

Inflation Targeting and Optimal Monetary Policy

Michael Woodford
Princeton University

October 8, 2003

PRELIMINARY

Prepared for the Annual Economic Policy Conference,
Federal Reserve Bank of St. Louis, October 16-17, 2003.

Since the early 1990s, an increasing number of countries have adopted explicit inflation targets as the defining principle that should guide the conduct of monetary policy. This development is often credited with having brought about substantial reductions in both the level and variability of inflation in the inflation-targeting countries, and is sometimes argued to have improved the stability of the real economy as well.¹

Inflation-forecast targeting, as a systematic decision procedure for the conduct of monetary policy, was developed at central banks like the Reserve Bank of New Zealand, the Bank of Canada, the Bank of England, and the Bank of Sweden on a trial-and-error basis, with little guidance from the academic literature on monetary policy rules. But the growing popularity of inflation targeting has more recently led to an active literature that seeks to assess the desirability of such an approach from the standpoint of theoretical monetary economics. This literature finds that an optimal policy regime — one that could have been designed on *a priori* grounds to achieve the highest possible degree of social welfare — might well be implemented through procedures that share important features of the inflation-forecast targeting that is currently practiced at central banks like those just mentioned. At the same time, the normative literature finds that one ought, in principle, to be able to do better through appropriate refinement of the practices developed at these banks.

Here I survey some of the most important conclusions of this literature. I shall begin by reviewing some of the respects in which inflation targeting as currently implemented represents a step toward what the theory of optimal monetary policy would recommend. In the final section of the paper, I then summarize some of the more important respects in which an optimal policy regime would go beyond current practice. Finally, as a concrete illustration of some of the general remarks that have been made about the form of an optimal policy rule, in an appendix I briefly discuss the quantitative character of optimal policy in the context of the small econometric model for the U.S. presented in Giannoni and Woodford (2003).

¹For surveys of early experiences with inflation targeting, see Leiderman and Svensson (1995) and Bernanke *et al.* (1999). King (2003) offers an optimistic assessment of the improvements made in the conduct of monetary policy in the U.K. under inflation targeting. For a more skeptical review of the lessons that can be gleaned from experience to date, see Ball and Sheridan (2003).

1 Advantages of an Explicit Target for Monetary Policy

Discussions of the desirability of inflation targeting for one country or another are often at cross purposes because of differing implicit assumptions about precisely what inflation targeting would mean. It is thus perhaps useful to be clear from the outset about what I regard to be the defining features of the approach to the conduct of policy with which I am concerned. Probably the most critical feature is the existence of a publicly announced, quantitative target which the central bank is committed to pursue, the pursuit of which structures both policy deliberations within the central bank and communications with the public. As should become clear from the discussion below, it is more important in my view that there should be an explicit *target* for policy than that it should be (in any strict sense) an *inflation* target. In my view, the most distinctive and most important achievement of the inflation-targeting central banks has not been the reorientation of the goals of monetary policy toward a stronger emphasis on controlling inflation — this has occurred, but it is has been a worldwide trend over the past two decades, neither limited to nor even necessarily most associated with the innovators in inflation targeting, and has hardly required a fundamental change in the traditional concerns of central bankers — but rather the development of an approach to the conduct of policy that focuses on a clearly defined target, that assigns an important role to quantitative projections of the economy's future evolution in policy decisions, and that is committed to a high degree of transparency as to the goals of policy, the decisions that are made, and the principles that guide those decisions.

It is useful to begin by discussing why it is desirable for a central bank to commit itself to an explicit target as the goal of its policy. The proposal that banks should do so runs contrary to a common instinct of central bankers, according to which it is wise to say as little as possible in advance about what one may do in the future. Because central banking is a complex task, the argument goes, any explicit target or policy rule would prove to be a straightjacket, preventing the full exercise of the judgment of central bankers on behalf of society when unanticipated circumstances arise, as they invariably do. Furthermore, even if a formula could be developed that would adequately describe what a good central banker should do, announcing it publicly would only invite second-guessing by the public and politicians of policy

decisions that are best left in the hands of professionals. The best approach, then, is to delegate the task to the best possible people, grant them full discretion, and require as little public comment as possible on the way in which they practice their arcane art.

But while it is true that central banking is complex, reasoning of this kind misses a fundamental point about the kind of problem that a central bank is called upon to solve. Central banking is not like steering an oil tanker, or even guiding a spacecraft, which follows a trajectory that depends on constantly changing factors, but that does *not* depend on the vehicle's own expectations about where it is heading. Because the key decisionmakers in an economy are forward-looking, central banks affect the economy as much through their influence on *expectations* as through any direct, mechanical effects of central bank trading in the market for overnight cash. As a consequence, there is good reason for a central bank to commit itself to a systematic approach to policy, that not only provides an explicit framework for decisionmaking within the bank, but that is also used to explain the bank's decisions to the public.

1.1 Central Banking as Management of Expectations

One important advantage of commitment to an appropriately chosen *policy rule* is that it facilitates public understanding of policy. It is important for the public to understand the central bank's actions, to the greatest extent possible, not only for reasons of democratic legitimacy — though this is an excellent reason itself, given that central bankers are granted substantial autonomy in the execution of their task — but also in order for monetary policy to be most effective. For not only do expectations about policy matter, but, at least under current conditions, very little *else* matters. Few central banks of major industrial nations still make much use of credit controls or other attempts to directly regulate the flow of funds through financial markets and institutions. Increases in the sophistication of the financial system have made it more difficult for such controls to be effective, and in any event the goal of improvement of the efficiency of the sectoral allocation of resources stressed above would hardly be served by such controls, which (if successful) inevitably create inefficient distortions in the relative cost of funds to different parts of the economy.

Instead, banks restrict themselves to interventions that seek to control the overnight interest rate in an interbank market for central-bank balances (for example, the fed-

eral funds rate in the U.S.). But the current level of overnight interest rates *as such* is of negligible importance for economic decisionmaking; if a change in the overnight rate were thought to imply only a change in the cost of overnight borrowing for that one night, then even a large change (say, a full percentage point increase) would make little difference to anyone's spending decisions. The effectiveness of changes in central-bank targets for overnight rates in affecting spending decisions (and hence ultimately pricing and employment decisions) is wholly dependent upon the impact of such actions upon other financial-market prices, such as longer-term interest rates, equity prices and exchange rates. These are plausibly linked, through arbitrage relations, to the short-term interest rates most directly affected by central-bank actions; but it is the expected future path of short-term rates over coming months and even years that should matter for the determination of these other asset prices, rather than the current level of short-term rates by itself.

Thus the ability of central banks to influence expenditure, and hence pricing, decisions is critically dependent upon their ability to influence market expectations regarding the *future path* of overnight interest rates, and not merely their current level. Better information on the part of market participants about central-bank actions and intentions should increase the degree to which central-bank policy decisions can actually affect these expectations, and so increase the effectiveness of monetary stabilization policy. Insofar as the significance of current developments for future policy are clear to the private sector, markets can to a large extent "do the central bank's work for it," in that the actual changes in overnight rates required to achieve the desired changes in incentives can be much more modest when expected future rates move as well.²

The importance of being able to influence expectations about future policy through

²There is evidence that this is already happening, as a result both of greater sophistication on the part of financial markets and greater transparency on the part of central banks, the two developing in a sort of symbiosis with one another. Blinder *et al.* (2001, p. 8) argue that in the period from early 1996 through the middle of 1999, one could observe the U.S. bond market moving in response to macroeconomic developments that helped to stabilize the economy, despite relatively little change in the level of the federal funds rate, and suggest that this reflected an improvement in the bond market's ability to forecast Fed actions before they occur. Statistical evidence of increased forecastability of Fed policy by the markets is provided by Lange *et al.* (2001), who show that the ability of Treasury bill yields to predict changes in the federal funds rate some months in advance has increased since the late 1980s.

means other than the announcement of a new operating target for the overnight interest rate becomes especially clear when the zero lower bound on nominal interest rates prevents further interest-rate cuts, in an environment where aggregate nominal expenditure is nonetheless too low. This is the situation that Japan has faced for more than four years now, and recently there has been considerable discussion in the U.S. as well as to whether the Fed is not nearly “out of ammunition” with which to fight a possible threat of deflation. The key to avoiding deflation and economic contraction under such circumstances, as Eggertsson and Woodford (2003) show, is to be able to credibly commit to looser monetary policy in the future.³ This requires explicit discussion of the way in which policy will be conducted in the future; furthermore, Eggertsson and Woodford show that the kind of commitment that is needed can be best expressed in terms of a commitment to a form of *price-level target*, which the central bank is committed to eventually hitting, even if the zero bound requires the target to be undershot for some period of time. While one might alternatively imagine a direct commitment regarding the length of time for which interest rates will remain low, the optimal continuation time will depend on how real conditions in the economy develop (that cannot yet be perfectly foreseen); it is thus easier to explain the kind of commitment that is actually appropriate by explaining the target that will have to be met in order for the zero interest-rate policy to be abandoned.

The existence of an explicit target for policy has similar advantages under more ordinary circumstances as well. An obvious consequence of the importance of managing expectations is that a transparent central-bank decisionmaking process is highly desirable. This has come to be widely accepted by central bankers over the past decade. (See Blinder *et al.*, 2001, for a detailed and authoritative discussion.) But it is sometimes supposed that the most crucial issues are ones such as the frequency of press releases or the promptness and detail with which the minutes of policy deliberations are published. Instead, from the perspective suggested here, what is important

³The basic point about the importance of commitment regarding future policy was first made by Krugman (1998); Eggertsson and Woodford present a fully dynamic analysis and characterize the optimal policy commitment in an optimizing model with staggered price-setting. The conclusion obtained by Auerbach and Obstfeld (2003) is not fundamentally different; in their analysis, it is actually only the expected money supply at the time of exit from the liquidity trap that matters, so that open market operations while in the trap are effective only to the extent that they are understood as implying a commitment to a higher money supply *after* the zero bound ceases to bind.

is not so much that the central bank's deliberations themselves be public, as that the bank give clear signals about what the public should expect it to do in the future. The public needs to have as clear as possible an understanding of the *rule* that the central bank follows in deciding what it does. Inevitably, the best way to communicate about this will be by offering the public an explanation of the decisions that have already been made; the bank itself would probably not be able to describe how it might act in all conceivable circumstances, most of which will never arise.

The *Inflation Reports* of the leading inflation-targeting central banks provide good practical examples of communication with the public about the central bank's policy commitments. These reports do not pretend to give a blow-by-blow account of the deliberations by which the central bank reached the position that it has determined to announce; but they do explain the *analysis* that justifies the position that has been reached. This analysis provides information about the bank's systematic approach to policy by illustrating its application to the concrete circumstances that have arisen since the last report; and it provides information about how conditions are likely to develop in the future through explicit discussion of the bank's own projections. Because the analysis is made public, it can be expected to shape future deliberations; the bank knows that it should be expected to explain why views expressed in the past are not later being followed. Thus a commitment to transparency of this sort helps to make policy more fully rule-based, as well as increasing the public's understanding of the rule.

It might be argued that it should be enough for a central bank to *follow* a systematic rule in its conduct of policy, without also needing to explain it to the public. If one assumes rational expectations on the part of the public, it would follow that the systematic pattern in the way that policy is conducted should be correctly inferred from the bank's observed behavior. Yet while it would be unwise to choose a policy the success of which depends on its *not* being understood by the public — which is the reason for choosing a policy rule that is associated with a desirable rational-expectations equilibrium — it is at the same time prudent not to rely too heavily on the assumption that the public will understand policy perfectly regardless of the efforts that are made to explain it. Insofar as explanation of the policy rule to the public does no harm under the assumption of rational expectations, but improves outcomes under the (more realistic) assumption that a correct understanding of the central bank's policy commitments does not occur automatically, then it is clearly

desirable for the central bank to explain the rule that it follows.

The advantages of a public target when the private sector must otherwise forecast future policy by extrapolating from experience are shown in a recent analysis by Orphanides and Williams (2003). In the Orphanides-Williams model, private agents forecast inflation using a linear regression model, the coefficients of which are constantly re-estimated using the most recent observations of inflation. The assumption of forecasting in this manner (on the basis of a finite time-window of historical observations) rather than a postulate of rational expectations worsens the tradeoff between inflation variability and output-gap variability that is available to the central bank. Allowing inflation variations in response to “cost-push” shocks for the sake of output-gap stabilization is more costly than it would be under rational expectations, because temporary inflation fluctuations in response to the shocks can be misinterpreted as indicating different inflation objectives on the part of the central bank. Orphanides and Williams then show that a credible commitment to a long-run inflation target — so that private agents do not need to estimate the long-run average rate of inflation, but only the dynamics of transitory departures from it — allows substantially better stabilization outcomes, though still not quite as good as if private agents were to *fully* understand the equilibrium dynamics implied by the central bank’s policy rule. This provides a nice example of theoretical support for the interpretation given by Mervyn King (2003) and others of practical experience with inflation targeting, which is that tighter anchoring of the public’s inflation expectations has made possible greater stability of *both* real activity and inflation.

1.2 Avoiding the Pitfalls of Discretionary Policy

There is also a further, somewhat subtler, reason why explicit commitment to a target or policy rule is desirable, given the forward-looking behavior of the people in the economy that one seeks to stabilize. It is not enough that a central bank have sound objectives (reflecting a correct analysis of social welfare), that it make policy in a systematic way, using a correct model of the economy and a staff that is well-trained in numerical optimization, and that all this be explained thoroughly to the public. A bank that approaches its problem as one of optimization under *discretion* — deciding afresh on the best action in each decision cycle, with no commitment regarding future actions except that they will be the ones that seem best in whatever circumstances

may arise — may obtain a substantially worse outcome, from the point of view of its own objectives, than one that commits itself to follow a properly chosen policy *rule*. As Kydland and Prescott (1977) first showed, this can occur even when the central bank has a correct quantitative model of the policy tradeoffs that it faces at each point in time, and the private sector has correct expectations about the way that policy will be conducted.

At first thought, discretionary optimization might seem exactly what one would want an enlightened central bank to do. All sorts of unexpected events constantly occur that affect the determination of inflation and real activity, and it is not hard to see that, in general, the optimal level of interest rates at any point in time should depend on precisely what has occurred. It is plainly easiest, as a practical matter, to arrange for such complex state-dependence of policy by having the instrument setting at a given point in time be determined only after the unexpected shocks have already been observed. Furthermore, it might seem that the dynamic programming approach to the solution of intertemporal optimization problems provides justification for an approach in which a planning problem is reduced to a series of independent choices at each of a succession of decision dates.

But standard dynamic programming methods are valid only for the optimal control of a system that evolves mechanically in response to the current action of the controller. The problem of monetary stabilization policy is of a different sort, in that the consequences of the central bank's actions depend not only upon the sequence of instrument settings up until the present time, but also upon private-sector expectations regarding future policy. In such a case, sequential (discretionary) optimization leads to a sub-optimal outcome because at each decision point, prior expectations are taken as *given*, rather than as something that can be affected by policy. Nonetheless, the predictable character of the central bank's decisions, taken from this point of view, do determine the (endogenous) expectations of the private sector at earlier dates, under the hypothesis of rational expectations; a commitment to behave differently, that is made credible to the private sector, could shape those expectations in a different way, and because expectations matter for the determination of the variables that the central bank cares about, in general outcomes can be improved through shrewd use of this opportunity.

The best-known example of a distortion created by discretionary optimization is the “inflation bias” analyzed by Kydland and Prescott (1977) and Barro and Gordon

(1983). In the presence of a short-run “Phillips curve” tradeoff between inflation and real activity (given inflation expectations), and a target level of real activity higher than the one associated with an optimal inflation rate (in the case of inflation expectations also consistent with that optimal rate), these authors showed that discretionary optimization leads to a rate of inflation that is inefficiently high on average, owing to neglect of the way that pursuit of such a policy raises inflation expectations (causing an adverse shift of the short-run Phillips curve). A commitment to an inflation target is one obvious way of eliminating the temptation of suboptimal behavior of this particular kind.

However, many central bankers would argue that they have absorbed the lesson of the Kydland-Prescott and Barro-Gordon models, and are able to avoid systematically higher inflation than is desirable, without any need for advance commitments regarding future policy. For example, they may view themselves as using their discretion to minimize a loss function that differs from the ones assumed by Kydland and Prescott or Barro and Gordon in a way that eliminates the high predicted average rate of inflation in the Markov equilibrium associated with discretionary policy.

In response to this, it is important to note that the distortions resulting from discretionary optimization go beyond simple bias in the average levels of inflation or other endogenous variables; this approach to the conduct of policy generally results in suboptimal *responses to shocks* as well. For example, various types of real disturbances can create temporary fluctuations in what Wicksell called the “natural rate of interest”, meaning that the level of nominal interest rates required to stabilize both inflation and the output gap varies over time (Woodford, 2003, chap. 4). However, the amplitude of the adjustment of short-term interest rates can be more moderate — and still have the desired size of effect on spending and hence on both output and inflation — if it is made more *persistent*, so that when interest rates are increased, they will not be expected to quickly return to their normal level, even if the real disturbance that originally justified the adjustment has dissipated. Because aggregate demand depends upon expected future short rates as well as current short rates, a more persistent increase of smaller amplitude can have an equal affect on spending. If one also cares about reducing the volatility of short-term interest rates, a more inertial interest-rate policy of this kind will be preferable; that is, the anticipation that the central bank will follow such a policy leads to a preferable rational-expectations equilibrium (Woodford, 1999a; 2003, chap. 7). But a central bank that optimizes

under discretion has no incentive to continue to maintain interest rates high once the initial shock has dissipated; at this point, prior demand has already responded to whatever interest-rate expectations were held then, and the bank has no reason to take into account any effect upon demand at an earlier date in setting its current interest-rate target.

This distortion in the dynamic response of interest-rate policy to disturbances cannot be cured by *any* adjustment of the way in which alternative possible future paths for the economy are ranked (assuming that the ranking depends only on the future paths of inflation and other welfare-relevant variables); instead, policy must be made *history-dependent*, *i.e.*, dependent upon past conditions even when they are no longer relevant to the determination of the current and future evolution of the variables that the bank cares about. In general, no *purely forward-looking* decision procedure — one that makes the bank’s action at each decision point a function solely of the set of possible paths for its target variables from that time onward — can bring about optimal equilibrium responses to disturbances. Discretionary optimization is an example of such a procedure, and it continues to be when the bank’s objective is modified, if the modified policy objective still involves only the future values of the welfare-relevant variables. A commitment to use policy to achieve a pre-specified *target*, instead, can solve this problem, if the target is defined in a way that incorporates the proper history-dependence.⁴

The advantages of an explicit target in solving this kind of problem are especially clear in the case of a binding zero lower bound on interest rates, as discussed by Eggertsson and Woodford (2003). In the case that the natural rate of interest is temporarily negative, the zero bound may prevent stabilization of inflation and the output gap at their desirable long-run average levels, as such an equilibrium would require a temporarily negative nominal interest rate. The key to preventing an undesirably sharp deflation and economic contraction is to convince people that the price level will eventually be raised, rather than being stabilized at whatever level it may fall to in the period during which the zero bound binds. A central bank that is expected to optimize under discretion will *not* be expected to subsequently undo the price decline that occurs during the “liquidity trap”, because the absolute level of prices is not welfare-relevant; it will therefore simply stabilize inflation and the out-

⁴The kind of modified inflation target that leads to optimal responses to the kinds of fluctuations in the natural rate of interest described above is derived in Giannoni and Woodford (2003, sec. 1.3).

put gap once this again becomes possible, accepting whatever level of prices happens to exist at that time. A commitment in advance to the achievement of a price level target — which target is not allowed to shift down even if actual prices undershoot it for many quarters in a row owing to the zero bound⁵ — will instead create expectations of the right sort. The farther prices fall while the economy is in the “trap,” the greater the expected future price increases will be; and this automatic increase in expected inflation will tend to prevent prices from falling very far, or demand from contracting very much, in the first place.

It is furthermore desirable, not simply that a central bank have a private intention of this sort, but that it be publicly committed to such a target. First, a public commitment is likely to make it easier for the central bank’s policy deliberations to remain focused on the right criterion — the one with the property that systematic conformity to it leads to an optimal equilibrium — rather than being tempted to “let bygones be bygones.” And second, the benefits associated with commitment to a history-dependent policy depend entirely on this aspect of policy being *anticipated* by the private sector; otherwise, it *would* be rational to “let bygones be bygones.” There is no point to a secret commitment to the future conduct of policy in accordance with a history-dependent rule, while the private sector continues to believe that the central bank will act in a purely forward-looking fashion; thus the target should be explained as clearly as possible to the public, and shown to be guiding the bank’s decisions.

1.3 Targeting Procedures as Policy Rules

It follows from the above discussion that there are important advantages to a central bank’s committing itself to conduct policy in accordance with a rule that can be explained in advance to the public. I turn now to the advantages of the particular type of rule that is followed by the inflation-targeting central banks. This is a rule under which the central bank’s commitment is defined by a *target* for certain variables at a certain distance in the future, together with a commitment to organize deliberations about policy actions around the question of whether the contemplated actions are consistent with the target.

⁵As Eggertsson and Woodford show, under an optimal policy the price-level target would actually shift *up* in response to the target misses during the period in which the zero bound is binding, and to a greater extent the greater the target misses and the longer they persist.

Much of the theoretical discussion of “rules versus discretion” since the seminal contribution of Kydland and Prescott (1977) has supposed that the conduct of policy in accordance with a “rule” would mean something rather different from this. On the one hand, an important branch of the literature on policy rules has emphasized the importance of limiting central bank discretion, in the sense of any scope for the exercise of judgment as to the nature of current conditions. A rule is then considered, by definition, to be a prescription of a fairly mechanical type, the dictates of which are unambiguous; it cannot pretend to allow optimal responses to all of the different types of shocks that an economy may face, and indeed it is often asserted that adherence to a rule means abandoning any concern for the stabilization of real variables. Inflation-forecast targeting as actually practiced is nowhere as rigid a framework as this; in particular, projections of the economy’s future evolution under alternative possible actions play a central role in policy deliberations, and these projections, even when disciplined by the use of a quantitative model, allow a rich range of information about current conditions to be taken into account in a way that could not be easily specified in advance by a computer program.

Alternatively, another branch of the literature identifies “commitment” with a once-and-for-all choice (at some initial date) of an optimal state-contingent plan for the central bank, which is implemented afterward by simply observing the state of the world each period and executing the instrument setting called for at that date and in that state. Under the conception of rule-based policy in this literature, the central bank may in principle pay attention to disturbances of all sorts; but there is no role in a specification of the policy commitment for any mention of *targets* for variables other than the instrument of policy itself (i.e., for anything besides a state-contingent *operating target* for the overnight interest rate).

The type of rule actually followed, at least in principle, by central banks like the Bank of England, is a policy rule of a different sort. Svensson (1999, 2003a) defines a *targeting rule* as a commitment to adjust the bank’s policy instrument as necessary so as to ensure that at each decision point the economy’s future evolution is still projected to satisfy a certain *target criterion*. For example, in the case of the Bank of England, the target criterion is that RPIX inflation should be projected to equal 2.5 percent per annum at a horizon 8 quarters in the future (Vickers, 1998; Goodhart, 2001). This is a “higher-level” specification of a policy rule than the kind generally considered in the two literatures just referred to, since it leaves unspecified

precisely what policy actions will be required in any given circumstance in order to conform to the rule. Implementation of the policy is only possible using a model of the economy (likely to be supplemented, in practice, by judgmental adjustments on the part of the monetary policy committee), with which projections of the economy's evolution under alternative hypothetical policy decisions can be constructed.

Commitment to a decision procedure of this kind has important advantages over both of the other two conceptions of a monetary policy rule.⁶ Achievement of the advantages of policy commitment — in particular, avoidance of the inflationary bias of discretionary policymaking — does not require one to give up on stabilization policy. Not only may policy adjust in response to disturbances, but it may adjust differently to each of an uncountable number of different types of disturbances, the nature of which need not even be specifiable in advance.

This is also true, in principle, under the conception of a policy rule as a commitment to a pre-specified state-contingent instrument path. But in practice one cannot imagine computing such an instrument rule in advance, and announcing one's commitment to it, unless one artificially assumes that the number of types of disturbances that can ever occur are extremely limited in number. This is a highly limiting assumption, given that in order to compute in advance the optimal dynamic response to a given shock, it is necessary to specify not simply which equations of one's model that it perturbs, but also to give a detailed quantitative specification of the dynamics of the shock — exactly how persistent it is expected to be, how far in advance it can be predicted, and so on. Shocks of a given type — for example, variations in government spending owing to the outbreak of war — that differ in the degree to which they are unanticipated or the length of time for which they are expected to last imply different optimal adjustments of the policy instrument, and so must be treated as *different shocks* in a complete specification of the optimal state-contingent instrument rule.

Of course, one can specify a quantitative model of the economy with a fairly small number of independent shocks (no more than the number of endogenous variables in the model), and estimate a joint stochastic process for those shocks using historical data. This method is often used, for example, in specifying the kind of stochastic model that is used for “stochastic simulation” exercises evaluating alternative simple

⁶For further discussion, see Svensson (1999, 2003a), Svensson and Woodford (2003), Giannoni and Woodford (2002), and Woodford (2003, chap. 7).

policy rules. And it may well be possible to calculate a complete specification of the optimal state-contingent instrument path for such a model. But a central bank would be highly unlikely to be willing to commit itself to follow a rule simply because it has been shown to be optimal in such an exercise.

For central bankers always have a great deal of highly specific information about the kind of disturbances that have just occurred, which are always somewhat different than those that have been faced at other times. Hence even if it is understood that “typically” disturbances to the level of military purchases have had a coefficient of serial correlation of 0.9 at the quarterly frequency, there will often be grounds to suppose that the particular conflict that is currently looming is likely to be either more persistent or less persistent than the “typical” one has been in the past. And it is unlikely that central bankers will be willing to commit themselves to stick rigidly to a rule that is believed to lead to outcomes that would be optimal in the case of “typical” disturbances, even in the case that they are aware of the economy’s instead being subjected to “atypical” disturbances. In order for an a proposed policy rule to be of practical interest, it must instead be believed that the rule is compatible with optimal (or at least fairly good) outcomes in the case of any of the extremely large number of possible types of disturbances that might be faced on different occasions. Yet if one were to try to write out the optimal state-contingent instrument path, allowing separate terms for each of the possible (finely-grained) types of disturbances that might actually be faced, such a description of optimal policy would be completely unwieldy.

Giannoni and Woodford (2002) instead show that if the central bank’s policy commitment is described in terms of a relation among endogenous variables that the bank is committed to bring about — rather than in terms of a mapping from exogenous states to the instrument setting — it is possible, in a large class of policy problems, to find a rule that is *robustly optimal*, in the sense that the same rule (with given numerical coefficients) continues to be optimal regardless of the assumed statistical properties of the (additive) disturbance terms in the model. Indeed, the target criterion that they derive characterizes optimal policy even if the disturbance terms in the model structural equations are actually composites of an extremely large (not necessarily finite) number of different types of real disturbances. This is possible because (as illustrated in the next section) the optimal target criterion is derived from certain first-order conditions that characterize an optimal evolution of the economy,

and these first-order conditions do not involve the additive disturbance terms in the structural relations.

A rule of this kind represents a policy commitment that a central bank could reasonably make, despite its awareness that it will constantly be receiving quite fine-grained information about current conditions. For a belief that the target criterion represents a sound basis for judging whether policy is on track does not require the central bank to believe that all shocks are alike, or even that all of the possible types of disturbances to which it may have to respond can all be listed in advance. At the same time, a public commitment to the target criterion tells the public in advance what it should expect with regard to the *outcome* to be achieved by policy. This is actually what the public most needs to be able to forecast well, and this is the aspect of the public's expectations that the central bank needs to influence, in order to achieve the benefits that are available in principle from policy commitment.

2 The Case for Price Stability

As noted above, the most important innovation of the inflation-targeting central banks, in my view, is the organization of policy deliberations around the achievement of an explicit target, quite apart from the type of target that happens to be chosen. But another distinctive feature of inflation targeting, of course, is the fact that the target is for some measure of *inflation*; while control of inflation has always been an important concern of central bankers, inflation targeting has given special, and sometimes exclusive, emphasis to this goal, and debates about the desirability of inflation targeting are often primarily discussions of the desirability of such a strong emphasis on inflation. Here I review what the theory of optimal monetary policy has to say about this.

First of all, the modern (micro-founded) literature on the real effects of monetary policy provides ample justification for the conventional wisdom of central bankers, according to which it is better for inflation to be both low and stable. It has been understood for some time that *expected* inflation creates distortions by increasing the opportunity cost of holding (non-interest-earning) money, leading to the inefficient use of real resources to economize on the use of money in transactions; this was the basis for the celebrated analysis of the optimal rate of inflation (which actually turned

out to be mild deflation) by Friedman (1969).⁷ However, models that incorporate some reason for prices to not adjust fully and instantaneously to changing market conditions — whether these involve infrequent price changes, or simply slow updating of the information on which prices are being set — imply that unanticipated variations in the inflation rate create real distortions as well, by causing prices that adjust at different times (or that are being set on the basis of different information sets) to become misaligned with one another.

So price stability has important advantages, in helping the market mechanism to work more effectively. Still, should this stabilization objective be given priority over others, such as stabilization of real economic activity or employment? I shall argue, below, that it should not be an *absolute* priority; but the recent literature on the welfare consequences of alternative monetary policies finds that there is less tension between inflation stabilization and properly defined real stabilization objectives than the traditional (non-welfare-theoretic) literature on monetary stabilization policy has often suggested. It is not a bad first approximation to say that the goal of monetary policy should be price stability.

2.1 When Full Price Stability is Optimal

Even when one grants that the economy is subject to exogenous real disturbances of many sorts — including various types of “supply shocks,” *i.e.*, disturbances that shift the “natural rate of output,” the level of output that would occur in an equilibrium with fully flexible prices — it is possible for the optimal monetary policy to be one that maintains completely stable prices in the face of these disturbances, and instead allows real activity to vary. In particular, this is true in a wide variety of “sticky price” or “sticky information” models (under varying assumptions about how many price-setters revise their prices or update their information in a given interval of time), as discussed in Woodford (2003, chap. 6), in the case that (i) the equilibrium fluctuations in the real allocation of resources *would* be optimal if only all prices were perfectly flexible and set on the basis of fully up-to-date information, and (ii) there are only aggregate shocks, so that in a flexible-price equilibrium all goods would have the same price. These hypotheses allow for the existence of a wide range of types

⁷Friedman’s argument remains correct in the case of a wide range of different ways of modeling the source of the demand for money; see, *e.g.*, Woodford (1990).

of real aggregate disturbances that should affect the natural rate of output — for example, exogenous variation in technology, in preferences regarding labor supply or impatience to consume, or in government purchases — though it does not allow for certain kinds of “supply shocks”, such as variations in the degree of market power in labor or product markets, or variations.

The basic intuition is fairly simple.⁸ The deadweight losses due to relative price distortions can be completely eliminated, in principle, by stabilizing the aggregate price level. For the aggregate price level is stabilized by creating an environment in which suppliers who choose a new price (under full information) have no desire at any time to set a price different from the average of existing prices. Then (because the average price level never changes), the price desired by any supplier that reconsiders its price is always the same, regardless of the number of future periods for which the price is expected to remain fixed, and regardless of how incomplete the supplier’s information may be about current market conditions. All new prices are then always chosen to equal the average of existing prices, and as a result the average price never changes. And all goods prices must eventually equal that same, constant value, so that inefficient relative-price dispersion due to price stickiness or information imperfections will not exist.

Furthermore, in such an environment, the equilibrium real allocation of resources will be the same as if all prices were fully flexible and set under perfect information. For by hypothesis, in that case suppliers would also all choose a common price equal to the current price index. Since they are able to charge this price at all times despite the infrequency of their reconsideration of their prices or the limitations of the information that they can use in adjusting prices, neither the stickiness of prices nor that of information has any effect on equilibrium behavior. Since by hypothesis the equilibrium allocation of resources would be optimal under full information and full price flexibility, it is optimal under the monetary policy that fully stabilizes prices.

⁸It is presented in the case of a model of staggered pricing by Goodfriend and King (1997). The fact that a similar conclusion is obtained in the case of “sticky information” is illustrated by the analysis of Ball *et al.* (2003).

2.2 Qualifications

The conditions under which full price stability can be shown to be an optimal policy are in some respects quite general; for example, the conclusion does not depend on fine details of how many prices are set a particular time in advance or left unchanged for a particular length of time. Nonetheless, the conditions assumed above are quite special — at least as an exact description of reality — in other respects, and it is likely that some degree of deviation from full price stability is warranted in practice. Some of the more obvious reasons for this are sketched here.

First of all, complete price stability may not be *feasible*. In the argument sketched above, I have supposed that it is possible to use monetary policy to maintain an environment in which a supplier with flexible prices and full information would never wish to change its price. Often there will exist a state-contingent path for short-term nominal interest rates consistent with such an equilibrium; it is shown in Woodford (2003, chap. 4) that this requires that the interest rate track the Wicksellian *natural rate of interest* — the real rate of return that would prevail in an equilibrium with flexible prices and full information — which varies in response to real disturbances. However, it is possible that at some times (as a result of exogenous real disturbances of a particular sort), the natural rate of interest is temporarily negative; if so, there cannot be an equilibrium in which the nominal rate of interest is equal at all times to the natural rate, and hence no equilibrium in which inflation is zero at all times. As a result, a policy will have to be pursued which involves less volatility of the short nominal interest rate in response to shocks, and some amount of price stability will have to be sacrificed for the sake of this.

Varying nominal interest rates as much as the natural rate of interest varies may also be desirable as a result of the “shoe-leather costs” involved in economizing on money balances. As argued by Friedman (1969), the size of these distortions is measured by the level of nominal interest rates, and they are eliminated only if nominal interest rates are zero at all times. Taking account of these distortions — from which we have abstracted thus far⁹ — provides another reason for the equilibrium with complete price stability, even if feasible, not to be fully efficient; for as Friedman argues,

⁹The hypothesis above that the equilibrium allocation of resources was efficient under flexible prices required, among other things, that transactions frictions of this kind be abstracted from. The economies referred to in the previous section are “cashless”, or at least near-cashless economies, in which transactions frictions are unimportant. See Woodford (2003) for details.

a zero nominal interest rate will typically require expected *deflation* at a rate of at least a few percent per year.

And taking account of these distortions affects more than the optimal *average* rate of inflation. As with distorting taxes, it is plausible that the deadweight loss resulting from a positive opportunity cost of holding money is a convex function of the relative-price distortion, so that temporary increases in nominal interest rates are more costly than temporary decreases of the same size are beneficial. In short, monetary frictions provide a further reason for it to be desirable to reduce the *variability* of nominal interest rates, even taking as given their average level. (At the same time, reducing their average level will require less variable rates, because of the zero floor.) Insofar as these costs are important, they too will justify a departure from complete price stability, in the case of any kinds of real disturbances that cause fluctuations in the natural rate of interest, in order to allow greater stability of nominal interest rates.

Yet while both of the factors just mentioned justify some departure from complete price stability, it is not clear that the volatility of inflation should be very great under an optimal policy, even when such factors are taken account of. For example, Rotemberg and Woodford (1997) characterize optimal policy for an estimated model of the U.S. economy, when a constraint that the mean federal funds rate must remain at least a certain number of standard deviations (greater than two) above zero is imposed as a substitute for the zero bound (that still allows a linear characterization of optimal policy). Even though the real disturbance processes in their model imply greater volatility of the natural rate of interest than many would assume (a standard deviation between three and four percentage points), they find that optimal policy involves an average rate of inflation only slightly greater than zero (11 basis points!), and not much variability of inflation (a standard deviation only 40 percent as large as the actual variability of U.S. inflation over their post-1980 sample period). Interest rates are smoothed considerably in the optimal policy, relative to what would be required to fully stabilize inflation, but this does not require too much variation in inflation, as their estimated model implies a variance tradeoff that is quite flat near the extreme of full inflation stabilization.

For the same reason, taking account of the distortions created by high nominal interest rates in a model with transactions frictions justifies only a relatively modest degree of inflation variation for the sake of greater stability of nominal interest rates, at least if the transactions frictions are calibrated at an empirically realistic magni-

tude. Woodford (2003, chap. 6) finds that when transactions frictions are calibrated to match facts about U.S. money demand, the penalty on nominal interest-rate variations that can be justified on welfare-theoretic grounds is a good bit smaller than the one assumed in Rotemberg and Woodford (1997). Hence in the case that the available tradeoff between interest-rate variability and inflation variability is the one estimated by Rotemberg and Woodford, the degree of inflation variability that could be justified on this ground would be even smaller than in their paper.

Even apart from these grounds for concern with interest-rate volatility, the class of models for which full price stability is optimal is a special one in several respects. One obvious restrictive assumption in the argument sketched above is that there are assumed to be no shocks that would require the relative prices of any goods to vary over time in an *efficient* equilibrium (*i.e.*, the shadow prices that would decentralize an optimal allocation of resources involve no variation in relative prices). If, instead, an efficient allocation of resources requires relative price changes, due to asymmetries in the way that different sticky-price commodities are affected by shocks, then full stabilization of a symmetric index of prices is not generally optimal, as shown by Aoki (2001) and Benigno (2003) in the context of two-sector models with asymmetric disturbances.

Nonetheless, it may still be possible to define an *asymmetric* price index, with the property that stabilization of this index is optimal, at least a good approximation to optimal policy, as these authors show.¹⁰ If the model is symmetric except for the frequency of adjustment of different types of prices, the optimal price index to stabilize puts more weight on the prices of the goods with “stickier” prices; this provides a theoretical justification for targeting an appropriately constructed measure of “core” inflation, rather than a standard consumer price index. But as long as the price index to be stabilized is appropriately chosen, complete stabilization of a price index is found (in calibrated examples) to be nearly optimal.

Similarly, the analysis sketched above assumed flexibility of wages. While this is a familiar assumption in sticky-price models used for pedagogical purposes, many

¹⁰Benigno (2003) applies this idea to an analysis of optimal stabilization objectives for a monetary union in which different regions are affected asymmetrically by real disturbances. In this application, the optimal inflation target for the monetary union does not necessarily put weights on the national inflation rates that are proportional to the shares of those country’s products in the union-wide consumption basket.

empirical models imply that wages are as sticky as are prices, and possibly more so.¹¹ But real disturbances almost inevitably require real wage adjustments in order for an efficient allocation of resources to be decentralized, and if *both* wages and prices are sticky, it will then not be possible to achieve all of the relative prices associated with efficiency simply by stabilizing the price level – specifically, the real wage will frequently be misaligned, as will be the relative wages of different types of labor if these are not set in perfect synchronization.

In such circumstances, complete price stability may not be a good approximation at all to the optimal policy, as Erceg *et al.* (2000) show. Nonetheless, one can show once again that stabilization of an appropriately weighted average of prices and wages may still be a good approximation to optimal policy; it is even fully optimal in special cases (Woodford, 2003, chap. 6). Thus concerns of this kind are not so much reasons not to pursue price stability as they are reasons why care in the choice of the index of prices (including wages) that one seeks to stabilize may be important.

Finally, even when wages are flexible (or there are efficient labor contracts) and all disturbances have symmetric effects on all sectors of the economy, the flexible-price equilibrium level of output need not be welfare-maximizing. Both market power and the existence of distorting taxes imply that in reality, the equilibrium level of economic activity is likely to be too low on average.¹² When this is true, not only is the flexible-price equilibrium level of output different from the (first-best) optimal level, but except in special cases, real disturbances will not shift these two quantities to quite the same extent (in percentage terms).¹³ This means that the gap between the level of output associated with a policy that maintains stable prices (which is the

¹¹See, *e.g.*, Amato and Laubach (2003), Christiano *et al.* (2001), Altig *et al.* (2002), Smets and Wouters (2002a, 2002b), and Giannoni and Woodford (2003).

¹²This does not occur in the model of Rotemberg and Woodford (1997) owing to the assumed presence of an output subsidy that offsets the consequences of the market power of the monopolistically competitive suppliers of differentiated goods.

¹³King and Wolman (1999) and Khan *et al.* (2002) analyze a model in which it is optimal to fully stabilize prices in response to technology shocks, despite the existence of an inefficiently low steady-state level of output. This result, however, depends on the assumption of special isoelastic functional forms for both preferences and technology, and also on the assumption of zero steady-state government purchases; deviations from any of these assumptions will result in full price stability no longer being optimal. Also, even under the assumed specification, other types of real disturbances imply that it will not be optimal to fully stabilize inflation, as Khan *et al.* show. See Woodford (2003, chap. 6) for further discussion.

same as the flexible-price equilibrium output, as explained above) and the optimal level of output will be time-varying. If we write the aggregate-supply relation as a relation between inflation and the welfare-relevant “output gap” (*i.e.*, the gap between the actual and efficient levels of output), an additional exogenous “cost-push” term appears. As a consequence, it will not be possible to simultaneously stabilize inflation and the welfare-relevant output gap.¹⁴

Yet even so, the degree of variability of inflation under an optimal policy may be quite modest. This is because the relative weight that should be placed on the goal of output-gap stabilization, relative to the weight on inflation stabilization, may not be large. (This is illustrated in the case of the welfare-based loss function for the model of Giannoni and Woodford, 2003, presented in the appendix.) There is a straightforward reason for this. In a variety of optimizing models with sticky prices, it is shown in Woodford (2003, chap. 6) that the loss function that corresponds to a quadratic approximation to expected utility involves a relative weight on output-gap stabilization that is proportional to the coefficient on the output gap in the short-run aggregate-supply relation. This means that the same underlying microeconomic factors that lead to a relatively flat aggregate-supply relation — and thus imply that fluctuations in nominal aggregate demand have large effects on output relative to their effects on prices — *also* imply that the welfare losses associated with fluctuations in the level of aggregate real activity are small relative to the welfare losses that result from the mis-alignment of prices that are not adjusted with perfect synchronization when inflation varies.

It follows that, while the welfare-theoretic loss functions derived for the estimated models of Rotemberg and Woodford (1997) and Giannoni and Woodford (2003) involve stabilization goals other than inflation stabilization, by far the largest coefficients are those on the inflation stabilization goal. Given this, optimal policy will still be focused to an important degree on inflation stabilization. While the considerations sketched in this section give one ample reason to consider the consequences of monetary policy for the evolution of variables other than inflation, it will nonetheless make sense to think of the optimal policy rule as a “flexible inflation targeting rule.”

¹⁴Even when the average level of output is efficient, the flexible-price level of output and the efficient level may be differently affected by certain kinds of real disturbances. As noted above, these include variations in market power or in the level of tax distortions.

3 Improving the Practice of Inflation-Forecast Targeting

I turn now to some ways in which an optimal forecast-targeting procedure for the conduct of monetary policy, from the perspective of the theoretical literature summarized above, would differ from inflation-forecast targeting as it is currently practiced by the central banks that have led the way in developing this approach to monetary policy. Of course, the precise details of an optimal procedure depend on the details of one's model of the monetary transmission mechanism, and it can hardly be argued that there is yet a consensus about the correct model to use for one country, let alone a model that can be claimed to apply equally to all countries. Nonetheless, it seems that one can draw at least a few broad lessons about the character of optimal policy rules from the analyses that have been undertaken thus far, and that these differ enough from current practice to allow some suggestions for improvement to be offered.

3.1 The Target Criterion Should Involve More than Inflation

The official target criterion of the Bank of England — ensuring that projected RPIX inflation eight quarters in the future should always equal 2.5 percent per annum — refers only to the projected future value of a particular measure of U.K. inflation. While other inflation-targeting central banks are often less explicit about the precise way in which current policy decisions are supposed to be determined by their inflation targets, it is very generally the case that there is an explicit target *only* for (some measure of) inflation, and no commitment to take into account the projected paths of any other variables. Hence debates about the desirability of inflation targeting in countries such as the U.S. often assume that such an approach to policy would mean a sole concern with inflation stabilization.

An optimal policy, instead, will not involve complete stabilization of inflation except under fairly special circumstances, as discussed in the previous section. In general, an optimal policy will involve some degree of temporary variation in the inflation rate in response to real disturbances, for the sake of greater achievement of other stabilization objectives. The degree to which this matters in practice will depend on the quantitative specification of one's model of the economy; but an identification of

inflation targeting with what Svensson (1999) calls “strict inflation targeting” makes it too easy for opponents of inflation targeting to argue that it would prevent the central bank from responding in appropriate ways to changing economic conditions.

It is sometimes argued that a coherent monetary policy requires “a single objective,” so that stabilization objectives in addition to inflation stabilization should play no role in the conduct of monetary policy, despite the admitted desirability of these ends.¹⁵ It is true that a simultaneous commitment to stabilize two different variables using a single policy instrument will, in general, represent a promise that cannot possibly be fulfilled. But a commitment to a single *target criterion*, on the basis of which the instrument of policy is to be adjusted, does not require that this criterion involve only a single *variable*. The target criterion may well be a linear combination of projections for several different variables (just as it may also involve inflation projections at more than one horizon). In general, an optimal target criterion will be of this form. For example, in the case of the Giannoni-Woodford model of the U.S. monetary transmission mechanism discussed in the appendix, the optimal target criterion involves not only projected inflation, but also real-wage and output-gap projections. A lower projection of real wage growth or of the (welfare-relevant) output gap will justify acceptance of a higher projected inflation rate. Nonetheless, there is a single well-defined measure at each point in time of whether policy remains on track.

It is not obvious, of course, that actual inflation-targeting central banks do not take into account other stabilization objectives in their policy decisions, despite their use of an official rhetoric that suggests a strict inflation target. Commentators such as Bernanke *et al.* (1999) and Svensson (1999) argue that all actual inflation-targeting central banks are “flexible inflation targeters,” that trade off inflation stabilization against other stabilization objectives. Furthermore, it is often argued that a particular advantage of inflation-forecast targeting as a policy rule is precisely that it allows monetary policy to be used to reduce the short-run effects of disturbances on real variables (such as the output gap), while retaining firmly anchored medium-term inflation expectations, and hence reducing the degree of inflation variability that is

¹⁵A related view asserts that other goals may be introduced only to the extent that they do not interfere with achievement of the inflation target. However, absolute priority of the inflation target would not seem to leave any room for stabilization of output or other variables, unless the inflation target is not understood to require stabilization of inflation to the greatest extent possible. Such formulations are thus hopelessly vague about what the policy commitment actually promises.

required in order to achieve a given degree of stability of the real variables.

I do not doubt that actual inflation-targeting central banks do take some account of real objectives. For example, one notes that the introductory summary of the Bank of England's *Inflation Report* always presents a chart of the Bank's current real GDP projection as well as its inflation projection — and the GDP projection is always discussed *first*, even if it is solely the inflation projection that is cited as showing that policy is on track. But it would be desirable for central banks to commit themselves to the pursuit of explicit target criteria that involve real variables as well as inflation. For one thing, if the criteria on which policy is actually based include projections for other variables, it would increase transparency, facilitating the public's ability to correctly anticipate future policy, to explain policy in this way. In addition, greater frankness about aspect of banks' policy commitments would help to dispel some of the resistance to the adoption of inflation targeting in countries like the United States. In particular, it would show that adoption of a targeting framework by the Federal Reserve need not imply any departure from the Fed's current legal mandate — which requires it to pursue full employment as well as price stability — and hence need not wait for Congressional authorization.¹⁶

3.2 A “Medium Term” Target is Not Enough

Many would argue that the reason that inflation targeting central banks have only an unqualified, time-invariant target for inflation — rather than a target criterion that takes account of output projections, or other variables, as well — is that the inflation target represents only a “medium-term” goal, that leaves unspecified the precise transition path by which the medium-term goal is to be reached. (This is explicit in the case of the Bank of England's official target criterion. Only the rate of inflation 8 quarters in the future must equal the time-invariant target rate; nearer-term inflation projections are allowed to vary.) The appropriate *medium-term* inflation target can be stated in an unqualified, time-invariant form, it is argued, because there is no substantial tradeoff between the inflation rate and real variables this far in the future. Other stabilization goals are instead appropriately taken into account in choosing among the possible nearer-term transition paths that are consistent with the medium-

¹⁶On the issue of whether the adoption of inflation targeting in the U.S. would require new legislative authority, see also Goodfriend (2003).

term target.

In fact, the sort of optimal target criteria that can be derived using the method of Giannoni and Woodford (2002) involve much nearer-term projections than those that are officially targeted by the Bank of England or other inflation-targeting central banks. For example, while the targeting criteria discussed in the appendix involve weighted averages of projections for many different future quarters, it is the projection for one or two quarters in the future that receives the greatest weight. Thus the optimal target criteria do not merely describe the state that one wishes to reattain once the effects of recent disturbances have worked themselves out; they also characterize the optimal transition dynamics following a disturbance.

A simple example may be useful in clarifying this. Suppose that the prices of individual goods are re-optimized at random intervals as proposed by Calvo (1983), but that all prices are fixed a quarter in advance, so that even those new prices that are chosen in quarter t take effect only beginning in quarter $t + 1$. Suppose furthermore that between the occasions on which the optimality of a given price is reconsidered, it is automatically indexed to an aggregate price index (but again, the aggregate price index of the quarter before the one in which the price will apply), as proposed by Christiano *et al.* (2001). In a simple model with fixed capital and no labor-market frictions, this results in an aggregate-supply relation of the form¹⁷

$$\pi_t - \pi_{t-1} = \kappa E_{t-1} x_t + \beta E_{t-1} (\pi_{t+1} - \pi_t) + u_{t-1}, \quad (3.1)$$

where π_t is the quarter t inflation rate, x_t is the (welfare-relevant) output gap, u_{t-1} is an exogenous (mean-zero) random disturbance at date $t - 1$, κ is a positive coefficient, and β is the discount factor of the representative household. Exogenous fluctuations in the “cost-push” term u_{t-1} as a result of various real disturbances then create a tension between the goals of inflation stabilization and output-gap stabilization.

Under the microeconomic foundations proposed for the aggregate-supply relation above, the appropriate welfare-theoretic stabilization objective corresponds to mini-

¹⁷This is essentially the form of aggregate-supply relation proposed by Fuhrer and Moore (1995), and a simplified version of the aggregate-supply blocks of the empirical optimizing models of Christiano *et al.* (2001), Altig *et al.* (2002), Smets and Wouters (2002a, 2002b), and Giannoni and Woodford (2003).

mization of a loss function of the form¹⁸

$$E_{t_0} \sum_{t=t_0}^{\infty} \beta^{t-t_0} [(\pi_t - \pi_{t-1})^2 + \lambda(x_t - x^*)^2], \quad (3.2)$$

where both the optimal output gap x^* (positive in the empirically realistic case) and the positive relative weight λ depend on model parameters. Because it is assumed that prices are automatically indexed to a lagged aggregate price index, inflation creates distortions in the model only to the extent that the aggregate inflation rate *differs* from that in the previous quarter; hence policy should aim to stabilize the *rate of change* of inflation, rather than its absolute level.¹⁹ As we shall see, however, this does not mean that it is not desirable for the central bank to commit to a fixed long-run inflation target.

Let us consider now the problem of conducting policy from some date t_0 onward so as to minimize (3.2), subject to a constraint

$$\pi_{t_0+1} = \bar{\pi}_{t_0}. \quad (3.3)$$

This last constraint prevents the policy authority from choosing a policy at date t_0 that fails to internalize the effects of policy at t_0 (insofar as it could have been forecasted in the previous quarter) on the inflation-output tradeoff faced in quarter $t_0 - 1$. Choosing a policy commitment from date t_0 onward in the absence of any such constraint would result in selection of a policy that is not *time-consistent*, for one *would* commit to a policy at all later dates that took account of these effects. If the constraint $\bar{\pi}_{t_0}$ is chosen (as a function of the state of the world in quarter t_0) in a “self-consistent” way, the optimization problem just posed can be solved by a time-invariant policy rule. Furthermore, if one reconsiders the desirability of following the policy rule at any later date, then (assuming that one’s model of the economy and policy objectives have not changed in the meantime) one would continue to find that the same time-invariant policy rule would continue to solve the corresponding constrained optimization problem looking forward from the later date.²⁰

¹⁸For details of the derivation, see Woodford (2003, chap. 6).

¹⁹The conclusion that the absolute level of inflation has no consequences for welfare is extremely special to the simple case considered here, and surely not realistic, as discussed in Woodford (2003, chap. 7). For similar analyses of the form of optimal target criteria in the case that there is no indexation, or only partial indexation, see Svensson and Woodford (2003) and Giannoni and Woodford (2003).

Finally, let us suppose that the component of aggregate real expenditure that is sensitive to interest rates is also determined a quarter in advance, so that the output gap x_t cannot be affected by monetary policy decisions later than quarter $t - 1$.²¹ It follows that monetary policy can affect only the evolution of inflation and the component of the output gap that is forecastable a quarter in advance, and that the possible stochastic paths for these variables that can be achieved by any monetary policy are the set of processes consistent with relation (3.1) for $t \geq t_0 + 1$.

The first-order conditions for the optimization problem just stated are then of the form

$$\pi_{t+1} - \pi_t + \varphi_t - \varphi_{t-1} = 0 \quad (3.4)$$

$$\lambda(E_t x_{t+1} - x^*) - \kappa \varphi_t = 0 \quad (3.5)$$

for each $t \geq t_0$, where φ_{t-1} is the Lagrange multiplier associated with constraint (3.1) for each $t > t_0$, and φ_{t_0-1} is a multiplier associated with the constraint (3.3). These conditions, together with the constraints, determine the optimal state-contingent evolution of inflation and the forecastable component of the output gap; the unforecastable component of the output gap, of course, is exogenously given.

How should monetary policy be conducted in order to ensure that this desired state-contingent evolution of inflation and output is realized? Applying the method of Giannoni and Woodford (2002), one can eliminate the Lagrange multiplier from equations (3.4) – (3.5), and show that one must have

$$(\pi_{t+1} - \pi_t) + \phi(E_t x_{t+1} - E_{t-1} x_t) = 0 \quad (3.6)$$

for each $t \geq t_0$, where $\phi \equiv \lambda/\kappa > 0$. This in turn implies that

$$\pi_{t+1} + \phi E_t x_{t+1} = \pi^* \quad (3.7)$$

for each $t \geq t_0$, for some constant π^* , the value of which will depend on the initial constraint $\bar{\pi}_{t_0}$. For any value of π^* , there exists a self-consistent specification of the

²⁰See Woodford (2003, chap. 7) for further discussion. A policy that solves a problem of this form is “optimal from a timeless perspective,” as discussed in Woodford (1999b).

²¹For models of aggregate demand with this property, see Woodford (2003, chap. 5). This kind of predetermination of interest-sensitive aggregate expenditure is a feature of many empirical optimizing models, such as Rotemberg and Woodford (1997), Amato and Laubach (2003), Christiano *et al.* (2001), Altig *et al.* (2002), and Giannoni and Woodford (2003), along with many *ad hoc* macroeconomic models.

initial constant under which optimal policy satisfies (3.7) for all $t \geq t_0$. Thus the optimal long-run inflation target π^* is not determined within this model.²²

Optimal policy, then, must arrange that (3.7) holds at each date, or equivalently, that

$$E_t[\pi_{t+1} + \phi x_{t+1}] = \pi^*. \quad (3.8)$$

(This alternative form emphasizes the fact that the terms in the target criterion can be affected only by monetary policy decisions in quarter t or earlier.) Conversely, one can show that if policy ensures that (3.8) is satisfied at each date $t \geq t_0$, the unique non-explosive rational-expectations equilibrium consistent with the policy commitment is the one that solves the optimization problem stated above. Hence (3.8) is an optimal *target criterion* for the central bank's policy decision in quarter t .

In the model sketched above, the central bank cannot expect to affect whether (3.8) holds in quarter t through adjustment of the interest rate i_t in that quarter, for the predetermination of the interest-sensitive component of expenditure implies that unforecastable interest-rate changes have no effect on aggregate demand. The central bank's period t policy decision should then be a commitment $i_{t+1,t}$ regarding its operating target for the interest rate in quarter $t + 1$. The value of $i_{t+1,t}$ should be chosen so as to lead the central bank to project that (3.8) is satisfied, conditional on the state of the economy in quarter t .²³ The expectation that $i_{t+1,t}$ will be chosen in this way in each quarter $t \geq t_0$, and that the central bank will then act to ensure that $i_{t+1} = i_{t+1,t}$ in the following quarter, will then imply the desired state-contingent evolution of inflation and output.

The proposed policy rule involves a constant long-run inflation target, since satisfaction of (3.8) each quarter implies that one must have

$$\lim_{T \rightarrow \infty} E_t \pi_T = \pi^* \quad (3.9)$$

at all times t . And it would surely be desirable for the central bank to emphasize to the public its commitment to a policy that implies (3.9), as this should help to

²²The addition of even small additional frictions can break the indeterminacy of the optimal long-run inflation target, as discussed in Woodford (2003, chap. 7). In practice, one can be certain that the optimal long-run inflation target is not far from zero; it could even be slightly below zero.

²³Note that in the model sketched here, both $E_t \pi_{t+1}$ and $E_t x_{t+1}$ should be affected by the bank's choice of $i_{t+1,t}$, assuming that the bank's announcement of its target for the following quarter is credible to the private sector.

anchor long-run inflation expectations (which would never be allowed to vary in an optimal equilibrium). Nonetheless, a commitment to ensure that (3.9) is satisfied at all times is not *sufficient* for optimality; many different sorts of transitory responses to disturbances would be equally consistent with it.

Nor is it clear what a central bank is committing itself to do if it pledges to ensure that (3.9) is satisfied at all times. Condition (3.9) does not place any restrictions on the behavior of interest rates over any finite horizon; hence it is not clear what one would be able to monitor about a central bank’s decisions in order to verify that it is indeed acting in conformity with its supposed commitment. Condition (3.8), instead — together with the expectation that (3.8) will also hold at all future dates — does imply a particular rational-expectations equilibrium value for $E_t i_{t+1}$, and so one could monitor, at least in principle, whether $i_{t+1,t}$ is chosen in accordance with it.

The same is true in the case of a “medium term” target that refers to a specific future date. Condition (3.8) implies that

$$E_t \left[\pi_{t+k} + \frac{\lambda}{\kappa} x_{t+k} \right] = \pi^* \tag{3.10}$$

must also hold at all times, for any $k \geq 1$, in an optimal equilibrium. So one might imagine that it would suffice for the central bank to commit to ensure that (3.10) holds at all times, where k might be 8 quarters in the future. But if $k > 1$, this condition does not suffice to determine a unique non-explosive rational-expectations equilibrium, in the context of the model set out above. For any commitment of the form

$$E_t \left[\pi_{t+1} + \frac{\lambda}{\kappa} x_{t+1} \right] = \pi^* + u_t, \tag{3.11}$$

where u_t is an exogenous random variable satisfying

$$E_t u_{t+k-1} = 0,$$

suffices to determine an equilibrium, for the same reason that (3.8) does, though the equilibrium will not be the optimal one (the one determined by (3.8) except when $u_t = 0$ at all times. Yet ensuring that (3.11) holds at all times is consistent with commitment (3.10); thus all of the different equilibria corresponding to different choices of the process $\{u_t\}$ are equally consistent with a commitment of the form (3.10). It follows that a commitment to ensure (3.10) fails to determine a unique

equilibrium, and indeed it fails to uniquely determine the required interest-rate policy on the part of the central bank.

A well-known argument for the desirability of a target criterion referring only to inflation two years in the future is provided by Svensson (1997). In the simple model used in that paper for illustrative purposes, the optimal target criterion (in the sense of Giannoni and Woodford) is shown to be of this form. But this results from the fact that in that model, an interest-rate decision by the central bank has no effect on inflation until two years later. It is also true in the case of the model sketched above, in which inflation can only be affected by monetary policy decisions in the previous quarter, that the optimal target criterion involves a forecast of inflation one quarter in the future; if the assumed delay were longer, the optimal target criterion would look farther into the future.

However, empirical models of the monetary transmission mechanism do *not* commonly imply delays of greater than a quarter before monetary policy is able to affect inflation, even if (because of various sorts of inertia in the transmission mechanism) the models imply that the effects of disturbances on the inflation rate are greatest only after several quarters. For example, the aggregate-supply relation (3.1) assumed above has the property that a demand disturbance (due to monetary policy or some other source) that raises output above its natural rate for several quarters will steadily increase inflation for several quarters, with the full effect on inflation being observed only after output has returned to its natural level. Nonetheless, optimal policy is described by a target criterion (3.8) that involves only a one-quarter-ahead inflation forecast. A similar result is obtained in the more complex model of Giannoni and Woodford (2003), discussed in the appendix: the optimal target criteria involve forecasts at many future horizons, but the weight is greatest on the forecasts for the nearest horizon at which the variables in question can still be affected by the current policy decision.

This should not be surprising; if the target criterion is to completely determine a policy decision at each date, it must specify what defines an acceptable outcome at the nearest date that can still be influenced policy, and not merely what must happen later, at dates that can be influenced by *later* policy decisions. The preference for “medium term” target criteria at inflation-targeting central banks represents a preference for *incomplete* specifications of the banks’ policy commitments. This probably reflects a greater degree of certainty about the desirability of the particular aspect of

policy about which the commitment is being made, and this is understandable. One can indeed state with greater confidence that it is desirable for medium-term inflation expectations to be highly stable (and to suggest a plausible value for the target π^*) than one can argue for the desirability of a particular criterion such as (3.8) that should be satisfied by the transition dynamics for inflation following a temporary disturbance.

Nonetheless, it is possible to make an argument for a particular near-term target criterion such as (3.8) that is surprisingly robust. For example, it might be thought better to leave the transition dynamics following disturbances unspecified on the ground that the optimal transition dynamics will look very different in the case of different types of disturbances. Yet Giannoni and Woodford (2002) show that it is possible quite generally to find a target criterion that applies regardless of the character of (additive) disturbances, yet which is sufficiently specific to uniquely determine the transition dynamics in response to any type of disturbance.

It is sometimes proposed, in discussions of inflation-forecast targeting, that a suitable form of central bank commitment, that is specific enough about the desired transition dynamics to determine an appropriate policy action, involves specification of a medium-term inflation target, together with a specification of the *rate* at which policy should seek to restore inflation to the target level when it deviates from it. A commitment of this form can be expressed in terms of a near-term target criterion of the form

$$E_t\pi_{t+1} = \pi^* + \mu(\pi_t - \pi^*), \quad (3.12)$$

where $0 < \mu < 1$ indicates the rate at which departures from the target should be eliminated; thus such a proposal amounts to a near-term target criterion, and not simply a medium-term inflation target. However, it is not generally possible to express a robustly optimal target criterion (in the sense of Giannoni and Woodford) in a form like this, that makes no reference to the projected path of any variable other than inflation. A robustly optimal criterion such as (3.8) implies a particular rate of convergence of $E_t\pi_{t+k}$ to π^* as k is made large, but this will differ depending on the recent history of disturbances; it is only the criterion (3.8), which involves the output-gap projection as well, that represents a robust criterion for optimality.

3.3 Constant-Interest-Rate Projections an Inappropriate Basis for Policy

One way that inflation-targeting central banks resolve the problem that their medium-term inflation target alone does not suffice to determine a particular current interest-rate decision — at least according to their official rhetoric — is by asking what *constant* interest-rate setting over the forecast horizon would result in a projection consistent with the medium-term target criterion, and then choosing a current interest-rate operating target at that level. For example, the Bank of England’s *Inflation Reports* justify current policy by showing that a projection based on the assumption that the interest rate will remain at the current level for the next two years indicates projected RPIX inflation equal to 2.5 percent 8 quarters from now.²⁴ This does not, however, mean that such banks constrain themselves to actually maintain a constant interest rate for two years at a time; instead, a new interest-rate setting is to be chosen each time the projection exercise is repeated.²⁵

This “solution” to the problem of the incompleteness of the policy commitment represented by the medium-term target has the advantage of being simple to explain to the public — as long as the public is not sophisticated enough to ask what it really means — but has a number of unappealing implications.²⁶ First of all, many optimizing models of the monetary transmission mechanism have the property first demonstrated by Sargent and Wallace (1975) for a rational-expectations IS-LM framework, namely, that the equilibrium path of the price level (and hence of the inflation rate) is *indeterminate* under the assumption of a fixed nominal interest rate (or indeed, any exogenously specified interest-rate process).²⁷ If such a model were to be used for the central bank’s projection exercise, the staff would be unable to

²⁴Former MPC member Charles Goodhart (2001) describes himself as having tried to set interest rates in this way, and says “This was, I thought, what the exercise was supposed to be” (p. 177). Heikensten (1999) describes the similar procedure used by the Bank of Sweden.

²⁵Indeed, Goodhart (2001) lists as an advantage of the constant-interest-rate projection-based procedure that “no one infers any commitment from the MPC to abide by that assumption in the future, nor is the credibility of the MPC damaged when, having made this assumption in a forecast one month, it decides to change interest rates even in the next month” (pp. 174-175).

²⁶Goodhart (2001) reviews what he calls “the prima facie case against” this approach before offering his defense of it. Other critical discussions include Leitemo (2003), Svensson (2003), and Honkapohja and Mitra (2003).

²⁷See Woodford (2003, chap. 4) for further discussion.

compute predicted paths for inflation or other variables under the hypothesis of any constant level of nominal interest rate, and so unable to assert that one particular level would imply satisfaction of the target criterion.²⁸

Alternatively, many backward-looking models (including optimizing models in which expectations are assumed to be based on extrapolation from past time series) have the property discussed by Friedman (1968), namely, that maintaining a constant nominal interest rate indefinitely will lead to explosive inflation dynamics, through a Wicksellian “cumulative process.”²⁹ Goodhart (2001) suggests that the Bank of England’s model has this latter property, and that as a result, “the rate of change of most variables visible at the two-year horizon in the Bank’s forecast generally (though not invariably) tends to persist, and on occasion to accelerate, in the third and subsequent years” (p. 171). In this case, it is possible to ask which constant interest rate would imply satisfaction of the target criterion at a certain finite horizon, but only at the expense of making it clear that hitting the target at (say) the 8-quarter horizon does *not* also imply expecting to hit it in subsequent quarters. Hence it cannot be the case that one expects to be content to maintain the constant-interest-rate policy indefinitely, even in the absence of any developments that cannot already be foreseen.³⁰

In fact, there is no reason to suppose that the constant interest-rate path represents the bank’s best current estimate of the future path of interest rates. This is at least implicitly conceded by the Bank of England in its published discussions of the accuracy of its projections.³¹ In these discussions, the Bank gives exclusive

²⁸Leitemo (2003) discusses possible interpretations of the constant-interest-rate projection exercise that would allow it to yield a policy recommendation even in the case of a forward-looking model of the transmission mechanism; but these do not eliminate the other unappealing features of such a procedure.

²⁹See Bullard and Mitra (2002) and Preston (2002) for analyses of forward-looking models with least-squares learning by the private sector.

³⁰If one’s model currently implies that inflation will depart significantly from the target rate at the three-year horizon if interest rates are maintained at their current level for that long, then it also implies that one should expect that a year from now — barring unforeseen developments — if interest rates have been maintained at their current level, it will then be forecasted that inflation will depart from the target at the *two-year* horizon if interest rates are not changed. Hence one cannot expect that interest rates should remain at their current level for an entire year, even in the absence of any “news”.

³¹See the *Inflation Reports* of August 2001 and August 2002.

attention to the projections that it also publishes in the *Inflation Report*, in which an interest-rate path is assumed that corresponds to current market expectations, rather than to the projections conditional on the constant interest-rate path, even though the latter ones are given primary emphasis in the justification of policy. It is evident that the Bank does *not* regard the constant interest rate assumption as the best available forecast of its behavior, for if it did, it would want to test the accuracy of the projections made under that assumption, rather than under whatever contrary assumptions might be made by traders in financial markets.

Thus the auxiliary assumption that is used to allow the forecast-targeting procedure to determine an interest-rate recommendation has the consequence that the targeting procedure is based on forecasts that are not actually believed, even in the Bank itself. Such a procedure has the paradoxical implication that the central bank may choose a policy under which it does *not* truly expect the target criterion to be satisfied, though it may believe that it would be under the counterfactual hypothesis of the constant interest rate.

Such a state of affairs can hardly be defended as conducive to transparency in the conduct of monetary policy. If policy is genuinely based on constant-interest-rate conditional projections, then one's policy decisions are not aimed at ensuring satisfaction of the target criterion that is announced to the public; and the projections published by the central bank are not accurate forecasts that should better help the private sector to correctly anticipate the economy's evolution. On the other hand, if the central bank genuinely does expect the target criterion to be satisfied, then policy is not actually determined in the way that the official rhetoric implies that it is; and if the forecasts are unbiased, then they are not the kind of forecasts that they are officially described as being.

The kind of forecast-targeting procedure recommended by Svensson and Woodford (2003) as a way of implementing optimal monetary policy is of a different sort. In this procedure, one projects the economy's future evolution under alternative contemplated policy decisions, assuming that in future decision cycles the central bank will *again act to ensure satisfaction of the target criterion*. This amounts to asking what action is needed in order to project that that the criterion should be satisfied in the current period, taking as given that it is expected to be satisfied in later periods (as a result of the policy actions to be taken in those periods). Such a calculation yields a determinate outcome as long as there is a determinate rational-expectations equi-

librium implied by the target criterion; this is always the case if the target criterion is selected according to the method of Giannoni and Woodford (2002).

Thus policy should be based on a projection exercise that includes a model of the central bank's own future behavior — one that is furthermore consistent with the procedure that it actually follows in making its policy decisions. This is the kind of projection exercise used as the basis for policy decisions at some central banks, notably the Reserve Bank of New Zealand, which also publishes some information about the non-constant interest-rate path implicit in its projections, along with its projections for inflation and other variables.

Goodhart (2001) objects that such a procedure is impractical, on the ground that it would be much more difficult for a monetary policy committee to reach agreement on an entire future path for interest rates, rather than allowing them to decide only about the current interest rate each time they meet. But the procedure described by Svensson and Woodford does not involve a multi-dimensional decision problem in each decision cycle. As with the constant-interest-rate projection method, one makes a decision for the current period only, on the basis of projections of the future that (necessarily) incorporate a hypothesis about future policy; the hypothesis about future policy is simply a more realistic one than the notion that interest rates will not change, regardless of how inflation and output evolve. And there is no greater need for agreement among the members of the policy committee about that particular aspect of the model specification than about the other assumptions involved in making projections for the future.

Goodhart also argues that revealing a projected non-constant path for interest rates is problematic, because “any indication that the MPC is formally indicating a future specific change in rates ... would be taken to indicate some degree of commitment” (p. 175). This is clearly a delicate issue in the proper explanation to the public of how the central bank's projections are to be interpreted. Yet the experience in New Zealand suggests that it is possible to reveal interest-rate projections to the public without being understood to have made an advance commitment about the path of the official cash rate. Moreover, a “fan chart” for the path of interest rates, like those that the Bank of England currently publishes for its inflation and output projections, ought to make it clear that the bank is not committing itself to a definite path; rather, the expected evolution will depend on a variety of contingencies that can at best be assigned probabilities.

If necessary for reasons of the sort to which Goodhart refers, it would be preferable to base policy on projections conditioned on predicted future policy, and to publish inflation and output projections of that sort, without any mention of the interest-rate path implicit in these projections, than to base policy on projections conditional upon a model of policy that one knows to be false. But there are likely to be advantages to publication of the interest-rate projections. One of the crucial ways in which central banks affect the economy is through the effects of their announcements on expectations regarding the future path of short-term interest rates, which expectations then determine longer-term bond yields, asset prices, and exchange rates, which in turn affect spending, employment, wage-setting and price-setting decisions. The current level of overnight interest rates is in itself of little importance for most economic decisions; the real significance of central-bank decisions about the overnight rate is what they are taken to signal about the likely path of interest rates months and years into the future. Given the importance to a central bank of steering expectations regarding future interest rates in a desirable way, it would seem that revealing to the public the expected future path of rates implied by the bank's policy commitments should help it to better achieve its goals.

3.4 Advantages of A History-Dependent Target Criterion

A notable feature of the kind of projection exercises upon which policy is currently based at banks like the Bank of England is that they are *purely forward-looking*. By this I mean that the decision made at any time is a function solely of the policy committee's judgment about the possible paths from now on for inflation and other variables (if any) relevant to its target criterion; past conditions are irrelevant except insofar as these have an effect on what it is possible to achieve from now on. Of course any projection-based decision procedure will be forward-looking; but under a procedure like the Bank of England's (at least as Goodhart, 2001, describes it), the past is irrelevant, because the target criterion is a time-invariant function of the projected future path of the target variable (RPIX inflation).

One might think that forward-looking behavior of this kind is a necessary feature of optimal policy — that “bygones should be bygones” for a rigorous optimizer. But as explained above, this is not correct in the case of the optimal control of a forward-looking system; if it were, there would be no flaw in the reasoning of a purely

discretionary policymaker. When the private sector is forward-looking, expectations regarding future policy matter for what can be achieved at any point in time, and outcomes can generally be improved through a judicious commitment regarding future policy. This requires, however, that policy be expected to be conducted at the later date in a way that is *history-dependent* — that is, in a way that depends on the earlier conditions (at the time at which it was desirable to alter expectations) as well as upon conditions at the time that the action is taken.

This history-dependence can be incorporated into a forecast-targeting procedure through the use of a history-dependent target criterion to evaluate whether the economy’s projected evolution from now on should be considered to be consistent with the bank’s general policy commitments. This means that the acceptable projections for the target variables looking forward should depend on recent *past* conditions.³² This is a further reason why, under an optimal regime, the *short-term* target for inflation will be time-varying, even though there is likely to be a constant long-run inflation target, around which the short-term target fluctuates. The way in which an optimal short-term target criterion is likely to be history-dependent is illustrated in the discussion in the appendix of the optimal target criteria in the case of the estimated model of Giannoni and Woodford (2003).

A particular clear example of the advantages of a history-dependent target criterion is the situation currently faced by the Bank of Japan, in which the zero lower bound on nominal interest rates has been reached yet deflation continues, so that further monetary stimulus is desirable. As I have already mentioned, the main lever by which monetary policy can still affect the economy under such circumstances is by changing expectations regarding the future conduct of policy: committing to a more expansionary policy *later* than would otherwise have been pursued. But this requires

³²The optimal target criterion (3.8) in the simple example above might seem not to confirm this principle, as it involves only forecasts of π_{t+1} and x_{t+1} . But in that model, inflation is not technically a “target variable,” as it is the rate of inflation acceleration, rather than the absolute rate of inflation, that enters the loss function (3.2). A purely forward-looking target criterion would then be one that involves only the projected future paths of the output gap and of *inflation acceleration*. The target criterion (3.8) is not of this form, as it implies a commitment to eventually reverse past increases in the inflation rate. We could alternatively adopt (3.6) as a target criterion, and this would also be optimal. This criterion involves only the projected acceleration of inflation in period $t + 1$, not the absolute rate of inflation. However, the criterion is history-dependent because of its dependence on the value of $E_{t-1}x_t$.

that policy not be expected to be conducted later in accordance with a purely forward-looking target criterion. For example, Eggertsson and Woodford (2003) show that the expectation that the central bank will remain committed to the forward-looking pursuit of a low (time-invariant) inflation target — and hence will adjust interest rates so as to be consistent with the target as soon as this can be done without violation of the zero lower bound — can lead to a disastrous outcome when real disturbances result in a temporarily negative natural rate of interest. This analysis suggests that the problem of the Bank of Japan at present is not that it is not understood to be committed to a non-negative inflation target, but that it is expected to pursue its tacit inflation target in a purely forward-looking manner, with the implication that the (unwanted) price declines that occur while the zero bound constrains policy will never be undone.

A commitment to a time-invariant inflation target would be more likely to avoid the problem caused by the zero bound, of course, if the target were set several percentage points above zero, as advocated by Summers (1991). But this would result in substantial losses of another sort, those created by chronic inflation. The optimal policy rule, as Eggertsson and Woodford show, would instead involve commitment to a history-dependent target criterion, as a result of which temporary inflation would be created following a period in which the zero bound has constrained policy — and a greater amount of inflation the longer the time for which the zero bound has continued to bind, and the greater the cumulative deflation that has occurred during that time — while a low inflation rate would again be targeted once a sufficient period of time has passed in which the zero bound has not prevented the central bank from hitting its target. Credible commitment to a history-dependent policy of this kind would create the desired kind of expectations while the economy is in the “liquidity trap” — so that the deflation and output contraction at that time should remain quite modest — without requiring chronic inflation during normal times, and creating an “inflation scare” during the period in which the economy is reflat as it exits from the trap.

4 Conclusions

Inflation-forecast targeting, as currently practiced at central banks such as the Bank of England, represents an important innovation in decision procedures with regard

to monetary policy, which has moved the actual practice of leading central banks closer to the ideal that would be recommended on the basis of economic theory. The organization of the decision process around the achievement of an explicit, quantitative target that is also communicated to the public, and a commitment to the explanation of policy decisions to the public in terms that allow verification of the central bank's commitment to its putative target are important improvements upon prior procedures, that can both help to safeguard a central bank against the trap of discretionary policymaking, and help the private sector to more accurately anticipate future policy, increasing the effectiveness of policy. The introduction of *targeting rules* as a way of specifying policy commitments is also an important conceptual advance, allowing commitments to be stated in a way that incorporates a kind of flexibility that is of considerable practical value, while being specific about the aspects of policy that are most critical for anchoring private-sector expectations.

At the same time, current practice falls short of the theoretical ideal sketched in this paper in some notable respects. Perhaps the most important of these is the exclusive emphasis on “medium term” targets that leave unspecified the basis on which a particular nearer-term path toward that target is to be preferred. At best, this represents a significant degree of vagueness about the criterion that is actually used to make policy decisions. It may also indicate that the choice among alternative near-term paths for the economy is still made on a discretionary basis that will ensure suboptimal policy even when decisions are made by an omniscient monetary policy committee with a perfect understanding of social welfare. The question whether it would be practical for central banks to commit themselves to more explicit nearer-term target criteria, of the form indicated by the theory of optimal monetary policy rules, should be an important issue for further study by central bankers and monetary economists alike.

A An Optimal Targeting Rule for the Model of Giannoni and Woodford (2003)

Here I summarize the quantitative form of the optimal targeting rule derived by Giannoni and Woodford (2003), in the context of a small, empirical optimizing model of the U.S. monetary transmission mechanism. The desirability of this precise rule depends, of course, on the details of the quantitative model, many of which are highly debatable. Nonetheless, it may be useful to consider this example of an optimal policy rule for an estimated model, as an illustration of some of the general points made in the text about the likely character of an optimal policy rule.

The model of Giannoni and Woodford incorporates both wage and price stickiness, with random intervals between the times at which both individual wages and prices are reconsidered, as in the theoretical analysis of Erceg *et al.* (2000). In addition, both wages and prices are allowed to be indexed to the previous quarter's index of prices between the occasions on which they are re-optimized, as proposed by Christiano *et al.* (2001). The degrees of indexation of both wages and prices are treated as free parameters to be estimated, as in Smets and Wouters (2002a, 2002b), but our parameter estimates indicate the best fit under the assumption of full indexation of both wages and prices, as assumed by Christiano *et al.* Both wages and prices are also determined a quarter in advance. On the demand side of the model, the preferences of the representative household are assumed to allow for habit persistence, and the best fit is obtained when the habit-persistence coefficient takes the largest allowable value, so that utility depends on the *change* in real expenditure rather than its level. In addition, real private expenditure is determined two quarters in advance. The several free parameters of the model are estimated by minimizing the distance between the predicted impulse responses of four variables (output, inflation, the real wage, and the short-term nominal interest rate) to a monetary policy shock and those implied by an

unrestricted VAR model of the same four time series. The parameter estimates are consistent with the sign restrictions implied by theory, and several restricted versions of the model — a restricted model with no indexation to the lagged price index, a restricted model with flexible wages, and a restricted model with no habit persistence — can each be statistically rejected.

Here I summarize the implications for optimal policy of treating the best-fitting parameter values as representing the literal truth. First of all, the estimated model implies that maximization of the expected utility of the representative household corresponds to minimization of a quadratic loss function of the form

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\lambda_p (\pi_t - \gamma_p \pi_{t-1})^2 + \lambda_w (\pi_t^w - \gamma_w \pi_{t-1}^w)^2 + \lambda_x (x_t - \delta x_{t-1} - \hat{x}^*)^2 \right], \quad (\text{A.1})$$

where π_t is an index of goods price inflation between quarter $t - 1$ and quarter t , π_t^w is an index of wage inflation, and x_t is the output gap (log real output relative to a “natural rate of output” that varies in response to several types of real disturbances). The discount factor is calibrated to equal 0.99 (to imply a realistic long-run average real rate of return), while the model’s estimated parameters imply the values $\lambda_p = 0.9960$, $\lambda_w = 0.0040$, $\lambda_x = 0.0026$, and $\delta = 0.035$ for the coefficients of the loss function.

The fact that prices are indexed to a lagged price index implies that it is inflation acceleration, rather than the rate of inflation as such, that creates distortions, as in the simpler model discussed in the text. The fact that wages are also sticky implies that wage inflation also creates distortions, even when the rate of goods price inflation is stable; because wages are indexed to the lagged price index, it is actually wage inflation relative to lagged price inflation that measures this distortion. Finally, because of habit persistence, the distortions associated with fluctuations in the output gap are not proportional simply to a sum of squared deviations of the output gap each period from its optimal level, but rather to a sum of squared deviations of the output

gap from an increasing function of the previous quarter's output gap. However, the weight δ on the lagged output gap turns out to be quite small, despite the existence of substantial habit persistence.

We also find that the estimated parameter values imply a very small relative weight on the wage-inflation stabilization objective relative to the price-inflation stabilization objective. This is not because wages are found to be flexible, but because other estimated parameters imply larger distortions resulting from misalignment of prices than from misalignment of wages. The relative weight on the output-gap stabilization objective implied by the parameter estimates is also quite small; this follows directly from the estimation of parameters that imply only weak responses of wage and price inflation to variations in the output gap, as discussed in the text.

An optimal policy for the estimated model — and one with the desirable property that it is optimal *regardless* of the assumed statistical properties of the disturbances, and not solely in the case of disturbance processes of the kind implied by the estimated model for the historical sample period — can be implemented by a targeting procedure of the following kind.³³ First, in each quarter t , the central bank intervenes in the money markets (through open-market operations, repurchases, standing facilities in the interbank market for central-bank balances, etc.) so as to implement the interest-rate target $i_{t,t-1}$ announced in quarter $t - 1$. As in the simpler model discussed in the text, the fact that wages, prices and spending are all predetermined for a quarter implies that nothing can be gained from allowing variations in interest rates that are not forecastable in the previous quarter.

Second, in the quarter t decision cycle the bank must choose an operating target

³³Because the empirical model is quarterly, it is simplest to discuss the policy process as if a policy decision is also made once per quarter, even though in reality most central banks reconsider their operating targets for overnight interest rates somewhat more frequently than this. The discussion should not be taken to imply that it is optimal for the policy committee to meet only once per quarter; this would follow only if (as in the model) all other markets were also open only once per quarter.

$i_{t+1,t}$ to announce for the following quarter. This is chosen in order to imply a projected evolution of (wage and price) inflation from quarter $t + 1$ onward that satisfies a target criterion of the form

$$F_t(\pi) + \phi_w[F_t(w) - w_t] = \bar{\pi}_t, \quad (\text{A.2})$$

where $\bar{\pi}_t$ is a target value that has been determined in quarter $t - 1$. Here for each of the variables $z = \pi, w$, the expression $F_t(z)$ refers to a weighted average of forecasts of the variable z at various future horizons, conditional on information at date t ,

$$F_t(z) \equiv \sum_{k=1}^{\infty} \alpha_k^z E_t z_{t+k}, \quad (\text{A.3})$$

where the weights α_k^z sum to one. Thus the coefficient ϕ_w is actually the sum of the weights on real-wage forecasts at different horizons k . We observe that the target criterion can be thought of as a wage-adjusted inflation target.

Third, it is also necessary, as part of the quarter t decision cycle, for the central bank to choose the target $\bar{\pi}_{t+1}$ for the following quarter. This is chosen so as to ensure that future policy will be conducted in a way that allows the bank to project (conditional on its current information) that another target criterion, of the form

$$F_t^*(\pi) + \phi_w^* F_t^*(w) + \phi_x^* F_t^*(x) = \pi_t^*, \quad (\text{A.4})$$

should be satisfied, where the expressions $F_t^*(z)$ are again weighted averages of forecasts at different horizons (but with relative weights α_k^{z*} that may be different in this case), and π_t^* is another time-varying target value, once again a predetermined variable. In this case the criterion specifies a target for a wage- and output-adjusted inflation projection.

In this last procedure, optimality requires that the target value be given by an expression of the form

$$\pi_t^* = (1 - \theta_\pi^*)\pi^* + \theta_\pi^* F_{t-1}^1(\pi) + \theta_w^* F_{t-1}^1(w) + \theta_x^* F_{t-1}^1(x), \quad (\text{A.5})$$

where the expressions $F_t^1(z)$ are still other weighted averages of forecasts at different horizons, with relative weights α_k^{z1} that again sum to one, and π^* is an arbitrary constant.³⁴ Note that the optimal target value depends on the previous quarter's forecasts of the economy's subsequent evolution; this is an example of the history-dependence of optimal target criteria, discussed generally in the text.

The estimated parameter values imply the following numerical coefficients in the optimal target criteria. In the case of the short-term criterion (A.2), the coefficient ϕ_w is equal to 0.565.³⁵ Thus if unexpected developments in quarter t are projected to imply a higher future level of real wages than had previously been anticipated, policy must ensure that projected future price inflation is correspondingly reduced. This is because of a desire to stabilize (nominal) wage inflation as well as price inflation, and under circumstances of expected real wage growth, inflation must be curbed in order for nominal wage growth to not be even higher.

The relative weights that this criterion places on projections at different future horizons are shown in Figure 1. The two panels plot the coefficients α_k^π, α_k^w respectively, as functions of the horizon k . Note that the quarter for which the projections receive greatest weight is one quarter in the future, in each case. This is also the first quarter in which it is possible for wage or price inflation to be affected by the choice of $i_{t+1,t}$, according to the estimated model. However, while the real-wage projection that matters is primarily the projected growth in real wages between the present quarter and the next one, substantial weight is also placed on projected inflation farther in the future; in fact, the mean lead $\sum_k \alpha_k^\pi k$ is between 10 and 11 quarters in the future in the case of the inflation projection $F_t(\pi)$. Thus the short-run target criterion is

³⁴Note that in the model considered here, as in the simpler model discussed in the text, there is no welfare significance to any absolute inflation rate, only to changes in the rate of inflation, and to wage growth relative to prices. There is therefore no particular inflation rate that could be justified as optimal from a timeless perspective.

³⁵Here and below, the coefficients are presented for a target criterion where the inflation rate is measured in annualized percentage points.

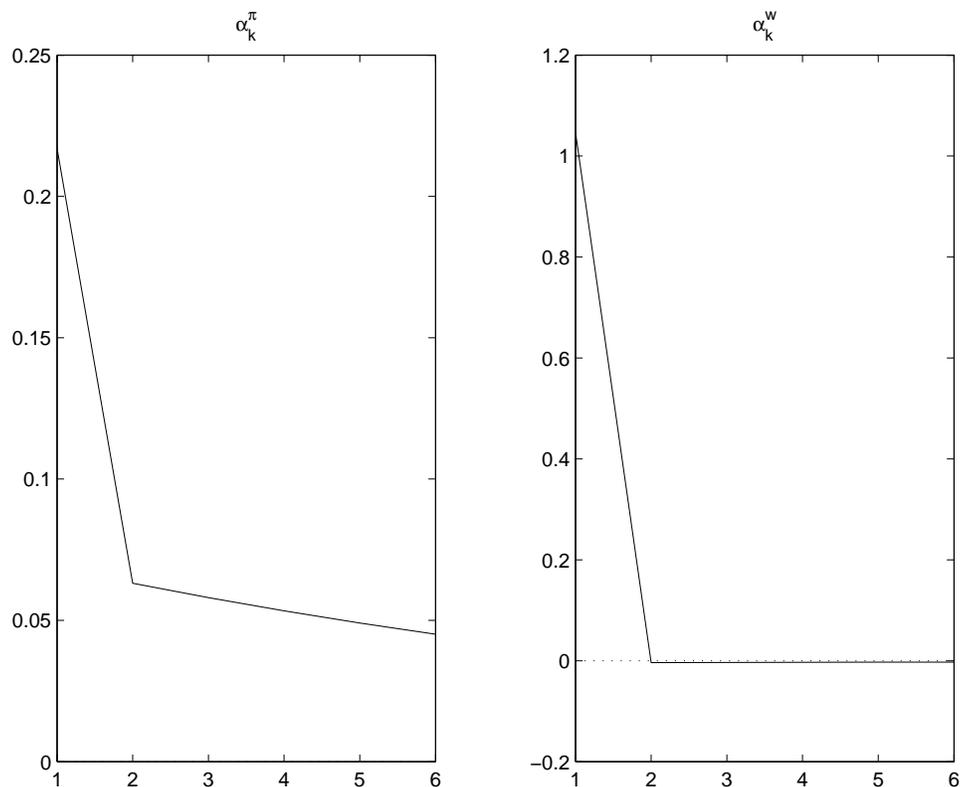


Figure A.1: Relative weights on projections at different horizons in the short-run target criterion (A.2). The horizontal axis indicates the horizon k in quarters.

a (time-varying) target for the average rate of inflation that is projected over the next several years, adjusted to take account of expected wage growth, mainly over the coming quarter. Roughly speaking, optimal policy requires the central bank to choose $E_t i_{t+1}$ in quarter t so as to head off any change in the projected average inflation rate over the next several years that is due to any developments not anticipated in quarter $t - 1$ (and hence reflected in the current target $\bar{\pi}_{t-1}$). This is a criterion in the spirit of inflation-forecast targeting as currently practiced at central banks such as the Bank of England, except that projected wage growth matters as well as price inflation, and that the target shifts over time.

In the case of the long-term criterion (A.4), instead, the numerical coefficients of

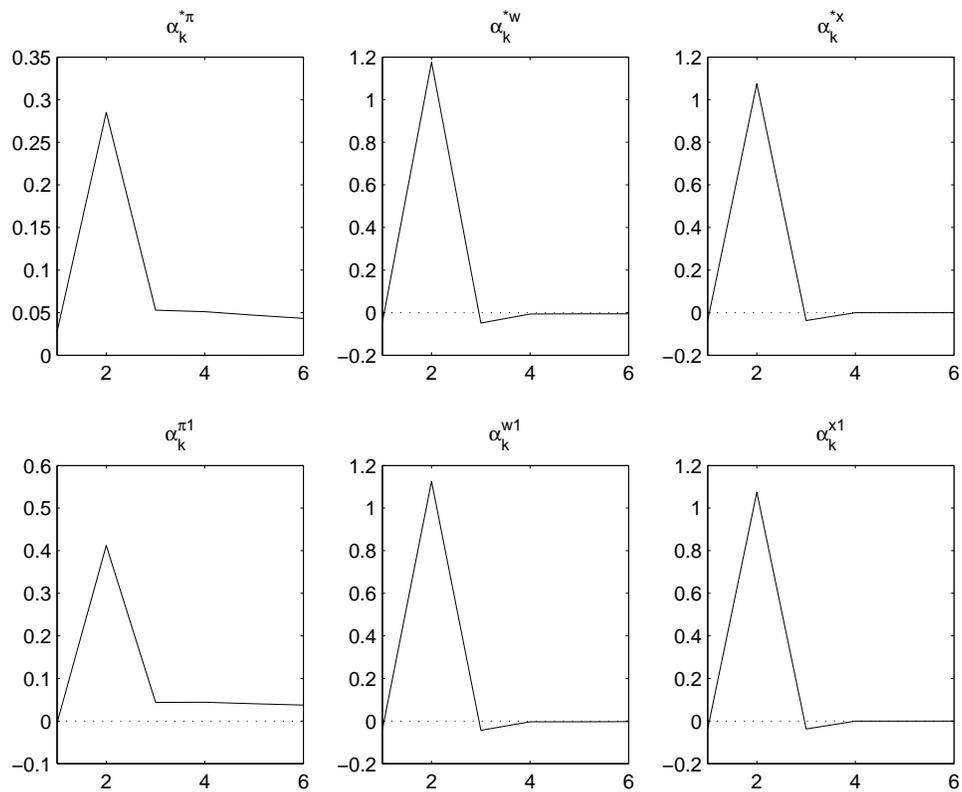


Figure A.2: Relative weights on projections at different horizons in the long-run target criterion. Panels in the first row indicate the projections in (A.4), while the second row indicates the projections from the previous quarter that define the target value π_t^* .

the target criterion are given by

$$\phi_w^* = 0.258, \quad \phi_x^* = 0.135.$$

In this case, output-gap projections matter as well; a higher projected future output gap will require a reduction in the projected future rate of inflation, just as will a higher projected future real wage. The numerical size of the weight placed on the output-gap projection may appear modest; but as we shall see in the next section, the degree of variability of output-gap projections in practice are likely to make this a quite significant correction to the path of the target criterion.

The relative weights on forecasts at different horizons in this criterion are plotted in the panels in the first row of Figure 2. We observe that in the case of this criterion, the projections that mainly matter are those for two quarters in the future; the criterion is nearly independent of projections regarding the quarter after the current one. Hence it makes sense to think of this criterion as the one that should determine the policy that the central bank plans on in periods two or more quarters in the future (and hence its choice in quarter t of the target $\bar{\pi}_{t+1}$ to constrain its choice in the following period of $i_{t+2,t+1}$), but not as a primary determinant of whether the bank's intended policy in period $t + 1$ is on track. The projections that receive the greatest weight under this criterion are those for the same quarter (quarter $t + 2$) that will receive the greatest weight in the targeting procedure for which $\bar{\pi}_{t+1}$ provides the target value.

Finally, the coefficients of the rule (A.5) determining the target value for the long-term criterion are given by

$$\theta_{\pi}^* = 0.580, \quad \theta_w^* = 0.252, \quad \theta_x^* = 0.125.$$

The weights in the projections (conditional on information in the previous quarter) at various horizons are plotted in the second row of Figure 2. Here too, it is primarily projections for two quarters in the future that matter in each case. Roughly speaking, then, the target value for the wage- and output-adjusted inflation projection two quarters in the future is high when a similar adjusted inflation projection (again for a time two quarters in the future) was high in the previous quarter.

Thus forecasting exercises, in which the central bank projects the evolution of both inflation and real variables many years into the future under alternative hypothetical policies on its own part, play a central role in a natural approach to the implementation of optimal policy. A forecast of inflation several years into the future is required in each (quarterly) decision cycle in order to check whether the intended

interest-rate operating target for the following quarter is consistent with the criterion (A.2). In addition, the time-varying medium-term inflation target $\bar{\pi}_t$ must be chosen each period on the basis of yet another forecasting exercise. While the long-run target criterion (A.4) primarily involves projections for a time only two quarters in the future, the choice of $\bar{\pi}_{t+1}$ requires that the central bank solve for a projected path of the economy in which (A.4) is satisfied not only in the current period, but in all future periods as well. Hence this exercise as well requires the construction of projected paths for inflation and real variables extending many years into the future. The relevant paths, however, will not be constant-interest-rate projections, but rather projections of the economy's future evolution given how policy is expected to evolve. Indeed, the projections are used to select constraints upon the bank's own actions in future decision cycles, by choosing both the interest-rate operating target $i_{t+1,t}$ and the adjusted inflation target $\bar{\pi}_{t+1}$ in period t .

References

- Altig, David, Lawrence J. Christiano, Martin S. Eichenbaum, and Jesper Linde, "Technology Shocks and Aggregate Fluctuations," unpublished, Federal Reserve Bank of Cleveland, 2002.
- Amato, Jeffery D., and Thomas Laubach, "Estimation and Control of an Optimization-Based Model with Sticky Prices and Wages," *Journal of Economic Dynamics and Control* 27: 1181-1215 (2003).
- Aoki, Kosuke, "Optimal Monetary Policy Responses to Relative Price Changes," *Journal of Monetary Economics* 48: 55-80 (2001).
- Auerbach, Alan J., and Maurice Obstfeld, "The Case for Open-Market Purchases in a Liquidity Trap," NBER working paper no. 9814, July 2003.
- Ball, Laurence, and Niamh Sheridan, "Does Inflation Targeting Matter?" in B.S. Bernanke and M. Woodford, eds., *Inflation Targeting*, Chicago: University of Chicago Press, forthcoming 2003.
- Ball, Laurence, N. Gregory Mankiw, and Ricardo Reis, "Monetary Policy for Inattentive Economies," NBER working paper no. 9491, February 2003.
- Barro, Robert J., and David B. Gordon, "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy* 91: 589-610 (1983).
- Benigno, Pierpaolo, "Optimal Monetary Policy in a Currency Area," *Journal of International Economics*, forthcoming 2003.
- Bernanke, Ben S., Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen, *Inflation Targeting*, Princeton: Princeton University Press (1999).
- Blinder, Alan, Charles Goodhart, Philipp Hildebrand, David Lipton, and Charles Wyplosz, *How Do Central Banks Talk?* Geneva Report on the World Economy no. 3, International Center for Monetary and Banking Studies, 2001.
- Bullard, James, and Kaushik Mitra, "Learning about Monetary Policy Rules," *Journal of Monetary Economics* 49: 1105-1129 (2002).
- Calvo, Guillermo, "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics* 12: 383-98 (1983).

- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," NBER Working Paper no. 8403, July 2001.
- Eggertsson, Gauti B., and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity* 2003:1, pp. 139-233.
- Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin, "Optimal Monetary Policy with Staggered Wage and Price Contracts," *Journal of Monetary Economics* 46: 281-313 (2000).
- Friedman, Milton, "The Optimum Quantity of Money," in *The Optimum Quantity of Money and Other Essays*, Chicago: Aldine, 1969.
- Fuhrer, Jeffrey C., and Geoffrey R. Moore, "Inflation Persistence," *Quarterly Journal of Economics* 110: 127-159 (1995).
- Giannoni, Marc P. and Michael Woodford, "Optimal Interest-Rate Rules: I. General Theory," NBER working paper no. 9419, December 2002.
- — and — —, "Optimal Inflation Targeting Rules," in B.S. Bernanke and M. Woodford, eds., *Inflation Targeting*, Chicago: University of Chicago Press, 2003.
- 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.
- Goodfriend, Marvin, "Inflation Targeting in the United States?" in B.S. Bernanke and M. Woodford, eds., *Inflation Targeting*, Chicago: University of Chicago Press, forthcoming 2003.
- —, and Robert G. King, "The New Neoclassical Synthesis and the Role of Monetary Policy," *NBER Macroeconomics Annual* 12: 231-283 (1997).
- Heikensten, Lars, "The Riksbank's Inflation Target – Classification and Evaluation," Bank of Sweden *Quarterly Review* 1/1999, pp. 5-17.
- Honkapohja, Seppo, and Kaushik Mitra, "Problems in Inflation Targeting Based on Constant Interest Rate Projections," unpublished, University of Helsinki, August 2003.

- Khan, Aubhik, Robert G. King, and Alexander L. Wolman, "Optimal Monetary Policy," NBER working paper no. 9402, December 2002.
- King, Mervyn, "What Has Inflation Targeting Achieved?" in B.S. Bernanke and M. Woodford, eds., *Inflation Targeting*, Chicago: University of Chicago Press, forthcoming 2003.
- King, Robert G., and Alexander L. Wolman, "What Should the Monetary Authority Do when Prices are Sticky?" in J.B. Taylor, ed., *Monetary Policy Rules*, Chicago: University of Chicago Press, 1999.
- Krugman, Paul, "It's Baaack! Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity* 1998-2, pp. 137-187.
- Kydland, Finn E., and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85: 473-491 (1977).
- Lange, Joe, Brian Sack, and William Whitesell, "Anticipations of Monetary Policy in Financial Markets," Federal Reserve Board, FEDS paper no. 2001-24, April 2001.
- Leiderman, Leonardo, and Lars E.O. Svensson, eds. *Inflation Targets*, London: Centre for Economic Policy Research, 1995.
- Leitemo, Kai, "Targeting Inflation by Constant-Interest-Rate Forecasts," *Journal of Money, Credit and Banking*, forthcoming 2003.
- Orphanides, Athanasios, and John C. Williams, "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in B.S. Bernanke and M. Woodford, eds., *Inflation Targeting*, Chicago: University of Chicago Press, forthcoming 2003.
- Preston, Bruce, "Learning about Monetary Policy Rules when Long-Horizon Forecasts Matter," unpublished, Princeton University, August 2002.
- Rotemberg, Julio J., and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual* 12: 297-346 (1997).
- Sargent, Thomas J., and Neil Wallace, "'Rational' Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule," *Journal of Political Economy* 83: 241-254 (1975).

- Smets, Frank, and Raf Wouters, “Monetary Policy in an Estimated Stochastic Dynamic General Equilibrium Model of the Euro Area,” unpublished, European Central Bank, May 2002a.
- — and — —, “Sources of Business Cycle Fluctuations in the U.S.: A Bayesian DSGE Approach,” seminar presentation, Princeton University, November 1, 2002b.
- Svensson, Lars E.O., “Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets,” *European Economic Review* 41: 1111-1146 (1997).
- —, “Inflation Targeting as a Monetary Policy Rule,” *Journal of Monetary Economics* 43: 607-654 (1999).
- —, “What is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules,” *Journal of Economic Literature*, June 2003a.
- —, “The Inflation Forecast and the Loss Function,” in P. Mizen, ed., *Central Banking, Monetary Theory and Practice: Essays in Honour of Charles Goodhart*, vol. 1, Edward Elgar, 2003b.
- — and Michael Woodford, “Implementing Optimal Policy through Inflation-Forecast Targeting,” in B.S. Bernanke and M. Woodford, eds., *Inflation Targeting*, Chicago: Univ. of Chicago Press, forthcoming 2003.
- Vickers, John, “Inflation Targeting in Practice: The U.K. Experience,” *Bank of England Quarterly Bulletin*, November 1998.
- Woodford, Michael, “The Optimum Quantity of Money,” in B.M. Friedman and F.H. Hahn, eds. *Handbook of Monetary Economics*, vol. II, Amsterdam: North-Holland, 1990.
- —, “Optimal Monetary Policy Inertia,” NBER Working Paper no. 7261, August 1999a.
- —, “Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?” in *New Challenges for Monetary Policy*, Kansas City: Federal Reserve Bank of Kansas City, 1999b.
- Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.