

COPYRIGHT NOTICE:

Michael Woodford: Interest and Prices

is published by Princeton University Press and copyrighted, © 2003, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

For COURSE PACK and other PERMISSIONS, refer to entry on previous page. For more information, send e-mail to permissions@pupress.princeton.edu

CHAPTER 1

The Return of Monetary Rules

If it were in our power to regulate completely the price system of the future, the ideal position . . . would undoubtedly be one in which, without interfering with the inevitable variations in the relative prices of commodities, the general average level of money prices . . . would be perfectly invariable and stable.

And why should not such regulation lie within the scope of practical politics? . . . Attempts by means of tariffs, state subsidies, export bounties, and the like, to effect a partial modification of the natural order of [relative prices] almost inevitably involve some loss of utility to the community. Such attempts must so far be regarded as opposed to all reason. Absolute prices on the other hand—money prices—are a matter in the last analysis of pure convention, depending on the *choice of a standard of price* which it lies within our own power to make.

—KNUT WICKSELL

Interest and Prices, 1898, p. 4

The past century has been one of remarkable innovation in the world's monetary systems. At the turn of the twentieth century, it was taken for granted by practical men that the meaning of a monetary unit should be guaranteed by its convertibility into a specific quantity of some precious metal. Debates about monetary policy usually concerned the relative advantages of gold and silver standards or the possibility of a bimetallic standard. But through fits and starts, the world's currencies have come progressively to be more completely subject to "management" by individual central banks. Since the collapse of the Bretton Woods system of fixed exchange rates in the early 1970s, the last pretense of a connection of the world's currencies to any real commodity has been abandoned. We now live instead in a world of pure "fiat" units of account, where the value of each depends solely upon the policies of the particular central bank with responsibility for it.

This has brought both opportunities and challenges. On the one hand, vagaries of the market for gold or some other precious metal no longer cause variations in the purchasing power of money, with their disruptions of the pattern of economic activity. The recognition that the purchasing

power of money need not be dictated by any “natural” market forces and is instead a proper subject of government regulation, as proposed by the monetary reformer Knut Wicksell a century ago, should in principle make possible greater stability of the standard of value, facilitating contracting and market exchange. At the same time, the responsibilities of the world’s central banks are more complex under a fiat system than they were when the banks’ tasks were simply to maintain convertibility of their respective national currencies into gold, and it was not immediately apparent how the banks’ new freedom should best be used. Indeed, during the first decade of the new regime, the policies of many industrial nations suffered from a tendency toward chronic inflation, leading to calls from some quarters in the 1980s for a return to a commodity standard.

This has not proven to be necessary. Instead, since the 1980s the central banks of the major industrial nations have been largely successful at bringing inflation down to low and fairly stable levels. Nor does this seem to have involved any permanent sacrifice of other objectives. For example, real GDP growth has been if anything higher on average, and certainly more stable, since inflation was stabilized in the United States. Somewhat paradoxically, this period of improved macroeconomic stability has coincided with a *reduction*, in certain senses, in the ambition of central banks’ efforts at macroeconomic stabilization. Banks around the world have committed themselves more explicitly to relatively straightforward objectives with regard to the control of inflation, and have found when they do so that not only is it easier to control inflation than previous experience might have suggested, but that price stability creates a sound basis for real economic performance as well.

What appears to be developing, then, at the turn of another century, is a new consensus in favor of a monetary policy that is disciplined by clear rules intended to ensure a stable standard of value, rather than one that is determined on a purely discretionary basis to serve whatever ends may seem most pressing at any given time. Yet the new monetary rules are not so blindly mechanical as the rules of the gold standard, which defined monetary orthodoxy a century ago. They are instead principles of systematic conduct for institutions that are aware of the consequences of their actions and take responsibility for them, choosing their policies with careful attention to what they accomplish. Indeed, under the current approaches to rule-based policymaking, more emphasis is given to explicit commitments regarding desired economic outcomes, such as a target rate of inflation, than to particular technical indicators that the central bank may find it useful to monitor in achieving that outcome.

The present study seeks to provide theoretical foundations for a rule-based approach to monetary policy of this kind. The development of such a theory is an urgent task, for rule-based monetary policy in the spirit

that I have described is possible only when central banks can develop a conscious and articulate account of what they are doing. It is necessary in order for them to know how to act systematically in a way that can serve their objectives, which are now defined in terms of variables that are much further removed from their direct control. It is also necessary in order for them to be able to communicate the nature of their systematic commitments to the public, despite the absence of such mechanical constraints as a commitment to exchange currency for some real commodity. As I explain below, the advantages of a sound monetary policy are largely dependent upon the policy's being *understood* and relied upon by the private sector in arranging its affairs.

There can be little doubt that the past decade has seen a marked increase in the self-consciousness of central banks about the way in which they conduct monetary policy and in the explicitness of their communication with the public about their actions and the considerations upon which they are based. A particularly important development in this regard has been the adoption of "inflation targeting" as an approach to the conduct of monetary policy by many of the world's central banks in the 1990s.¹ As I subsequently discuss in more detail, this approach (best exemplified by the practices of such innovators as the Bank of England, the Bank of Canada, the Reserve Bank of New Zealand, and the Swedish Riksbank) is characterized not only by public commitment to an explicit target, but also by a commitment to explain the central bank's policy actions in terms of a systematic decisionmaking framework that is aimed at achieving this target. This has led to greatly improved communication with the public about the central bank's interpretation of current conditions and the outlook for the future, notably through the publication of detailed *Inflation Reports*. It has also involved fairly explicit discussion of the approach that they follow in deliberating about policy actions and, in some cases, even publication of the model or models used in producing the forecasts that play a central role in these deliberations. As a consequence, these banks in particular have found themselves in need of a clear theory of how they can best achieve their objectives and have played an important role in stimulating reflection on this problem.

It is true that the conceptual frameworks proposed by central banks to deal with their perceived need for a more systematic approach to policy were, until quite recently, largely developed without much guidance from the academic literature on monetary economics. Indeed, the central questions of practical interest for the conduct of policy—how should central banks decide about the appropriate level of overnight interest rates? how should monetary policy respond to the various types of unexpected disturbances that occur?—had in recent decades ceased to be considered

1. See, e.g., Bernanke et al. (1999) for a thorough discussion of this development.

suitable topics for academic study. Reasons for this included the trenchant critique of traditional methods of econometric-policy evaluation by Lucas (1976); the critique of the use of conventional methods of optimal control in the conduct of economic policy by Kydland and Prescott (1977); and the development of a new generation of quantitative models of business fluctuations (“real-business-cycle theory”) with more rigorous microeconomic foundations, but which implied no relevance of monetary policy for economic welfare.

Nonetheless, recent developments, to be discussed in detail in this volume, have considerably changed this picture. The present study seeks to show that it is possible to use the tools of modern macroeconomic theory—intertemporal equilibrium modeling, taking full account of the endogeneity of private-sector expectations—to analyze optimal interest-rate setting in a way that takes the concerns of central bankers seriously, while simultaneously taking account of the “New Classical” critique of traditional policy-evaluation exercises. In this way, the basic elements are presented of a theory that can provide the groundwork for the kind of systematic approach to the conduct of monetary policy that many central banks are currently seeking to develop. In the present chapter, I review some of the key features of this theory, as preparation for the more systematic development that begins in Chapter 2.

1 The Importance of Price Stability

A notable feature of the new rule-based approaches to monetary policy is the increased emphasis given to a particular policy objective: maintaining a low and stable rate of inflation. This is most obvious in the case of countries with explicit inflation targets. But it also seems to characterize recent policy in the United States as well, where the past decade has seen unusual stability of the inflation rate and many econometric studies have found evidence of a stronger Fed reaction to inflation variations. (See further discussion of recent U.S. policy in Section 4.1.)

Yet the justification of such an emphasis from the standpoint of economic theory may not be obvious. Standard general-equilibrium models—and the earliest generation of quantitative equilibrium models of business fluctuations, the real-business-cycle models of the 1980s—indicate that the absolute level of prices should be irrelevant for the allocation of resources, which depends only on *relative* prices. Traditional Keynesian macroeconomic models, of course, imply otherwise: Variations in the growth rate of wages and prices are found to be associated with substantial variations in economic activity and employment. Yet the existence of such “Phillips-curve” relations has typically been held to imply that monetary policy should be

used to achieve output or employment goals, rather than giving priority to price stability.

The present study argues instead for a different view of the proper goals of monetary policy. Its use to stabilize an appropriately defined price index is in fact an important end toward which efforts should be directed—at least to a first approximation, it should be the primary aim of monetary policy. But this is *not*, as proponents of inflation targeting sometimes argue, because variations in the rate of inflation have no real effects. Rather, it is exactly because instability of the general level of prices causes substantial real distortions—leading to inefficient variation both in aggregate employment and output and in the sectoral composition of economic activity—that price stability is important.

Moreover, the existence of predictable real effects of shifts in monetary policy need not imply that policy should be based primarily on a calculation of its effects on output or employment. For the efficient aggregate level and sectoral composition of real activity is likely to vary over time, as a result of real disturbances of a variety of types. The market mechanism performs a difficult computational task—much of the time, fairly accurately—in bringing about a time-varying allocation of resources that responds to these changes in production and consumption opportunities. Because of this, variation over time in employment and output relative to some smooth trend cannot in itself be taken to indicate a failure of proper market functioning. Instead, instability of the general level of prices is a good indicator of inefficiency in the real allocation of resources—at least when an appropriate price index is used—because a general tendency of prices to move in the same direction (either all rising relative to their past values or all falling) is both a cause and a symptom of systematic imbalances in resource allocation.

This general vision is in many respects an attempt to resurrect a view that was influential among monetary economists prior to the Keynesian revolution. It was perhaps best articulated by the noted Swedish economic theorist Knut Wicksell at the turn of the previous century, along with his followers in the “Stockholm school” of the interwar period (such as Erik Lindahl and Gunnar Myrdal) and others influenced by Wicksell’s work, such as Friedrich Hayek. However, these authors developed their insights without the benefit of either modern general-equilibrium theory² or macroeconometric modeling techniques, so that it may be doubted whether Wicksellian theory can

2. Of course, Wicksell and his followers were quite familiar with Walrasian general-equilibrium theory and used it as a starting point for their own thought. But at the time, general-equilibrium theory meant a *static* model of resource allocation, not obviously applicable to the problems of intertemporal resource allocation with which they were primarily concerned. See, e.g., Myrdal (1931, chap. 2, sec. 4, and chap. 3, sec. 5).

provide a basis for the kind of quantitative policy analysis in which a modern central bank must engage—and which has become only more essential given current demands for public justification of policy decisions. This book seeks to provide theoretical foundations for the view just sketched that meet modern standards of conceptual rigor and are capable of elaboration in a form that can be fit to economic time series.

1.1 Toward a New “Neoclassical Synthesis”

The approach to monetary policy proposed here builds upon advances in the analysis of economic fluctuations and, in particular, of the monetary transmission mechanism over the past few years.³ The models analyzed in this volume differ in crucial respects from the first two generations of equilibrium business-cycle models, namely, the New Classical models that Lucas (1972) took as a starting point and the real-business-cycle (RBC) models pioneered by Kydland and Prescott (1982) and Long and Plosser (1983). Neither of these early illustrations of the possibility of rigorous intertemporal general-equilibrium analysis of short-run fluctuations contained elements that would make them suitable for the analysis of monetary policy. While the Lucas model allows for real effects of unexpected variations in monetary policy (modeled as stochastic variation in the growth rate of the money supply), it implies that any real effects of monetary policy must be purely transitory and that monetary disturbances should have *no* real effects to the extent that their influence on aggregate nominal expenditure can be forecast in advance. Yet, as shown Chapter 3, VAR evidence on the effects of identified monetary policy shocks is quite inconsistent with these predictions. Rather, these effects on aggregate nominal expenditure are forecastable at least 6 months in advance on the basis of federal funds rate movements, whereas the (similarly delayed) effects on real activity are substantial and persist for many quarters. Nor is this empirical failure of the model one of minor import for the analysis of monetary policy. The conclusion that only unanticipated monetary policy can have real effects leads fairly directly to the skeptical conclusions of Sargent and Wallace (1975) about the necessary ineffectiveness of any attempt to use monetary policy to stabilize real activity.

The RBC models of the 1980s offered a very different view of the typical nature of short-run fluctuations in economic activity. But the classic models in this vein similarly imply no scope at all for monetary stabilization policy because real variables are modeled as evolving in complete independence of any nominal variables. Monetary policy is thus (at least implicitly) assumed

3. Useful surveys of recent developments include Goodfriend and King (1997) and Gali (2001).

to be of no relevance as far as fluctuations in real activity are concerned. Since neither the empirical evidence from VAR studies nor the practical experience of central bankers supports this view, I am reluctant to discuss the nature of desirable monetary policy rules using models of this kind.

Chapters 3 through 5 review a more recent literature that has shown, instead, how one can develop models with equally rigorous foundations in intertemporal optimizing behavior that allow a more realistic account of the real effects of monetary disturbances. These models also imply that systematic monetary policy can make a substantial difference for the way that an economy responds to real disturbances of all sorts, and this is actually the prediction that is of greatest importance for the concerns here. VAR models typically do *not* imply that a large part of the variance of fluctuations in real activity should be attributed to monetary policy shocks—that is, to the purely random component of central-bank interest-rate policy. Moreover, in any event, one does not really need to understand exactly what the effects of such shocks are, since under almost any view it is desirable to eliminate such shocks (i.e., to render monetary policy predictable) to the greatest extent possible. (Here, I discuss the ability of the models being proposed to account for evidence with regard to the effects of such shocks only because this is the aspect of the effects of monetary policy that can be *empirically identified* under relatively weak, and hence more convincing, identifying assumptions.) On the other hand, I am very interested in what a model implies about the way in which alternative systematic monetary policies determine the effects of real disturbances. The question of practical importance in central banking is never “should we create some random noise this month?,” but rather “does this month’s news justify a change in the level of interest rates?” To think about this, one needs to understand the consequences of different types of possible monetary responses to exogenous disturbances.

The key to obtaining less trivial consequences of systematic monetary policy in the models proposed here is the assumption that prices and/or wages are not continually adjusted, but remain fixed for at least short periods (a few months, or even a year) at a level that was judged desirable at an earlier time. However, this postulate does not mean accepting the need for mechanical models of wage and price adjustment of the kind that were at the heart of the Keynesian macroeconometric models of the 1960s. Rather than postulating that prices or wages respond mechanically to some measure of market disequilibrium, they are set optimally, that is, so as to best serve the interests of the parties assumed to set them, according to the information available at the time they are set. The delays involved before the next time that prices are reconsidered (or perhaps, before a newly chosen price takes effect) are taken here to be an institutional fact, just like the available production technology. However, the resulting constraints are taken account of by the decisionmakers who set them. Thus the assumed “stickiness” of

prices implies that when they are reconsidered, they are set in a *forward-looking* manner, on the basis of expectations regarding future demand and cost conditions, and not simply in response to current conditions. As a result, expectations turn out to be a crucial factor in the equilibrium relation between inflation and real activity (as argued by Phelps and Friedman in the 1960s). Under certain special assumptions, described in Chapter 3, the relation is of exactly the form assumed in the New Classical literature: Deviations of output from its “natural rate” are proportional to the unexpected component of inflation. However, this is not true more generally. Other models, which I would judge to be more realistic, also lead to “expectations-augmented Phillips-curve” relations of a sort, but not of the precise sort that implies that anticipated monetary policy cannot have real effects.

It is also important to note that the emphasis here upon nominal rigidities does not in any way mean ignoring the real factors in business fluctuations stressed by RBC theory. One important achievement of the RBC literature has been to show that the equilibrium level of output can easily be disturbed by real disturbances of many sorts—variations in the rate of technical progress, variations in government purchases, changes in tax rates, or different kinds of shifts in tastes. I do not want to abstract from the existence of such disturbances in the proposed models. After all, it is only the existence of real disturbances (i.e., disturbances other than those originating from randomness in monetary policy itself) that gives rise to nontrivial questions about monetary policy, and I strive to obtain results that remain valid for as broad a class of possible disturbances as possible. Of course, the predicted effects of real disturbances are not necessarily the same in the models presented here as in RBC theory, which, in its classic form, assumes complete flexibility of both wages and prices. Rather, in these models, the predicted effects of real disturbances depend on the nature of monetary policy.

Nonetheless, the predicted evolution of real variables under complete wage and price flexibility—the topic studied in RBC theory—represents an important benchmark in the theory developed here. The level of output that *would* occur in an equilibrium with flexible wages and prices, given current real factors (tastes, technology, government purchases)—what is called the “natural rate” of output, following Friedman (1968)—turns out to be a highly useful concept, even if the present theory does not imply that this is the *actual* level of output, regardless of monetary policy. It is the gap between actual output and this natural rate, rather than the level of output as such (or output relative to trend) that is related to inflation dynamics in a properly specified Phillips-curve relation, as I show in Chapter 3. It is also this concept of the output gap to which interest rates should respond if a “Taylor rule” is to be a successful approach to inflation stabilization, as I discuss in Chapter 4; it is this concept of the output gap that monetary policy

should aim to stabilize in order to maximize household welfare, as shown in Chapter 6; and it is this concept of the output gap to which optimal interest-rate rules and/or optimal inflation targets should respond, as shown in Chapter 8. From the point of view of any of these applications, the fact that the natural rate of output may vary at business-cycle frequencies, as argued in the RBC literature, is of tremendous practical importance. I am also quite interested in the consequences of time variation in what Wicksell (1898) called the “natural rate of interest”—the equilibrium real rate of interest in the case of flexible wages and prices, given current real factors.⁴ Once again, RBC theory has a great deal to tell us about the kind of factors that should cause the natural rate of interest to vary. Hence RBC theory, when correctly interpreted, constitutes an important building block of the theory to be developed here.

It is for this reason that Goodfriend and King (1997) speak of models of this kind as representing a “new neoclassical synthesis,” in the spirit of the synthesis between Keynesian short-run analysis and neoclassical long-run analysis proposed by Hicks and Samuelson. In the modern, more explicitly dynamic version of such a synthesis, the neoclassical theory (i.e., RBC theory) defines not a static “long-run equilibrium” but rather a dynamic path that represents a sort of *virtual* equilibrium for the economy at each point in time—the equilibrium that one *would* have if wages and prices were not in fact sticky. The evolution of the virtual equilibrium matters because the gaps between actual quantities and their virtual equilibrium values are important measures of the incentives for wage and price adjustment and hence determinants of wage and price dynamics.

At the same time, the stickiness of prices and/or wages implies that short-run output determination can be understood in a manner reminiscent of Keynesian theory. Indeed, the basic analytical framework in this study has the structure of a simple model consisting of an “IS equation,” a monetary policy rule, and an “AS equation.” (The monetary policy rule—which I often suppose is something similar to a Taylor rule—replaces the “LM equation” of Hicksian pedagogy, since for the most part I am not interested

4. This is of course the origin of the natural rate terminology—Friedman’s (1968) concept of a “natural rate of unemployment” appealed to an analogy with Wicksell’s “natural rate of interest,” a concept with which his readers were presumed to be familiar. Nowadays, many readers are more familiar with Friedman’s concept and find the natural rate of interest most easy to understand as an analogy with Friedman’s natural rate. In fact, Wicksell was also an early proponent of the “natural rate hypothesis” enunciated in Friedman’s address. “Those people who prefer a continually upward moving to a stationary price level forcibly remind one of those who purposely keep their watches a little fast so as to be more certain of catching their trains. But to achieve their purpose they must not be conscious . . . of the fact that their watches are fast; otherwise they become accustomed to take the extra few minutes into account and so after all, in spite of their artfulness, arrive too late” (Wicksell, 1898, pp. 3–4).

here in the consequences of monetary targeting.) Nonetheless, even for purposes of “short-run” analysis, the proposed model is less static than an old-fashioned Keynesian model; in particular, *expectations* will be crucial elements in its structural relations (e.g., its “intertemporal IS relation”), so that anything that causes a change in expectations should shift them.

The inclusion of significant forward-looking terms in the key structural relations has substantial consequences for the analysis of the character of optimal policy, just as Lucas (1976) argued, even if the consequences are not necessarily the ones suggested in the New Classical literature. For example, estimated IS equations in traditional macroeconometric models often indicate an effect of *lagged* rather than current interest rates on aggregate demand, since the coefficients on lagged rates are found to be more significant than those on a current interest rate in the case of a regression seeking to explain aggregate real expenditure in terms of observable variables. On the other hand, in the optimization-based model estimated by Rotemberg and Woodford (1997), the observed delay in the effects of an interest-rate innovation on real GDP is explained by an assumption that the interest-sensitive component of private spending is predetermined, though chosen in a forward-looking way. Thus current aggregate demand is assumed to depend on *past expectations* of *current and future* interest rates, rather than on past interest rates.

Econometrically, the two hypotheses are not easily distinguished, given the substantial serial correlation of observed interest rates. Yet the second hypothesis, I would argue, has a much simpler logic in terms of the optimal timing of expenditure, once one grants the hypothesis of predetermination of spending decisions (just as with pricing decisions). Moreover, the specification assumed matters greatly for one’s conclusions about the conduct of policy. If expenditure is really affected solely by lagged interest rates, it becomes important for the central bank to adjust interest rates in response to its *forecast* of how it would like to affect aggregate demand at a later date; “preemptive” actions are essential. If instead only past expectations of current and future interest rates matter, then *unforecastable* interest-rate movements will not affect demand, so that immediate responses to news will serve no purpose. It will then be important for interest rates to continue to respond to the outlook that *had* been perceived in the past, even if more recent news has substantially modified the bank’s forecasts. This *inertial* character of optimal interest-rate policy is discussed further in Chapters 7 and 8.

1.2 Microeconomic Foundations and Policy Analysis

I consider the development of a model of the monetary transmission mechanism with clear foundations in individual optimization to be important for two reasons: It allows us to evaluate alternative monetary policies in a way

that avoids the flaw in policy evaluation exercises using traditional Keynesian macroeconomic models stressed by Lucas (1976); and the outcomes resulting from alternative policies can be evaluated in terms of the preferences of private individuals that are reflected in the structural relations of one's model.

Lucas (1976) argued that traditional policy evaluation exercises using macroeconomic models were flawed by a failure to recognize that the relations typically estimated—a “consumption equation,” a “price equation,” and so on—were actually (at least under the hypothesis of optimizing behavior by households and firms) reduced-form rather than truly structural relations. In particular, in the estimated equations, expectations regarding future conditions (future income in the case of consumers and future costs and future demand in the case of pricesetters) were proxied for by current and lagged observable state variables. But the correlation of expectations with those observables ought to be expected to change in the case of a change in the government's policy rule, as contemplated in the policy-evaluation exercise.

This problem can be addressed by making use of structural relations that explicitly represent the dependence of economic decisions upon expectations regarding future endogenous variables. The present study illustrates how this can be done, deriving the structural relations that are to be used in the calculation of optimal policy rules from the first-order conditions (Euler equations) that characterize optimal private-sector behavior. These conditions explicitly involve private-sector expectations about the future evolution of endogenous variables, and often they only *implicitly* define private-sector behavior, rather than giving a consumption equation or price equation in closed form. My preference for this form of structural relations is precisely that they are ones that should remain invariant (insofar as the proposed theory is correct) under changes in policy that alter the stochastic laws of motion of the endogenous variables.

Of course, the mere fact that the structural relations derived here follow from explicit optimization problems for households and firms is no guarantee that they are correctly specified; the (fairly simple) optimization problems that I consider here may or may not be empirically realistic. (Indeed, insofar as I illustrate the principles of my approach in the context of very simple examples, one can be certain that they are not very precise representations of reality.) But this is not an objection to the method that I advocate here. It simply means that there is no substitute for careful empirical research to flesh out the details of a quantitatively realistic account of the monetary transmission mechanism. While the present study does include some discussion of the extent to which the simple models presented here are consistent with empirical evidence, in order to motivate the introduction of certain model elements, no attempt is made to set out a model that

is sufficiently realistic to be used for actual policy analysis in a central bank. Nonetheless, the basic elements of an optimizing model of the monetary transmission mechanism, developed in Chapters 3 through 5, are ones that I believe are representative of crucial elements of a realistic model; and indeed, the illustrative models discussed here have many elements in common with rational-expectations models of the monetary transmission mechanism that are already being used for quantitative policy evaluation at a number of central banks.

A second advantage of proceeding from explicit microeconomic foundations is that in this case, the welfare of private agents—as indicated by the utility functions that underlie the structural relations of one’s model of the transmission mechanism—provides a natural objective in terms of which alternative policies should be evaluated. In taking this approach, the present study seeks to treat questions of monetary policy in a way that is already standard in other branches of public economics, such as the analysis of optimal tax policy. Nonetheless, the approach is not common in the literature on monetary policy evaluation, which instead typically evaluates alternative policies in terms of ad hoc stabilization objectives for various macroeconomic indicators.

Until recently, welfare-theoretic analyses of monetary policy have been associated exclusively with the problem of reducing the transactions frictions (sometimes called “shoe-leather costs”) that account for the use of money in purchases.⁵ This is because this was for a long time the only sort of inefficiency present in general-equilibrium monetary models, which typically assumed perfectly flexible wages and prices and perfect competition. Here I show how welfare analysis of monetary policy is also possible in settings that incorporate nominal rigidities. Allowing for these additional frictions—crucial to understanding the real effects of alternative monetary policies—provides a welfare-theoretic justification for additional policy goals.

As shown in Chapter 6, taking account of delays in the adjustment of wages and prices provides a clear justification for an approach to monetary policy that aims at price stability. It might seem more obvious that allowing for real effects of monetary policy provides a justification for concern with output stabilization. The stickiness of prices explains why actual output may differ from the natural rate, and so justifies a concern for the stabilization of the “output gap,” that is, the discrepancy between the actual and natural levels of output. But price stickiness also justifies a concern with price stability. For when prices are not constantly adjusted, instability of the general level of prices creates discrepancies between relative prices owing to the absence of perfect synchronization in the adjustment of the prices of different

5. For reviews of that traditional literature, see Woodford (1990) and Chari and Kehoe (1999).

goods. These relative-price distortions lead in turn to an inefficient sectoral allocation of resources, even when the *aggregate* level of output is correct.

Moreover, the present theory implies not only that price stability should matter *in addition* to stability of the output gap, but also that, at least under certain circumstances, inflation stabilization eliminates any need for further concern with the level of real activity. This is because, at least under the conditions described more precisely in Chapter 6, the time-varying efficient level of output is the *same* (up to a constant, which does not affect the basic point) as the level of output that eliminates any incentive for firms on average to either raise or lower their prices. It then follows that there is no conflict between the goal of inflation stabilization and output-gap stabilization, once the welfare-relevant concept of the output gap is properly understood. Furthermore, because of the difficulty involved in measuring the efficient level of economic activity in real time—depending as this does on variations in production costs, consumption needs, and investment opportunities—it may well be more convenient for a central bank to concern itself simply with monitoring the stability of prices.

The development of an explicit welfare analysis of the distortions resulting from inflation variations has advantages beyond the mere provision of a justification for central bankers' current concern with inflation stabilization. For the theory presented here also provides guidance as to *which* price index it is most desirable to stabilize. This is a question of no small practical interest. For example, the stock-market booms and crashes in many industrial nations in the late 1990s led to discussion of whether central banks ought not target an inflation measure that took account of "asset-price inflation" as well as goods prices.⁶

The answer provided by the theory developed here is *no*. The prices that monetary policy should aim to stabilize are the ones that are infrequently adjusted and that consequently can be expected to become misaligned in an environment that requires these prices to move in either direction. Large movements in frequently adjusted prices—and stock prices are among the most flexible—can instead be allowed without raising such concerns, and if allowing them to move makes possible greater stability of the sticky prices, such instability of the flexible prices is desirable.⁷ In Chapter 6, I show how such a conclusion can be justified from the point of view of welfare

6. For examples of scholarly attention to the question, see Goodhart (1999) and Bryan et al. (2002).

7. The basic point was already evident to authors of the Stockholm school: "If one desires the greatest possible diminution of the business cycle . . . then one must try to stabilize an index of those prices which are sticky in themselves. . . . Stability of the level of the sticky prices permits a certain freedom for all other price levels, including capital values. . . . It is evident that [the price of capital goods] is the last price that one should try to stabilize in a capitalist society. . . . The same is naturally true for all indices of flexible commodity prices" (Myrdal,

economics. I further show how to develop a quantitative measure of the deadweight loss resulting from stabilization of alternative price indices, so that more subtle distinctions between the relative stickiness of different prices can be dealt with.

In addition to implying that an appropriate inflation target ought not involve asset prices, the present theory suggests that not all goods prices are equally relevant. Rather, central banks should target a measure of “core” inflation that places greater weight on those prices that are stickier. Furthermore, insofar as wages are also sticky, a desirable inflation target should take account of wage inflation as well as goods prices. The empirical results discussed in Chapter 3 suggest that wages and prices are sticky to a similar extent, suggesting (as I show in Chapter 8) that a desirable inflation target should put roughly equal weight on wage and price inflation.

2 The Importance of Policy Commitment

Thus far, I have summarized a theoretical justification for the concern of the inflation-targeting central banks with price stability. But why should it follow that there is a need for public commitment to a target inflation rate, let alone for commitment to a systematic procedure for determining appropriate instrument settings? Why is it not enough to appoint central bankers with a sound understanding of the way the economy works and then grant them complete discretion to pursue the public interest in the way that they judge best? Should it not follow from my analysis that this would result in price stability, to the extent that it is possible given the instruments available to the central bank and the information available at the time that policy decisions must be made?

I argue instead that 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. There are two important advantages of commitment to an appropriately chosen *policy rule* of this kind. One is that the effectiveness of monetary policy depends as much on the public’s expectations about future policy as upon the bank’s actual actions. Hence it is important not only that a bank manage to make the right decision as often as possible, but that its actions be *predictable*.

The second, and subtler, reason is that even if the public has no difficulty in correctly perceiving the pattern in the central bank’s actions—as assumed

1931, pp. 192–193). Simons (1948) and Meade (1951) advocated targeting a domestic-goods price index rather than a consumer price index on similar grounds; the CPI would include the prices of imported goods, that could vary freely through exchange-rate movements. Clarida et al. (2001) also obtain the latter conclusion, in the context of a fully articulated open-economy sticky-price model.

under the hypothesis of rational expectations—if a bank acts at each date under the assumption that it cannot commit itself to any future behavior (and is not bound by any past commitments), it will choose a systematic pattern of behavior that is suboptimal. I take up each of these arguments in turn.

2.1 Central Banking as Management of Expectations

The first advantage of commitment to a 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 successful monetary policy is not so much a matter of effective control of overnight interest rates as it is of shaping market *expectations* of the way in which interest rates, inflation, and income are likely to evolve over the coming year and later. On the one hand, optimizing models imply that private sector behavior should be forward looking; hence expectations about future market conditions should be important determinants of current behavior. It follows that, insofar as it is possible for the central bank to affect expectations, this should be an important tool of stabilization policy. Moreover, given the increasing sophistication of market participants about central banking over the past two decades, it is plausible to suppose that a central bank's commitment to a systematic policy will be factored into private-sector forecasts—at least insofar as the bank's actions are observed to match its professed commitments.

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 previously 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 (e.g., the federal funds rate in the United States). 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.⁸

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 is 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.⁹

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

8. An effect of the same kind is obtained in the basic "neo-Wicksellian" model developed in Chapter 4, insofar as the short-run real rate of interest determines not the absolute level of desired private-sector expenditure, but rather the current level relative to the expected future level of expenditure, as a result of an Euler equation for the optimal timing of expenditure. Expected future expenditure, relative to expected expenditure even farther in the future, similarly depends upon expected future short rates, and so on for expectations regarding the still farther future.

9. There 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.

minutes of policy deliberations are published. Instead, from the perspective suggested here, what is important 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 is 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.

Some good practical examples of communication with the public about the central bank's policy commitments are provided by the *Inflation Reports* of the leading inflation-targeting banks. 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, with the bank knowing that it should be expected to explain why views expressed in the past are not being followed later. Thus a commitment to transparency of this sort helps to make policy more fully rule based and also increases the public's understanding of the rule.

It is perhaps worth clarifying further what I intend by "rule-based" policy. I do not mean that a bank should commit itself to an explicit state-contingent plan for the entire foreseeable future, specifying what it would do under every circumstance that might possibly arise. That would obviously be impractical, even under complete unanimity about the correct model of the economy and the objectives of policy, simply because of the vast number of possible futures. But it is also not necessary. To obtain the benefits of commitment to a systematic policy, it suffices that a central bank commit itself to a systematic way of determining an appropriate response to future developments, without having to list all of the implications of the rule for possible future developments.¹⁰

Nor is it necessary to imagine that commitment to a systematic rule means that once a rule is adopted it must be followed forever, regard-

10. I show in Chapter 8 how policy rules can be designed that can be specified without any reference to particular economic disturbances, but that nonetheless imply an optimal equilibrium response to additive disturbances of an arbitrary type. The targeting rules advocated by Svensson (1999, 2003b) are examples of rules of this kind.

less of subsequent improvements in understanding the effects of monetary policy on the economy, including experience with the consequences of implementing the rule. If the private sector is forward looking, and it is possible for the central bank to make the private sector aware of its policy commitments, then there *are* important advantages of commitment to a policy other than discretionary optimization—that is, simply doing what seems best at each point in time, with no commitment regarding what may be done later. This is because there are advantages to having the private sector able to anticipate *delayed* responses to a disturbance that may not be optimal *ex post* if one reoptimizes taking the private sector's past reaction as given. But one can create the desired anticipations of subsequent behavior—and justify them—without committing to following a fixed rule in the future no matter what may happen in the meantime.

It is enough that the private sector have no grounds to forecast that the bank's behavior will be *systematically* different from the rule that it pretends to follow. This will be the case if the bank is committed to choosing a rule of conduct that is justifiable on certain *principles*, given its model of the economy. (An example of the sort of principles that I have in mind is given in Chapter 8.) The bank can then properly be expected to continue to follow its current rule, as long as its understanding of the economy does not change. Furthermore, as long as there is no *predictable* direction in which its future model of the economy should be different from its current one, private-sector expectations should not be different from those in the case of an indefinite commitment to the current rule. Yet changing to a better rule remains possible in the case of improved knowledge (which is inevitable); and insofar as the change is justified in terms of both established principles and a change in the bank's model of the economy that can itself be defended, this need not impair the credibility of the bank's professed commitments.

It follows that rule-based policymaking necessarily means a decision process in which an explicit *model* of the economy (albeit one augmented by judgmental elements) plays a central role, both in the deliberations of the policy committee and in explanation of those deliberations to the public. This too has been a prominent feature of recent innovations in the conduct of monetary policy by the inflation-targeting central banks. While there is undoubtedly much room for improvement in both current models and current approaches to the use of models in policy deliberations, one can only expect the importance of models to policy deliberations to increase in a world of increasingly sophisticated financial markets.

2.2 Pitfalls of Conventional Optimal Control

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 trade-offs 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, as in the kind of industrial problems of typical interest in engineering control theory. 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 suboptimal 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, does 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 trade-off 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 variety of solutions to the problem of inflation bias have been proposed. One influential idea is that this bias can be eliminated by assigning the central bank targets for inflation and output that differ from those reflected in the true social welfare function (i.e., the central-bank objective assumed by Kydland and Prescott or Barro and Gordon), without otherwise constraining the central bank's discretion in the selection of policies to achieve its objective. This is one of the primary reasons for the popularity of inflation targeting, which involves commitment of a central bank to the pursuit of an assigned target rather than being left to simply act as seems best for society at any point in time, while leaving the bank a great deal of flexibility as to the way in which the assigned goal is to be pursued.

However, 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, as shown in Chapter 7. For example, various types of real disturbances can create temporary fluctuations in what Wicksell called the “natural rate of interest,” meaning (as shown in Chapter 4) that the level of nominal interest rates required to stabilize both inflation and the output gap varies over time. 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 on 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 is preferable; that is, the anticipation that the central bank will follow such a policy leads to a preferable rational-expectations equilibrium. 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 that 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 targets that the bank is directed to aim for regardless of what disturbances may occur. Rather, policy must be made *history dependent*, that is, 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. Indeed, 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 modification does not introduce any history dependence. But other popular proposals are often purely forward looking as well. Thus the classic “Taylor rule” (Taylor, 1993) prescribes setting an interest-rate operating target at each decision point as a function of current estimates of inflation and the output gap only (see below), and Taylor (1999b) expresses skepticism about the desirability of partial-adjustment dynamics of the kind that characterize most estimated central-bank reaction functions. Popular descriptions of inflation-forecast targeting are typically purely forward looking as well. The interest-rate setting at each decision point is to be determined purely as a function of the forecast from that date forward for inflation (and possibly other target variables). Thus the intuition that optimal policy *should* be purely forward looking seems to be fairly commonplace; but when the private sector is forward looking, any purely forward-looking criterion for policy is almost invariably suboptimal.

Obtaining a more desirable pattern of responses to random disturbances therefore requires commitment to a systematic policy rule, and not just a (one-time) adjustment of the bank's targets. The primary task of this study is to provide principles that can be used in the design of such rules. By saying that a *policy rule* is necessary, I mean to draw a distinction with two other conceptions of optimal policy. One is discretionary optimization, as just discussed; specifying a rule means a more detailed description of the way in which a decision is to be reached than is involved in a simple commitment to a particular objective. But I also mean to distinguish the approach advocated here from the usual understanding of what an *optimal commitment* involves.

In the literature that contrasts policy commitment with discretionary policymaking, following Kydland and Prescott, “commitment” is generally taken to mean a specification, once and for all, of the state-contingent action to be taken at each subsequent date. An optimal commitment is then a choice of such a state-contingent plan so as to maximize the ex ante expected value of the policymaker's objective, as evaluated at the initial date t_0 at which the commitment is chosen. This leads to a description of optimal

policy in terms of a specification of the instrument setting as a function of the history of exogenous shocks since date t_0 .

But the solution to such an optimization problem is not an appealing policy recommendation in practice. For it is generally not *time consistent*: Solving the same optimization problem at a later date t_1 to determine the optimal commitment from *that* date onward does not result in a state-contingent plan from date t_1 onward that continues the plan judged to be optimal at date t_0 . This is because the commitment chosen at date t_0 takes account of the consequences of the commitments made for dates t_1 and later for expectations between dates t_0 and t_1 , while at date t_1 these expectations are taken as historical facts that cannot be changed by the policy chosen from then on. (This is just the reason why discretionary optimization does not lead to the same policy as an optimal commitment.) Hence this policy proposal cannot be regarded as proposing a decision procedure that can be used at each date to determine the best action at that date. Rather, a state-contingent plan must be determined once and for all, for the rest of time, and thereafter simply implemented, whether it continues to appear desirable or not.

Such a proposal is not a practical one for two reasons. First, enumeration in advance of all of the possible subsequent histories of shocks is not feasible—the kinds of situations that the central bank may face at a given date are too various to possibly be listed in advance. Second, the arbitrariness of continuing to stick to a particular specified policy simply because it looked good at a particular past date—the date t_0 at which one happened to make the commitment—is sufficiently unappealing that one cannot imagine a central bank binding itself to behave in that way or the private sector believing that it had. Here my argument is not that central bankers are *incapable* of commitment to a systematic rule of conduct, so that they are inevitably discretionary optimizers; it is rather that their commitment must be based upon an understanding of the rational justification of the rule, and not on the mere fact that it happens to have been chosen (even by themselves) on a past occasion.

Both problems can be avoided by commitment to a systematic rule for determining their policy action at each decision point that does not reduce to a once-and-for-all specification of the instrument setting as a function of the history of shocks. In Chapter 8, it is shown that one can design rules for setting the central bank's interest-rate operating target that lead to optimal dynamic responses to shocks, without the rule specification having to refer to the various disturbances that may have occurred. The disturbances affect the instrument setting, of course. But they affect it either as a result of having influenced endogenous variables, such as inflation and output, to which the instrument setting responds or as a result of being factored into the central bank's projections of the future evolution of the economy under alternative

possible instrument settings. Such a rule can result in optimal equilibrium responses to disturbances of any of a vast number of possible types, so that the potential disturbances need not even be listed in advance in order to describe the rule and evaluate its desirability.

The optimal rules derived in accordance with the principles set out in Chapter 8 are also time invariant in form. This means that the optimal rule that would be derived at date t_0 , on the basis of a particular structural model of the monetary transmission mechanism and a particular understanding of the central bank's stabilization objectives, is also derived at date t_1 , assuming that the bank's model and objectives remain the same. A commitment to conduct policy in accordance with a rule that is judged optimal by this criterion is thus *time consistent*, in the sense that reconsideration of the matter at a later date on the basis of the same principle leads to a decision to continue the same course of action as had been intended earlier.¹¹

Because of this, adherence to a policy rule need not be taken to mean adoption of a rule at some initial date, after which the rule is followed blindly, without ever again considering its desirability. Instead, *rule-based policymaking* as the term is intended here means that at each decision point an action is taken that conforms to a policy rule, which rule is itself one that is judged to be optimal (from a "timeless perspective" that is explained in Chapter 7) given the central bank's understanding of the monetary transmission mechanism *at the time that the decision is made*. The desire to follow a rule (and so to avoid the trap of discretionary optimization) does not mean that the bank must refrain from asking itself whether adherence to the rule is consistent with its stabilization objectives. It simply means that whenever this question is taken up, the bank should consider what an optimal *rule* of conduct would be, rather than asking what an optimal *action* is on the individual occasion, and that it should consider the desirability of alternative rules from an impartial perspective that does not amount to simply finding a rationalization for the action it would like to take on this particular occasion.¹² A central bank might reconsider this question as often as it likes,

11. Note that "time consistency" in the sense that I use the term here does *not* mean that the policymaker does not believe at any time that it is possible to achieve a higher expected value for its objective by deviating from its intended rule. Time consistency does not require this, because this is not the criterion according to which the central bank's action is judged to be optimal at *any* time, including the initial date t_0 .

12. The distinction between these two perspectives is similar to the distinction that is made in ethical theory between "rule utilitarianism" and "act utilitarianism" (Brandt, 1959; Harsanyi, 1982). Act utilitarianism is the view that the right act on any occasion is the one that will maximize social utility in the situation that the actor is in at that time. Rule utilitarianism instead maintains that a right act is one that conforms to the correct rule for this sort of situation, where a correct rule is one that would maximize social utility if *always* followed in all situations of this type.

without being led into the kind of suboptimal behavior that results from discretionary optimization. Moreover, when considering the desirability of a policy rule, it is correct for the bank to consider the effects of its being expected to follow the rule indefinitely, even though it does not contemplate *binding* itself to do so; for as long as its view of the policy problem does not change (which it has no reason to expect), a commitment to rule-based policymaking should guarantee that it will continue to act according to the rule judged to be optimal.

Rule-based policymaking in this sense avoids the sorts of rigidity that are often associated with commitment to a “rule” and that probably account for much of the resistance that central bankers often display toward the concept of a policy rule. A commitment to rule-based policymaking does not preclude taking account of all of the information, from whatever sources, that the central bank may have about current economic conditions, including the recognition that disturbances may have occurred that would not have been thought possible a few months earlier. For a policy rule need not specify the instrument setting as a function of a specified list of exogenous states, and indeed it is argued in Chapter 7 that an optimal rule should in general *not* take this form. Nor does it preclude changing the form of the policy rule when the bank’s view of the monetary transmission changes, as it surely will, owing both to institutional change in economies themselves and to the progress of knowledge in economics. Hence it allows the sort of flexibility that is often associated with the term “discretion,” while at the same time eliminating the systematic biases that follow from policy analysis that naively applies dynamic-programming principles.

3 Monetary Policy without Control of a Monetary Aggregate

Thus far I have discussed the desirability of a monetary policy rule without saying much about the precise form of rule that is intended. To be more concrete, the present study considers the design of a rule to be used in determining a central bank’s operating target for a short-term nominal interest rate. This target will ordinarily be revised at intervals of perhaps once a month (as at the ECB) or eight times a year (as in the United States).¹³

My focus on the choice of an *interest-rate rule* should not surprise readers familiar with the current practice of central banks. Monetary policy decisionmaking almost everywhere means a decision about the operating target for an overnight interest rate, and the increased transparency about policy in recent years has almost always meant greater explicitness about

13. The question of the optimal frequency of reconsideration of the interest-rate target is one of obvious practical interest. But I shall not take it up in this study, as I consider optimal policy in the context of a discrete-time model of the transmission mechanism with “periods” corresponding to the length of the central bank’s decision cycle.

the central bank's interest-rate target and the way in which its interest-rate decisions are made. In such a context, it is natural that adoption of a policy rule should mean commitment to a specific procedure for deciding what interest-rate target is appropriate.

Nonetheless, theoretical analyses of monetary policy have until recently almost invariably characterized policy in terms of a path for the money supply, and discussions of policy rules in the theoretical literature have mainly considered money-growth rules of one type or another. This curious disjunction between theory and practice predates the enthusiasm of the 1970s for monetary targets. Goodhart (1989) complains of “an unhelpful dichotomy, between the theory and the reality of Central Bank operations,” which equally characterized the work of John Maynard Keynes and Milton Friedman:

When either of these two great economists would discuss practical policy matters concerning the level of short-term interest rates, they had no doubts that these were normally determined by the authorities, and could be changed by them, and were not freely determined in the market. . . . But when they came to their more theoretical papers, they often reverted to the assumption that the Central Bank sets the nominal money stock, or alternatively fixes the level of the monetary base, [with] the demand and supply of money . . . equilibrated in the short run . . . by market-led developments in nominal interest rates. (pp. 330–331)

The present study seeks to revive the earlier approach of Knut Wicksell and considers the advantages of systematic monetary policies that are described in terms of rules for setting a nominal interest rate. While the implied evolution of the money supply is sometimes discussed, the question is often ignored. Some of the time, I do not bother to specify policy (or an economic model) in sufficient detail to determine the associated path of the money supply, or even to tell if one can be uniquely determined in principle. Some readers may fear as a result that I consider an ill-posed question—that the “policy rules” studied here may not represent sufficiently complete descriptions of a policy to allow its consequences to be determined or may not represent states of affairs that the central bank is able to bring about. Hence some general remarks may be appropriate about why it is possible to conceive of the problem of monetary policy as a problem of interest-rate policy before turning to examples of the specific types of interest-rate rules that I wish to consider.

3.1 Implementing Interest-Rate Policy

An argument that is sometimes advanced for specifying monetary policy in terms of a rule for base-money growth rather than an interest-rate rule is that central banks do not actually fix overnight interest rates. Even when

banks have an operating target for the overnight rate, they typically seek to implement it through open-market operations in Treasury securities or their equivalent—that is, by adjusting the supply of central-bank liabilities to a level that is expected to cause the market for overnight cash to clear near the target rate. Thus it may be argued that the action that the central bank actually takes each day is an adjustment of the nominal magnitude of the monetary base, so that a complete specification of policy should describe the size of this daily adjustment.

But even when banks implement their interest-rate targets entirely through quantity adjustments, as is largely correct as a description of current U.S. arrangements, this conclusion hardly follows. Central banks like the U.S. Federal Reserve determine their quantity adjustments through a two-step procedure: first the interest-rate target is determined by a monetary policy committee (the Federal Open Market Committee in the United States) without consideration of the size of the implied open-market operations, and then the appropriate daily open-market operations required to maintain the funds rate near the target are determined by people closer to the financial markets (mainly the Trading Desk at the New York Fed). The higher-level policy decision about the interest-rate target is the more complicated one, made much less frequently because of the complexity of the deliberations involved,¹⁴ and it is accordingly this decision with which the present study is concerned.

Nor is it the case that a central bank's interest-rate target *must* be implemented through choice of an appropriate supply of central-bank liabilities. A central bank can also influence the interest rate at which banks lend overnight cash to one another through adjustment of the interest rate paid on overnight balances held at the central bank and/or the interest rate at which the central bank is willing to lend overnight cash to banks that run overdrafts on their clearing accounts at the central bank. These are important policy tools outside the United States, and in some countries are the primary means through which the central bank implements its interest-rate targets.

As is discussed in more detail in Woodford (2001c), countries like Canada, Australia, and New Zealand now implement monetary policy through a “channel system.” In a system of this kind, the overnight interest rate is kept near the central bank's target rate through the provision of standing

14. The comparative simplicity of the decision about each day's open-market operation is not so much because each day's demand for Fed balances is highly predictable as because the Fed learns immediately how much it has misjudged market demand each day and can act the following day in response to the previous day's gap between the actual funds rate and the target rate. Owing to intertemporal substitution in the demand for reserves under U.S. regulations, a credible commitment by the Fed to respond the following day is enough to keep the funds rate from deviating too much from the target most of the time. See Taylor (2001) for further discussion of the Trading Desk's reaction function.

facilities by the central bank, with interest rates determined by the central bank's current target interest rate $\bar{i}_t$. In addition to supplying a certain aggregate quantity of clearing balances (adjusted through open-market operations), the central bank offers a lending facility through which it stands ready to supply an arbitrary amount of additional overnight balances at an interest rate determined by a fixed spread over the target rate (i.e., $i_t^l = \bar{i}_t + \delta$). In the countries just mentioned, the spread δ is generally equal to 25 basis points, regardless of the level of the target rate. Finally, depository institutions that settle payments through the central bank also have the right to maintain excess clearing balances overnight with the central bank at a deposit rate $i_t^d = \bar{i}_t - \delta$, where δ is the same fixed spread.

The lending rate, on the one hand, and the deposit rate, on the other, then define a channel within which overnight interest rates should be contained.¹⁵ Because these are both standing facilities (unlike the Fed's discount window in the United States), no bank has any reason to pay another bank a higher rate for overnight cash than the rate at which it could borrow from the central bank. Similarly, no bank has any reason to lend overnight cash at a rate lower than the rate at which it can deposit with the central bank. The result is that the central bank can control overnight interest rates within a fairly tight range regardless of what the aggregate supply of clearing balances might be; frequent quantity adjustments accordingly become less important.

Woodford (2001c) describes a simple model of overnight interest-rate determination under such a system. In this model, the daily demand for clearing balances by depository institutions depends only on the location of the interbank market rate relative to the channel established by the two standing facilities, rather than on the absolute level of this interest rate. The interbank market then clears at an interest rate

$$i_t = i_t^d + F(-S_t/\sigma_t)(i_t^l - i_t^d), \quad (3.1)$$

where S_t is the aggregate supply of clearing balances (determined by the central bank's open-market operations), σ_t is a factor measuring the degree of uncertainty about payment flows on a given day, and F is a cumulative distribution function that increases monotonically from 0 (when its argument is $-\infty$) to 1 (as the argument approaches $+\infty$).

15. It is arguable that the actual lower bound is somewhat above the deposit rate because of the convenience and lack of credit risk associated with the deposit facility, and similarly that the actual upper bound is slightly above the lending rate because of the collateral requirements and possible stigma associated with the lending facility. Nonetheless, market rates are observed to stay within the channel established by these rates (except for occasional slight breaches of the upper bound during the early months of operation of Canada's system—see Figure 1.1), and typically near its center.

As noted, the market overnight rate is necessarily within the channel: $i_t^d \leq i_t \leq i_t^l$. Its exact position within the channel should be a decreasing function of the supply of central-bank balances. The model predicts an equilibrium overnight rate at exactly the target rate (the midpoint of the channel) when the supply of clearing balances is equal to

$$S_t = -F^{-1}(1/2) \sigma_t. \quad (3.2)$$

If the probability distribution of unexpected payment flows faced by each institution is roughly symmetric, so that $F(0)$ is near one-half, then the aggregate supply of clearing balances required to maintain the overnight rate near the target rate should not vary much with changes in σ_t . Even if this is not quite true, the adjustments of the supply of clearing balances required by (3.2) are unrelated to changes in the target level of interest rates.

Thus achievement of the central bank's operating target does not require any quantity adjustments through open-market operations in response to deviations of the market rate from the target rate; nor are any changes in the supply of central-bank balances required when the bank wishes to change the level of overnight interest rates. The target level of clearing balances in the system (3.2) has to be adjusted only in response to "technical" factors (e.g., changes in the volume of payments on certain days that can be expected to affect the σ_i), but *not* on occasions when it is desired to "tighten" or "loosen" monetary policy. Instead, changes in the level of overnight rates, when desired, are brought about through the shifts in the deposit rate and lending rate that automatically follow from a change in the target rate (and constitute the operational meaning of such a change), without any need for quantity adjustments.

This type of system has proven highly effective in Canada, Australia, and New Zealand in controlling the level of overnight interest rates. For example, Figure 1.1 plots the overnight rate in Canada since the adoption of the Large-Value Transfer System for payments in February 1999, at which time the standing facilities described previously were adopted.¹⁶ One observes that the channel system has been quite effective, at least since early in 2000, in keeping the overnight interest rate not only within the Bank's 50-basis-point operating band or channel, but usually within about 1 basis point of the target rate. Australia and New Zealand similarly now achieve considerably tighter control of overnight interest rates than is realized under the current operating procedures employed in the United States.¹⁷ For purposes of comparison, Figure 1.2 plots the federal funds rate together

16. A system of the kind described here has been used in Australia since June 1998, and in New Zealand since March 1999.

17. See Woodford (2001c) for corresponding plots for the other two countries, and for discussion of the differences in the four countries' ability to respond to the Y2K panic without loss of control of short-term interest rates.




Figure 1.1 The channel or operating band and the market overnight rate since introduction of the LVTs system in Canada. Source: Bank of Canada.

with the Fed's operating target over the same time period. Note that the standard deviation of the gap between the target and actual federal funds rate is much larger in the United States than in the countries using channel systems (about 1 basis point for Canada, but over 10 basis points for the United States), even under normal circumstances. At times of particular uncertainty about money demand (the Y2K panic, the 9/11/01 attacks on the United States), the deviations from the target funds rate in the United States have been much larger, whereas the same events had little effect on overnight rates in countries like Canada.

Thus the quantity adjustments of the supply of central-bank balances¹⁸ that are involved in implementation of interest-rate policy are quite different

18. I refer here to adjustments of the supply of central-bank balances rather than adjustments of the monetary base because in all of the countries under discussion, changes in the public's demand for currency are automatically accommodated by open-market operations that change the monetary base while seeking to insulate the supply of central-bank balances from the effects of such developments. Thus despite the emphasis of the academic literature on monetary-base rules, in practice a quantity-targeting rule that is intended to directly specify the central bank's daily open-market operation would have to specify a target supply of central-bank balances rather than a target value for the monetary base.




Figure 1.2 The U.S. federal-funds rate and the Fed's operating target. Source: Federal Reserve Board.

under a channel system as opposed to the system used in the United States. In the latter, policy can be tightened *only* by restricting the supply of Fed balances, so that the equilibrium spread between the return available on interbank lending and that available on Fed balances increases; in Canada, on the other hand, there need be no change in supply, as there is no desire to change the spreads $i_t^l - i_t$ or $i_t - i_t^d$. Yet there is no reason to believe that these institutional details have any important consequences for the effects of interest-rate policy on these economies, and hence for the way in which it makes sense for these different central banks to determine their interest-rate operating targets. It follows that these conclusions would be of less universal validity if I were to formulate them in terms of a rule for determining the appropriate size of open-market operations, assuming American institutional arrangements.

Furthermore, for a country with a channel system, it would *not be possible* to formulate this advice in terms of a quantity-targeting rule. On the occasions upon which it is appropriate for the central bank to tighten or loosen policy, this does not imply any change in the appropriate target for

the supply of central-bank balances; yet this does not at all mean that the central bank should not act! Because the crucial policy instruments in these countries are in fact the interest rates associated with the two standing facilities, which are in turn directly based on the time-varying interest-rate target, a policy rule for such countries must necessarily be formulated as an interest-rate rule. In fact, this way of specifying monetary policy is equally convenient for a country like the United States, and is the one that I use in this study.

3.2 Monetary Policy in a Cashless Economy

Another case in which a monetary policy prescription would *have* to be specified in terms of an interest-rate rule would be if the foregoing advice were to be applicable to a “cashless” economy, by which I mean an economy in which there are *no monetary frictions* whatsoever. In a hypothetical economy of this kind, central-bank liabilities have no special role to play in the payments system that results in a willingness to hold them despite the fact that they yield a lower return than other, equally riskless short-term claims. Consideration of this extreme case is of interest for two reasons.

First, it is possible to imagine that in the coming century the development of electronic payments systems could not only substitute for the use of currency in transactions, but also eliminate any advantage of clearing payments through accounts held at the central bank, as discussed by King (1999). This prospect is highly speculative at present; most current proposals for variants of “electronic money” still depend upon the final settlement of transactions through the central bank, even if payments are made using electronic signals rather than old-fashioned instruments such as paper checks.¹⁹ Yet it is possible that in the future central banks will face the problem of what their role should be in such a world. Moreover, the question of how the development of electronic money should be regulated will face them much sooner. If one takes the view that monetary policy can be implemented only by rationing the supply of something that fulfills an essential function in the payments system, it is likely to be judged important to *prevent* the development of alternatives to payments using central-bank money in order to head off a future in which the central bank is unable to do anything at all on behalf of macroeconomic stabilization—in which it becomes “an army with only a signal corps,” in the evocative phrase of Benjamin Friedman (1999).

19. Charles Freedman (2000), for one, argues that the special role of central banks in providing for final settlement is unlikely ever to be replaced, owing to the unimpeachable solvency of these institutions, as government entities that can create money at will. Some, such as Goodhart (2000), equally doubt that electronic media can ever fully substitute for the use of currency.

A second reason why it is useful to consider policy implementation in this hypothetical case is that if I can show that effective interest-rate control is possible even in the complete absence of monetary frictions, it may well simplify the analysis basic issues in the theory of monetary policy to start from the frictionless case, just as a physicist does when analyzing the motion of a pendulum or the trajectory of a cannonball. The appeal of this analytical approach was clear already to Wicksell (1898), who famously began his analysis (though writing at the end of the nineteenth century!) by considering the case of a “pure credit economy,” defined as

A state of affairs in which money does not actually circulate at all, neither in the form of coin (except perhaps as small change) nor in the form of notes, but where all domestic payments are effected by means of . . . bookkeeping transfers. (p. 70)

This is the approach that is taken in the chapters to follow. The basic model (developed beginning in Chapter 2) is one that abstracts from monetary frictions, in order to focus attention on more essential aspects of the monetary transmission mechanism, such as the way that spending decisions depend on expected future interest rates as well as current ones or the way in which fluctuations in nominal expenditure affect real activity. I then pause at various points to consider the modifications of the analysis that are required in order to take account of the monetary frictions that evidently exist, given the observation that non-interest-earning currency continues to be held. It is shown that, as a quantitative matter, these modifications are of relatively minor importance.

In the discussion of interest-rate determination under a present-day channel system, I have supposed that there is a demand for at least a small quantity of central-bank balances for clearing purposes, and these are held despite the existence of a small opportunity cost (25 basis points on average). But once the idea has been accepted that the central bank can vary the overnight interest rate without ever having to vary the size of this return spread, the functioning of the system no longer depends on the existence of a clearing demand. Consider instead the supposition that balances held with the central bank cease to be any more useful to commercial banks than any other equally riskless overnight investment. In this case, the demand for central-bank balances is zero for all interest rates higher than the deposit rate i_t^d . But banks should still be willing to hold arbitrary balances at the central bank as long as the market overnight rate is no higher than the rate paid by the central bank. In this case, it would no longer be possible to induce the overnight cash market to clear at a target rate higher than the rate paid on overnight balances at the central bank; for equation (3.1) reduces to $i_t = i_t^d$ in the case of any positive supply of central-bank balances.

But the central bank could still control the equilibrium overnight rate, by using the deposit rate as its policy instrument.²⁰ Such a system would differ from current channel systems in that an overnight lending facility would no longer be necessary, so that there would no longer be a “channel.”²¹ Furthermore, the rate paid on central-bank balances would no longer be set at a fixed spread δ below the target overnight rate, but rather at exactly the target rate.

Perfect control of overnight rates should still be possible through adjustments of the rate paid on overnight central-bank balances, and changes in the target overnight rate would not have to involve any change in the target supply of central-bank balances, just as is true under current channel systems. Indeed, in this extreme case, any variations that did occur in the supply of central-bank balances would cease to have any effect at all upon the equilibrium overnight rate.

But how can interest-rate variation be achieved without any adjustment at all of the supply of central-bank balances? Informal discussions often treat interest-rate control by the central bank like a species of price control. Certainly, if a government decides to peg the price of some commodity, it may be able to do so, but only by holding stocks of the commodity that are sufficiently large relative to its world market, and by standing ready to vary those holdings by large amounts as necessary. If the market in question is a large one (more to the point, if either supply or demand in the market is relatively price elastic) relative to the size of the balance sheet of the government entity seeking to control the price, one doubts that such efforts will be effective. What is different about controlling short-term nominal interest rates?

The difference is that there is no inherent “equilibrium” level of interest rates to which the market would tend in the absence of central-bank intervention and against which the central bank must therefore exert a significant countervailing force in order to achieve a given operating

20. Grimes (1992) makes a related point, showing that variation of the interest rate paid on central-bank balances would be effective in an environment in which central-bank reserves are no more useful for carrying out transactions than other liquid government securities, so that open-market purchases or sales of such securities are completely ineffective. Hall (1983, 2002) has also proposed this as a method of price-level control in the complete absence of monetary frictions. Hall speaks of control of the interest yield on a government “security,” without any need for a central bank at all. But because of the special features that this instrument would need to possess, which are not possessed by privately issued securities—it is a claim only to future delivery of more units of the same instrument, and society’s unit of account is defined in terms of this instrument—it seems best to think of it as still taking the same institutional form that it does today, namely, balances in an account with the central bank.

21. This presumes a world in which no payments are cleared using central-bank balances. Of course, there would be no harm in continuing to offer such a facility as long as the central-bank clearing system were still used for at least some payments.

target.²² This is because there is no inherent value (in terms of real goods and services) for a fiat unit of account such as the “dollar,” except insofar as a particular exchange value results from the monetary policy commitments of the central bank. The basic point was clear to Wicksell (1898, pp. 100–101), who compares relative prices to a pendulum that always returns to the same equilibrium position when perturbed, while the money prices of goods in general are compared to a cylinder resting on a horizontal plane, which can remain equally well in any location on the plane to which it may happen to be moved.²³ Alternative price-level paths are thus equally consistent with market equilibrium in the absence of any intervention that would vary the supply of any real goods or services to the private sector; and associated with these alternative paths for the general level of prices are alternative paths for short-term nominal interest rates.

Of course, this analysis might suggest that while central banks can bring about an arbitrary level of *nominal* interest rates (by creating expectations of the appropriate rate of inflation), they should not be able to significantly affect *real* interest rates, except through trades that are large relative to the economy that they seek to affect. It may also suggest that they should be able to move nominal rates only by altering inflation expectations; yet banks generally do not feel that they can easily alter expectations of inflation over the near term, so that one might doubt that they should be able to affect *short-term* nominal rates through such a mechanism.

However, once one recognizes that many prices (and wages) are fairly sticky over short time intervals, the arbitrariness of the path of nominal prices (in the sense of their underdetermination by real factors alone) implies that the path of real activity and the associated path of equilibrium real interest rates are equally arbitrary. It is equally possible, from a logical standpoint, to imagine allowing the central bank to determine, by arbitrary fiat, the path of aggregate real activity, or the path of real interest rates, or the path of nominal interest rates as it is to imagine allowing it to determine the path of nominal interest rates.²⁴ In practice, it is easiest for central

22. This does not mean that Wicksell’s notion of a natural rate of interest determined by real factors is of no relevance to the consideration of the policy options facing a central bank. It is indeed, as argued in Chapter 4. But the natural rate of interest is the rate of interest required for *an equilibrium with stable prices*; the central bank nonetheless can arbitrarily choose the level of interest rates (within limits), because it can choose the degree to which prices will increase or decrease.

23. This is the grounds for his argument—in the quotation from the introduction to his book that begins this chapter—that control of the general level of prices involves no interference with the market mechanism of the kind that is required if some relative price is to be controlled.

24. This does not mean, of course, that absolutely any paths for these variables can be achieved through monetary policy; the chosen paths must be consistent with certain constraints

banks to exert relatively direct control over overnight nominal interest rates, and so they generally formulate their short-run objectives (their operating target) in terms of the effect that they seek to bring about in this variable rather than in one of the others.

Even recognizing the existence of a very large set of rational-expectations equilibria—equally consistent with optimizing private-sector behavior and with market clearing, in the absence of any specification of monetary policy—one might nonetheless suppose, as Fischer Black (1970) once did, that in a fully deregulated system the central bank should have no way of using monetary policy to select among these alternative equilibria. The path of money prices (and similarly nominal interest rates, nominal exchange rates, and so on) would then be determined solely by the self-fulfilling expectations of market participants. Why should the central bank play any special role in determining which of these outcomes should actually occur if it does not possess any monopoly power as the unique supplier of some crucial service?

The answer is that the unit of account in a purely fiat system is *defined* in terms of the liabilities of the central bank.²⁵ A financial contract that promises to deliver a certain number of U.S. dollars at a specified future date is promising payment in terms of Federal Reserve notes or clearing balances at the Fed (which are treated as freely convertible into one another by the Fed). Even in the technological utopia imagined by the enthusiasts of electronic money—where financial-market participants are willing to accept as final settlement transfers made over electronic networks in which the central bank is not involved—if debts are contracted in units of a national currency, then clearing balances at the central bank still define the thing to which these other claims are accepted as equivalent.

This explains why the nominal interest yield on clearing balances at the central bank can determine overnight rates in the market as a whole. The central bank can obviously define the nominal yield on overnight deposits in its clearing accounts as it chooses; it is simply promising to increase the nominal amount credited to a given account, after all. It can also determine

implied by the conditions for a rational-expectations equilibrium, e.g., those presented in Chapter 4. But this is true even in the case of the central bank's choice of a path for the price level. Even in a world with fully flexible wages and prices, for instance, it would not be possible to bring about a rate of deflation so fast as to imply a negative nominal interest rate.

25. See Hall (2002) and White (2001) for expressions of similar views. White emphasizes the role of legal-tender statutes in defining the meaning of a national currency unit. But such statutes do not represent a restriction upon the means of payment that can be used within a given geographical region—or at any rate, there need be no such restrictions upon private agreements for the point to be valid. What matters is simply what contracts written in terms of a particular unit of account are taken to mean, and the role of law in stabilizing such meanings is essentially no different than, say, in the case of trademarks.

this independently of its determination of the quantity of such balances that it supplies. Commercial banks may exchange claims to such deposits among themselves on whatever terms they like. But the market value of a dollar deposit in such an account cannot be anything other than a dollar—*because this defines the meaning of a “dollar”!*

This places the Fed in a different situation than any other issuer of dollar-denominated liabilities.²⁶ Citibank can determine the number of dollars that one of its jumbo CDs will be worth at maturity, but must then allow the market to determine the current dollar value of such a claim; it cannot determine *both* the quantity that it wishes to issue of such claims and the interest yield on them. Yet the Fed can, and does so daily—though at present it chooses to fix the interest yield on Fed balances at zero and only to vary the supply. The Fed’s current position as monopoly supplier of an instrument that serves a special function is necessary in order for variations in the quantity supplied to affect the equilibrium spread between this interest rate and other market rates, but *not* in order to allow separate determination of the interest rate on central-bank balances and the quantity of them in existence.

Yes, some may respond, a central bank would still be able to determine the interest rate on overnight deposits at the central bank, and thus the interest rate in the interbank market for such claims, even in a world of completely frictionless financial markets. But would control of *this* interest rate necessarily have consequences for other market rates, the ones that matter for critical intertemporal decisions such as investment spending? The answer is that it must—and all the more so in a world in which financial markets have become highly efficient, so that arbitrage opportunities created by discrepancies among the yields on different market instruments are immediately eliminated. Equally riskless short-term claims issued by the private sector (say, shares in a money-market mutual fund holding very short-term Treasury bills) would not be able to promise a different interest rate than the one available on deposits at the central bank; otherwise, there would

26. Costa and De Grauwe (2001) instead argue that “in a cashless society . . . the central bank cannot ‘force the banks to swallow’ the reserves it creates” (p. 11), and speak of the central bank being forced to “liquidate . . . assets” in order to redeem the central-bank liabilities that commercial banks are “unwilling to hold” in their portfolios. This neglects the fact that the definition of the U.S. dollar allows the Fed to honor a commitment to pay a certain number of dollars to account holders the next day by simply crediting them with an account of that size at the Fed—there is no possibility of demanding payment in terms of some other asset valued more highly by the market. Similarly, Costa and De Grauwe argue that “the problem of the central bank in a cashless society is comparable to [that of a] central bank pegging a fixed exchange rate” (footnote 15). But the problem of a bank seeking to maintain an exchange-rate peg is that it promises to deliver a foreign currency in exchange for its liabilities, not liabilities of its own that it freely creates. Costa and De Grauwe say that they imagine a world in which “the unit of account remains a national affair . . . and is provided by the state” (p. 1), but seem not to realize that this means defining that unit of account in terms of central-bank liabilities.

be an excess supply or demand for the private-sector instruments. Moreover, determination of the overnight interest rate would also have to imply determination of the equilibrium overnight holding return on longer-lived securities, up to a correction for risk; and so determination of the expected future path of overnight interest rates would essentially determine longer-term interest rates.

The special feature of central banks, then, is simply that they are entities whose liabilities happen to be used to define the unit of account in a wide range of contracts that other people exchange with one another. There is perhaps no deep, universal reason why this need be so; it is certainly not essential that there be one such entity per national political unit. Nonetheless, the provision of a well-managed unit of account—one in terms of which the equilibrium prices of many goods and services are relatively stable—clearly facilitates economic life. Furthermore, given the evident convenience of having a single unit of account used by most of the parties with whom one wishes to trade, one may well suppose that this function should properly continue to be taken on by the government, even in a world of highly efficient information processing. I assume here a world in which central banks (whether national or supranational, as in the case of the ECB) continue to fulfill this function, and in which they are interested in managing their fiat currency in the public interest. The present study aims to supply a theory that can help them to do so.

4 Interest-Rate Rules

I have argued that the central problem of the theory of monetary policy is to provide principles that can be used in selecting a desirable rule for setting a central bank's interest-rate operating target. It is perhaps worth saying a bit more at this point about exactly what form of rules I have in mind and what sort of questions I would like to answer about them. This will provide a more concrete background for the analysis to be developed in the chapters to come.

Probably the earliest example of a prescription for monetary policy in terms of an interest-rate rule is due to Wicksell (1898, 1907). Although writing at a time when the leading industrial nations remained committed to the gold standard, with even most scholars assuming the necessity of a commodity standard of one sort or another, Wicksell foresaw the possibility of a pure fiat standard and indeed argued that it was essential for the development of "a rational monetary system." The soundness of his advocacy of price-level targeting in the context of a pure fiat monetary standard was shown during the 1930s, when Sweden abandoned the gold standard that had become an engine of worldwide deflation. The Swedish Riksbank adopted a price-level target as a substitute and used its interest-rate policy to achieve

this target as recommended by Wicksell. The policy was quite successful, especially in comparison with the price-level fluctuations suffered by many other countries during this period (Jonung, 1979). Nonetheless, the experiment was radical for its time and was not to be repeated for another 50 years.²⁷

Wicksell advocated not only price-level targeting, but a specific form of interest-rate rule for the management of such a system. His original (1898) statement of the proposed rule was as follows:

So long as prices remain unaltered the [central] banks' rate of interest²⁸ is to remain unaltered. If prices rise, the rate of interest is to be raised; and if prices fall, the rate of interest is to be lowered; and the rate of interest is henceforth to be maintained at its new level until a further movement of prices calls for a further change in one direction or the other (p. 189, italicized in original).

Wicksell's proposal can be represented mathematically as a commitment to set the central bank's interest-rate operating target i_t according to a relation of the form²⁹

$$i_t = \bar{i} + \phi p_t, \quad (4.1)$$

where p_t is the log of some general price index (the one that the policy aims to stabilize) and ϕ is a positive response coefficient, or alternatively by a rule of the form

$$\Delta i_t = \phi \pi_t \quad (4.2)$$

where $\pi_t \equiv \Delta p_t$ is the inflation rate.³⁰ I discuss price-level determination under policy rules of this kind in Chapter 2 and argue that such a rule

27. For an example of favorable academic comment on the Swedish experiment at the time, see Fisher (1934, pp. 399–410).

28. In the passage quoted here, Wicksell proposes a coordinated policy by the world's central banks in order to establish a stable world unit of account. Elsewhere he offers similar advice for the stabilization of the value of a single national unit of account.

29. Wicksell proposes nothing so specific as a log-linear relation of this kind, of course; he only describes a monotonic relationship. The log-linear specification is useful for the simple calculations of the next section. In Chapter 2, I discuss the usefulness of this sort of log-linear approximation of what is necessarily not a globally log-linear rule. The specification (4.1) cannot be maintained for all possible price levels, owing to the requirement that the nominal interest rate be nonnegative.

30. Note that a commitment to set the interest rate according to (4.2) from some date t_0 forward is equivalent to a commitment to set it according to a rule of the form (4.1), where the intercept $\bar{i}$ corresponds to $i_{t_0-1} - \phi p_{t_0-1}$. Fuhrer and Moore (1995c) propose a more complicated interpretation of Wicksell's proposal, in which the interest change is instead a

should indeed succeed in stabilizing the price index around a constant level. The principles that determine the equilibrium price level under such a regime are briefly sketched in Section 4.3 of this introduction.

A simple Wicksellian rule such as this has certain advantages as well, at least in comparison to other equally simple rules, as discussed by Giannoni (2000). Nonetheless, I do not confine my attention to rules of this kind. My primary interest in this study is in the analysis of proposals that are closer in form to the policies currently followed by many central banks. These rules involve an (explicit or implicit) target for the *inflation rate*, rather than for the *price level*; nor are they expressible solely in terms of interest-rate *changes*, so that they are equivalent to a rule that responds to the price level, as in the case of (4.2). The rules typically allow for “base drift” in the price level as a result—even if the inflation rate is kept within a narrow interval at all times, there is no long-run mean reversion in the price level, or even in the price level deflated by some deterministic target path. Moreover, as I eventually conclude in Chapter 8, optimal interest-rate rules are likely to have this property.

4.1 Contemporary Proposals

The best-known example of a proposed rule for setting interest rates is probably the one proposed by John Taylor (1993), both as a rough description of the way that policy had actually been made by the U.S. Federal Reserve under Alan Greenspan’s chairmanship and as a normative prescription (on the basis of stochastic simulations using a number of econometric models). According to the Taylor rule, as it has come to be known, the Fed’s funds-rate operating target i_t is set as a linear function of measures of the current inflation rate and the current gap between real output and potential:

$$i_t = 0.04 + 1.5(\bar{\pi}_t - 0.02) + 0.5(y_t - y_t^p), \quad (4.3)$$

where $\bar{\pi}_t$ is the rate of inflation (the change in the log GDP deflator over the previous four quarters in Taylor’s illustration of the rule’s empirical fit), y_t is log output (log real GDP in Taylor’s plot), and y_t^p is log “potential” output (log real GDP minus a linear trend in Taylor’s plot). The constants in Taylor’s numerical specification indicate an implicit inflation target of 2 percent per annum, and an estimate of the long-run real federal funds rate of 2 percent per annum as well, so that a long-run average inflation rate at the target requires a long-run average funds rate of 4 percent. A slightly

function of the price *level*. While their rule is slightly more difficult to analyze, it does not lead to substantially different conclusions about the consequences of commitment to a Wicksellian rule.

simpler rule in the same vein was proposed for the United Kingdom by Charles Goodhart (1992), according to which “there should be a presumption” that the nominal interest rate would satisfy an equation of the form

$$i_t = 0.03 + 1.5 \bar{\pi}_t,$$

and “the Governor should be asked, say twice a year, to account for any divergence from that ‘rule’ ” (p. 324).

The coefficients 1.5 and 0.5 in the Taylor rule are round figures argued to approximately characterize U.S. policy between 1987 and 1992 and that were found to result in desirable outcomes (in terms of inflation and output stability) in simulations.³¹ In Taylor’s discussions of the rule, he places particular stress upon the importance of responding to inflation above the target rate by raising the nominal interest-rate operating target by *more than the amount* by which inflation exceeds the target; the importance of this “Taylor principle” is considered in detail in Chapters 2 and 4. Taylor (1999c) argues that the Fed did not adhere to this principle before 1979 (at which time Fed chairman Paul Volcker instituted a radical shift in policy), and this failure may well have been responsible for the greater U.S. macroeconomic instability during the 1960s and 1970s. Taylor illustrates the change in policy by estimating simple Fed reaction functions of the form

$$i_t = \bar{i} + \phi_\pi(\bar{\pi}_t - \bar{\pi}) + \phi_x x_t \quad (4.4)$$

for two different sample periods, using ordinary least squares; his coefficient estimates are shown in Table 1.1. (Here I introduce the notation x_t for the *output gap*, again equated with deviations of log real GDP from trend in Taylor’s empirical work.) Nelson (2001) finds that estimates of Taylor-type rules for the United Kingdom tell a similar story. Prior to the adoption of inflation targeting in 1992, U.K. interest rates rose less than one-for-one with increases in inflation (and in the mid-1970s responded little or not at all), but since 1992, the long-run inflation response coefficient is estimated to have been nearly 1.3.

Estimates of empirical central-bank reaction functions typically find that a dynamic specification fits the data better, whatever the validity may be of Taylor’s (1999b) preference for a purely contemporaneous specification on normative grounds.³² For example, Judd and Rudebusch (1998) estimate

31. A similar form of policy rule was advocated, also on the basis of simulation studies, at around the same time by Henderson and McKibbin (1993).

32. This is also true of the estimates for the United Kingdom reported in Nelson (2001). This is why I refer to Nelson’s “long-run inflation response coefficient” in the previous paragraph, rather than to a contemporaneous response coefficient of the kind estimated by Taylor (1999c).

TABLE 1.1 Alternative estimates of Fed reaction functions^a

	ϕ_π	(s.e.)	ϕ_x	(s.e.)	γ	ρ_1	(s.e.)	ρ_2	(s.e.)
Taylor (1999c)									
1960–1979	0.81	(.06)	0.25	(.05)					
1987–1997	1.53	(.16)	0.77	(.09)					
Judd-Rudebusch (1998)									
1979–1987	1.46	(.26)	1.53	(.80)	1	0.56	(.12)		
1987–1997	1.54	(.18)	0.99	(.13)	0	0.72	(.05)	0.43	(.10)
Clarida et al. (2000)									
1960–1979	0.83	(.07)	0.27	(.08)		0.68	(.05)	*	
1979–1996	2.15	(.40)	0.93	(.42)		0.79	(.04)	*	
Orphanides (2003)									
1966–1979	1.49	(.38)	0.46	(.13)		0.68	(.07)	0.26	(.14)
1979–1995	1.89	(.64)	0.18	(.20)		0.77	(.10)	0.08	(.19)

^a Asterisks indicate values not constrained to be zero in the estimation, but not reported in the published account of regression results.

Fed reaction functions according to which the funds-rate operating target adjusts in response to changes in an implicit desired level of the funds rate $\bar{i}_t$ according to partial-adjustment dynamics of the form³³

$$i_t = (1 - \rho_1)\bar{i}_t + \rho_1 i_{t-1} + \rho_2(i_{t-1} - i_{t-2}). \quad (4.5)$$

The desired level of the funds rate in turn depends upon inflation and the output gap in a manner similar to that postulated by Taylor,

$$\bar{i}_t = \bar{i} + \phi_\pi(\bar{\pi}_t - \bar{\pi}) + \phi_x(x_t - \gamma x_{t-1}), \quad (4.6)$$

except that the allowance for nonzero γ means that the desired funds rate may respond to the rate of change of the output gap as well as (or instead of) its level. The Judd-Rudebusch estimated coefficients for two different sample periods, corresponding to the Fed chairmanships of Paul Volcker and Alan Greenspan, respectively, are also reported in Table 1.1.³⁴ Taylor's view of the nature of policy in the Greenspan period is largely confirmed, with the exception that Judd and Rudebusch estimate partial-adjustment

33. Here as in all other regressions reported in this section, periods are assumed to be quarters, and quarterly data are used in the estimation.

34. In their preferred estimates, the value of γ is imposed rather than estimated. The extreme values assumed for the separate periods, however, are suggested by preliminary regressions in which the value of γ is unconstrained.

dynamics implying substantial persistence. They give a similar characterization of policy in the Volcker period, except that the desired funds rate is found to depend on the rate of change of the output gap, rather than its level.³⁵

Many recent discussions of central-bank behavior, both positive and normative, argue instead for specifications in which a bank's operating target depends on forecasts. For example, Clarida et al. (2000) estimate Fed reaction functions of the form³⁶

$$\bar{i}_t = \bar{i} + \phi_\pi E[\pi_{t+1} - \bar{\pi} | \Omega_t] + \phi_x E[x_t | \Omega_t], \quad (4.7)$$

where Ω_t is the information set assumed to be available to the Fed when setting i_t , and the actual operating target is again related to the desired funds rate $\bar{i}_t$ through partial-adjustment dynamics. Like Taylor, these authors find an important increase in the degree to which the Fed's desired level for the funds rate responds to inflation variations since 1979, though in their specification the Fed responds to an inflation *forecast* rather than inflation that has already occurred.³⁷

Finally, it should be noted that the view that the Fed has responded more vigorously to inflation variations since 1979 has not gone unchallenged. Orphanides (2003) argues that the findings of Taylor and the other authors just cited are distorted by the use of inflation and output-gap data (especially the output-gap estimates), which were not available to the Fed at the time that its interest-rate decisions were made. When he estimates Fed reaction functions of the kind assumed by Clarida et al. using the forecasts actually produced by Fed staff at the time rather than econometric projections using the data available now, he obtains much more similar estimates for the pre-Volcker and post-Volcker periods, as shown in Table 1.1. The inflation-response coefficient ϕ_π is well above one in both periods, according to Orphanides's estimates; he instead emphasizes the reduction in the size of ϕ_x as the crucial policy change after 1979, and the key to U.S. macroeconomic

35. Judd and Rudebusch also estimate a reaction function for the period (1970–1978) corresponding to the chairmanship of Arthur Burns. Like Taylor, in this period they estimate an inflation-response coefficient ϕ_π less than one, though not significantly so in their case.

36. Clarida et al. also estimate variants of the rule in which the forecast horizon is assumed to be more than one quarter in the future. The policies of inflation-targeting central banks have often been represented by rules in which the interest-rate operating target responds to a forecast of inflation as many as 8 quarters in the future. See, e.g., Black et al. (1997b), Batini and Haldane (1999).

37. Like Taylor, they also suggest that this change has led to greater macroeconomic stability in the later period. They provide a theoretical analysis of why this could have been so, in terms of the vulnerability of an economy to instability due to self-fulfilling expectations in the case of a policy rule of the kind that they estimate for the period 1960–1979. Reasons for this are discussed in Chapter 4.

stability since the mid-1980s. I do not seek to resolve this debate about historical Fed policy here, but simply note that much current debate about both the explanation of recent U.S. policy successes and the reason for past policy failures turns upon claims concerning the desirability of particular coefficients in Taylor-type rules.

These alternative characterizations suggest a number of questions about the form of a desirable interest-rate rule. One obvious question is whether the variables to which the Fed is described as responding in the Taylor rule and the estimated reaction functions just discussed—some measures of inflation and the output gap—are ones that make sense. Is it desirable for interest rates to be adjusted in response to variations in these variables, and with the signs proposed by Taylor? Are there any grounds for thinking it more important to respond to variations in these variables than in others? Is responding to variations in these variables an adequate substitute for attempting to respond to the underlying disturbances that are perceived to be currently affecting the economy?

If it does make sense to respond to these variables, how exactly should they be defined? Which sort of price index is most appropriately used in the inflation measure? Relative to what concept of potential output should the output gap measure be defined? And how strongly is it desirable to respond to variations in these variables? Is a value of ϕ_π greater than one essential, as argued by Taylor? Is a large value of ϕ_x dangerous, as argued by Orphanides?

I am also interested in the most desirable dynamic specification of such an interest-rate rule. Are purely contemporaneous responses, as prescribed by Taylor, preferable? Is there any justification for the more inertial interest-rate dynamics indicated by the estimated reaction functions? If so, how inertial is it desirable for interest-rate policy to be? Is it preferable to respond to forecasts rather than to current or past values of inflation and the output gap? If so, how far in the future should the forecasts look?

Another type of policy rule that has figured prominently both in recent descriptions of actual central-bank behavior and in normative prescriptions is an *inflation-forecast targeting* rule. A classic example is the sort of rule that is often used to explain the current procedures of the Bank of England (e.g., Vickers, 1998). According to the formula, the Bank should be willing to adopt a given operating target i_t for the overnight interest rate at date t if and only if the Bank's forecast of the evolution of inflation over the next 2 years, conditional upon the interest rate remaining at the level i_t , implies an inflation rate of 2.5 percent per annum (the Bank's current inflation target) 2 years after date t . This is an example of what Svensson (1999, 2003b) calls a "targeting rule" as opposed to an instrument rule. No formula is specified for the central bank's interest-rate operating target. Rather, it is to be set at whatever level may turn out to be required in order for the bank's projections to satisfy a certain target criterion. That criterion need

not involve only future inflation; for example, Svensson (1999) advocates a flexible inflation-targeting rule in which the interest rate is adjusted at date t so as to ensure that

$$E_t \pi_{t+j} + \lambda E_t x_{t+k} = \bar{\pi}, \quad (4.8)$$

where $\bar{\pi}$ is the average inflation target, the coefficient $\lambda > 0$ depends on the relative importance of output-gap stabilization, and the horizons j and k are not necessarily the same distance in the future.

I also wish to consider the desirable specification of a target criterion in the case of a policy rule of this sort. Again, a basic question is whether it makes sense to define the target criterion in terms of projections for these particular variables, inflation and the output gap, rather than others, such as monetary aggregates. If so, what should determine the relative weight, if any, to be placed on the output-gap forecast? How far into the future should the forecasts in the target criterion look? And is a desirable criterion *purely* forward looking, as in the case of the two examples just mentioned, or should the inflation target be history dependent, in addition to (possibly) depending on projected future output gaps?

This study will seek to elaborate a methodology that can be used to give quantitative answers to questions of this sort about optimal policy rules. Of course, the answers obtained depend on the details of what one assumes about the nature of the monetary transmission mechanism, and I do not propose to argue for a specific quantitative rule. The aim of the present study is more to suggest a way of approaching the problem than to announce the details of its solution. However, certain model elements recur in many of the models currently used in studies of the effects of monetary policy, both in the academic literature and in central banks; and given the likelihood that a reasonable model is judged to include these features, I may obtain some tentative conclusions as to the likely form of reasonable policy rules.

4.2 General Criticisms of Interest-Rate Rules

Before taking up specific questions of these kinds about the form of desirable interest-rate rules, it is first necessary to address some more basic issues. Would *any* form of interest-rate rule represent a sensible approach to monetary policy? Proponents of monetary targeting have often argued against interest-rate control *as such*—asserting not that skill is required in the choice of an interest-rate operating target, but that it is a serious mistake to have one at all.

One famous argument, mentioned previously, is that of Sargent and Wallace (1975). These authors consider a general class of money-supply rules on the one hand, and a general class of interest-rate rules on the other,

and argue that while *any* of the money-supply rules leads to a determinate rational-expectations equilibrium (in the context of a particular rational-expectations IS-LM model), *none* of the interest-rate rules do. By *determinacy* of the equilibrium I mean that there is a unique equilibrium satisfying certain bounds, made precise in Chapter 2. Sargent and Wallace argue that interest-rate rules lead to *indeterminacy*, meaning that even if one restricts one's attention to bounded solutions to the equilibrium relations (as I largely do in this study), there is an extremely large set of equally possible equilibria. These include equilibria in which endogenous variables such as inflation and output respond to random events that are completely unrelated to economic "fundamentals" (i.e., to the exogenous disturbances that affect the structural relations that determine inflation and output) and also equilibria in which "fundamental" disturbances cause fluctuations in equilibrium inflation and output that are arbitrarily large relative to the degree to which the structural relations are perturbed. Thus in such a case, macroeconomic instability can occur owing purely to self-fulfilling expectations. This is plainly undesirable, if one's objective is to stabilize inflation and/or output.³⁸ Hence Sargent and Wallace argue that interest-rate rules can be excluded from consideration as a class; the problem of optimal monetary policy is then properly framed as a question of what the best money-supply rule would be.

However, as McCallum (1981) notes, the Sargent-Wallace indeterminacy result applies, even in the context of their own model, only in the case of interest-rate rules that specify an *exogenous* evolution for the nominal interest rate. This includes the possibility of rules that specify the nominal interest rate as a function of the history of exogenous disturbances, but not rules that make the nominal interest rate a function of *endogenous* variables, such as inflation or output. Yet the Taylor rule, and the other interest-rate rules discussed previously, are all rules of the latter sort, so that the Sargent-Wallace result need not apply. The same is shown to be true, in Chapters 2 and 4, in the case of the optimizing models of inflation and output determination considered here. Indeed, I find that, at least in the case of the simple model of the monetary transmission mechanism that is most extensively analyzed here, *either* the type of feedback from the general price level to the interest rate (or from changes in the price level to changes in the interest rate) advocated by Wicksell *or* the type of feedback from inflation and output to the central bank's interest-rate operating target prescribed by Taylor would suffice to imply a determinate rational-expectations equilibrium. In

38. The point remains valid if one's objective is, as I argue that it should be, to stabilize output relative to its natural rate, rather than output stabilization as such. For the fluctuations in output owing purely to self-fulfilling expectations will imply fluctuations in the output gap as well.

the case of a level-to-level (or change-to-change) Wicksellian specification, it is only necessary that the sign of the response be the one advocated by Wicksell. In the case of a change-to-level specification such as that proposed by Taylor, the Taylor principle mentioned earlier—the requirement that a sustained increase in inflation eventually result in an increase in nominal interest rates that is even larger in percentage points—turns out to be the critical condition that determines whether equilibrium should be determinate or not.³⁹

A related criticism of interest-rate targeting also maintains that such a policy is dangerous because it allows instability to be generated by self-fulfilling expectations, but is not based on the possibility of multiple rational-expectations equilibria. Friedman (1968) argues that attempting to control nominal interest rates is dangerous on the basis of Wicksell's (1898, 1907) famous analysis of the "cumulative process." With a nominal interest rate that is fixed at a level below the natural rate, inflation is generated that increases inflation expectations, which then stimulates demand even further owing to the reduction in the real rate, generating even faster inflation, further increasing inflation expectations, and so on without bound.⁴⁰ The same process should occur with the opposite sign if the interest rate happens to be set above the natural rate; thus any attempt to fix the nominal interest rate would seem almost inevitably to generate severe instability of the inflation rate. (In Friedman's analysis, there is no indeterminacy of the path of inflation, as inflation expectations are assumed to be a specific function of previously observed inflation.)

As is discussed in Chapter 4, this analysis can be formalized in the context of an optimizing model in which inflation and income forecasts are based on extrapolation from past data (e.g., using empirical time-series models). But once again, the classic analysis applies only in the case of a policy that exogenously specifies the path of nominal interest rates. If instead a surge in inflation and output leads to increases in nominal interest rates large enough to raise real rates, then demand should be damped, tending to lower inflation as well—so that there should be no explosive instability of either inflation or output dynamics under adaptive learning. Indeed, the analyses of Bullard and Mitra (2002) and Preston (2002a) find that conformity

39. Clarida et al. (2000) argue on this basis that the macroeconomic instability of the 1970s in the United States may have been increased by self-fulfilling expectations, given that their estimates (see Table 1.1) imply that the Taylor principle has been satisfied by post-1979 policy but not by previous policy.

40. This summarizes Friedman's account, rather than Wicksell's original discussion. Wicksell does not discuss endogenous inflation expectations and so concludes that the price level rather than the inflation rate should rise without bound. Lindahl (1939) was the first to introduce endogenous inflation expectations into the analysis and so to conclude that the inflation rate could rise without bound.

to the Taylor principle is both a necessary and sufficient condition (at least within certain simple classes of policy rules) for adaptive learning dynamics to converge to a stationary rational-expectations equilibrium, in which inflation and output fluctuate only in response to “fundamentals.”

Thus it is important to realize that these well-known criticisms of interest-rate targeting assume that under such a policy the interest-rate target would remain *fixed*, regardless of the path of inflation. The analyses are quite inapplicable in the case of policy rules such as Wicksell’s rule, the Taylor rule, or typical inflation-forecast targeting rules, which require that interest rates be raised sharply if inflation is either observed or forecasted to exceed the target rate consistently for a substantial period of time. In fact, in these conventional arguments for monetary targeting, the reason for control of money growth is precisely that this is a policy commitment that ensures that an excessive rate of inflation will lead to interest-rate increases sufficient to curb the growth of nominal expenditure. A fixed target path for the money supply (or more generally, a path that is contingent only upon exogenous state variables, not upon the path of the price level) implies that if the price level grows more rapidly, the private sector will be forced to operate with a lower level of real money balances. This will require interest-rate increases and/or a reduction in real activity sufficient to reduce desired real money balances to the level of the real money supply.

But the same kind of automatic increase in interest rates, curbing expenditure, can be arranged through a simple commitment of the central bank to raise interest rates in response to deviations of the general level of prices from its desired path, as first proposed by Wicksell. Moreover, once one recognizes that quantity control is not necessary for such a system to work, it is hard to see why one should wish to be encumbered by it. Over the course of the twentieth century, it came to be accepted that no convertibility of national currencies into a real commodity such as gold was necessary in order for central banks to act in a way that controlled the value of their currencies. Moreover, once this was accepted, it quickly became evident that nations were better off *not* relying upon such a crude mechanism as a gold standard, which left the value of the national unit of account vulnerable to fluctuations in the market for gold. Similarly, once one accepts that the adjustment of interest rates to head off undesired price-level variation can be managed by central banks without any need for so mechanical a discipline as is provided by a money-growth target, it should be clear that a properly chosen interest-rate rule can be more efficient than monetary targeting, which has the unwanted side effect of making interest rates (and hence the pace of aggregate expenditure) vulnerable to variations in the relation between desired money balances and the volume of transactions.

A more subtle criticism of interest-rate rules as an approach to systematic monetary policy would argue that even if such rules lead to well-defined,

well-behaved equilibria, the description of policy in this way may still not be useful to a central bank that wishes to understand, and thus to accurately calibrate, the consequences of its actions. It is often supposed that the key to understanding the effects of monetary policy on inflation must always be the quantity theory of money, according to which the price level is determined by the relation between the nominal money supply, on the one hand, and the demand for real money balances, on the other. It may then be concluded that what matters about *any* monetary policy is the implied path of the money supply, whether this is determined through straightforward monetary targeting or in some more indirect manner.⁴¹ From such a perspective, it might seem that a clearer understanding of the consequences of a central bank's actions would be facilitated by an explicit focus on what evolution of the money supply the bank intends to bring about—that is, by monetary targeting—rather than by talk about interest-rate policy that, even if it does imply a specific path for the money supply, does not make the intended path entirely transparent.

The present study aims to show that the basic premise of such a criticism is incorrect. One of the primary goals of Part I of this book is the development of a theoretical framework in which the consequences of alternative interest-rate rules can be analyzed, which does not require that they first be translated into equivalent rules for the evolution of the money supply. Indeed, much of the time I analyze the consequences of interest-rate rules without having to solve for the implied path of the money supply, or even having to specify the coefficients of a “money-demand” relation. In the case of an economy without monetary frictions—a case that I argue is an analytically convenient approximation for many purposes, and that may well represent the future, as discussed earlier—there will not even be any meaningful money supply or demand to be defined. If instead one takes account of the sort of frictions that evidently still exist in an economy like that of the United States at present, then the models imply an equilibrium path for the money supply along with other endogenous variables. But the factors determining the equilibrium paths of both inflation and output continue to be nearly the same as in the frictionless economy, so that it does not seem at all natural or useful to try to *explain* the predicted paths of inflation and output as *consequences* of the implied path of the money supply. Instead, it proves to be possible to discuss the determinants of inflation and output in a fairly straightforward way in terms of the coefficients of an interest-rate rule. Thus the characterizations of central-bank policy offered previously are found to be quite convenient for analysis of the consequences of one quantitative specification or another. I further show, in Chapter 8, that optimal policy

41. This is, e.g., the perspective taken in the *Monetary History* of Friedman and Schwartz (1963).

can be conveniently represented in terms of specifications of exactly this sort, leading to answers to the very specific questions about interest-rate policy posed in the previous section.

4.3 Neo-Wicksellian Monetary Theory

The non-quantity-theoretic analytical framework developed here develops several important themes from the monetary writings of Knut Wicksell (1898, 1906, 1907, 1915). Wicksell argued that even the variations in the price level observed in his own time, under the international gold standard, were not primarily due to variations in the world gold supply, but rather to two other factors—the policies followed by central banks, adjusting the “bank rate” at which they were willing to discount short-term bills, on the one hand, and real disturbances, affecting the natural rate of interest, on the other. In Wicksell’s view, price stability depended on keeping the interest rate controlled by the central bank in line with the natural rate determined by real factors (such as the marginal product of capital). Inflation occurred whenever the central banks lowered interest rates without any decline in the natural rate having occurred to justify it or whenever the natural rate of interest increased (due, e.g., to an increase in the productivity of investment opportunities) without any adjustment of the interest rates controlled by central banks in response. Deflation occurred whenever a disparity was created of the opposite sign.

Whatever the validity of such a non-quantity-theoretic approach for the analysis of price-level determination under the gold standard, Wicksell’s approach is a particularly congenial one for thinking about our present circumstances—a world of purely fiat currencies in which central banks adjust their operating targets for nominal interest rates in response to perceived risks of inflation, but pay little if any attention to the evolution of monetary aggregates—to say nothing of the one toward which we may be headed, in which monetary frictions become negligible.⁴² In such a world, where the concepts of money supply and demand become inapplicable, what is there to determine an equilibrium value for the general level of money prices? One possible answer is the role of *past* prices in determining

42. As noted earlier, Wicksell’s basic exposition of his theory is for the case of a “pure credit economy.” On the non-quantity-theoretic character of Wicksell’s explanation of price-level determination, see in particular Haavelmo (1978), Leijonhufvud (1981), and Niehans (1990). Others have insisted on the compatibility of Wicksell’s analysis with a quantity-theoretic framework; see, e.g., Humphrey (2002). I do not deny the *consistency* of a neo-Wicksellian account of a price-level determination under an interest-rate rule with a standard quantity-theoretic model; see, e.g., Section 3 of Chapter 4. But I show that the existence of a well-defined demand for money is not *essential* to the logic of such an account, which is consequently equally applicable to a world in which money-demand relations are undefined or unstable.

current equilibrium prices, due either to wage or price stickiness or to the effect of past prices on expectations regarding future prices (the critical factor in Wicksell's own analysis). Thus once prices have been at a certain level (for whatever arbitrary reason), this historical initial condition ties down their subsequent evolution, though they may subsequently drift arbitrarily far from that level. But probably the most important factor, in general, is *the interest-rate policy of the central bank*, insofar as this responds to the evolution of some price index. A state of affairs in which all wages and prices were 10 percent higher than they presently are would not be equally possible as an equilibrium, if the observation of such a jump in the price level would trigger an increase in interest rates, as called for under either a Wicksellian rule or the Taylor rule.

The way in which the equilibrium price level can be determined by the central bank's interest-rate response to price-level variations, without any reference to the associated fluctuations in any monetary aggregate, can be illustrated very simply. Suppose that the equilibrium real rate of interest is determined by real factors (such as time preference and the productivity of capital), in complete independence of how nominal quantities may evolve,⁴³ and let $\{r_t\}$ be an exogenous stochastic process for this real rate. It then follows that the short-term nominal interest rate i_t and the log price level p_t must at all times satisfy the Fisherian relation

$$p_t = E_t p_{t+1} + r_t - i_t, \quad (4.9)$$

assuming rational expectations on the part of the private sector. Because r_t is a certain function of exogenous real factors, rather than the measured real rate of return, this is an *equilibrium* relation—the condition required for equality between aggregate saving and investment—rather than an identity. This “flexible-price IS equation” indicates how the price level that clears the goods market—or equivalently, that equates saving and investment—depends on the expected future price level, real factors affecting saving and investment, and the nominal interest rate controlled by the central bank.

Now suppose that the central bank sets the short-term nominal interest rate according to the Wicksellian rule

$$i_t = \bar{i}_t + \phi p_t, \quad (4.10)$$

which generalizes (4.1) in allowing for a time-varying intercept, indicating possible shifts over time in monetary policy. Suppose further that $\{\bar{i}_t\}$ is another exogenous stochastic process (i.e., determined independently of the

43. In Chapter 2, I present assumptions under which this is true in an explicit intertemporal equilibrium model with flexible prices.

evolution of prices), which may or may not be correlated with the exogenous fluctuations in the equilibrium real rate of interest. Then substituting (4.10) into (4.9) to eliminate i_t , one obtains a relation of the form

$$p_t = \alpha E_t p_{t+1} + \alpha(r_t - \bar{i}_t) \quad (4.11)$$

to determine the equilibrium evolution of the price level, given the exogenous processes $\{r_t, \bar{i}_t\}$, where $\alpha \equiv 1/(1 + \phi)$ is a coefficient satisfying $0 < \alpha < 1$.

In the case that $\{r_t, \bar{i}_t\}$ are bounded processes, equation (4.11) has a unique bounded solution, obtained by “solving forward,” namely,

$$p_t = \sum_{j=0}^{\infty} \alpha^{j+1} E_t(r_{t+j} - \bar{i}_{t+j}). \quad (4.12)$$

Thus the equilibrium price level fluctuates in a bounded interval around the long-run average value

$$\bar{p} \equiv \phi^{-1}(\bar{r} - \bar{i}),$$

where $\bar{r}, \bar{i}$ are the long-run average values of r_t and $\bar{i}_t$, respectively. This analysis shows how a policy rule that involves no targets for any monetary aggregate can nonetheless control the long-run price level. It also shows how the determinants of equilibrium inflation can be understood without any reference to the determinants of either the money supply or money demand—indeed, it does not matter for the analysis just presented whether there is any well-defined demand function for the monetary base.

The account of price-level determination implied by this theory has a strongly Wicksellian flavor. One observes from (4.12) that the equilibrium price level at any date t is increased by either a loosening of monetary policy—represented by a reduction of the intercept term $\bar{i}_t$ —not justified by any decline in the equilibrium real rate or by an increase in the equilibrium real rate r_t that is not matched by a tightening of policy. The proposed forward-looking model also implies that any news that allows the private sector to forecast the future occurrence of either of these things should stimulate inflation immediately.

In the simple model sketched here, there is no distinction of the sort that Wicksell makes between the actual real rate of return and the natural rate that would occur in an intertemporal equilibrium with flexible prices. In Chapter 4, however, I show how one can usefully introduce such a distinction, in the context of a model with sticky prices. When prices are temporarily sticky, the real rate of return at which borrowing and lending occurs can differ from the natural rate of interest, just as the level of output can differ

from its natural rate; and the degree to which both occur depends on the degree of instability of the overall price level, as it is only when the general level of prices is changing that price rigidity creates distortions. Equilibrium condition (4.9) must then be replaced by a more general one, of the form

$$i_t - E_t \pi_{t+1} = r_t^n + \delta(\pi_t, \dots), \quad (4.13)$$

where r_t^n is the natural rate of interest (still assumed to depend only on exogenous real factors) and the discrepancy $\delta(\cdot)$ is a function of both current and expected future inflation.⁴⁴ The system consisting of conditions (4.10) and (4.13) can again be solved for a unique bounded process for the price level, and the solution is of the form

$$p_t = \sum_{j=0}^{\infty} \psi_j E_t (r_{t+j}^n - \bar{r}_{t+j})$$

for certain coefficients $\{\psi_j\}$. Thus in the more general case, it is variation in the *natural rate* of interest due to real disturbances of various sorts, to the extent that such variation is not matched by corresponding adjustment of the central bank's reaction function, that causes inflation variation. Just as in Wicksell's theory, real disturbances affecting desired saving and investment are predicted to be important sources of price-level variations; and as in that theory, the implied variation in the natural rate of interest is a useful summary statistic for the way in which a variety of real disturbances should affect the rate of inflation.

In Chapters 2 and 4, I show how a similar analysis of equilibrium inflation determination is possible in the case of a rule like the Taylor rule. In this case a positive response of the interest rate to fluctuations in the inflation rate is not sufficient to guarantee a determinate equilibrium (a unique nonexplosive equilibrium path for the inflation rate, rather than the price level). It is instead necessary that the response coefficient be greater than one, in accordance with the Taylor principle mentioned earlier. But in that case similar results are obtained; equilibrium inflation is a function of current and expected future gaps between the natural rate of interest and the intercept term in the Taylor rule.

One finds, then, that it is possible to determine the consequences for inflation dynamics of a given monetary policy rule when it is expressed in terms of an interest-rate rule, without any need to first translate the rule into an implied state-contingent path for the money supply. Hence the

44. In the case of the basic neo-Wicksellian model developed in Chapter 4, δ is a function of π_t , $E_t \pi_{t+1}$, and $E_t \pi_{t+2}$.

terms used to describe both actual central-bank policies and simple policy prescriptions in the literature summarized previously are not inappropriate ones. The consequences of systematic policies of these types can conveniently be analyzed as functions of exactly the coefficients appearing in Table 1.1.

While the usefulness of the neo-Wicksellian framework sketched here is perhaps most evident in the case of an economy without monetary frictions, so that the familiar quantity-theoretic apparatus is plainly inapplicable, it is also equally useful in the case of an economy in which monetary frictions still exist, at least of the modest sort that are indicated by the observed willingness to hold non-interest-earning currency in an advanced economy like that of the United States today. In what I have written previously, I have not actually relied upon any assumed absence of monetary frictions, except in assuming that the equilibrium real rate of interest (or more generally, the natural rate of interest) is independent of the evolution of nominal variables. But even in the presence of transactions frictions resulting in a demand for base money despite its below-market rate of return, it is unlikely that the natural rate of interest is much affected, as a quantitative matter, by variations in the rate of inflation. (The accuracy of the approximation involved in neglecting such effects is considered numerically in Chapters 2 and 4.) Hence the approach proposed here is also appropriate for analysis of the effects of a Taylor rule in an economy like that of the United States, where changes in the Fed's interest-rate operating target are implemented through adjustments in the supply of (non-interest-earning) Fed balances. The monetary frictions that create a demand for such balances are important for the size of quantity adjustment required to achieve a given change in the funds rate, but are of little significance for the effects upon output and inflation of any given change in the path of the funds rate.

Nor are the predictions of the neo-Wicksellian theory really any different from those of a standard quantity-theoretic analysis, despite the apparent dissimilarity of approach. In a quantity-theoretic analysis of inflation determination, central importance is given to the money-demand relation that describes desired real money balances as a function of interest rates and other variables. However, in the case of an interest-rate rule like any of those described earlier, the money supply varies passively so as to satisfy this relation; hence the relation places no restrictions upon the equilibrium evolution of interest rates and goods prices. Thus in such a system, the equilibrium paths of interest and prices are determined by solving equations (4.9) and (4.10), or equations (4.10) and (4.13) in the case of sticky prices, just as previously. Once one knows the equilibrium paths of interest and prices, the money-demand relation can then be used to determine the implied evolution of the money supply as well. But this last relation does not play an important role in determining equilibrium inflation under such an

analysis. Moreover, insisting upon first solving for the state-contingent path of the money supply implied by the policy rule and then deriving the equilibrium path of inflation from this along quantity-theoretic lines would be an unnecessarily roundabout procedure, given that one must first solve for the path of prices and interest rates in order to determine the path of the money supply.

The neo-Wicksellian approach is thus clearly preferable, even granting the existence of a well-defined, econometrically stable money-demand relation, if one wishes to analyze the consequences of interest-rate rules such as the Taylor rule. But, it might be asked, is it clear that *desirable* policy rules should belong to this class, regardless of the current popularity of such prescriptions? Might not a money-growth rule be preferable, in which case a more traditional quantity-theoretic approach would also be necessary in order to explain its effects?

The results of this study suggest that the answer is no. As is shown in Chapter 8, it is possible to derive optimal policy rules that indicate how a short-term nominal interest-rate operating target should be set, as a function of the projected evolution of inflation and the output gap, without any reference to the paths of monetary aggregates. It is further argued that this form of prescription has the advantage, relative to other possible characterizations of optimal policy, of being invariant under a larger class of possible exogenous disturbances. For example, in the case of an economy with a well-defined demand for base money, it is possible to compute both the state-contingent evolution of overnight interest rates and the state-contingent evolution of the monetary base in an optimal equilibrium. However, the desired evolution of the monetary base, even when well defined, depends upon factors that are of little or no relevance to the desired evolution of interest rates, and this makes it simpler to characterize optimal policy in terms of an interest-rate operating target.

One such factor is the dependence of the optimal path of the monetary base on changes in the transactions technology—for example, available opportunities for substitution among alternative means of payment. Such changes may have significant effects on money demand in the presence of a given interest differential between base money and other riskless assets, but have little effect on the relations between interest rates and the incentives for intertemporal substitution of expenditure that determine the desired evolution of interest rates. (In an economy where the financial system is already highly efficient, one expects further innovations to represent movements from one highly efficient system to another, so that the relations between interest rates and the real allocation of resources remain near those predicted by a model with no financial frictions; but money demand may be greatly affected in percentage terms, as it ceases to be defined in the frictionless limit.) Another factor is the dependence of the optimal evolution of the monetary base on the details of monetary policy implementation.

The desired path of money-market interest rates is largely independent of the rate of interest paid on the monetary base. This depends instead on the intertemporal marginal rates of substitution, the marginal products of real investment implied by the desired allocation of resources, and the desired path for inflation. But the desired path of the monetary base depends greatly on whether it is assumed for institutional reasons that zero interest is paid on base money or whether instead the interest paid on money varies when other short-term interest rates vary; for the demand for base money depends not on the absolute level of nominal interest rates, but on the *spread* between the interest rate paid on base money and that available on other assets.

Thus even when the desired evolution of the monetary base is well defined, it is more dependent on special “technical” factors than is the desired evolution of short-term nominal interest rates. This makes a description of optimal monetary policy in terms of a state-contingent money growth rate less convenient. Moreover, if, as some forecast, monetary frictions are largely eliminated in the coming century owing to the development of electronic payments media, a description of optimal policy in terms of the desired evolution of a monetary aggregate is likely to become awkward if not altogether impossible. Yet a description of optimal policy in terms of the principles that should regulate the adjustment of an interest-rate operating target should still be possible. Indeed, increasing efficiency of the financial system should only simplify the description of optimal policy in these terms, insofar as the arbitrage relations that connect the overnight interest rate directly targeted by the central bank to other interest rates and asset prices become simpler and more reliable. Hence the neo-Wicksellian framework proposed here directs attention to precisely those elements of the monetary transmission mechanism that are likely to remain of fundamental importance for the design of effective monetary policies in a world of increasingly efficient financial markets and institutions.

5 Plan of the Book

Part I develops a theoretical framework that can be used to analyze the consequences of alternative monetary policy rules in a way that takes full account of the implications of forward-looking private sector behavior. Chapter 2 begins by considering price-level determination when monetary policy is specified by an interest-rate rule in the case of a model where, for simplicity, prices are completely flexible and the supply of goods is given by an exogenous endowment. This chapter demonstrates the possibility of a coherent theory of price-level determination even in the complete absence of monetary frictions—a special case that is considered repeatedly in what follows, in order to direct attention more closely to the economic relations that are considered to be of more fundamental importance for

the characterization of optimal policy. But it also considers price-level determination under an interest-rate rule in a standard optimization-based monetarist framework, allowing a comparison between the consequences of monetary targeting and those of commitment to an interest-rate rule, and an analysis of the extent to which the presence of monetary frictions changes one's conclusions about the effects of such a rule.

Chapter 3 then introduces endogenous goods supply and nominal price and wage rigidities, so that monetary policy can affect the level of real activity as well as the inflation rate. Considerable attention is given to the microeconomic foundations of the aggregate-supply relations that result from delays in the adjustment of prices or wages, in order to select specifications (from among those that might appear similarly consistent with econometric evidence) with clear behavioral interpretations that thereby allow one to take account of the "Lucas critique." At the same time, attention is also paid to the need to find a specification of the dynamic relations between real activity and inflation that is consistent with econometric evidence regarding the effects of identified monetary disturbances. A series of modifications of a basic sticky-price model are introduced that can improve the model's fit with estimated responses on various dimensions.

Chapter 4 then integrates the analysis in Chapters 2 and 3, considering the effects of interest-rate rules in a framework where monetary policy can affect real activity and where feedback from measures of real activity to the central bank's interest-rate operating target matter for the predicted effects of such rules. In the neo-Wicksellian framework developed here, inflation dynamics results from the interaction between real disturbances on the one hand and the central bank's interest-rate rule on the other. Wicksell's natural rate of interest is shown to play a central role, summarizing the effects of a variety of real disturbances that are relevant for inflation and output-gap determination in the case of a class of policy rules that may be thought of as generalized Taylor rules. The chapter also includes a first analysis of the consequences of such a framework for the design of desirable policy rules by considering the conditions under which a Taylor-like rule should be able to stabilize inflation and the output gap.

Chapter 5 discusses further elaborations of the basic neo-Wicksellian model, mainly with regard to aggregate-demand determination, that allow the model to better match econometric evidence regarding the effects of monetary policy. Various reasons are discussed for delayed effects of interest-rate changes on aggregate demand to differ from the immediate effect, and then some simple examples are presented of small optimizing models of the U.S. monetary transmission mechanism that have been fit to U.S. time series. This discussion of empirical models is intended to motivate the kinds of model specifications that are considered in the remainder of this study.

Part II then considers the optimal conduct of monetary policy in the light of the theoretical framework introduced in the earlier chapters. Chapter 6 begins by considering appropriate stabilization goals for monetary policy. An advantage of the derivation of the proposed model's structural relations from explicit microeconomic foundations in Part I is that it is possible to ask what sort of monetary policy would best serve economic welfare, given the objectives and constraints of the agents whose decisions account for the observed effects of monetary policy. Chapter 6 considers the connection between the obvious measure of economic welfare in such a model—the expected utility of the representative household—and the stabilization of macroeconomic aggregates such as inflation and the output gap.

It is shown that a quadratic approximation to expected utility, which suffices (under certain conditions) for the derivation of a linear approximation to an optimal policy rule, can be expressed in terms of the expected value of squared deviations of certain aggregate variables from target values for those variables. The variables that are relevant and the details of the quadratic loss function that can be justified on welfare-theoretic grounds depend on the microeconomic foundations of one's model of the monetary transmission mechanism. In particular, it is shown that different assumptions regarding price and/or wage stickiness imply that price and/or wage inflation should enter the central bank's loss function in different ways. Nonetheless, it is argued that price stability, suitably interpreted—for example, quite possibly in terms of an index that includes wages as well as the prices of final goods and services—should be an important consideration, though not necessarily the only one, in the selection of a monetary policy rule. Grounds for inclusion of output-gap stabilization and interest-rate stabilization objectives in the loss function as well are considered; but it is argued that in practice, these additional concerns are not likely to justify either an average rate of inflation much greater than zero or substantial variability in the rate of inflation in response to shocks.

Chapter 7 then considers the optimal state-contingent evolution of inflation, output, and interest rates in response to real disturbances of various sorts, from the point of view of the sort of loss function argued for in Chapter 6 and in the context of a forward-looking model of the monetary transmission mechanism of the kind developed in Part I. An important general issue treated in this chapter is the way in which optimal control techniques must be adapted in the case of control of a forward-looking system. The responses to shocks under an optimal commitment are distinguished from the equilibrium responses under discretionary optimization by the central bank. Particular attention is given to the fact that in general, optimal responses will be *history dependent* in a way that is inconsistent with any purely forward-looking decision procedure for monetary policy. Alternative approaches to introducing the desired sort of history dependence into the conduct of policy

are surveyed. The one that is emphasized in this study is the possibility of commitment to a *policy rule* that is history dependent in the desired way. Commitment to a *targeting rule*, such as a flexible inflation target, is shown to be an especially desirable way of formulating such a policy commitment.

Finally, Chapter 8 considers the problem of choice of a policy rule to implement the desired state-contingent evolution of inflation, output, and interest rates, as derived in Chapter 7. From the standpoint taken here, this problem of *implementation* of the desired equilibrium is a nontrivial part of the characterization of optimal policy. The mapping from the history of exogenous disturbances to the desired overnight interest rate at any point in time is not a suitable description of an optimal policy rule, for reasons taken up in this chapter. Instead, it is argued that a more suitable policy prescription should relate the instrument setting to the evolution (observed or projected) of endogenous variables such as inflation and the output gap, as in the proposals mentioned earlier. A general method is described for constructing optimal policy rules of this form in the case of a fairly general class of log-linear structural models and quadratic loss functions. The method is then applied to several of the simple optimizing models of the monetary transmission mechanism developed in previous chapters.

It is shown that optimal rules can easily take the form of generalized Taylor rules or of target criteria for a forecast-targeting procedure like that used at the Bank of England. However, even in the case of fairly simple models of the transmission mechanism, the optimal rules are somewhat different from the proposals described earlier. While there is a fairly clear logic for rules that respond to (and perhaps only to) variations in inflation and in the output gap, the theoretically appropriate measures of inflation and of the output gap may not be the ones used in the characterizations above of current central-bank behavior. Moreover, the optimal rules that I obtain are also typically different in their dynamic specifications. Optimal rules are history dependent in ways other than those of the classic Taylor rule or familiar descriptions of inflation-forecast targeting; and while they may well be more forward looking than the classic Taylor rule, in all of the calibrated examples they are considerably less forward looking than the procedures currently used at the inflation-targeting central banks.

Conclusions about the precise content of an optimal policy rule depend, of course, on the details of one's model of the transmission mechanism, and I do not attempt here to reach final conclusions in that regard. The answer is likely to be somewhat different for different countries in any event. My more important goal is to provide a method that individual central banks can use in order to choose sensible systematic policies on the basis of their own research on the nature of the transmission mechanism in their respective economies. It is hoped that the present volume can provide useful guidelines for such an inquiry.