

Principled Policymaking in an Uncertain World

Michael Woodford
Columbia University

Revised July 2012

Revised text of a presentation at the Conference on Microfoundations for Modern Macroeconomics, Columbia University, November 2010. I would like to thank Amar Bhidé, Roman Frydman, and Andy Haldane for helpful comments, and the Institute for New Economic Thinking for research support.

Rule-Based Policy or Discretion?

A crucial legacy of Phelps et al (1970) has been recognition of the importance of economic agents' *anticipations* as a determinant of macroeconomic outcomes. This has had many profound consequences for macroeconomic analysis. Among them is the fact that the subsequent theoretical literature on monetary policy has focused on the analysis of *monetary policy rules*, rather than on decisions about individual policy *actions*. The present essay considers the reasons for this development, and the extent to which such a focus continues to be appropriate in the light of subsequent events --- changes in central banks' approach to monetary policy in the decades since the publication of the Phelps volume, and even more crucially the reconsideration of macroeconomic theory and policy that is necessary in the wake of the global financial crisis.

A. Policy Rules as an Object of Study

There are at least two important reasons why recognition of the importance of expectations led to an emphasis on policy rules in the theoretical literature.¹ First of all, if one is to specify agents' anticipations in one's economic model using the common hypothesis of *rational expectations* (RE), one *can't answer* questions about the predicted effect of a given policy action, even in the context of a particular economic model, except if one specifies *expected future policy* as well, over an indefinite future and under all potential future contingencies. One can't solve for the consequences of a given action (say, purchases of a certain quantity of Treasury securities by the Fed this month) without specifying what people expect about future outcomes (for example, future inflation), both with and without the action in question. Under the hypothesis of "rational" (or, more properly, model-consistent) expectations, what people expect in either case should be what one's model predicts will occur; but that will depend on what is assumed about future policy, and in what respects it does or does not change as a result of the policy change that one wishes to analyze.

Hence the object of analysis must always be a complete specification of current and future policy, including a specification of how policy can be expected to respond to all possible future developments; in other words, the only possible object of analysis is a complete policy *strategy*.² This does not mean that such an approach cannot be used to analyze the consequences of approaches to policy other than ones under which policymakers consciously follow a rule; the analyst might postulate a systematic pattern of conduct --- which it is furthermore assumed that the public should also be able to predict --- even if it is neither announced nor consciously formulated by policymakers themselves. But once the object of study is defined as the comparative advantages of alternative systematic patterns of conduct, the only goal of *normative* policy analysis must be to propose a systematic pattern of conduct that it would be desirable for the policy authority to follow, in a sufficiently faithful way for the pattern to be

¹ See Woodford (2003, chapter 1, section 2) for further discussion.

² See, for example, Sargent (1993) for discussion of this point.

predictable. In other words, even if positive analyses of policy in particular times and places do not necessarily assume that policymakers consciously follow a rule, a normative analysis would necessarily recommend a rule that should be followed systematically, rather than an individual action that is appropriate to some particular situation.

There is a second reason for the recent literature's focus on policy rules, under which a "rule" has a more specific meaning, namely, a prescription that constrains the policymaker to behave in a way other than the way that would be judged desirable under a sequential optimization procedure of the kind criticized by Kydland and Prescott (1977). Kydland and Prescott criticize "discretionary" policy, by which they mean sequential optimization each time an action must be taken, under no constraints resulting from prior commitments. This *sequence* of optimizing decisions --- under which no more is decided at any one time than is necessary to determine the action that must be taken at that time --- can be contrasted with an alternative form of optimization, under which an overall pattern of conduct that will be followed forever after is chosen once and for all. Note that the complaint about "discretionary" policy is not that it is not systematic or predictable --- as conceived by Kydland and Prescott, it involves a clear objective that is pursued consistently over time, and in their analysis of the consequences of such behavior, they assume (following RE methodology) that it is completely predictable.

The Kydland-Prescott critique of such a sequential approach to decisionmaking is rather that it fails to internalize the consequences of people's anticipation of systematic patterns in the policymaker's conduct. At each time that a decision must be made about a specific action (say, the level at which the federal funds rate should be maintained for the next six weeks), people's prior expectations about that action are a fact about the past that can no longer be affected, and so an analysis of the consequences of the action for the policymaker's objectives assumes no possibility of influencing those expectations --- even though a different systematic approach to choice regarding this action could have given people a reason to have had different expectations, and shaping expectations is relevant to the achievement of the policymaker's own objectives. In general, a superior outcome can be achieved (under the RE analysis) through *commitment* by the policy authority to behave in a systematically different way than a discretionary policymaker would wish to behave *ex post*; this requires commitment to follow a *rule*. The key feature of the Kydland-Prescott conception of a policy rule is thus the element of advance commitment, which is contrasted with ad hoc decisionmaking at the time when action is necessary.

Even a brief review of these familiar arguments raises an important question. Does not a recognition of the possibility (indeed, the inevitability, eventually) of non-routine change undermine the desirability of commitment to a policy rule? In a theoretical exposition of the advantages of policy commitment --- such as the examples presented by Kydland and Prescott --- it is easy to assume that the possible future states in which the policymaker may find herself can be enumerated in advance, and that a commitment can be chosen *ex ante* that specifies what will be done in each state if it is reached. In practice, this will not be possible, for reasons that go beyond a mere assertion that the number of possible future states is very large (the elements of some infinite-dimensional space). There are often developments that are not simply elements in a large space of possibilities *the dimensions of which* were conceptualized in advance, but that instead were inconceivable previously --- so that policymakers are confronted not simply with the question whether it is now desirable to behave differently than they were expected to

behave in such a situation, but with a need to think afresh about a type of situation to which they have given little prior thought.³ The experience of policymakers after the unexpected eruption of the global financial crisis in the summer of 2007 underlines the relevance of this possibility, if further proof were needed.

It is fairly obvious that the existence of non-routine change of this sort undermines the desirability of *a certain conception* of a policy rule: one where a rule is understood to mean a fully explicit formula that prescribes a precise action for any possible circumstance. Nonetheless, it does little to reduce the relevance of the considerations mentioned above as reasons that the literature of recent decades on monetary policy has focused on the evaluation of policy rules. It does not eliminate the need to assess policy *strategies*, rather than individual decisions considered in isolation, even if such strategies cannot realistically be supposed to represent complete specifications of behavior in all possible future circumstances. Nor does it eliminate the potential benefits from requiring policy decisions to be based on general *principles*, rather than making an ad hoc decision about what will achieve the best outcome under current circumstances.

B. Policy Analysis without Rational Expectations

Strategies and principles would be irrelevant only if one were to view decisionmakers as responding mechanically to the current economic environment, and not on the basis of anticipations that can be influenced in any way by the announced policy commitments of a central bank --- that is, only if one were to deny the relevance of the “modern” turn advocated by Phelps. In fact, a variety of approaches to dynamic economic analysis have been proposed that still allow a role for anticipations that should take into account what is known about central-bank policy commitments, without imposing the strong form of expectational coordination implied by the postulate of rational expectations.

One example is the concept of “calculation equilibrium” proposed by Evans and Ramey (1992). Evans and Ramey propose that individuals make decisions that are optimal under a particular anticipated future evolution of the economy (extending, in principle, indefinitely into the future); they also propose that individuals possess a correct model of the economy, in the sense that they are able to correctly predict the evolution of the variables that they wish to forecast under a particular conjecture about *the way that others expect the economy to evolve*. People’s expectations can then be disciplined by requiring them to result from a calculation using the economic model, starting from an expectation about others’ expectations. They relax, however, the RE assumption that everyone must forecast a future evolution that is predicted by the commonly agreed-upon model under the assumption that others predict precisely that same evolution. Instead, they propose that individuals start with some initial conjecture about the future path of economic variables, and progressively refine this forecast by calculating (at each stage in an iterative process) the evolution that should be forecasted if *others* are

³ On the importance for policy analysis of confronting the occurrence of non-routine change, see Frydman and Goldberg (2011).

expected to forecast using the output of the previous stage's calculation. (The thought process that this involves is like the one described by Keynes, 1936, in his famous analysis of the "beauty contest.")

If this iterative calculation were pursued to the point of convergence⁴ --- so that a forecast were eventually obtained with the property that expecting others to forecast that way would lead to the same forecast --- the resulting forecast would correspond to an RE equilibrium of the model used by decisionmakers. (Of course, this would still only be an equilibrium *relative to the model* that they happen to believe in, since the iterative calculation is merely a check on the *internal* consistency of their forecasting, and not a proof that it must correctly describe how the world will actually evolve. Thus such a conception of how people forecast could still allow for surprises, at which points in time there might be an abrupt change in the model that people believe and hence in the way that they forecast.) But Evans and Ramey assume instead (like Keynes) that in practice decisionmakers will truncate such calculations after a finite number of iterations; they propose that calculation costs limit the number of iterations that it is reasonable for a decisionmaker to undertake (and propose a particular stopping rule, that need not concern us). Given the truncation of the expectation calculations, dynamic phenomena are possible --- even assuming that people's model of the economy is actually correct --- that would not occur under an RE analysis. These include asset "bubbles" that last for a period of time (though not indefinitely), and are sustained by beliefs that are consistent with the economic model, under a belief about others' beliefs that is also consistent with others' understanding of the model, and so on *for a finite number of iterations* --- but that ultimately depend on higher-order beliefs that will be disconfirmed.

The "eductive stability" analysis proposed by Guesnerie (2005) similarly assumes that individuals make decisions that are optimal under a particular anticipated future evolution of the economy, and that they each possess a correct model of the economy. It further imposes the stronger restriction that both of these things are *common knowledge* in the sense that that term is used in game theory: each individual's beliefs are consistent with knowledge that all others know that all others know [and so on *ad infinitum*] that these things are true. Nonetheless, as Guesnerie stresses, only under rather special circumstances are RE beliefs the only ones consistent with such a postulate. (It is in this case that Guesnerie refers to the REE as "eductively stable," and hence a reasonable prediction of one's model.) Under more general circumstances, he proposes that one should consider the entire set of possible paths for the economy that can be supported by beliefs consistent with common knowledge of rationality (the analog of the "rationalizable" outcomes considered by Bernheim, 1984, and Pearce, 1984). This includes paths along which fluctuations in asset prices occur that are sustained purely by changing conjectures about how others will value the assets in the future --- conjectures that must be consistent with similar rationalizability of the conjectured future beliefs. Guesnerie proposes that policies should be selected with an eye on the entire set of rationalizable outcomes associated with a given policy; for example, it may be desirable to eliminate the risk of fluctuations due to arbitrary

⁴ This assumes that the process *would* converge, if pursued far enough. In the examples considered by Evans and Ramey, this is the case, and their interest is in the alternative forecasts that remain possible when the calculation is instead truncated after a finite number of iterations. But such an algorithm need not converge at all, nor need there be a unique limiting forecast independent of the initial conjecture, as Guesnerie (2005) emphasizes.

changes in expectations, by choosing a policy under which a unique REE is “eductively stable” --- but this is a criterion for policy design, rather than something that can be taken for granted.

Under the approach proposed by Woodford (2010), a given policy is again associated with an entire set of possible outcomes, rather than with a unique prediction, and it is argued that one should seek a policy which ensures the greatest possible lower bound for the average level of welfare, over the set of outcomes associated with the policy. The set of possible outcomes corresponds to a set of possible (not perfectly model-consistent) beliefs about the economy’s future evolution that people may entertain. Under this approach, however, the set of possible beliefs that are to be entertained is disciplined not by a requirement that the evolution in question be rationalizable under a *theory of others’ behavior* (more generally, be consistent with knowledge of the correct model of the economy), by rather by a requirement that subjective beliefs not be grossly out of line with actual probabilities --- an assumption of “near-rational expectations.” For example, events that occur with complete certainty (according to the policy analyst’s model) are assumed to be correctly anticipated, though events that occur with probabilities strictly between zero and one may be assigned somewhat incorrect probabilities; a parameter (that indexes the analyst’s degree of concern for robustness of policy to departures from model-consistent expectations) determines how large of discrepancies between subjective and model-implied probabilities⁵ are to be contemplated. This approach requires policymakers to contemplate equilibrium outcomes that differ from the REE prediction to a greater or lesser extent, depending on policy and other aspects of the economic structure. For example, for a given value of the robustness parameter, equilibrium valuations of long-lived risky assets can depart to a greater extent from their “fundamental” (REE) value when the short-term riskless rate of return is lower, so that the anticipated future sale price of the asset accounts for a larger share of its current valuation.

Each of these concepts assumes less perfect coordination of expectations than does the hypothesis of RE, and so may provide a more plausible basis for policy analysis following structural change. Yet in each case, the central bank’s commitments regarding future policy will influence the set of possible subjective forecasts consistent with the hypothesis. Under the proposals of Evans and Ramey (1992) or of Guesnerie (2005), this is because the mapping from given conjectures about others’ forecasts to what one should oneself forecast (using the model of the economy) is influenced by the central bank’s public commitments regarding its conduct of policy; under the proposal of Woodford (2010), this is because the degree of discrepancy between given subjective beliefs and model-consistent beliefs will depend on policy commitments. Hence under any of these approaches, a comparative evaluation of alternative monetary policies will require a specification of the entire (state-contingent) future path of policy, and not simply a current action, just as in the case of REE analysis. Similarly, there will be potential benefits from commitment relative to the outcome under discretionary policy. Indeed, Woodford (2010) finds that when the policymaker wishes to choose a policy that is robust to departures from fully model-consistent expectations, the advantages of commitment over discretionary policy are even greater than when one assumes that agents in the economy will necessarily have RE.

⁵ Note that what is relevant is the discrepancy between the subjective beliefs and what the model predicts should happen *if people hold those beliefs*, and *not* the discrepancy between subjective beliefs and REE beliefs. These may be quite different, if the model’s prediction for the economy’s evolution is highly sensitive to subjective beliefs.

Yet regardless of the degree of expectational coordination that is assumed, is it not still unreasonable for a central bank to commit itself in advance to a rule, chosen under one view of what possible future contingencies may arise, despite the fact that the future situations may well arise that were not contemplated at all?

I believe that the argument that rule-based policymaking is necessarily foolhardy in a world where non-routine change occurs depends on too narrow a conception of what is involved in following a rule. In particular, it is important to recognize that there are several different *levels* at which it is possible to describe the process through which policy decisions are made. Judgment may be exercised in the application of policy to particular circumstances, at a more concrete level of description of the policy, even though the judgment is used to determine the implications of a rule, at a more general level of description, that has been stated explicitly in advance. The judgment that is exercised at the more concrete level will appropriately take account of unforeseen changes in circumstances, while the rule that has been formulated at the more general level might nonetheless remain applicable despite the occurrence of the novel situation.

I illustrate this idea with a more detailed discussion of possible types of commitments in the case of monetary policy.

I. Alternative Levels of Policy Commitment: The Case of Monetary Policy

One might imagine a rule for the conduct of monetary policy being specified at any of four distinct levels of description of the policy in question.⁶ These involve increasing degrees of abstraction as one proceeds to “higher-level” descriptions.

The lowest level is what I shall call the *operational* level. At the most concrete level of description, monetary policy (under routine conditions, rather than those during the recent crisis) involves a decision about a quantity of bank reserves to inject or withdraw each day, typically through open-market purchases or repo transactions. One might imagine that a monetary policy rule should be a specific formula that would tell the Trading Desk of the Federal Reserve Bank of New York (or the corresponding branch of another central bank) which trades to execute each day, as a function of various observable conditions. McCallum (1988, 1999) argues for a policy rule that is “operational” in this sense, and so proposes rules that specify equations for the adjustment of the monetary base.

The literature on monetary policy rules has instead often discussed specifications at a higher level, that I shall call the *instrument* level. At most central banks, the key decision of the policy committee (again, under routine conditions) is the choice of an operating target for a particular overnight interest rate --- the federal funds rate, under the operating procedures of the Federal Reserve since at least the mid-1980s. The decision about how to achieve this target through market trades is then delegated to staff members with these operational responsibilities, or is at any rate determined without having to convene the central’s policy committee (i.e., the committee that chooses the

⁶The first three of these levels are distinguished in Woodford (2007, sec. 1), which also discusses the possible specification of policy rules at the different levels.

operating target for the instrument of policy: the Federal Open Market Committee, in the case of the US).⁷ One might imagine that a monetary policy rule should be a specific formula that determines what the correct target for the federal funds rate should be at each point in time, as a function of observable variables. The celebrated “Taylor rule” (Taylor, 1993) is of this form, and so are most empirical characterizations of policy through estimation of a central-bank reaction function, as well as most of the normative proposals considered in the theoretical literature.

A still higher-level description of policy is possible, however, at least in the case of those central banks which base their policy decisions on a clear intellectual framework, that remains constant over the course of many meetings of the policy committee, and that is articulated with some degree of explicitness in the public communications of the central bank; this can be referred to as the *policy targets* level.⁸ A central bank may determine the correct instrument setting (i.e., the operating target for the policy rate) at each meeting of the policy committee on the basis of previously specified targets for other macroeconomic variables that are expected to be indirectly influenced by the path of the policy rate.

In particular, a *forecast targeting* regime (Svensson, 1999, 2005; Woodford, 2007) involves choosing a target for the policy rate at each meeting that is consistent with the anticipated path for the policy rate that is required (according to the policy committee’s analysis) in order for the economy’s projected evolution to satisfy a quantitative *target criterion*. A policy rule might be specified by announcing the particular target criterion that will guide such deliberations; Svensson calls such a prescription a “targeting rule” (as opposed to an “instrument rule” like the Taylor rule). This is in fact the level at which central banks have most been willing to commit themselves to explicit criteria for the conduct of policy. Many central banks now have explicit quantitative targets for some measure of medium-run inflation; a few have also been fairly explicit about the criteria used to judge that nearer-term economic projections are acceptable, and about the way in which projections for variables other than inflation are taken into account (Qvigstad, 2006).⁹

Finally, it is possible, at least in principle, for a central bank’s policy commitments to be formulated at a still higher level, that I shall call the *policy design* level. At this level, one would specify the principles on the basis of which policy targets are chosen, given a particular model of the way that monetary policy affects the economy. A commitment to specified principles at this level could be maintained in the face of a change in either the structure of the economy or in policymakers’ understanding of that structure, even though it might well be appropriate to modify a central bank’s policy targets in the light of such change. I do not think that any central banks have yet made explicit statements committing themselves to principles of policy design at this level of abstraction. But the formulation of useful principles at this level has been a goal of at least a part of the research literature on normative monetary policy, and I believe that the quest for useful principles at this level becomes more important the more seriously one takes the likelihood of non-routine change.

⁷For the distinction between instrument choice and the decisions involved in implementation of that decision, see for example Friedman and Kuttner (2011).

⁸The distinction between policy prescriptions that are specified at the instrument level (“instrument rules”) and those specified at the policy targets level (“targeting rules”) has been stressed in particular by Svensson (2003).

⁹See Woodford (2007, sec. 2) for further discussion of the Norges Bank procedures as a particularly explicit example of a forecast targeting approach.

A. At Which Level of Specification is Policy Commitment Appropriate?

Note that these four distinct levels are mutually compatible ways of describing a central bank's policy; the same policy might simultaneously be correctly described at each of these levels. Hence when one contrasts possible specifications of monetary policy "rules" of these four distinct types, one is not necessarily talking about policies that are different, in terms of the actions that they would require a central bank to take under particular circumstances. But the levels of description differ in the degree to which it is useful to imagine specifying a rule for policy in advance.

At each successively lower level of the specification, one comes closer to saying precisely what the central bank ultimately must do. At each lower level, finer institutional details about the precise mechanism through which monetary policy affects the economy become relevant. And finally, at each lower level, it is appropriate for the central bank to be prepared to adjust course more frequently on the basis of more recent information. In practice, decisions are reviewed more frequently, the lower the level. For example, in the case of the Federal Reserve, decisions at the operational level are adjusted daily, and sometimes more often, during periods of particular market turmoil; instead, decisions at the instrument level are scheduled for review only 8 times a year, though occasionally inter-meeting changes in the funds rate target are judged necessary. The policy committees of inflation-targeting central banks reconsider the appropriateness of the planned path for the policy rate 8 or 12 times per year, but the inflation target itself remains unchanged for years. Yet even inflation targets change from time to time; for example, the UK changed the Bank of England's official target at the end of 2003, and the ECB slightly changed its "definition of price stability" after a review of its monetary policy strategy in 2003. The Bank of Canada's inflation target has been modified several times since its introduction in 1991, and is reviewed at five-year intervals. And surely the possibility of such changes in the light of changing knowledge is entirely appropriate.

The degree to which it is either possible or useful to articulate the principles on the basis of which decisions are to be made also differs greatly depending on the level of specification. I think that few monetary economists or central bankers, even among those who are strong proponents of rule-based policy and of central-bank transparency, would argue that there is a need for explicit policy commitments at the *operational* level. The literature that compares the consequences of alternative policy rules generally takes it as given that any non-negative target for the policy rate can be implemented with a high degree of accuracy over time scales (a day or two) that are quite short compared to those that matter for the effects of interest rates on the basis of which the policy is to be judged, and that the details of the required open-market operations have few if any consequences for the objectives of policy. Moreover, while in principle the same policy prescription (say, adherence to the Taylor rule) should have an equivalent formulation at the operational level, it would be complex to describe this in detail (by which I mean, to give a precise algorithm that would allow the correct operational decision to be computed under all possible circumstances). A simplified description at the operational level might instead be practical, but then, if this is regarded as an actual commitment about how policy will be conducted, it would be less successful at achieving the desired outcome with regard to the state-contingent evolution of the policy instrument, and hence less successful at achieving the

central bank's higher-level stabilization objectives as well. Hence insistence on an "operational" policy commitment would have clear costs.

Deviations from a bank's routine approach to the implementation of its interest-rate target are often necessary at times of crisis, as increased uncertainty leads to a sudden increase in demands for liquidity. For example, Sundaesan and Wang (2009) describe the special measures introduced by the Fed to deal with unusual liquidity needs around the time of the millennium date change (the so-called "Y2K" scare), in such a way as to minimize the consequences of this unusual behavior for money-market interest rates. An inability to respond in this way, owing to the existence of a rigid policy commitment at the operational level, would likely have meant greater disruption of financial markets. At the same time, the benefits of so low-level a commitment seem minimal. The most important argument for the desirability of a lower-level commitment is that accountability of the central bank to the public is increased by specifying exactly what must be done in terms that can be verified by outside observers. But one cannot really say that a commitment at the instrument level, without specifying in advance the precise operational decisions that this will require, reduces accountability to any significant extent, given central banks' degree of success at achieving their interest-rate targets over quite short time horizons in practice.

It is less obvious that a description of policy at a level of abstraction higher than the instrument level should suffice, and indeed the literature on the quantitative evaluation of policy rules has almost exclusively focused on rules specified as formulas to determine the value of a policy instrument that is under the relatively direct control of the central bank.¹⁰ Nonetheless, I think there are important advantages to considering rules that are specified by target criteria that need not involve any variable over which the central bank has direct control.

A first question is whether a mere specification of a target criterion suffices to *fully determine* outcomes under the policy, so that one can compare the outcomes associated with alternative policies. This is undoubtedly an issue of practical relevance. For example, a criterion that only involves projected outcomes two or more years in the future (as is true of the explicit commitments of many inflation-targeting central banks) is one that is unlikely to imply a determinate solution; there will be alternative paths by which the economy could reach a situation consistent with the criterion, and in such a case the target criterion fails to fully determine policy. In my view, it is important to adopt a target criterion that *does* fully determine (but not over-determine) a particular equilibrium. But this is a property that one can analyze given a specification of the target criterion alone; one need not specify the policy at a lower level in order to check this. Giannoni and Woodford (2010) illustrates how this kind of calculation can be undertaken under the assumption of rational expectations, using a structural model of the economy that specifies the constraints on feasible equilibrium paths of the target variables, but that need not even include the additional model equations that would be needed to determine the required evolution of the central bank's policy instrument; this paper also describes a general approach to the derivation of target criteria which guarantees, among other desiderata, that the target criterion necessarily determines a unique bounded REE. And one should recall that there is also a question whether a given interest-rate feedback rule determines a unique REE or not; one argument for the importance of choosing a rule that conforms to the "Taylor Principle" is that in many models, rules with weaker

¹⁰ See, for example, the review of this literature by Taylor and Williams (2011).

feedback from realized inflation to the interest-rate operating target have been found to result in indeterminacy of equilibrium (e.g., Woodford, 2003, chap. 4).

It is true that the literature on this topic typically assumes rational expectations, and one might wonder instead how precisely the predicted evolution of variables such as inflation and real activity is pinned down if one admits that people in the economy may not all anticipate precisely the evolution that the policy analyst's own model predicts. I believe that a consideration of this issue is another important part of an analysis of the desirability of a proposed target criterion. But this question as well can be analyzed without any need to specify the associated evolution of the central bank's interest-rate instrument, as illustrated by Woodford (2010). Moreover, specification of a policy commitment at the level of an instrument rule (or central-bank reaction function), rather than through a target criterion, only increases the degree of uncertainty about the equilibrium outcome that results from doubts about whether people in the economy will have model-consistent expectations. This is because the relation between the interest-rate reaction function and the evolution of the variables of interest (the "target variables" in terms of which the central bank's stabilization objectives are expressed) is more indirect than the relation between the target criterion and the paths of these variables, and the number of ways in which departures from model-consistent expectations can skew the outcome implied by the policy is correspondingly larger.

The analysis by Evans and Honkapohja (2003) illustrates this point, though they do not express the matter in this way, in the context of a standard New Keynesian model (the one analyzed under the assumption of rational expectations in Clarida *et al.*, 1999) and a particular hypothesized alternative to rational expectations, namely least-squares learning dynamics. Evans and Honkapohja compare predicted outcomes under the learning dynamics under two different policy specifications that, under the assumption of rational expectations, would be regarded as equivalent. (Both imply a determinate REE, and the RE equilibrium evolution of the endogenous variables that each determines is the same.) The two rules are each described in their paper as interest-rate reaction functions. However, the one that they call the "expectations-based policy rule" is the equation for the interest-rate instrument setting that they obtain by inverting their model's structural equations to determine the interest rate --- as a function of *observed private-sector expectations*, whether those correspond to the model-consistent expectations or not --- that would imply that the evolution of inflation and the output gap satisfy a particular target criterion (a type of "flexible inflation target") that could be expressed in terms of those two variables alone. Systematic adherence to this rule is equivalent to commitment to the target criterion, rather than to any particular equation specifying the instrument as a function of "objective" factors without reference to people's expectations. The alternative rule that they consider (the "fundamentals-based policy rule") is instead a formula for the instrument setting as a function of the exogenous state of the world at each point in time. The two rules are chosen so that they would both determine precisely the same equilibrium evolution for the economy, under the assumption of rational expectations. Yet Evans and Honkapohja show that under least-squares learning dynamics, a commitment to the target criterion (i.e., "expectations-based policy") leads to convergence to the REE, while commitment to the instrument rule results in unstable dynamics.

A second question is whether specification of a target criterion, rather than a reaction function for the instrument, is a useful way of providing a *guideline for policymakers* in their deliberations. Of course, a monetary policy committee has to decide on the level of overnight interest rates, so the target

criterion alone does not provide them sufficient information to discharge their duty. Nonetheless, a target criterion relating the paths of some of the variables that the policy committee wishes to stabilize seems the appropriate level of detail for a prescription that a policy committee can agree to use to structure its discussions, that can be explained to new members of the committee, and that can ensure some degree of continuity in policy over time. Special factors are likely to be important at each meeting, in deciding upon the level of interest rates consistent with fulfillment of the target criterion; hence it is difficult to impose too much structure on this kind of deliberation, without the committee members feeling that their procedures are grossly inadequate to dealing with the complexity of the situation in which they find themselves. The considerations involved in a judgment that a particular target criterion is sensible are instead less likely to constantly change.

Indeed, there are important theoretical reasons to expect that a desirable target criterion will depend on fewer details about the current economic environment than would a desirable specification of a reaction function. Giannoni and Woodford (2010) show how to construct robustly optimal target criteria, which implement an optimal response to shocks regardless of which types of shocks are more important, or of the degree of persistence, forecastability, and so on of the shocks that occur. The coefficients of an optimal reaction function will instead depend on the statistical properties of the shocks.¹¹ Since each shock to the economy that occurs is always somewhat different from any other, there will always be new information about the particular types of disturbances that have most recently occurred, making advance commitment to a particular reaction function inconvenient. The types of structural change that imply a change in the form or coefficients of the desirable target criterion instead occur more infrequently, though they certainly also occur.

As an example of the undesirability of strict commitment to an instrument rule, consider the consequences of the disruption of financial markets during the global financial crisis of 2007-09. Prior to the crisis, other U.S. dollar money-market interest rates moved closely with changes in the federal funds rate, so that adjustment of the FOMC's target for the funds rate (which in turn resulted in actions that kept the effective funds rate very close to that target, on virtually a daily basis) had direct implications for other rates as well. During the crisis, instead, many other short-term rates departed substantially from the path of the funds rate. For example, one closely monitored indicator, the U.S. dollar LIBOR rate --- to which the lending terms available to many non-financial borrowers are automatically linked --- had remained always close to the market forecast of the average funds rate over the corresponding horizon (as indicated by the overnight interest-rate swap rate); but after the summer of 2007, a spread that had previously been extremely stable and of 10 basis points or less became highly volatile and at certain times several percentage points in magnitude.

The same kind of "Taylor rule" for the federal funds rate as a function of general macroeconomic conditions (inflation and real activity) as might be appropriate at other times should not be expected to remain an equally reliable indicator when the relation between the funds rate and other market interest rates (others of which are more directly relevant to many people's economic decisions) changes. For example, in the simulations of Cúrdia and Woodford (2010a), an inflexible commitment to the standard "Taylor rule" lead to policy that is too tight in the case of a financial disturbance that increases spreads between the funds rate and the rates faced by other borrowers. Yet commitment to a

¹¹ This is illustrated in Woodford (2003, chap. 7) in the context of a simple example.

target criterion is not subject to the same critique. Even if the central bank's target criterion involves only the projections for inflation and some measure of aggregate real activity, if the bank correctly takes account of the consequences of financial disruptions for the monetary transmission mechanism in the forecast targeting exercise, it will necessarily be sensitive to changing financial conditions in choosing a path for its policy rate; and it will similarly modify its implementation procedures, if necessary, in order to more effectively keep the policy rate close to that path.

One might counter that such an example shows only that a central bank must be willing to consider modifications of its commitment to an instrument rule occasionally, under sufficiently unusual circumstances. Indeed, John Taylor himself (Taylor, 2008) proposed a modification of his celebrated rule during the crisis, under which the funds rate target (given any particular values for the current inflation rate and output gap) should be adjusted downward one-for-one with any increase in the LIBOR-OIS spread. However, even this proposed modification is unlikely to provide as accurate a guideline as would be provided by commitment instead to a target criterion, and the use of many indicators to determine the instrument setting required to implement it. Taylor's quest for a simple reaction function that can be stated publicly in advance of the decisions made under it requires him to choose a *single* indicator of financial conditions --- the LIBOR-OIS spread at one particular term --- when in fact there are many market rates and asset prices that influence economic decisions, and the different possible spreads and other indicators of financial conditions behave differently, especially during periods of financial instability (see, e.g., Hatzius *et al.*, 2010). Commitment to a target criterion rather than to a specific reaction function automatically allows a large number of indicators to be taken into account in judging the instrument setting that is required for consistency with the target criterion, and the indicators that are considered and the weights given to them can easily be changed when economic conditions change.

Hence I would argue [see Woodford, 2007, for detailed discussion] that the level at which it is most valuable for a central bank to make an explicit commitment --- one that can be expected to guide policy for at least some years into the future --- is at the level of a target criterion. This can then be used to guide policy decisions at the instrument level through commitment to a forecast targeting procedure for making instrument decisions; and these can in turn be used to guide policy decisions at the operational level by making the staff in charge of operations responsible for achieving the operating target over a fairly short time horizon, without any need to pre-specify the market interventions that will be required to do so. Note that the process by which the instrument path and the concrete market transactions required to implement it are derived should take account of changes in market conditions, including ones that may not have been foreseeable at the time of adoption of the target criterion.

But while I believe that it is useful for policymakers to articulate a policy commitment at the level of a target criterion, the kind of commitment that I have in mind does not preclude reconsideration of the appropriateness of the target criterion in the event of a significant change in the policy authority's view of the conditions under which it must act, due for example to progress in knowledge about how the economy works. The benefits that are expected to be obtained from an explicit policy commitment are not vitiated by allowing for occasional reconsideration of the target criterion, if the authority remains committed to choosing the new target criterion on any such occasion in accordance with its higher-level commitment to particular *principles of policy design*.

These highest-level principles will include, of course, a specification of the ultimate goals that the policy targets are intended to serve. (In the theory of monetary policy expounded in Woodford, 2003, for example, the ultimate goal is assumed to be the maximization of the expected utility of a representative household.) But there are other important principles that deserve to be articulated as well. For example, I have proposed that whenever policy targets are reconsidered, they should be chosen from what I have called a *timeless perspective* (Woodford, 1999).

B. Policy Design from a "Timeless Perspective"

By this I mean the rule of conduct that is chosen is the one that the policy authority *would have wished to commit itself to make* --- had it then had the knowledge of the economy's structure that it has now --- at a time far enough in the past for all possible consequences of the public's anticipation of the bank's systematic pattern of conduct to be taken into account.¹² I argue that this is a desirable criterion for choice even though, at the time that the new target criterion is actually adopted, the public has already anticipated whatever it has anticipated up until that point, and these past expectations can no longer be affected by the current decision --- they can only be fulfilled or disappointed.

This proposal is somewhat in the spirit of John Rawls' (1971) interpretation of social contract theory, according to which citizens should accept as binding upon them principles of justice to which they have not *actually* voluntarily submitted themselves, on the ground that these principles are ones that they *should have* been willing to choose in a hypothetical "original position," from which --- not yet knowing anything about the actual situation that they will occupy in society --- they would not make choices that seek to take advantage of the particular circumstances of the individual that they actually become. The doctrine of the timeless perspective similarly argues that a central bank should accept to be bound by principles for making monetary policy decisions that it would have wished to be bound by, if considering the matter before reaching the situation that it is actually in at the time that the action must be taken, though considering the *possibility* of reaching that situation among others.

A commitment to always choose new policy targets from a timeless perspective means that the occasion of a reconsideration of the policy targets can never be used as an excuse for renegeing on previous policy commitments simply because the policymaker's incentives are different *ex post* (when the effects of the anticipation of her actions need no longer be taken into account) than they were *ex ante* (when such effects were internalized). In the absence of a commitment to this principle --- if, instead, the policy authority simply chooses the new target criterion at each date that is associated with the best possible equilibrium from the current date onward, from the standpoint of its stabilization objective --- the need to reconsider policy targets from time to time raises similar difficulties to those discussed in the critique of discretionary policy by Kydland and Prescott (1977). In fact, this approach would reduce precisely to discretionary policy in the sense of Kydland and Prescott, if the policy target were reconsidered each time a policy action must be taken. The problem is less severe if reconsiderations are less frequent; but the question of why frequent reconsiderations should not be justifiable would itself have to be faced. Strictly speaking, the state of knowledge will constantly be

¹² See Woodford (2008; 2011, sec. 1) for further discussion of this issue in general terms.

changing, so that if reconsideration is justified whenever the policy authority's model of the economy changes, there is no obvious limit to the frequency of possible reconsiderations. Moreover, a policy authority that is not committed to the choice of targets from a timeless perspective would have a constant incentive to use any pretext, however minor, to call for reconsideration of the policy targets that it has previously announced, in order to have an excuse to renege on prior commitments that it wished to be expected to follow *ex ante*, but that it does not wish to actually fulfill *ex post*.

With a commitment to choose the target criterion from a timeless perspective, it is no longer essential to pre-specify the kinds of situations in which it will be legitimate to reconsider the target criterion. When this principle is followed, a reconsideration will always lead the policy authority to re-affirm precisely the same target criterion as it chose on the previous occasion *if there has been no change in its model of the economy*.¹³ In this case it will, as a practical matter, not make sense to go through the necessarily laborious process of debating the appropriateness of the target criterion except when there has been a substantial change in the authority's view of the economy's functioning. Hence one will expect reconsiderations of the target criterion to occur much less frequently than reconsiderations of the operating target for the policy rate, as stated above.

II. The Theory of Monetary Policy After the Global Financial Crisis

A thorough discussion of the kind of target criterion that it is appropriate for a central bank to adopt is beyond the scope of this paper. However, some brief remarks may nonetheless be appropriate about an issue that is likely to be on the minds of many at present: to what extent have the dramatic complications facing central banks during the recent global financial crisis shown that ideas about the possibility of rule-based policymaking that were popular (at least in academic circles, but also at some central banks) prior to the crisis must be thoroughly reconsidered? And to the extent that they must be, should this not cast doubt on the very wisdom of proposing that central banks undertake to articulate policy commitments on the basis of economic models that must always be regarded as, at best, provisional attempts to comprehend a complex and ever-changing reality?¹⁴

While reassessments of the theory of monetary policy in the light of the crisis have only begun, a few conclusions are already clear. Disruption of the normal functioning of financial markets, of the kind observed during the crisis, certainly affects the connection between central-bank market interventions and the bank's policy rate, and the connection between that policy rate and other equilibrium rates of return, and hence the bank's stabilization objectives. It follows that the appropriate policy decisions, at least at the operational and instrument levels, will surely be affected.

This makes commitment to a mechanical rule specified at one of these lower levels unwise under such circumstances. For example, an inflexible commitment to the standard "Taylor rule" will lead to policy that is too tight in the case of financial disturbances, as illustrated by the simulations of Curdia

¹³ A general method for deriving an optimal target criterion given the policy authority's stabilization objective and its economic model, in a way that conforms to this principle, is explained in Giannoni and Woodford (2010).

¹⁴ See Woodford (2010b) for further discussion of these issues.

and Woodford (2010). But as argued in the previous section, monetary policy recommendations that are expressed in the form of a target criterion are not so obviously problematic. In fact, except under quite special circumstances, taking account of financial market imperfections should also have consequences for the form of a desirable target criterion. For example, in the presence of financial distortions, there are additional appropriate stabilization goals for policy, that could safely be neglected if the financial system could be relied upon to function efficiently. (The minimization of financial distortions becomes an additional stabilization goal, in addition to the traditional concerns for price stability and an efficient aggregate level of resource utilization, because of the implications of financial intermediation for the efficiency of the *composition* of expenditure and of production, and not just for their aggregate levels.) These additional concerns almost certainly imply that an ideal target criterion should involve additional variables, beyond those that would suffice in a world with efficient financial intermediation. Nonetheless, the severity of the distortions resulting from neglect of such refinements is probably not as great in the case of commitment to a target criterion as in the case of commitment to an instrument rule for the federal funds rate; this is, at any rate, what the simulations reported in Cúrdia and Woodford (2010a) suggest.

Another of the special problems raised for many central banks by the crisis is that the zero lower bound on short-term nominal interest rates became a binding constraint on the degree to which traditional interest-rate policy could be used to achieve the desired degree of monetary stimulus. A situation in which the constraint binds is theoretically possible, but practically unlikely, in the absence of substantial disruption of the financial system; hence the issue was ignored in many analyses of optimal monetary policy rules prior to the crisis.

This constraint certainly changes what can be achieved by interest-rate policy and must be taken into account when choosing an appropriate state-contingent path for the policy rate. However, it does not mean that an appropriate criterion for choosing the path for the policy rate is necessarily much different from the kind that would have been recommended by the standard literature.¹⁵ Eggertsson and Woodford (2003) show that even when the zero lower bound is expected sometimes to bind, an optimal policy commitment can still be characterized by commitment to a particular target criterion, and while the optimal target criterion in this case is slightly more complex than any that had been recommended in the literature that assumed the bound would never be a binding constraint, they also show that a particular type of simpler target criterion that had already been advocated by the theoretical literature continues to provide a fairly good approximation to optimal policy (at least in the numerical example that they analyze) even in the case of a crisis that causes the zero lower bound to bind for a substantial number of quarters.¹⁶

¹⁵ It is also worth noting that this is not an issue that had been neglected in the theoretical literature on optimal monetary policy prior to the crisis. Thanks to Japan's experience since the late 1990s, the consequences of a binding zero lower bound had already been the topic of fairly extensive analysis prior to the current crisis.

¹⁶ Eggertsson and Woodford (2003) analyze the issue in the context of a DSGE model with perfectly functioning financial markets, but Cúrdia and Woodford (2010b) show how the same analysis applies to a model with credit frictions in which the zero lower bound comes to bind as a result of a disruption of financial intermediation.

The key feature that is required in order for a targeting regime to have desirable properties when the zero lower bound binds is for the target criterion to involve a *price-level target path*, rather than only a target for the rate of inflation looking forward. A purely forward-looking approach to inflation targeting can lead to a very bad outcome when the zero lower bound constrains policy, as shown by Eggertsson and Woodford (2003), because the central bank may be unable to prevent undershooting of its target while the constraint binds, and yet will permanently lock in any unwanted price declines that occur, by continuing to target inflation in a purely forward-looking way once it regains control of aggregate expenditure again. An expectation that this will occur leads to expectations of a deflationary bias to policy (to the extent that people correctly understand how the regime will work), which make the zero lower bound on nominal interest rates an even tighter constraint; under such a regime, expectations of deflation and contraction become self-fulfilling, amplifying the effects of the original disturbance. In the case of commitment to a price-level target path, instead, any undershooting of the target path implies a greater degree of future inflation that will be required to “catch up” to the target path; hence (again, to the extent that people correctly understand how the regime will work) undershooting should create inflationary expectations that, by lowering the anticipated real rate of return associated with a zero nominal interest rate, will tend to automatically limit the degree of undershooting that occurs.¹⁷

The simple target criterion proposed by Eggertsson and Woodford is actually one that has already been recommended as an optimal target criterion in a variety of simple New Keynesian models that abstracted from the ZLB. In fact, target criteria that involve a target path for the price level, and not simply a target for the rate of inflation going forward, have been found to be more robust in the sense of reducing the extent to which economic stabilization suffers as a result of errors in achieving the target due to imperfect knowledge on the part of the central bank, whether due to poor estimates of parameters of the bank’s structural model or to mistaken judgment of the economy’s current state.¹⁸ The greater robustness of this form of target criterion to difficulties caused by a failure to achieve the target owing to the zero lower bound is closely related to those other robustness results.

Financial disruptions also require reconsideration of the traditional doctrine according to which interest-rate policy is the sole tool that a central bank should use for macroeconomic stabilization purposes, and policy can be conducted while maintaining a balance sheet made up solely of short-term Treasury securities. I would argue that the traditional doctrine is a sound one, as long as financial markets operate with a high degree of efficiency; but disruption of the ability of private parties to effectively arbitrage between different markets, as during the recent crisis, creates a situation in which targeted asset purchases by the central bank and/or special credit facilities serving particular classes of institutions become additional relevant dimensions of central-bank policy.

Cúrdia and Woodford (2011) analyze the effects and illustrate the potential usefulness of these additional dimensions of policy in the context of a DSGE model with credit frictions. They find, however,

¹⁷ While no central bank has yet adopted a target of this kind, there has recently been some discussion within the Federal Reserve System of the advantages of doing so, at least as a temporary measure in response to the current crisis. See in particular Evans (2010).

¹⁸ See Woodford (2007; 2011, sec. 1) for further discussion of this topic.

that the existence of potential additional dimensions of policy does not greatly change the principles according to which it is appropriate to choose a target path for the policy rate, and hence does not call into question the desirability of a forecast-targeting framework for addressing that issue, or even provide a reason for departure from a conventional form of target criterion. The extent to which the central bank is able to limit anomalous behavior of credit spreads through unconventional policies will matter, of course, for the appropriate path of the policy rate, as Cúrdia and Woodford show through numerical examples. But this kind of modification in the character of interest-rate policy will automatically occur under the forecast-targeting procedure, just as the procedure requires the central bank to take account of changes in the anticipated behavior of credit spreads due to financial disturbances; it does not require a change in the target criterion.

The effective use of unconventional dimensions of policy also requires that policy be conducted within a systematic framework that involves some degree of advance commitment of policy actions, rather than in a purely discretionary fashion, for reasons similar to those advanced in discussions of conventional interest-rate policy. Once again, the effects of policy depend not only upon current actions (say, the quantity and type of assets that the Fed purchases this month) but also upon expectations about future policy (whether this month's purchases are only the start of an intended sequence of further purchases, how long it intends to hold these assets on its balance sheet, and so on). Given this, a purely discretionary approach to policy, which chooses a current action in order to achieve some immediate effect without internalizing the consequences of having been anticipated to act in that way, is likely to be quite suboptimal. In particular, the introduction of unconventional measures ought to be accompanied by an explanation of the anticipated "exit strategy" from these measures.

The crisis has also led to much discussion of the extent to which the stance of monetary policy (of the Fed in particular) during the real estate boom contributed to the occurrence or severity of the crisis, which raises the question whether even during times that financial markets appear to be functioning well, monetary policy decisions do not need to take into account potential consequences of monetary policy for risks to financial stability. This is not a topic that is yet well understood, but surely an important topic for central bankers to study. In Woodford (2012), I consider some standard arguments for trying to separate this issue from monetary policy deliberations, and conclude that the arguments do not justify avoiding this inquiry.

To the extent that financial crisis risk is endogenous and influenced by monetary policy, this is a concern that is not present in traditional analyses of optimal monetary policy rules (for example, the analyses in Woodford, 2003, or even in Cúrdia and Woodford, 2010a, b, 2011), and the target criteria for the conduct of monetary policy proposed in the traditional literature are not necessarily appropriate when one takes this additional consideration into account.¹⁹ This is an example of a circumstance under which it might be justifiable for a central bank to modify its existing policy commitment at the "policy targets level" --- i.e., the target criterion on the basis of which it makes decisions about the correct path

¹⁹ Of course, the degree to which the problem is significant will depend both on the degree to which one concludes that financial crisis risk is predictably influenced by monetary policy, and the extent to which such risk cannot be adequately controlled using other policy tools, such as improved regulation, "macro-prudential" supervision, and the like. But I do not believe that we can confidently conclude at this time that the problem is negligible.

for the policy rate --- though in a way that is consistent with its existing higher-level commitments at the “policy design level.”

In Woodford (2012), I give an example of how this might be done. In the simple model proposed there, the optimal target criterion for interest-rate policy involves not only the projected path of the price level and of the output gap, but also the projected path of a “marginal crisis risk” variable, which measures the degree to which marginal adjustments of the policy rate are expected to affect the risk of occurrence of a financial crisis (weighted by the expected welfare loss in the event of such a crisis). In periods when the marginal crisis risk is judged to be negligible, the recommended procedure would reduce to “flexible price-level targeting” of the kind discussed in Woodford (2007). But in periods when this is no longer true, the target criterion would require the central bank to tolerate some degree of undershooting of the price level target path, of output relative to the natural rate, or both, in order to prevent a greater increase in the marginal crisis risk.

While the adoption of such a procedure would require a departure from recent conventional wisdom in that it would allow some sacrifice of conventional policy targets in order to reduce crisis risk, it would maintain many salient characteristics of the kind of policy regime advocated in the pre-crisis literature. It would still be a form of inflation targeting regime (more precisely, a form of price-level targeting regime); it would not only ensure relative constancy of the inflation rate that people would have reason to expect in the “medium run” (i.e. a few years in the future), but it would in fact ensure constancy of the long-run price level path, regardless of the occurrence either of occasional financial crises or of (possibly more frequent) episodes of non-trivial marginal crisis risk.

I do not mean to suggest that the current crisis provides no ground for reconsideration of recently popular doctrines about central banking. To the contrary, it raises many new issues, some of which are already the topics of an active literature. However, I will be surprised if confronting these issues requires wholesale abandonment of the lessons for policy emphasized by the literature on policy rules. Among the insights that I think most likely to be of continuing relevance is the recognition that suitably chosen policy commitments can substantially improve upon the macroeconomic outcomes that should be expected from purely discretionary policy, even when those chosen to exercise the discretion are policymakers of superb intelligence and insight into current economic conditions.

REFERENCES

- Bernheim, Douglas, "Rationalizable Strategic Behavior," *Econometrica* 52: 1007-1028 (1984).
- Clarida, Richard, Jordi Gali, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature* 37: 1661-1707 (1999).
- Cúrdia, Vasco, and Michael Woodford, "Credit Spreads and Monetary Policy," *Journal of Money, Credit and Banking* 42(s1): 3-35 (2010a).
- Cúrdia, Vasco, and Michael Woodford, "Conventional and Unconventional Monetary Policy," Federal Reserve Bank of St. Louis *Review*, July/August 2010b, 92(4), pp. 229-264.
- Cúrdia, Vasco, and Michael Woodford, "The Central-Bank Balance Sheet as an Instrument of Monetary Policy," *Journal of Monetary Economics* 58: 54-79 (2011).
- Eggertsson, Gauti, and Michael Woodford, "The Zero Interest-Rate Bound and Optimal Monetary Policy," *Brookings Papers on Economic Activity* 2003(1): 271-333.
- Evans, Charles, "Monetary Policy in a Low-Inflation Environment: Developing a State-Contingent Price-Level Target," speech at the 55th Economic Conference, Federal Reserve Bank of Boston, October 16, 2010. [URL: http://www.chicagofed.org/webpages/publications/speeches/2010/10_16_boston_speech.cfm]
- Evans, George W., and Garey Ramey, "Expectation Calculation and Macroeconomic Dynamics," *American Economic Review* 82: 207-224 (1992).
- Evans, George W., and Seppo Honkapohja, "Expectations and the Stability Problem for Optimal Monetary Policy," *Review of Economic Studies* 70: 807-824 (2003).
- Friedman, Benjamin M., and Kenneth N. Kuttner, "Implementation of Monetary Policy: How Do Central Banks Set Interest Rates?" in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3B, Amsterdam, Elsevier, 2011.
- Frydman, Roman, and Michael Goldberg, *Beyond Mechanical Markets*, Princeton: Princeton University Press, 2011.
- Giannoni, Marc P. and Michael Woodford, "Optimal Target Criteria for Stabilization Policy," NBER Working Paper no. 15757, February 2010.
- Guesnerie, Roger, *Assessing Rational Expectations 2: Eductive Stability in Economics*, Cambridge: MIT Press, 2005.

Hatzius, Jan, Peter Hooper, Frederic S. Mishkin, Kermit L. Schoenholtz, and Mark W. Watson, "Financial Conditions Indexes: A Fresh Look After the Financial Crisis," NBER Working Paper no. 16150, July 2010.

Kydland, Finn E., and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy* 85: 473-491 (1977).

McCallum, Bennett T., "Robustness Properties of a Rule for Monetary Policy," *Carnegie-Rochester Conference Series on Public Policy* 29: 173-203 (1988).

McCallum, Bennett T., "Issues in the Design of Monetary Policy Rules," in J.B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, vol. 1C. Amsterdam: Elsevier, 1999.

Pearce, David, "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica* 52: 1029-1050 (1984).

Phelps, Edmund S., et al., *Microeconomic Foundations of Employment and Inflation Theory*, New York: W.W. Norton, 1970.

Qvigstad, Jan F., "When Does an Interest Rate Path 'Look Good'? Criteria for an Appropriate Future Interest Rate Path: A Practitioner's Approach," Norges Bank staff memo 2006/5, April 2006.

Rawls, John, *A Theory of Justice*, Cambridge, Mass.: Harvard University Press, 1971.

Sargent, Thomas J., "Rational Expectations and the Reconstruction of Macroeconomics," in *Rational Expectations and Inflation*, 2d ed., New York: HarperCollins, 1993.

Sundaresan, Suresh, and Zhenyu Wang, "Y2K Options and the Liquidity Premium in Treasury Markets," *Review of Financial Studies* 22: 1021-1056 (2009).

Svensson, Lars E.O., "Inflation Targeting as a Monetary Policy Rule," *Journal of Monetary Economics* 43: 607-654 (1999).

Svensson, Lars E.O., "What Is Wrong with Taylor Rules? Using Judgment in Monetary Policy through Targeting Rules," *Journal of Economic Literature* 41: 426-477 (2003).

Svensson, Lars E.O., "Monetary Policy with Judgment: Forecast Targeting," *International Journal of Central Banking* 1: 1-54 (2005).

Taylor, John B., "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214 (1993).

Taylor, John B., "Monetary Policy and the State of the Economy," testimony before the Committee on Financial Services, U.S. House of Representatives, February 26, 2008. [URL: http://www.stanford.edu/~johntayl/Onlinepaperscombinedbyyear/2008/Monetary_Policy_and_the_State_of_the_Economy.pdf]

Taylor, John B., and John C. Williams, "Simple and Robust Rules for Monetary Policy," in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3B, Amsterdam, Elsevier, 2011.

Woodford, Michael, "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, 1999.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.

Woodford, Michael, "Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice," NBER Working Paper no. 13716, December 2007.

Woodford, Michael, "Principles and Public Policy Decisions: The Case of Monetary Policy," paper for the Yale Law School Legal and Economic Organization seminar, March 2008. [URL: www.columbia.edu/~mw2230/Principles and Policy_YLS.pdf]

Woodford, Michael, "Robustly Optimal Monetary Policy with Near-Rational Expectations," *American Economic Review* 100: 274-303 (2010).

Woodford, Michael, "Optimal Monetary Stabilization Policy," in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3B, Amsterdam: Elsevier, 2011.

Woodford, Michael, "Inflation Targeting and Financial Stability," *Sveriges Riksbank Economic Review* 2012(1): 7-32.