technical issues have been finessed,¹⁵ there remains the more fundamental problem of the intertemporal consistency of the procedure. Here it is important to realize that the mere use of a *consistent criterion* over time to rank alternative projected paths for the endogenous variables — not just a criterion that provides a transitive ordering of outcomes within each decision cycle, but one that ranks different possible paths the same way, regardless of the date at which the decision is being made — is not enough to ensure intertemporal consistency, in the sense defined above. Thus the problems of choosing a forward path for policy are not resolved simply by asking the members of the policy committee to agree on a loss function that they will then use (for an entire sequence of meetings) to rank alternative possible outcomes, as proposed by Svensson (2007).

Even in the case of a single decisionmaker who minimizes a well-defined loss function that remains the same over time, using a correct economic model that also remains the same over time, and who never makes any calculation errors, the choice of a new optimal path for policy each period will not general lead to intertemporal consistency. For in the case of a forward-looking model of the transmission mechanism, the procedure will lead to the choice of a forward path for policy that one will not be lead by the same procedure to continue in subsequent decision cycles, even if there have been no unexpected developments in the meantime. The reason is the same as in the celebrated argument of Kydland and Prescott (1977) for the “time inconsistency of optimal plans”: the forward path chosen at one time will take account of the benefits at earlier dates of certain expectations about policy at the later dates, but as the later dates approach (and the earlier expectations are now historical facts), there will no longer be a reason to take into account any effect of the policy chosen for those dates on earlier expectations. This problem does not arise solely in connection with the bias in the average rate of inflation chosen by a sequential optimizer, as in the example of Kydland and Prescott (1977). One may

¹⁵For example, one might specify future policy by a policy rule, such a Taylor rule, with some number of free parameters that are optimized, in each decision cycle, so as to result in projections that are acceptable to the monetary policy committee. If only rules that are considered that imply a determinate equilibrium, the first problem is avoided. And since the rule that is chosen would make the interest rate endogenous, an assumption about the distribution of shocks in each future period would result in a probability distribution for future interest rates, just as for the future inflation rate.

solve the problem of “inflationary bias” by assigning the central bank a loss function in which the target level of the output gap is not higher than the level consistent on average with its inflation target, but the optimal dynamic responses to shocks are still not generally the ones that would be chosen under sequential (or discretionary) optimization.¹⁶

The problem can be illustrated using the familiar discussion of optimal monetary policy in Clarida *et al.* (1999). The available trade-off between inflation and real activity is assumed to be characterized by a “New Keynesian Phillips curve,”

$$\pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t, \quad (2.1)$$

where π_t is the rate of inflation between period $t - 1$ and period t , x_t is the output gap in period t , u_t is an exogenous “cost-push shock” (a stationary process with mean zero), and $\kappa > 0, 0 < \beta < 1$. The goal of policy is assumed to be the minimization of a quadratic loss function of the form

$$E_0 \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda x_t^2], \quad (2.2)$$

for some relative weight $\lambda > 0$. Here the specification of a target value for the output gap of zero means that the target output gap is no different from the level that is consistent with hitting the inflation target (also zero here) on average, and as a result there is no inflationary bias associated with discretionary optimization.

The optimal dynamic response to a cost-push shock is nonetheless *not* the one that would occur in an equilibrium¹⁷ with sequential optimization on the part of

¹⁶In the literature on inflation targeting, it is sometimes supposed instead that there is no problem with allowing a central bank complete discretion in its choice of the instrument settings that will minimize its loss function, as long as the loss function involves an output-gap target that is consistent with the inflation target; hence inflation targeting is argued to differ from purely discretionary policy *only* in the fact that policy is made on the basis of a loss function with this property. King (1997) obtains a formal result to this effect, but in the context of a model where the aggregate-supply relation is assumed to be of the “New Classical” form assumed by Kydland and Prescott (1977). The result is in fact dependent on extremely special properties of that form of aggregate-supply relation; see Woodford (2003, chap. 7) for further discussion.

¹⁷Here I mean more precisely a Markov equilibrium, in which equilibrium outcomes depend only on state variables that affect either the policymaker’s objective or the set of possible outcomes from the current period onward, as is common in the literature on the suboptimality of discretionary policy.




Figure 2: Impulse responses to a purely transitory cost-push shock, under discretionary policy (dashed lines) and under an optimal commitment (solid lines).

the central bank. The difference is illustrated in Figure 2, where the responses to a purely transitory positive cost-push shock in period zero are shown in the case of the two alternative policies.¹⁸ Because the effect of the shock on the location of the inflation/output tradeoff is not expected to last beyond period zero, under discretionary policy (shown by the dashed line), the central bank will choose policy that brings about zero inflation and a zero output gap from period 1 onward. (This minimizes the bank's objective from period 1 onward, and in the absence of any prior commitment, it is the policy that an optimizing policy committee would choose.) In period zero, instead, it is necessary to accept either a positive inflation rate, a negative output gap, or both; the optimal policy accepts some of both so as not to have losses that are too great of either kind. The equilibrium response under an optimal policy commitment (shown by the solid line) is instead more complex: instead of allowing

¹⁸This figure reproduces Figure 7.3 from Woodford (2003), where the calculations are explained.

the output gap to return to zero immediately following the dissipation of the shock, it remains negative, returning to zero only gradually. The continued tight policy results in disinflation following the dissipation of the shock, so that the price level returns asymptotically to the level it would have been at had the cost-push shock never occurred. The expectation of subsequent disinflation helps to offset the shift in the inflation-output tradeoff in period zero due to the shock; as a consequence, both less inflation and a less negative output gap are possible during the period of the cost-push shock.

Suppose, now, that a central bank were to choose a forward path for policy each period, given the economy's state at that time, under the assumption that it can choose any forward path represents a possible rational-expectations equilibrium, though it does not commit itself to subsequently act in accordance with the paths chosen at earlier dates. When the cost-push shock occurs, the policy committee should choose the forward path that results in responses for inflation and output shown by the solid lines in Figure 2, since these responses minimize the loss function, among those paths which, if anticipated, are consistent with the structural relation (2.1). Suppose that in period zero the committee chooses the nominal interest rate associated with period zero in this equilibrium; and that the same decision process is repeated in period one. Since the effects of the shock have dissipated by period one, the policy committee now judges (correctly) that it is possible to achieve an equilibrium from then on in which $\pi_t = x_t = 0$ in all periods. This is obviously the best possible equilibrium from the standpoint of the criterion (2.2), and so *this* is the forward path for the policy that will be chosen in period one and later.¹⁹ However, because the path chosen in period one does not continue the path chosen in period zero, even if nothing unexpected occurs between the two meetings, the procedure is not intertemporally consistent, and the projections on the basis of which policy is chosen in period zero will not represent good forecasts of the economy's future evolution.

Moreover, if the private sector correctly understands how policy will be conducted under this procedure, the equilibrium outcome will not only fail to be optimal; it

¹⁹This assumes that no further shocks occur in subsequent periods. But even if there are subsequent shocks, since they are distributed independently of the shock that occurred in period zero, they are equally likely to result in departures from the zero-inflation steady state in one direction as in the other; so the *mean* expected outcome from the bank's policy choices in period one and later is one in which $\pi_t = x_t = 0$ for all $t \geq 1$.

will be even worse than the equilibrium associated with discretionary policy. In response to a shock of the kind just considered, the private sector will expect no effect on inflation or the output gap in any future periods. Under these expectations (which are not the ones associated with the optimal equilibrium), the best possible outcome is the one shown by the dashed lines in Figure 2. (In that equilibrium, the central bank optimally exploits the tradeoff that exists in period zero, given that it recognizes that expected inflation does not shift in response to the shock.) But if the central bank instead supposes that it can choose any forward path for policy that it likes (instead of recognizing that its future behavior will predictably bring about $\pi_t = x_t = 0$ in every period), and chooses the path corresponding to the solid lines, then it will not choose the nominal interest rate in period zero associated with the equilibrium under discretionary policy; instead, it will raise the nominal interest rate much less than it would under discretionary policy.²⁰ As a result, the point on the inflation/output tradeoff that will actually be realized will involve more inflation (though a less negative output gap) than the optimal point; and as a consequence the shock will increase the loss function (2.2) to a greater extent than it would under discretion.

2.3 Using a Sequence of Target Criteria to Determine the Path of Interest Rates

An alternative approach, that avoids this problem, is to determine the forward path of policy as that path which results in projections that satisfy a *sequence* of quantitative target criteria, one for each of a sequence of future horizons. It is true that a *single* criterion — say, involving the projections for 8 quarters in the future only — can determine only a single dimension of policy, and thus can only determine an entire path if one is constrained to consider only a one-parameter family of possible paths (such as constant-interest-rate paths). But a sequence of similar criteria can independently determine the stance of policy at each of a sequence of dates, and

²⁰This follows both from the fact that it does not believe that the real interest rate needs to rise as much as in the discretionary equilibrium — it does not seek to contract current output as much, and it expects the return of output to the natural rate to be more gradual — and from the fact that it expects a given nominal interest rate to correspond to a higher real interest rate than it would in the discretionary equilibrium, because it has chosen a forward path for policy that involves disinflation in the following period.




Figure 3: The inflation and output-gap projections of the Norges Bank superimposed. Source: Norges Bank, *Inflation Report 3/2006*, Chart 1.13.

thus can determine the entire forward path of policy. Moreover, if the sequence of target criteria for different horizons are of the *same form* — *i.e.*, if the target criterion is independent of the horizon — then the forecast-targeting procedure will be intertemporally consistent.

As a practical example, consider the targeting procedure used by the Norges Bank beginning in 2005. Through the end of 2006,²¹ each issue of the Bank’s *Inflation Report* included a box labeled “Criteria for an appropriate future interest rate path.”²² According to the first of the criteria listed, “inflation should be stabilized near the target [*i.e.*, 2.5 percent per year] within a reasonable time horizon, normally 1-3 years,” and moving *toward* that target rate even sooner. This criterion alone would sound

²¹Beginning with the 2007/1 issue of the Bank’s *Monetary Policy Report*, the description of the criterion used to select the forward path of policy has been less explicit; presumably the criterion used at present is more complex than the simple one discussed in the text. (Note that the name of the report was also changed with the 2007/1 issue.

²²The criteria are discussed in more detail in Qvigstad (2006).

similar to the Bank of England target criterion mentioned above, except with greater vagueness about the horizon. But there is then a second criterion: that “the inflation gap [the amount by which actual inflation exceeds the medium-run target rate] and the output gap should be in reasonable proportion to each other until they close,” and in particular that the two gaps “should normally not be positive or negative at the same time.”²³ The second criterion indicates not only what the projections should look like in some medium run, but also what the *transition path* should look like: there should be an inverse relation between the inflation gap and the output gap, with the two gaps shrinking to zero together. In order to allow visual inspection of the extent to which the projections satisfy this criterion, the Norges Bank presents a figure in which the projections for its preferred measures of inflation²⁴ and of the output gap are superimposed, as shown in Figure 3. A criterion of this kind can determine the entire forward path for policy. And with such a criterion, it is not necessary to specify independently the rate at which the inflation rate should be projected to approach the target rate; the appropriate rate is exactly the rate that allows the output gap to remain in the desired proportion to the inflation gap. (Under such a criterion, the inflation gap *will* be projected to close eventually, as long as it is not possible to have a non-zero permanent output gap.)

The criterion just cited applies to each of a sequence of future horizons. It can be represented formally as the requirement that

$$(\pi_{t+h,t} - \pi^*) + \phi x_{t+h,t} = 0 \quad (2.3)$$

for each horizon $h \geq \underline{h}$, for some coefficient $\phi > 0$. Here $y_{t+h,t}$ denotes the projected value at date t of some variable y , at a horizon h periods in the future; $\underline{h} \geq 0$ indicates the shortest horizon at which it is still possible for policy to affect the projections, and I suppose that a sequence of criteria (2.3) for $h \geq \underline{h}$ suffices to uniquely determine the acceptable projections (including an implied forward path for policy).²⁵ Suppose also that the central bank’s forecast of its own forecasts in future decision cycles satisfy the principle that one should expect one’s future forecasts to be the same as one’s

²³See, for example, the box on pp. 10 of Norges Bank, *Inflation Report* 3/2006.

²⁴The inflation measure emphasized by the Norges Bank in its targeting procedure, CPI-ATE, is a consumer price index that is adjusted for tax changes and energy prices.

²⁵See Svensson and Woodford (2005) for algebraic analysis of a specific example. In the case considered there, prices and spending decisions are each predetermined a period in advance, so that $\underline{h} = 1$.

current forecasts (except, of course, as a result of developments that cannot currently be foreseen), so that

$$[y_{t+h_2,t+h_1}]_t = y_{t+h_2,t}$$

for any horizons $h_2 \geq h_1 \geq 0$. Then if at date t a forward path for policy is chosen that leads to projections satisfying (2.3) for each $h \geq \underline{h}$, it should also be projected at that time that at any later date $t + h_1$, the continuation of that same path should lead to projections satisfying a corresponding sequential criterion, since at date t the bank should project that

$$[(\pi_{t+h_2,t+h_1} - \pi^*) + \phi x_{t+h_2,t+h_1}]_t = 0$$

for all horizons $h_2 \geq h_1 + \underline{h}$. This makes the procedure of choosing a forward path for policy on such a basis intertemporally consistent.

I believe that this kind of targeting procedure provides the most appealing solution to the problem of intertemporal consistency. The way in which the target criterion is used to determine an appropriate forward path for policy is essentially the same as under the procedure used by the Bank of England prior to 2004, as discussed above, except without either the arbitrary emphasis on a single horizon or the arbitrary restriction to forward paths for policy involving a constant interest rate. Since forecast-targeting central banks already publish charts showing their projections for each of a sequence of future horizons, rather than only presenting a set of numerical forecasts for a specific horizon, discussion of a target criterion that should apply at each horizon is fairly straightforward within the existing frameworks for deliberation and communication about policy, as the example of the Norges Bank shows. Moreover, both the Norges Bank and the Riksbank now discuss quite explicitly the fact that their targeting procedures involve the choice of a forward path for policy, and publish “fan charts” for the paths of short-term nominal interest rates implicit in their projections. Hence this aspect of the recommended approach is entirely possible within the context of existing procedures as well.

The main “practical” obstacle to such an approach, I believe, is that it would require a central bank to adopt a highly structured approach to policy deliberations, and to describe that approach rather explicitly to the public. It would require the bank to be more open about its own view of the likely future evolution of policy than even some forecast-targeting central banks have been willing to be thus far. And it would require the bank to discuss explicitly the nature of the trade-offs that

determine an acceptable transition path following a disturbance, and not merely the nature of the “medium-run” targets that one hopes to reach some years in the future. The latter goal will almost surely require that a bank be explicit about the ways in which projections for variables other than a single measure of inflation are relevant to judgments about the appropriate stance of policy. Even though all inflation-targeting central banks appear to care about projections for real variables as well as inflation,²⁶ most have been quite cautious about discussing the way in which this may factor into their policy decisions. But this would have to be different if forecast targeting were to be adopted by an institution with a “dual mandate” like the U.S. Federal Reserve (at least, in the absence of a substantial modification of the Federal Reserve Act by Congress). And even in the case of other central banks, I believe that it would greatly enhance the transparency of policymaking — and ultimately, the credibility of their commitments to inflation control, by making clearer the extent to which temporary failures to return inflation immediately to its medium-run target level are nonetheless consistent with a systematic approach to policy that does indeed guarantee stability of inflation over the medium run.

3 Which Target Criterion is Appropriate?

The desirability of a forecast-targeting regime depends crucially, of course, on what kind of target criterion is used as the focus of policy deliberations. I have mentioned above one desideratum, which is that the target criterion suffice to determine a unique (or at least a unique non-explosive) rational-expectations equilibrium, and hence to uniquely determine the forward path of policy that should be anticipated in each decision cycle. But of course one wants, not merely to determine an equilibrium, but to determine one with desirable properties. I have suggested above that the relative flexibility of a forecast-targeting approach, by comparison with previous conceptions of policy rules, should in principle allow a greater degree of success at using monetary policy to stabilize the economy in the face of exogenous disturbances, while also maintaining a reasonable degree of confidence in the future value of the currency.

²⁶For example, the summary justification of current policy in the introduction to each issue of the Bank of England’s *Inflation Report* always begins by discussing the projection for real GDP growth before turning to the inflation projection, despite the apparent concern with the inflation projection alone in the simple target criterion discussed above.

But what sort of target criterion should be adopted, in order to take best advantage of this flexibility?

It might be feared that a commitment to pursue a specific quantitative target criterion at each future horizon, rather than allowing the policy committee to choose the forward path associated with the most appealing projections at each point in time without any such straightjacket, means choosing a rule that will frequently involve undesirable responses to disturbances, simply for the sake of making policy predictable and transparent. But a commitment to a specific target criterion need not imply any departure from an optimal stabilization policy at all; it is possible to choose a target criterion such that fulfillment of the criterion implements precisely the optimal equilibrium responses to shocks. Moreover, it is possible for this to be true, not only in the case of some particular kind of stochastic disturbance, but for any member of a broad class of possible disturbances; this result makes the commitment to a target criterion a particularly useful way of specifying a central bank's policy commitment.

3.1 Example: Optimal Policy in the Basic New Keynesian Model Revisited

Giannoni and Woodford (2002) show that it is possible quite generally to find a target criterion that implements optimal responses to any of a broad class of disturbances, and Giannoni and Woodford (2005) illustrate the application of this method to a variety of models. Here I shall review only a very simple example, that illustrates why a specification of policy in terms of a target criterion is particularly robust. Consider again the problem of optimal stabilization policy treated by Clarida *et al.* (1999). The problem of finding the state-contingent evolution of inflation and of the output gap that minimize (2.2) subject to the constraint that (2.1) hold each period can be solved by minimizing a Lagrangian. This results in first-order conditions of the form

$$\pi_t + \varphi_t - \varphi_{t-1} = 0 \tag{3.1}$$

$$\lambda(x_t - x^*) - \kappa\varphi_t = 0 \tag{3.2}$$

each period, where φ_t is a Lagrange multiplier associated with the constraint (2.1) in period t .²⁷

Manipulation of these constraints so as to eliminate the Lagrange multipliers then implies that in the optimal equilibrium, a linear relation of the form

$$\pi_t + \frac{\lambda}{\kappa} (x_t - x_{t-1}) = 0 \quad (3.3)$$

should hold each period between the paths of inflation and of the output gap. Note that this can be interpreted as a target criterion. In fact, it is similar in form to the criterion of the Norges Bank discussed above (and represented by equation (2.3)): there is a constant long-run inflation target (here, zero), short-run departures from which are justified if they are associated with large enough projected short-run output gaps of the right sign. (And the more negative the projected output gap in a given period, the greater the positive inflation gap that should be accepted in the projection for that period. The crucial difference is that in (3.3) it is the *change* in the output gap over a given period that is relevant for determining the inflation that should be accepted over that period, rather than the absolute *level* of the output gap as in Figure 3.

This is not only a relation that holds each period in the optimal equilibrium, but one that suffices to define the optimal equilibrium dynamics. For the system of equations consisting of (2.1) together with (3.3), if both assumed to hold each period from some date onward, determine a unique non-explosive REE.²⁸ Moreover, in that equilibrium, there exists a system of Lagrange multipliers $\{\varphi_t\}$ such that the first-order conditions (3.1)–(3.2) are satisfied each period; hence the equilibrium responses to shocks are indeed optimal, if it can be ensured that the target criterion (3.3) is satisfied each period.

More precisely, the equilibrium in which (3.3) is satisfied in each period $t \geq 0$ is the one in which the stochastic evolution of inflation and output minimize the

²⁷These conditions must be satisfied in each period $t \geq 0$, in each possible state of the world at each date, if policy is being optimally chosen for periods $t \geq 0$. In the event that there is no additional pre-commitment of the form indicated by (3.4) below — *i.e.*, in the conventional “Ramsey policy” problem — there is no constraint prior to the constraint (2.1) for $t = 0$, and so $\varphi_{-1} = 0$ in the first-order condition (3.1) for $t = 0$. If there is instead an initial pre-commitment of the kind discussed below, this multiplier can be non-zero as well, and indicates the value of relaxing that constraint.

²⁸For further discussion of this and other assertions made in this section, see Woodford (2003, chap. 7).

objective (2.2), among all processes consistent with the AS relation (2.1) each period, with an initial inflation rate

$$\pi_0 = \bar{\pi}(x_{-1}; \xi_0). \quad (3.4)$$

Here ξ_t is the (exogenous) state of the world in period t , and $\bar{\pi}(x_{t-1}; \xi_t)$ is a function that indicates the equilibrium inflation rate each period in the unique non-explosive equilibrium determined by the equations (2.1) and (3.3). Thus conformity to the target criterion each period implements a policy that is optimal subject to an initial pre-commitment (a constraint on what policy can bring about, but that restricts outcomes *only* in the initial period); this initial pre-commitment is to behave in the same way in period zero as one will in fact wish to commit oneself to behave in the subsequent periods.²⁹ In the absence of any such pre-commitment, an optimal commitment regarding policy from time zero onward can be characterized by fulfillment of a different target criterion in period zero only, but fulfillment of (3.3) from period 1 onward. Thus in either case, an optimal policy commitment can be implemented through a commitment to a decision procedure under which (at least in every period after the initial one) policy is chosen so as to ensure that the target criterion (3.3) is satisfied, and is furthermore expected to be satisfied in all future periods.³⁰

Of course, one could bring about the same equilibrium through a commitment to use policy to bring about the inflation rate $\bar{\pi}(x_{t-1}; \xi_t)$ each period. This would also be an example of an intertemporally consistent targeting procedure, that results in a unique REE, which again involves responses to shocks that are optimal in the sense just discussed. However, this way of formulating an optimal target criterion is not robust to changes in the assumed character of the disturbances that make up the vector ξ_t . If one writes $\bar{\pi}$ as a function of x_{t-1} and some number of lags of u_t , for example, then the optimal dependence of the function on the various lags of u_t will

²⁹Policy that is optimal subject to a self-consistent initial pre-commitment of this kind is called “optimal from a timeless perspective” in Giannoni and Woodford (2002); see also Woodford (2003, chap. 7).

³⁰It is of course the fact that Ramsey optimal policy involves a different target criterion in the initial period that accounts for the fact that choosing a forward path for policy in each decision cycle so as to minimize the loss function (2.2) does not lead to the same policy choices as the choice of a forward path that is expected to satisfy (3.3) at each horizon — so that the former procedure is not intertemporally consistent. Under sequential re-optimization, the current period is always treated as “the initial period”, resulting in a failure to choose to continue the path of policy chosen in the previous decision cycle.

depend on the autoregressive dynamics of the process $\{u_t\}$. The target criterion (3.3) instead implements the optimal responses to disturbances *regardless* of the assumed properties of the disturbances; note that we did not need to make any assumption about the process $\{u_t\}$ in order to derive the criterion (3.3), though we did need to make a very specific assumption (namely, that u_t is an i.i.d. random variable) in order to derive the dynamic response coefficients shown in Figure 2. The target criterion (3.3) also has the advantage that it can be explained to the public without having to make any reference to particular types of structural disturbances; instead, the target criterion only involves “target variables,” *i.e.*, variables that are relevant to the central bank’s stabilization objectives.

3.2 Robustness of the Optimal Target Criterion

The above example relates, of course, to a single very simple model — one that is obviously too simple to provide the basis for policy deliberations in an actual central bank. Nonetheless, the possibility of implementing optimal responses to disturbances using a target criterion is not a special feature of this example. In fact, Giannoni and Woodford (2002) show that an optimal target criterion can be derived for any stabilization policy problem belonging to a fairly general family of linear-quadratic problems.³¹ In any such case, the target criterion is a linear relationship among the various target variables (arguments of the quadratic loss function) such that there is a unique non-explosive REE consistent with satisfaction of the target criterion at all times; and this REE always implements the responses to disturbances associated with an optimal *ex ante* commitment, regardless of the statistical properties of the disturbances (as long as these are purely additive terms in the model’s linear structural relations).

Hence it is possible for any central bank that is willing to base its policy on a particular quantitative model of the transmission mechanism to derive an appropriate target criterion. (The best target criterion, of course, might be different for different economies.) Nor is there any need for the bank to reconsider its commitment to the target criterion each time a shock of some unusual type occurs. Since novel types

³¹One can furthermore find a linear target criterion such that the equilibrium determined by the criterion agrees with the optimal equilibrium responses, to a linear approximation, in the case of a fairly general class of sufficiently differentiable policy problems.

of disturbances will occur relatively frequently (in fact, no two disturbances are ever of exactly the same type), while structural change in the economy (or a substantial change in a central bank’s understanding of the structure of the economy) occurs much less often, a target criterion that is equally suitable regardless of the nature of the disturbance is one to which a central bank should be able to provisionally commit itself, even if from time to time it would still be appropriate to reconsider the target criterion in the light of new knowledge or new circumstances.³²

But while a fully optimal target criterion of this type exists in principle, it is likely to be too complicated to be useful as the basis for communication with the public, in the case of any model with even a modest claim to empirical realism.³³ The fully optimal criterion is also likely to depend on the precise parameterization of a number of fine details of one’s model of the economy — parameter values (or even model features) about which there will in fact be considerable uncertainty. As a practical matter, what one should really be interested in is a fairly simple target criterion that provides a reasonable approximation to optimal policy, even if it is not fully optimal; and one that continues to lead to an equilibrium that is not too far from optimal, even if the true model of the economy is slightly different from the one for which the rule was designed.

What we need, then, is a study of economic performance under alternative quantitative target criteria in the spirit of the Taylor program — one should study alternative target criteria in precisely the same way that alternative interest-rate feedback rules have been studied over the past fifteen years. This means that one should not only address the level of welfare associated with alternative parameterizations in the context of a single model, but also investigate the degree to which a given quantitative target criterion may be robust, in the sense of achieving relatively good performance

³²In principle, frequent reconsideration of the target criterion is possible, without this eliminating the advantages of policy commitment, if the central bank always chooses a new policy rule that is “optimal from a timeless perspective,” rather than one that seeks to exploit already existing expectations at the time of the policy change. But in practice, public understanding of the systematic character of policy, and verifiability of the bank’s adherence to its alleged commitments, will be better served if reconsiderations of the target criterion do not occur every year.

³³The complexity of the optimal target criterion derived by Giannoni and Woodford (2005) for a small empirical model of the U.S. economy — a four-equation model that is still extremely simple relative to the models used for policy analysis in most central banks — illustrates this point. While it is possible to give a precise statement of the optimal target criterion, I shall not attempt it here!

across a range of model specifications. In fact, the methods required to do this are exactly the same ones that have already been employed in the study of interest-rate rules; one simply uses the equation that states the target criterion (*i.e.*, an equation such as (3.3) as the “missing equation” that specifies monetary policy, rather than adding an interest-rate equation (such as the Taylor rule) to one’s stochastic simulation model.

Relatively little work of this kind has already been undertaken, so I cannot survey its conclusions. But there are good reasons to suppose that relatively robust target criteria can be found. A number of results that can be obtained from the study of optimal target criteria in relatively simple examples suggest that the description of desirable policy in terms of a target criterion is likely to be more robust than other ways of characterizing the same policy.

First of all, there is the fact that it is possible to state a target criterion that describes optimal policy (and that suffices to determine the optimal equilibrium — the description of optimal policy is not robust at the cost of being incomplete), regardless of the statistical properties of the disturbances to which the central bank may need to respond. This does not mean that the central bank does not need to determine the specific properties of the disturbances that have just occurred; determining what adjustment of its interest-rate operating target is required in order to *conform to* the target criterion will require an assessment of the nature of current shocks. But the target criterion that is to be achieved can be stated more generally. A description of the optimal policy at the level of an appropriate interest-rate reaction function is not similarly robust.

Second, certain kinds of structural change will not require any change in the appropriate target criterion. In the policy problem discussed above, a change in the equations of the structural model that determine how interest rates affect spending will not have any effect on the above derivation if it does not change either the aggregate-supply relation (2.1) or the appropriate stabilization objective (2.2); and so (3.3) would continue to be the optimal target criterion. Again, this does not mean that structural change of this sort can be ignored by the central bank; it will have to be taken account of if the bank’s interest-rate decisions are to have the desired effect. But it would remain possible to organize policy deliberations around the goal of satisfying an unchanged quantitative target criterion. (Again, a description of optimal policy at the level of an interest-rate reaction function would not be similarly

invariant.)

But perhaps more surprisingly, even types of structural change that *do* change the quantitative specification of the aggregate-supply relation do not always require a change in the quantitative specification of the target criterion, in order for policy to continue to be optimal. This is because the appropriate loss function (in terms of the paths of prices and output), in order for the central bank's stabilization objective to correspond to household welfare maximization, *also* changes in the case of a change in those aspects of economic structure that determine available tradeoff between output and inflation, indicated by the aggregate-supply relation. In fact, there is a whole range of changes that would imply a different quantitative specification of aggregate supply, but that do *not* change the optimal target criterion, because the changes in the aggregate-supply relation and in the welfare-based loss function precisely cancel one another.

The aggregate-supply relation (2.1) assumed by Clarida *et al.* (1999) represents a log-linear approximation to the relation implied by a particular model of price-setting, namely a model in which price changes are staggered in the way proposed by Calvo (1983). The loss function (2.2) is an appropriate objective in the case of this model, as it can be shown to be (inversely related to) a quadratic approximation to the expected utility of the representative household in the DSGE model. In the case that we assume an efficient steady-state level of output,³⁴ the value of λ corresponding to welfare maximization is given by

$$\lambda = \frac{\kappa}{\theta} > 0, \quad (3.5)$$

where θ is the elasticity of substitution among alternative differentiated goods (and hence the elasticity of demand faced by each monopolistically competitive producer). Hence the optimal target criterion (3.3) can alternatively be written

$$\pi_t + \theta^{-1}(x_t - x_{t-1}) = 0. \quad (3.6)$$

³⁴This assumption simplifies the analysis of the welfare-based loss function, because it is not necessary to consider the effects of stabilization policy on the average level of output in order to obtain a welfare measure that is accurate to second order, as discussed in Woodford (2003, chap. 6). All of the results reviewed in this section are derived under that simplification. However, a welfare-based quadratic loss function can also be derived without this assumption, using the method illustrated in Benigno and Woodford (2005). For example, in the Calvo model the welfare-based loss function is of the form (2.2) even without the assumption of an efficient steady-state, but the correct value of λ is in general more complicated than the simple formula (3.5).

Note that the quantitative specification of the optimal target criterion turns out to depend on only *one* underlying structural parameter of the model: the value of θ . (This parameter is important because it determines the extent to which the dispersion of prices that results from variability of the inflation rate causes inefficient non-uniformity in the composition of the national product.) A variety of other parameters — parameters describing household preferences, the production technology, and the frequency of adjustment of prices — all affect the slope of the aggregate-supply curve, but have *no* effect on the quantitative specification of the optimal target criterion (3.6). Thus the optimal target criterion is not only more robust than a representation of optimal policy in terms of a reaction function, it is also a more robust specification of the appropriate goal of policy than the loss function (2.2). This is a further reason why it is more desirable for a central bank to commit itself to a specific target criterion than for it to publicly specify its loss function.

Other, even more substantial modifications of the model are also possible, without affecting the optimality of the target criterion (3.6). The basic new Keynesian model considered by Clarida *et al.* assumes that the utility of the representative household is additively separable over time, but empirical models often find that the data are better explained by preferences that incorporate habit persistence, and Giannoni and Woodford (2005) consider the consequences for optimal monetary policy of this generalization. The most important consequence of habit persistence, of course, is a change in the relation between real interest rates and the dynamics of expenditure; but this in itself is no reason for a change in the optimal target criterion, though it certainly affects the interest-rate path required by the target criterion. However, habit persistence also changes the relation between current output (or expenditure) and the marginal utility of income; as a consequence, the way that output enters both the aggregate-supply relation and the welfare-based loss function must change. Yet the two changes cancel one another in their effects on the optimal target criterion, for Giannoni and Woodford show that (3.6) continues to characterize optimal policy, regardless of the degree of habit persistence. Once again, the optimal target criterion is more robust to changes in the economic structure than is the welfare-based quadratic loss function.

All of the examples discussed thus far assume price stickiness of the form introduced by Calvo (1983). But the “new Keynesian Phillips curve” (2.1) is often criticized for its implication that inflation determination is purely forward-looking, and

empirical models often incorporate some degree of structural inflation inertia. One way to motivate such inertia is to assume (contrary to the simple Calvo specification) that the probability of revision of a given firm’s price is an increasing (rather than constant) function of the time since the price was set; this results in a generalization of (2.1) in which lagged inflation rates appear as determinants of current inflation, along with expectations of future inflation, as in common (“hybrid NKPC”) empirical specifications. But as shown by Sheedy (2005), the appropriate welfare-based loss function changes as well, since the relation between the dynamics of aggregate inflation and the degree of cross-sectional dispersion of prices at a given point in time changes; and optimal policy continues to be characterized by the target criterion (3.6), regardless of the values of the parameters that determine the hazard rate for price changes as a function of price duration.

An alternative interpretation of apparent inflation inertia is provided by the “sticky information” model of Mankiw and Reis (2002). In this model, each firm charges at all times the price that it expects at that time to maximize its current profits, but firms do not continuously update their information about market conditions; inflation inertia thus results from inertia in the average *estimate* of the current inflation rate. This kind of model leads to a very different kind of aggregate-supply relation, but again Ball *et al.* (2005) show that optimal policy is characterized by a target criterion of the form (3.6).³⁵ More precisely, optimal policy requires that the quantity on the left-hand side of (3.6) must be a *deterministic sequence*, that can be perfectly forecasted arbitrarily far in advance; it need not be zero (or even constant), since in this model there are no real effects of changes in the general price level that are expected by all firms, no matter how long it has been since they last updated their information. But a rule of the form (3.6) would be *one* example of an optimal policy rule in this model, and since this particular rule is also optimal in the case of sticky-price models, it is therefore a more robust policy choice than would result from some other deterministic sequence.

I do not mean to suggest that (3.6) should be an optimal target criterion, regardless of the structure of the economy; it is easy to give examples where the optimal target criterion is not of this form (see, e.g., Giammoni and Woodford, 2005). All of the examples just reviewed have important elements in common — for example, they

³⁵Note that (3.6) implies that $\log P_t + \theta^{-1}x_t$ is equal to a constant, which is how Ball *et al.* describe the optimal policy.

are all models in which only prices are sticky (not wages), and they are all models that are perfectly symmetric across sectors (different sectors have identical production technologies, the same kind and degree of nominal rigidities or informational frictions, etc.) — and I do not mean to suggest that these idealizations are adequate for practical monetary policy analysis. But these examples do indicate that even the *optimal* target criterion need not depend on all aspects of the model specification; this gives one reason to hope that it may be possible to find target criteria that are fairly robust across (even if not fully independent of) aspects of the model specification about which policymakers remain uncertain. In particular, they suggest that a description of desirable policy in terms of a target criterion may be more robust to model uncertainty than other descriptions that would be equally valid in the absence of such uncertainty.

4 Implementing the Target Criterion

Thus far, I have spoken about the choice of a target criterion as if one can take it for granted that there is a policy under which the target criterion will in fact be satisfied each period, and that selecting a target criterion means committing oneself to a decision procedure that implements that policy. (I have pointed out that one must ask whether there *exists* any equilibrium consistent with a given target criterion; but if there does, I have not treated the achievement of that equilibrium as problematic.) But this is hardly obvious; and a discussion of the practicality of forecast targeting as a monetary policy strategy would be incomplete without any consideration of how one should determine the actual actions to be taken by the central bank under such a strategy. Certainly one cannot compare a forecast targeting strategy to a “lower-level” specification of a policy rule, such as the Taylor rule, without also describing what forecast targeting means for the way in which the policy instrument should be adjusted over time.

I have mentioned above that, in the case that the target criterion is one that implies a determinate rational-expectations equilibrium, it therefore implies a particular state-contingent evolution of the policy rate. But that does not mean that there are no further issues to address with regard to the policy actions that should be chosen in order to implement the target criterion. For even when a particular rule for setting the central bank’s operating target for the policy rate is known to

be consistent with the rational-expectations equilibrium that the target criterion is intended to bring about, that need not mean that this equilibrium is the *only* one that is consistent with that rule. Thus we again face a question about whether a particular rule determines a unique rational-expectations equilibrium, distinct from the one discussed above — and one that the theorem of Giannoni and Woodford (2002) does not address. Even if there is only one stationary REE under which the target criterion is satisfied at all times, it does not follow that a reaction function that is *consistent* with fulfillment of the target criterion — because it is consistent with that REE — is not also consistent with other REE; because it may also be consistent with equilibria in which the target criterion is *not* fulfilled.

Even if this is not a problem, one has only verified that the rule for setting interest rates should have the desired consequences under the assumptions that (i) the economy will indeed reach a rational-expectations equilibrium, and (ii) the central bank's structural model is correct. It remains important to ask how *robust* the policy is to possible departures from these ideal cases. These are issues that have been extensively addressed in the literature on simple Taylor-type rules, but they need also to be addressed for interest-rate reaction functions that are intended to implement a target criterion. Here again, the issues are ones that can (and should) be addressed at each of *two* levels, in the case of a forecast-targeting strategy: one should ask, first, whether fulfillment of the target criterion will still lead to a fairly desirable outcome, in the case of a certain departure from rational expectations (for example); but then one should also ask whether the interest-rate rule intended to implement the target criterion still comes close to fulfilling the target criterion, in the case of the departure from rational expectations. A similar question about robustness can be posed at each of the two levels, for any of a variety of types of model mis-specification with which one might be concerned.

Once one concedes this, it might be asked whether there is any advantage to formulations of policy commitments in terms of a target criterion at all. If one still needs to address the questions of determinacy, learnability of the rational-expectations equilibrium, and robustness to model uncertainty all over again at the level of the specification of an interest-rate reaction function, even when the target criterion (if it could be assumed to be necessarily satisfied) has been shown to be desirable on all of those grounds, why not simply search for desirable interest-rate rules without the discussion of desirable target criteria at all? The answer is that I believe that there

will turn out to be advantages of a rule for the adjustment of interest rates that *explicitly uses a target criterion* in determining what the interest rate should be — that is, of a forecast-targeting procedure — over other interest-rate rules that might *happen to support* an equilibrium in which the target criterion is satisfied. (In the latter case, one would be justified in simply saying that the rule is good because it implements a *desirable equilibrium*, without any need to discuss the target criterion.) And the advantages that I believe such rules are likely to have are precisely the ability to exclude undesired alternative equilibria, the ability to facilitate the convergence of out-of-equilibrium expectations to the REE, and greater robustness of alternative model specifications. However, these advantages depend on a particular understanding of what it means to use the target criterion to determine the appropriate interest-rate operating target; in particular, it is necessary for the central bank to *monitor private-sector expectations*, and to respond to them in a particular way.

4.1 The Example Once Again

This can be illustrated using the analysis in Preston (2004) of alternative approaches to implementation of the optimal target criterion (3.3) in the basic new-Keynesian model. Preston shows that under arbitrary (subjective) private-sector expectations, the aggregate-supply relation takes the form

$$\pi_t = \kappa x_t + \hat{E}_t \sum_{T=t+1}^{\infty} (\alpha\beta)^{T-t} \left[\kappa x_T + \frac{1-\alpha}{\alpha} \pi_T + u_T \right] + u_t, \quad (4.1)$$

where the notation is the same as in (2.1) above, but now the operator $\hat{E}_t[\cdot]$ indicates the conditional expectation under the (common) probability beliefs of private agents at date t , which may or may not correspond to the true conditional expectation (*i.e.*, “rational expectations”). This inflation equation results from log-linearization of the optimal price-setting rule for firms in a model with monopolistic competition and Calvo staggering of price changes; $0 < \alpha < 1$ is the probability that any firm’s price will not be revised in a given period. Note that under the assumption of rational expectations — so that subjective forecasts of future inflation are consistent with inflation’s being determined by this equation in all future periods — this equation implies the relation (2.1) used by Clarida *et al.* (1999); but more generally, subjective forecasts of both inflation and output for more than one period in the future affect the location of the current inflation/output tradeoff.

In order to determine the interest rate required in order to implement the target criterion, one must adjoin an equilibrium relation between interest rates and private expenditure. Preston shows that under arbitrary private-sector expectations, the “intertemporal IS relation” of the standard new-Keynesian model takes the form

$$x_t = -\sigma(i_t - r_t^n) + \hat{E}_t \sum_{T=t+1}^{\infty} \beta^{T-t} [(\beta^{-1} - 1)x_T - \sigma(i_T - \beta^{-1}\pi_T - r_T^n)], \quad (4.2)$$

where r_t^n indicates exogenous variation in the Wicksellian “natural rate of interest”.³⁶ This is essentially a permanent-income model of expenditure by the representative household, augmented to take account of the effects of the anticipated path of real interest rates on the optimal timing of expenditure; again, under the assumption of rational expectations, it would reduce to the simpler relation (involving only expectations of variables one period in the future) used by Clarida *et al.* (1999).

We may now compare two of the approaches to implementation of the target criterion discussed by Preston. One approach would be to compute the REE implied by satisfaction of the target criterion (discussed in section 3), and determine the state-contingent path of nominal interest rates in that equilibrium. In the case of the target criterion (3.3) and structural relations (4.1)–(4.2), this leads to a solution of the form

$$i_t = \bar{i}(x_{t-1}; s_t) \quad (4.3)$$

for the interest rate, where s_t is the economy’s (exogenous) state in period t (in general, a vector), including a full description of all information in period t about the current and future values of the disturbances $\{u_T, r_T^n\}$.³⁷ One approach to implementation would be to compute this solution, and then use (4.3) to determine the central bank’s operating target i_t each period, after observing the current state s_t . Note that

³⁶This is a function of exogenous disturbances, such as the rate of time preference and the state of technology, that indicates what the short-run real interest rate would be each period in an equilibrium where a zero output gap is maintained at all times. For further discussion of the natural rate of interest, and of the microfoundations of this model, see Woodford (2003, chap. 4).

³⁷Note that in the three relations (3.3), (4.1), and (4.2) that we solve to determine the REE evolution of the endogenous variables, the only lagged endogenous variable that appears is x_{t-1} , appearing in (3.3). This ensures that if there is a determinate equilibrium, it must be described by a function of the form (4.3). Note also that the precise form of the function $\bar{i}(x; s)$ depends on the specification of the joint stochastic process for the exogenous disturbances, and not simply on the coefficients of the three relations just listed.

under this approach, the target criterion would not actually be used in the routine conduct of policy; it would only be used (once) to compute the function (4.3). Moreover, the target criterion is not really needed even in that calculation; equation (4.3) is simply a (partial) description of the optimal equilibrium, and so it is really only computation of the optimal equilibrium that is required under this approach. Thus this approach would not really deserve to be called a “forecast targeting” procedure, even if it is a policy consistent with satisfaction of the target criterion.

But this approach to implementation of optimal policy has undesirable properties. For many possible values of the parameters, rational-expectations equilibrium is *indeterminate* under commitment to an interest-rate rule of the form (4.3).³⁸ The reason, essentially, is that such a policy would not conform to the “Taylor Principle” that interest rates should be raised in response to inflation above the target rate, and sharply enough to make the *real* interest rate an increasing function of inflation. While the rule (4.3) is *consistent* with the desired equilibrium, policy under this rule is too “passive” to ensure that this is the only possible REE.³⁹ Moreover, a “passive” rule of this kind is also ineffective in steering expectations closer to the REE in the case that people initially have other beliefs, but update their forecasting rules on the basis of the data that they observe.⁴⁰ Preston (2004) analyzes learning dynamics in this model, under the assumption that people forecast future inflation, output and interest rates as linear functions of the exogenous state and the lagged output gap, and that the coefficients of their forecasting rules are determined by OLS regressions, constantly updated as additional data are observed. He shows that often a rule of the form (4.3) implies *unstable* learning dynamics, that will diverge from rational expectations with probability 1 even if people happen to start with forecasting rules close to rational expectations. There is also considerable reason to doubt the robustness of this kind of rule under model mis-specification, even if one is not troubled by the above results (perhaps because one presumes that the economy will necessarily settle

³⁸See the analysis in Svensson and Woodford (2005). Svensson and Woodford discuss a more elaborate version of the model, in which both pricing and spending decisions are predetermined, but the conditions on parameter values that control the determinacy of equilibrium are identical in the case discussed here.

³⁹On the role of the “Taylor Principle” for the determinacy of equilibrium under certain simple families of interest-rate rules, see e.g., Clarida *et al.* (2000), Woodford (2003, chap. 4).

⁴⁰The essential intuition that a passive interest-rate policy will fail to stabilize inflation expectations if these extrapolate observed past inflation is contained in an argument of Friedman (1968).

into the “minimum-state-variable” REE associated with a given policy rule, through a process that need not be explained). For example, the coefficients on the state variables in (4.3) will depend on the degree of persistence of the disturbances, so that the rule will not continue to achieve the optimal equilibrium if this has been wrongly specified, even though the target criterion (3.3) *would* continue to characterize the optimal equilibrium in this case.

Another approach is to use the structural relations (presumed to be correctly understood by the central bank) to determine the interest-rate operating target required in order to ensure that realized inflation and output satisfy the target criterion (3.3), *given the expectations that people actually have*, whether these correspond to the expectations that would be “rational” under this policy or not. This approach requires that the central bank monitor the expectations of market participants, in order to adjust interest rates in a way that *offsets* the effects of any departures of people’s actual expectations from rational expectations. In order to specify concretely what this involves, Preston assumes that the private sector’s period- t forecasts ($\hat{E}_t\pi_T$, etc.) are determined prior to the central bank’s decision about its period- t operating target i_t ,⁴¹ though they may depend on the exogenous state s_t , and that they are observed by the central bank (along with the exogenous state) prior to its choice of i_t . Given this information, the central bank can use (4.1) together with (3.3) to determine the values of π_t and x_t that it should seek to bring about, in order for the target criterion to be satisfied in the current period. The implied current output-gap target is given by

$$\begin{aligned}\hat{x}_t \equiv & \frac{\lambda}{\kappa^2 + \lambda}x_{t-1} - \frac{\kappa}{\kappa^2 + \lambda}u_t \\ & - \frac{\kappa}{\kappa^2 + \lambda}\hat{E}_t \sum_{T=t+1}^{\infty} (\alpha\beta)^{T-t} [\kappa x_T + (1 - \alpha/\alpha)\pi_T + u_T].\end{aligned}\quad (4.4)$$

It can then use (4.2) to determine the interest rate i_t that should achieve this target

⁴¹Note that this need not be the case, though it is a property of the particular kind of learning process that Preston specifies. While the assumption is a restrictive one, it does not rule out the possibility of rational expectations, since in the REE, the central bank’s choice of i_t will be a perfectly predictable function of variables that have already been determined at the time that people form their period- t forecasts.

for the output gap. This is given by

$$i_t = r_t^n - \sigma^{-1} \hat{x}_t + \hat{E}_t \sum_{T=t+1}^{\infty} \beta^{T-t} [\sigma^{-1} (\beta^{-1} - 1) x_T - (i_T - \beta^{-1} \pi_T - r_T^n)], \quad (4.5)$$

where $\hat{x}_t$ is defined in (4.4). Equation (4.5) represents an alternative interest-rate rule that, like (4.3), is consistent with the REE determined by the target criterion. But it describes a policy in which the target criterion plays a crucial role in the conduct of policy; indeed, while it is *possible* to describe the policy using the interest-rate rule, it is much more compactly described by the prescription to use one's model of the economy, together with one's estimates of current fundamentals and the current state of private-sector expectations, to determine the interest rate needed in order to ensure that the target criterion (3.3) is satisfied.

While (4.3) might seem to represent a simpler approach to policy — both one that is simpler to execute and one that makes outside verification of the central bank's compliance with the rule simpler — the rule (4.5) is a more reliable way of bringing about the desired outcome. First of all, there is necessarily a unique stationary REE consistent with rule (4.5).⁴² For regardless of what the private sector may be expecting, (4.5) implies that $x_t = \hat{x}_t$ and hence that π_t and x_t satisfy (3.3). As this will be true each period, if the private sector has rational expectations, they must be expecting (and their actions will generate) the unique REE consistent with the target criterion. Moreover, Preston (2004) shows that if instead private-sector expectations evolve through least-squares learning, in the case of rule (4.5) the learning dynamics converge to the REE.⁴³ This is because this rule involves responses to variations in private-sector expectations that neutralize the effects of the errant expectations on aggregate outcomes, thus short-circuiting the process through which out-of-equilibrium expectations might become self-fulfilling. Finally, while this has been less thoroughly

⁴²Thus there is a value to a commitment to respond to potential departures from rational expectations, even if one is confident that the economy must end up in a REE; for responding in this way ensures that there are not *other* REE besides the one in which the target criterion is always satisfied.

⁴³More precisely, Preston shows that the REE is “E-stable” in the sense of Evans and Honkapohja (2001). This implies convergence to the REE in expectation, and convergence with probability 1 in the case that certain bounds are imposed on how far people's forecasting rules can depart from the REE rules.

studied, it is likely that performance under rule (4.5) is more robust to variations in the model specification. For example, rule (4.5) implements the optimal dynamic responses to disturbances regardless of the degree of persistence of the disturbances, unlike the rule (4.3). (Note that no assumptions about the disturbance processes have been used in deriving (4.5). It would be desirable to study the robustness of a rule like this to other forms of model misspecification, as has often been done in the case of simple Taylor rules. As discussed in section 3, it is known that for a variety of alternative models of price adjustment, (3.3) continues to be an optimal target criterion. Hence the prescription “act so as to ensure that (3.3) holds” is a robust policy, within this set of alternative models, even if the more specific formula (4.5), which depends on the structural relations (4.1)–(4.2), will not be equally robust.

4.2 Comparison with the Taylor Rule

To what extent is an analysis of the kind proposed here — the search for a robust approach to implementation of a robust target criterion — likely to lead to recommendations for policy that differ from those obtained from the Taylor research program, namely, the search for a simple interest-rate reaction function with robust properties? I suspect that many of the most important conclusions from Taylor’s research will also be supported by an inquiry of this kind. That is to say, a desirable forecast-targeting framework will involve a conduct of policy that conforms to many of the broad guidelines for sound policy emphasized by Taylor.⁴⁴

This is illustrated by the presentation above of an optimal targeting procedure in the context of the basic new-Keynesian model. The interest-rate reaction function (4.5) that implements the optimal target criterion shares a number of features with the simple rule advocated by Taylor (1993). In fact, it could be called a sort of “forward-looking Taylor rule,” insofar as the prescribed short-term interest rate i_t is an increasing function of expected future inflation, expected future output gaps, and the current and expected future levels of the natural rate of interest. One respect in which (4.5) differs from more familiar examples of forward-looking Taylor rules is that it also involves forecasts of future *interest rates*. Because it is the expected

⁴⁴The observed conduct of actual forecast-targeting central banks has often been found to look fairly similar to the policy that would follow from a Taylor rule. For example, Nelson (2003) finds that UK policy has been much more similar to the kind prescribed by Taylor (1993) after the introduction of inflation-forecast targeting in 1992 than was true in earlier periods.

path of interest rates that determines aggregate demand, rather than the current short rate alone, higher future short rates can substitute for a higher short rate now; so the required current short rate is a decreasing function of expected future short rates, given expectations regarding inflation, the output gap, and the natural rate of interest.

Moreover, this reaction function conforms to the “Taylor principle”, insofar as an average of current and expected future nominal interest rates is specified to be an increasing function of an average of expected future inflation rates (for given expectations regarding the output gap and the natural rate), with a response coefficient that is greater than one. We can alternatively write (4.5) in the form

$$\hat{E}_t \bar{i}_t = \phi_\pi \hat{E}_t \bar{i}_{t+1} + \dots,$$

where $\bar{i}_t$ is a weighted average of the interest rates i_T in periods $T \geq t$, $\bar{\pi}_t$ is a weighted average of the inflation rates π_T in periods $T \geq t$, and the omitted terms involve (both actual and expected values of) the output gap and the exogenous disturbances. When the rule is expressed in this way,

$$\phi_\pi = 1 + \frac{\kappa \sigma^{-1}}{\kappa^2 + \lambda} \frac{\beta(1 - \beta)(1 - \alpha)}{1 - \alpha\beta} > 1,$$

in accordance with the Taylor Principle. This aspect of the reaction function is probably not unrelated to its success in eliminating instability due to self-fulfilling expectations.

Nonetheless, while this policy is similar to the policy prescribed by the Taylor rule in certain important respects, it does not follow that a central bank would therefore be best advised to simply try to follow the Taylor rule, or to explain its decisions to the public in those terms. There is an alternative, fairly simple description of the policy prescription represented by (4.5) that is also more precise: it is a prescription to set interest rates as required in order for inflation and output to satisfy the target criterion (3.3).

Moreover, this alternative prescription is more sophisticated in important respects. If the basic new-Keynesian model were correct, and the economy were to follow the REE associated with the policy, it would lead to optimal dynamic responses to all disturbances. Even more generally, some of the ways in which it differs from the simple Taylor rule are likely to be desirable. For example, (4.5) prescribes that the

interest rate should be an increasing function of expectations of the future output gap *relative to its recent past level*, rather than the absolute level of the gap as in the simple Taylor rule. A rule of this kind will be less vulnerable to the possibility that a persistent error in estimating the natural rate of output can lead to persistent inflation above the central bank's long-run target.⁴⁵ Because it is much easier to imagine a substantial, persistent error in a central bank's estimate of the *level* of the natural rate than in its estimate of the growth rate, a rule that depends only on the projected *change* in the output gap is likely to be much more robust to possible errors of this kind.⁴⁶

The targeting procedure also differs from the simple Taylor rule in the way that it requires the central bank to respond to measures of private-sector expectations. Under the simple Taylor rule, private-sector expectations can be ignored; one trusts that a rule with desirable stabilization properties will stabilize expectations as well. The targeting procedure instead directs the central bank to respond to variations in private-sector expectations, so as to offset their projected effects on the particular combination of inflation and output growth that constitutes the target criterion. To the extent that sufficient information about expectations exists to allow such a response, such a response should tend to stabilize the economy regardless of the vagaries of expectations — and should therefore stabilize expectations themselves more reliably as well.

5 Conclusion

I have argued that forecast-targeting procedures provide a useful way of implementing a rule-based approach to the conduct of monetary policy, and that the normative study of monetary policy rules would do well to pay more attention to proposed rules

⁴⁵Orphanides (2003) warns that this may be a serious risk under a Taylor rule that involves the absolute level of the output gap. He shows that US policy in the 1970s was fairly consistent with a forward-looking Taylor rule of the same kind as also describes US policy in the Greenspan period, if one assumes that the FOMC responded to estimates of the output gap that were available at the time; in this interpretation, the policy led to a decade of inflation as a result of a failure to recognize the slowdown in trend productivity growth.

⁴⁶Orphanides (2003) shows that real-time errors in estimates of the quarter-to-quarter *change* in the output gap were quite modest during the 1970s, if one assumes that Fed's estimates by the 1990s were roughly correct.

of this type. This would involve analysis of two distinct, though obviously complementary issues: on the one hand, the desirability of satisfying a given target criterion (specified as a relation among endogenous “target variables”), on the assumption that this can be arranged; and on the other hand, the efficacy of alternative procedures for adjusting the instrument of policy as ways of ensuring that the target criterion should be satisfied. In the case of each of these issues, it is important not only to analyze the question for a given model specification that is assumed to be precisely correct, and under the assumption that the state variables referred to in the model are precisely observable by the central bank, but also to consider the robustness of a given prescription to possible errors in either the model specification or in the evaluation of current conditions.

Forecast-targeting procedures are especially desirable classes of policy rules for two reasons. On the one hand, this way of specifying a rule of conduct seems especially likely to provide a useful basis for explaining the systematic character of policy to the public, while at the same time allowing policy to respond in a flexible way to unexpected developments, including changes in the central bank’s understanding of the character of disturbances and of the structure of the economy. And at the same time, it is a form of rule-based policy that is fairly close to the kind of procedures already being adopted by a number of central banks; hence it is a kind of policy rule that, if found to have desirable properties in theory, it is particularly plausible for central banks to consider. Further study of forecast-targeting procedures would thus serve the laudable goal, that has animated much of the research program of John Taylor, of providing a “translation” of the conclusions of normative economic theory into terms that can guide actual policy decisions.

While more work on rules of this class is needed, some provisional conclusions are already possible that may allow improvement upon the current practices of forecast-targeting central banks. The projections that are central to forecast targeting should be based on intertemporally consistent assumptions about future monetary policy. The most practical way to do this is to seek to determine an endogenous path for future policy that ensures the satisfaction of a particular target criterion at each of an entire sequence of future horizons, rather than some single horizon. Furthermore, the target criterion must apply to relatively near horizons and not only to outcomes several years in the future. It is important to specify not merely the inflation rate that should be expected in the “medium run” (though it is crucial that a central bank

be clear about that particular implication of the targeting procedure), but also the nature of an acceptable transition path by which an economy currently away from the medium-run target inflation rate is expected to approach it. This in turn will inevitably require that a central bank be more explicit than most forecast-targeting central banks have thus far been about the way in which projections for variables other than a single overall measure of inflation are taken into account in the policy decision. Finally, a robust approach to the implementation of the target criterion requires that the central bank monitor private-sector expectations (through the variety of direct and indirect sources of information about expectations that is available to it), with a view to correcting the effects on the macroeconomy of departures from the beliefs that the bank itself would regard as correct.

No central bank currently conforms very fully to all of these precepts, and so I believe that there is ample room for further productive dialogue between students of the theory of monetary policy and its practitioners. But forecast targeting as it has developed over the past fifteen years already represents an important step toward more systematic and more transparent policy. This is highly desirable, for reasons that monetary theory has made clear, and the further refinement of forecast targeting procedures represents the best hope at present for putting the insights provided by that theory into practice.

References

- [1] Ball, Laurence, N. Gregory Mankiw, and Ricardo Reis, “Monetary Policy for Inattentive Economies,” *Journal of Monetary Economics* 52: 703-725 (2005).
- [2] Benigno, Pierpaolo, and Michael Woodford, “Inflation Stabilization and Welfare: The Case of a Distorted Steady State,” *Journal of the European Economics Association* 3: 1185-1236 (2005).
- [3] Calvo, Guillermo, “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics* 12: 383-98 (1983).
- [4] Clarida, Richard, Jordi Gali and Mark Gertler, “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature* 37: 1661-1707 (1999).
- [5] Clarida, Richard, Jordi Gali and Mark Gertler, “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics* 115: 147-180 (2000).
- [6] Evans, George W., and Seppo Honkapohja, *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press, 2001.
- [7] Giannoni, Marc P., and Michael Woodford, “Optimal Interest-Rate Rules: I. General Theory,” NBER working paper no. 9419, December 2002.
- [8] Giannoni, Marc P., and Michael Woodford, “Optimal Inflation Targeting Rules,” in Ben S. Bernanke and Michael Woodford, eds., *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2005.
- [9] Goodhart, Charles A.E., “Monetary Transmission Lags and the Formulation of the Policy Decision on Interest Rates,” Federal Reserve Bank of St. Louis *Review*, July/August 2001, pp. 165-181.
- [10] Goodhart, Charles A.E., “The Interest Rate Conditioning Assumption,” Financial Markets Group discussion paper no. 547, London School of Economics, October 2005.

- [11] Honkapohja, Seppo, and Kaushik Mitra, “Performance of Inflation Targeting Based on Constant Interest Rate Projections,” *Journal of Economic Dynamics and Control* 29: 1867-1892 (2005).
- [12] Jansson, Per, and Anders Vredin, “Forecast-Based Monetary Policy: The Case of Sweden,” *International Finance* 6: 349-380 (2003).
- [13] King, Mervyn A., “Changes in UK Monetary Policy: Rules and Discretion in Practice,” *Journal of Monetary Economics* 39: 81-97 (1997).
- [14] Kydland, Finn E., and Edward C. Prescott, “Rules Rather than Discretion: The Inconsistency of Optimal Plans,” *Journal of Political Economy* 85: 473-491 (1977).
- [15] Leitemo, Kai, “Targeting Inflation by Constant-Interest-Rate Forecasts,” *Journal of Money, Credit and Banking* 35: 609-26 (2003).
- [16] Mankiw, N. Gregory, and Ricardo Reis, “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics* 117: 1295-1328 (2002).
- [17] Nelson, Edward, “UK Monetary Policy 1972-97: A Guide Using Taylor Rules,” in Paul Mizen, ed., *Central Banking, Monetary Theory and Practice: Essays in Honour of Charles Goodhart*, Volume 1. Cheltenham, UK: Edward Elgar, 2003.
- [18] Orphanides, Athanasios, “The Quest for Prosperity without Inflation,” *Journal of Monetary Economics* 50: 633-663 (2003).
- [19] Preston, Bruce, “Adaptive Learning and the Use of Forecasts in Monetary Policy,” unpublished, Princeton University, January 2004.
- [20] Qvigstad, Jan F., “When Does an Interest Rate Path ‘Look Good’? Criteria for an Appropriate Future Interest Rate Path: A Practitioner’s Approach,” Norges Bank staff memo 2006/5, April 2006.
- [21] Rosenberg, Irma, “Riksbank to Introduce Own Path for the Repo Rate,” a speech at Danske Bank, Stockholm, January 17, 2007. [Text available at www.riksbank.com.]

- [22] Sheedy, Kevin D., “Resistance to Persistence: Optimal Monetary Policy Commitment,” unpublished, Cambridge University, November 2005.
- [23] Svensson, Lars E.O., “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets,” *European Economic Review* 41: 1111-1146 (1997).
- [24] Svensson, Lars E.O., “What is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules,” *Journal of Economic Literature* 41: 426-477 (2003).
- [25] Svensson, Lars E.O., “Optimal Inflation Targeting: Further Developments of Inflation Targeting,” in Frederic Mishkin and Klaus Schmidt-Hebbel, eds., *Monetary Policy under Inflation Targeting*, Santiago: Central Bank of Chile, 2007.
- [26] Svensson, Lars E.O., and Michael Woodford, “Implementing Optimal Policy through Inflation-Forecast Targeting,” in Ben S. Bernanke and Michael Woodford, eds., *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2005.
- [27] Taylor, John B., “Estimation and Control of a Macroeconometric Model with Rational Expectations,” *Econometrica* 47: 1267-1286 (1979).
- [28] Taylor, John B., “Discretion Versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214 (1993a).
- [29] Taylor, John B., “The Use of the New Macroeconometrics for Policy Formulation,” *American Economic Review: Papers and Proceedings* 83(2): 300-305 (1993b).
- [30] Taylor, John B., “Applying Academic Research on Monetary Policy Rules: An Exercise in Translational Economics,” Harry G. Johnson Lecture, *The Manchester School* 66 (supp.): 1-16 (1998).
- [31] Taylor, John B., “Introduction,” in John B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999a.
- [32] Taylor, John B., “A Historical Analysis of Monetary Policy Rules,” in John B. Taylor, ed., *Monetary Policy Rules*, Chicago: Univ. of Chicago Press, 1999b.

- [33] Taylor, John B., “How the Rational Expectations Revolution Has Changed Macroeconomic Policy Research,” in Jacques Dréze, ed., *Advances in Macroeconomics*, Palgrave, 2000.
- [34] Vickers, John, “Inflation Targeting in Practice: The U.K. Experience,” Bank of England *Quarterly Bulletin*, November 1998.
- [35] Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.
- [36] Woodford, Michael, “Central-Bank Communication and Policy Effectiveness,” in *The Greenspan Era: Lessons for the Future*, Kansas City: Federal Reserve Bank of Kansas City, 2005.