

GAUTI B. EGGERTSSON

International Monetary Fund

MICHAEL WOODFORD

Princeton University

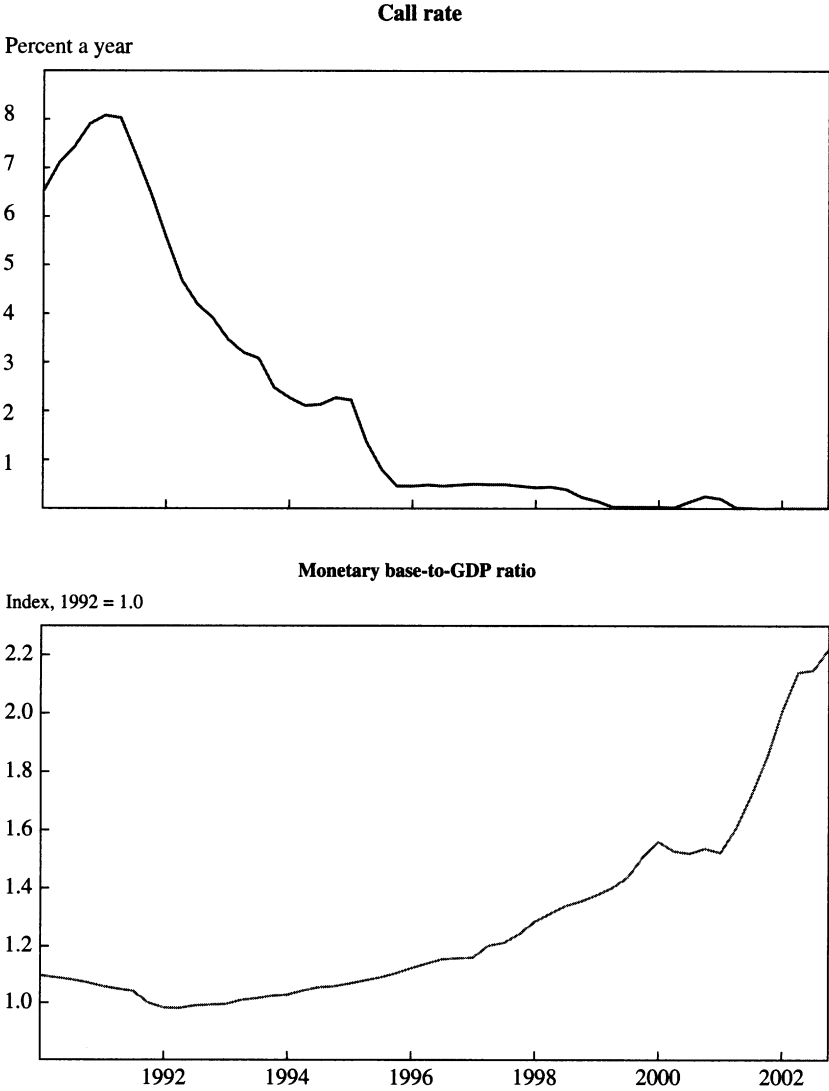
The Zero Bound on Interest Rates and Optimal Monetary Policy

THE CONSEQUENCES FOR THE PROPER conduct of monetary policy of the existence of a lower bound of zero for overnight nominal interest rates has recently become a topic of lively interest. In Japan the call rate (the overnight cash rate analogous to the federal funds rate in the United States) has been within 50 basis points of zero since October 1995, and it has been essentially equal to zero for most of the past four years (figure 1). Thus the Bank of Japan has had little room to further reduce short-term nominal interest rates in all that time. Meanwhile Japan's growth has remained anemic, and prices have continued to fall, suggesting a need for monetary stimulus. Yet the usual remedy—lower short-term nominal interest rates—is plainly unavailable. Vigorous expansion of the monetary base has also seemed to do little to stimulate demand under these circumstances: as figure 1 also shows, the monetary base is now more than twice as large, relative to GDP, as it was in the early 1990s.

In the United States, meanwhile, the federal funds rate has now been reduced to only 1 percent, and signs of recovery remain exceedingly fragile. This has led many to wonder if this country might not also soon find itself in a situation where interest rate policy is no longer available as a

We would like to thank Tamim Bayoumi, Ben Bernanke, Robin Brooks, Michael Dotsey, Benjamin Friedman, Stefan Gerlach, Mark Gertler, Marvin Goodfriend, Kenneth Kuttner, Maurice Obstfeld, Athanasios Orphanides, Kenneth Rogoff, David Small, Lars Svensson, Harald Uhlig, Tsutomu Watanabe, and Alex Wolman for helpful comments, and the National Science Foundation for research support through a grant to the National Bureau of Economic Research. The views expressed in this paper are those of the authors and do not necessarily represent those of the International Monetary Fund or IMF policy.

Figure 1. Japan: Call Rate on Overnight Loans and Ratio of Monetary Base to GDP, 1990–2002



Source: Nomura database, Bank of Japan.

tool for macroeconomic stabilization. A number of other countries face similar questions. John Maynard Keynes first raised the question of what can be done to stabilize the economy when it has fallen into a liquidity trap—when interest rates have fallen to a level below which they cannot be driven by further monetary expansion—and whether monetary policy can be effective at all under such circumstances. Long treated as a mere theoretical curiosity, Keynes's question now appears to be one of urgent practical importance, but one with which theorists have become unfamiliar.

The question of how policy should be conducted when the zero bound is reached—or when the possibility of reaching it can no longer be ignored—raises many fundamental issues for the theory of monetary policy. Some would argue that awareness of the possibility of hitting the zero bound calls for fundamental changes in the way policy is conducted even before the bound has been reached. For example, Paul Krugman refers to deflation as a “black hole,”¹ from which an economy cannot expect to escape once it has entered. A conclusion often drawn from this pessimistic view of the efficacy of monetary policy in a liquidity trap is that it is vital to steer far clear of circumstances in which deflationary expectations could ever begin to develop—for example, by targeting a sufficiently high positive rate of inflation even under normal circumstances. Others are more sanguine about the continuing effectiveness of monetary policy even when the zero bound is reached. For example, it is often argued that deflation need not be a black hole, because monetary policy can affect aggregate spending, and hence inflation, through channels other than central bank control of short-term nominal interest rates. Thus there has been much recent discussion, with respect to both Japan and the United States—of the advantages of vigorous expansion of the monetary base even without any further reduction in interest rates, of the desirability of attempts to shift longer-term interest rates through central bank purchases of longer-maturity government securities, and even of the desirability of central bank purchases of other kinds of assets.

Yet if these views are correct, they challenge much of the recent conventional wisdom regarding the conduct of monetary policy, both within central banks and among academic monetary economists. That wisdom has stressed a conception of the problem of monetary policy in terms of

1. Paul Krugman, “Crisis in Prices?” *New York Times*, December 31, 2002, p. A19.

the appropriate adjustment of an operating target for overnight interest rates, and the prescriptions formulated for monetary policy, such as the celebrated Taylor rule,² are typically cast in these terms. Indeed, some have argued that the inability of such a policy to prevent the economy from falling into a deflationary spiral is a critical flaw of the Taylor rule as a guide to policy.³

Similarly, concern over the possibility of entering a liquidity trap is sometimes presented as a serious objection to another currently popular monetary policy prescription, namely, inflation targeting. The definition of a policy prescription in terms of an inflation target presumes that there is in fact some level of the nominal interest rate that can allow the target to be hit (or at least projected to be hit, on average). But, some argue, if the zero interest rate bound is reached under circumstances of deflation, it will not be possible to hit any higher inflation target, because further interest rate decreases are not possible. Is there, in such circumstances, any point in having an inflation target? The Bank of Japan has frequently offered this argument as a reason for resisting inflation targeting. For example, Kunio Okina, director of the Institute for Monetary and Economic Studies at the Bank of Japan, was quoted as arguing that “because short-term interest rates are already at zero, setting an inflation target of, say, 2 percent wouldn’t carry much credibility.”⁴

We seek to shed light on these issues by considering the consequences of the zero lower bound on nominal interest rates for the optimal conduct of monetary policy, in the context of an explicitly intertemporal equilibrium model of the monetary transmission mechanism. Although our model is extremely simple, we believe it can help clarify some of the basic issues just raised. We are able to consider the extent to which the zero bound represents a genuine constraint on attainable equilibrium paths for inflation and real activity, and the extent to which open-market purchases of various kinds of assets by the central bank can mitigate that constraint. We are also able to show how the existence of the zero bound changes the character of optimal monetary policy, relative to the policy rules that would be judged optimal in its absence or in the case of real disturbances small enough for the bound never to matter under an optimal policy.

2. Taylor (1993).

3. Benhabib, Schmitt-Grohé, and Uribe (2001).

4. “Japan BOJ Official: Hard to Set Inflation Targets,” Dow Jones News, August 11, 1999.

To preview our results, we find that the zero bound does represent an important constraint on what monetary stabilization policy can achieve, at least when certain kinds of real disturbances are encountered in an environment of low inflation. We argue that the possibility of expanding the monetary base through central bank purchases of a variety of types of assets does little if anything to expand the set of feasible paths for inflation and real activity that are consistent with equilibrium under some (fully credible) policy commitment.

Hence the relevant trade-offs can correctly be studied by simply considering what alternative anticipated state-contingent paths of the short-term nominal interest rate can achieve, taking into account the constraint that this rate must be nonnegative at all times. Doing so, we find that the zero interest rate bound can indeed be temporarily binding, and when it is, it inevitably results in lower welfare than could be achieved in the absence of such a constraint.⁵

Nonetheless, we argue that the zero bound restricts possible stabilization outcomes under sound policy to a much more modest degree than the deflation pessimists presume. Even though the set of feasible equilibrium outcomes corresponds to those that can be achieved through alternative interest rate policies, monetary policy is far from powerless to mitigate the contractionary effects of the kind of disturbances that would make the zero bound a binding constraint. The key to dealing with this sort of situation in the least damaging way is to create the right kind of expectations regarding how monetary policy will be used *after* the constraint is no longer binding, and the central bank again has room to maneuver. We use our intertemporal equilibrium model to characterize

5. We do not explore here the possibility of relaxing the constraint by taxing money balances, as originally proposed by Gesell (1929) and Keynes (1936), and more recently by Buiter and Panigirtzoglou (2001) and Goodfriend (2000). Although this represents a solution to the problem in theory, it presents substantial practical difficulties, not the least of which is the political opposition that such an institutional change would be likely to generate. Our consideration of the problem of optimal policy also abstracts from the availability of fiscal instruments, such as the time-varying tax policy recommended by Feldstein (2002). We agree with Feldstein that there is a particularly good case for state-contingent fiscal policy to deal with a liquidity trap, even if fiscal policy is not a very useful tool for stabilization policy more generally. Nonetheless, we consider here only the problem of the proper conduct of monetary policy, taking as given the structure of tax distortions. As long as one does not think that state-contingent fiscal policy can (or will) be used to eliminate even temporary declines in the natural rate of interest below zero, the problem for monetary policy that we consider here remains relevant.

the kind of expectations regarding future policy that it would be desirable to create, and we discuss a form of price-level targeting rule that—if credibly committed to—should bring about the constrained-optimal equilibrium. We also discuss, more informally, how other types of policy actions could help increase the credibility of the central bank's announced commitment to this kind of future policy.

Our analysis will be recognized as a development of several key themes in Paul Krugman's treatment of the same topic in these pages a few years ago.⁶ Like Krugman, we give particular emphasis to the role of expectations regarding future policy in determining the severity of the distortions that result from hitting the zero bound. Our primary contribution, relative to Krugman's earlier treatment, will be the presentation of a more fully dynamic analysis. For example, our assumption of staggered pricing, rather than Krugman's simple hypothesis of prices that are fixed for one period, allows for richer (and at least somewhat more realistic) dynamic responses to disturbances. In our model, unlike in Krugman's, a real disturbance that lowers the natural rate of interest can cause output to remain below potential for years (as shown in figure 2 later in the paper), rather than only for a single "period," even when the average frequency of price adjustments is more than once a year. These richer dynamics are also important for a realistic discussion of the kind of policy commitment that can help to reduce economic contraction during a liquidity trap. In our model a commitment to create subsequent inflation involves a commitment to *keep interest rates low* for some time in the future, whereas in Krugman's model a commitment to a higher future price level does not involve any reduction in future nominal interest rates. We are also better able to discuss such questions as how the creation of inflationary expectations while the zero bound is binding can be reconciled with maintaining the credibility of the central bank's commitment to long-run price stability.

Our dynamic analysis also allows us to further clarify the several ways in which the central bank's management of private sector expectations can be expected to mitigate the effects of the zero bound. Krugman emphasizes the fact that increased expectations of inflation can lower the real interest rate implied by a zero nominal interest rate. This might suggest, however, that the central bank can affect the economy only insofar

6. Krugman (1998).

as it affects expectations regarding a variable that it cannot influence except quite indirectly; it might also suggest that the only expectations that should matter are those regarding inflation over the relatively short horizon corresponding to the term of the nominal interest rate that has fallen to zero. Such interpretations easily lead to skepticism about the practical effectiveness of the expectations channel, especially if inflation is regarded as being relatively “sticky” in the short run. Our model is instead one in which expectations affect aggregate demand through several channels.

First of all, it is not merely short-term real interest rates that matter for current aggregate demand; our model of intertemporal substitution in spending implies that the entire expected future path of short-term real rates should matter, or alternatively that very long term real rates should matter.⁷ This means that the creation of inflation expectations, even with regard to inflation that should not occur until at least a year into the future, should also be highly relevant to aggregate demand, as long as it is not accompanied by correspondingly higher expected future nominal interest rates. Furthermore, the expected future path of nominal interest rates matters, and not just their current level, so that a commitment to keep nominal interest rates low for a longer period of time should stimulate aggregate demand, even when current interest rates cannot be lowered further, and even under the hypothesis that inflation expectations would remain unaffected. Because the central bank can clearly control the future path of short-term nominal interest rates if it has the will to do so, any failure of such a commitment to be credible will not be due to skepticism about whether the central bank is *able* to follow through on its commitment.

The richer dynamics of our model are also important for the analysis of optimal policy. Krugman mainly addresses the question of whether monetary policy is completely impotent when the zero bound binds, and he argues for the possibility of increasing real activity in the liquidity trap by

7. In the simple model presented here, this occurs solely as a result of intertemporal substitution in private expenditure. But there are a number of reasons to expect long-term rates, rather than short-term rates, to be the critical determinant of aggregate demand. For example, in an open-economy model, the real exchange rate becomes an important determinant of aggregate demand. But the real exchange rate should be closely linked to a very long domestic real rate of return (or alternatively to the expected future path of short-term rates) as a result of interest rate parity, together with an anchor for the expected long-term real exchange rate (deriving, for example, from long-run purchasing power parity).

creating expectations of inflation. Although we agree with this conclusion, it does not answer the question of whether, or to what extent, it would be desirable to create such expectations, given the well-founded reasons that the central bank should have to not prefer inflation at a later time. Nor is Krugman's model well suited to address such a question, insofar as it omits any reason for even an extremely high subsequent inflation to be deemed harmful. Our staggered-pricing model instead implies that inflation (whether anticipated or not) does create distortions, justifying an objective function for stabilization policy that trades off inflation stabilization and output gap stabilization in terms that are often assumed to represent actual central bank concerns. We characterize optimal policy in such a setting and show that it does indeed involve a commitment to history-dependent policy of a sort that should result in higher inflation expectations in response to a binding zero bound. We can also show to what extent it should be optimal to create such expectations, assuming that this is possible. We find, for example, that it is not optimal to commit to so much future inflation that the zero bound ceases to bind, even though this is one possible type of equilibrium; this is why the zero bound does remain a relevant constraint, even under an optimal policy commitment.

Is Quantitative Easing a Separate Policy Instrument?

A first question we wish to consider is whether expansion of the monetary base represents a policy instrument that should be effective in preventing deflation and an associated output decline, even under circumstances where overnight interest rates have fallen to zero. According to Keynes's famous analysis,⁸ monetary policy ceases to be an effective instrument to head off economic contraction in a "liquidity trap," which can arise if interest rates fall so low that further expansion of the money supply cannot drive them lower. Others have argued that monetary expansion should increase nominal aggregate demand even under such circumstances, and the supposition that this is correct lies behind Japan's explicit adoption, since March 2001, of a policy of "quantitative easing" in addition to the zero interest rate policy that continues to be maintained.⁹

8. Keynes (1936).

9. See Kimura and others (2002) for a discussion of this policy as well as an expression of doubts about its effectiveness.

Here we consider this question in the context of an explicitly intertemporal equilibrium model, which models both the demand for money and the role of financial assets (including the monetary base) in private sector budget constraints. The model we use for this purpose is more detailed in several senses than that used in subsequent sections to characterize optimal policy. We do this to make it clear that we have not excluded a role for quantitative easing simply by failing to model the role of money in the economy.¹⁰

Our key result is an irrelevance proposition for open-market operations in a variety of types of assets that the central bank might acquire, under the assumption that the open-market operations do not change the expected future conduct of monetary or fiscal policy (in senses that we specify below). It is perhaps worth noting at the outset that our intention in stating such a result is not to vindicate the view that a central bank is powerless to halt a deflationary slump, and hence to absolve the Bank of Japan, for example, of any responsibility for the continuing stagnation in that country. Although our proposition establishes that there is a sense in which a liquidity trap is possible, this does not mean that the central bank is powerless under the circumstances we describe. Rather, our intent is to show that the key to effective central bank action to combat a deflationary slump is the management of expectations. Open-market operations should be largely ineffective to the extent that they fail to change expectations regarding future policy; the conclusion we draw is not that such actions are futile, but rather that the central bank's actions should be chosen with a view to signaling the nature of its policy commitments, and not for the purpose of creating some sort of "direct" effects.

A Neutrality Proposition for Open-Market Operations

Our model abstracts from endogenous variations in the capital stock and assumes perfectly flexible wages (or some other mechanism for efficient labor contracting), but it assumes monopolistic competition in goods markets and sticky prices that are adjusted at random intervals in the manner assumed by Guillermo Calvo, so that deflation has real effects.¹¹ We

10. Woodford (forthcoming, chapter 4) discusses the model in more detail and considers the consequences of various interest rate rules and money growth rules under the assumption that disturbances are not large enough for the zero bound to bind.

11. Calvo (1983).

assume that the representative household seeks to maximize a utility function of the form

$$E_t \sum_{T=t}^{\infty} \beta^{T-t} \left\{ u(C_t, M_t/P_t; \xi_t) - \int_0^1 v[H_t(j); \xi_t] dj \right\},$$

where C_t is a Dixit-Stiglitz aggregate of consumption of each of a continuum of differentiated goods,

$$C_t \equiv \left[\int_0^1 c_t(i)^{\frac{\theta}{\theta-1}} di \right]^{\frac{\theta-1}{\theta}},$$

with an elasticity of substitution $\theta > 1$; M_t measures end-of-period household money balances,¹² P_t is the Dixit-Stiglitz price index,

$$(1) \quad P_t \equiv \left[\int_0^1 p_t(i)^{1-\theta} di \right]^{\frac{1}{1-\theta}},$$

and $H_t(j)$ is the quantity supplied of labor of type j . Real balances are included in the utility function,¹³ as a proxy for the services that money balances provide in facilitating transactions.¹⁴ Each industry j employs an industry-specific type of labor, with its own wage.

For each value of the disturbances ξ_t , $u(\cdot, \cdot; \xi_t)$ is a concave function, increasing in the first argument and increasing in the second for all levels of real balances up to a satiation level $\bar{m}(C_t; \xi_t)$. The existence of a satiation level is necessary in order for it to be *possible* for the zero interest rate bound ever to be reached; we regard Japan's experience over the past several years as having settled the theoretical debate over whether such a level of real balances exists. Unlike many papers in the literature, we do not assume additive separability of the function u between the first two

12. We do not introduce fractional-reserve banking into our model. Technically, M_t refers to the monetary base, and we represent households as obtaining liquidity services from holding this base, either directly or through intermediaries (not modeled).

13. Following Sidrauski (1967) and Brock (1974, 1975).

14. We use this approach to modeling the transactions demand for money because of its familiarity. As shown in Woodford (forthcoming, appendix section A.16), a cash-in-advance model leads to equilibrium conditions of essentially the same general form, and the neutrality result that we present below would hold in essentially identical form were we to model the transactions demand for money after the fashion of Lucas and Stokey (1987).

arguments; this (realistic) complication allows a further channel through which money can affect aggregate demand, namely, by an effect of real money balances on the current marginal utility of consumption. Similarly, for each value of ξ_t , $v(\cdot; \xi_t)$ is an increasing convex function. The vector of exogenous disturbances ξ_t may contain several elements, so that no assumption is made about correlation of the exogenous shifts in the functions u and v .

For simplicity we assume complete financial markets and no limit on borrowing against future income. As a consequence, a household faces an intertemporal budget constraint of the form

$$E_t \sum_{T=t}^{\infty} Q_{t,T} [P_T C_T + \delta_T M_T] \leq W_t + E_t \sum_{T=t}^{\infty} Q_{t,T} \left[\int_0^1 \Pi_T(i) di + \int_0^1 w_T(j) H_T(j) dj - T_T^h \right],$$

looking forward from any period t . Here $Q_{t,T}$ is the stochastic discount factor that the financial markets use to value random nominal income at date T in monetary units at date t ; δ_t is the opportunity cost of holding money and is equal to $i_t/(1 + i_t)$, where i_t is the riskless nominal interest rate on one-period obligations purchased in period t , in the case that no interest is paid on the monetary base; W_t is the nominal value of the household's financial wealth (including money holdings) at the beginning of period t ; $\Pi_t(i)$ represents the nominal profits (revenue in excess of the wage bill) in period t of the supplier of good i ; $w_t(j)$ is the nominal wage earned by labor of type j in period t , and T_t^h represents the net nominal tax liabilities of each household in period t . Optimizing household behavior then implies the following necessary conditions for a rational expectations equilibrium. Optimal timing of household expenditure requires that aggregate demand Y_t for the composite good satisfy an Euler equation of the form¹⁵

$$(2) \quad u_c(Y_t, M_t/P_t; \xi_t) = \beta E_t \left[u_c(Y_{t+1}, M_{t+1}/P_{t+1}; \xi_{t+1})(1 + i_t) \frac{P_t}{P_{t+1}} \right],$$

15. For simplicity, we abstract from government purchases of goods. Our equilibrium conditions directly extend to the case of exogenous government purchases, as shown in Woodford (forthcoming, chapter 4).

Optimal substitution between real money balances and expenditure leads to a static first-order condition of the form

$$\frac{u_m(Y_t, M_t/P_t; \xi_t)}{u_c(Y_t, M_t/P_t; \xi_t)} = \frac{i_t}{1+i_t},$$

under the assumption that zero interest is paid on the monetary base, and that preferences are such as to exclude the possibility of a corner solution with zero money balances. If both consumption and liquidity services are normal goods, this equilibrium condition can be solved uniquely for the level of real balances $L(Y_t, i_t; \xi_t)$ that satisfy it in the case of any positive nominal interest rate. The equilibrium relation can then equivalently be written as a pair of inequalities:

$$(3) \quad \frac{M_t}{P_t} \geq L(Y_t, i_t; \xi_t)$$

$$(4) \quad i_t \geq 0,$$

together with the “complementary slackness” condition that at least one must hold with equality at any time. Here we define $L(Y, 0; \xi) = \bar{m}(Y; \xi)$, the minimum level of real balances for which $u_m = 0$, so that the function L is continuous at $i = 0$. Household optimization similarly requires that the paths of aggregate real expenditure and the price index satisfy the bounds

$$(5) \quad \sum_{T=t}^{\infty} \beta^T E_t [u_c(Y_T, M_T/P_T; \xi_T) Y_T + u_m(Y_T, M_T/P_T; \xi_T) (M_T/P_T)] < \infty$$

$$(6) \quad \lim_{T \rightarrow \infty} \beta^T E_t [u_c(Y_T, M_T/P_T; \xi_T) D_T/P_T] = 0$$

looking forward from any period t , where D_t measures the total nominal value of government liabilities (monetary base plus government debt) at the end of period t under the monetary-fiscal policy regime. (The condition in expression 5 is required for the existence of a well-defined intertemporal budget constraint, under the assumption that there are no limitations on households’ ability to borrow against future income, whereas the transversality condition in equation 6 must hold if the house-

hold hits its intertemporal budget constraint.) The conditions in expressions 2 through 6 also suffice to imply that the representative household chooses optimal consumption and portfolio plans (including its planned holdings of money balances) given its income expectations and the prices (including financial asset prices) that it faces, while making choices that are consistent with financial market clearing.

Each differentiated good i is supplied by a single, monopolistically competitive producer. There are assumed to be many goods in each of an infinite number of industries; each industry j uses a type of labor that is specific to that industry, and all goods in an industry change their prices at the same time. Each good is produced in accordance with a common production function

$$y_t(i) = A_t f[h_t(i)],$$

where A_t is an exogenous productivity factor common to all industries; $f(\cdot)$ is an increasing, concave function; and $h_t(i)$ is the industry-specific labor hired by firm i . The representative household supplies all types of labor and consumes all types of goods.¹⁶

The supplier of good i sets a price for that good at which it satisfies demand in each period, hiring the labor inputs necessary to meet that demand. Given households' allocation of demand across goods in response to firms' pricing decisions, on the one hand, and the terms on which optimizing households are willing to supply each type of labor, on the other, we can show that nominal profits (sales revenue in excess of labor costs) in period t of the supplier of good i are given by the function

$$\begin{aligned} \Pi \left[p_t(i), p_t^j, P_t, Y_t, M_t/P_t, \tilde{\xi}_t \right] &\equiv p_t(i) Y_t [p_t(i)/P_t]^{-\theta} \\ &- \frac{v_h \left\{ f^{-1} \left[Y_t (p_t^j/P_t)^{-\theta} / A_t \right]; \xi_t \right\}}{u_c(Y_t, M_t/P_t; \xi_t)} P_t f^{-1} \left\{ Y_t [p_t(i)/P_t]^{-\theta} / A_t \right\}, \end{aligned}$$

16. We might alternatively assume specialization across households in the type of labor supplied; in the presence of perfect sharing of labor income risk across households, household decisions regarding consumption and labor supply would all be as assumed here.

where p_i^j is the common price charged by the other firms in industry j .¹⁷ (We introduce the notation $\tilde{\xi}_t$ for the complete vector of exogenous disturbances, including variations in technology as well as in preferences.) If prices were fully flexible, $p_t(i)$ would be chosen each period to maximize this function.

Instead we suppose that prices remain fixed in monetary terms for a random period of time. Following Calvo, we suppose that each industry has an equal probability of reconsidering its prices each period, and we let $0 < \alpha < 1$ be the fraction of industries whose prices remain unchanged each period. In any industry that revises its prices in period t , the new price p_t^* will be the same. This price is implicitly defined by the first-order condition

$$(7) \quad E_t \left\{ \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \Pi_1(p_t^*, p_t^*, P_T; Y_T, M_T/P_T, \tilde{\xi}_T) \right\} = 0.$$

We note furthermore that the stochastic discount factor used to price future profit streams will be given by

$$(8) \quad Q_{t,T} = \beta^{T-t} \frac{u_c(C_T, M_T/P_T; \xi_T)}{u_c(C_t, M_t/P_t; \xi_t)}.$$

Finally, the definition in equation 1 implies a law of motion for the aggregate price index of the form

$$(9) \quad P_t = \left[(1 - \alpha) p_t^{*1-\theta} + \alpha P_{t-1}^{1-\theta} \right]^{\frac{1}{1-\theta}}.$$

Equations 7 and 9, which jointly determine the path of prices given demand conditions, represent the aggregate supply block of our model. It remains to specify the monetary and fiscal policies of the government.¹⁸

17. In equilibrium, all firms in an industry charge the same price at any given time. But we must define profits for an individual supplier i in the case of contemplated deviations from the equilibrium price.

18. The particular specification of monetary and fiscal policy proposed here is not intended to suggest that either monetary or fiscal policy *must* be expected to be conducted according to rules of the sort assumed here. Indeed, in later sections we recommend policy commitments on the part of both the monetary and the fiscal authorities that *do not* conform to the assumptions made here. The point is to define what we mean by the qualification that open-market operations are irrelevant if they do not change expected future monetary or fiscal policy. To make sense of such a statement, we must define what it would mean for

To address the question of whether quantitative easing represents an additional tool of policy, we suppose that the central bank's operating target for the short-term nominal interest rate is determined by a feedback rule in the spirit of the Taylor rule,¹⁹

$$(10) \quad i_t = \phi(P_t/P_{t-1}, Y_t; \tilde{\xi}_t),$$

where now $\tilde{\xi}_t$ may also include exogenous disturbances in addition to the ones listed above, to which the central bank happens to respond. We assume that the function ϕ is nonnegative for all values of its arguments (otherwise the policy would not be feasible, given the zero lower bound), but that there are conditions under which the rule prescribes a zero interest rate policy. Such a rule implies that the central bank supplies the quantity of base money that happens to be demanded at the interest rate given by this formula; hence equation 10 implies a path for the monetary base, so long as the value of ϕ is positive. However, under those conditions in which the value of ϕ is zero, the policy commitment in equation 10 implies only a lower bound on the monetary base that must be supplied. In these circumstances we may ask whether it matters whether a greater or a smaller quantity of base money is supplied. We assume that the central bank's policy in this regard is specified by a base-supply rule of the form

$$(11) \quad M_t = P_t L\left[Y_t, \phi(P_t/P_{t-1}, Y_t; \tilde{\xi}_t); \xi_t\right] \psi(P_t/P_{t-1}, Y_t; \tilde{\xi}_t),$$

where the multiplicative factor ψ satisfies the following two conditions:

$$\begin{aligned} \psi(P_t/P_{t-1}, Y_t; \tilde{\xi}_t) &= 1 \quad \text{if } \phi(P_t/P_{t-1}, Y_t; \tilde{\xi}_t) > 0, \text{ otherwise} \\ \psi(P_t/P_{t-1}, Y_t; \tilde{\xi}_t) &\geq 1. \end{aligned}$$

for all values of its arguments. (The second condition implies that $\psi = 1$ whenever $i_t > 0$.) Note that a base-supply rule of this form is consistent with both the interest rate operating target specified in equation 10 and the equilibrium relations in expressions 3 and 4. The use of quantitative

these policies to be specified in a way that prevents them from being affected by past open-market operations. The specific classes of policy rules discussed here show not only that our concept of "unchanged policy" is logically possible, but indeed that it could correspond to a policy commitment of a fairly familiar sort, one that would represent a commitment to "sound policy" in the views of some.

19. Taylor (1993).

easing as a policy tool can then be represented by a choice of a function ψ that is greater than 1 under some circumstances.

It remains to specify which sort of assets should be acquired (or disposed of) by the central bank when it varies the size of the monetary base. We allow the asset side of the central bank balance sheet to include any of k different types of securities, distinguished from each other by their state-contingent returns. At the end of period t , the vector of nominal values of central bank holdings of the various securities is given by $M_t \omega_t^m$, where ω_t^m is a vector of central bank portfolio shares. These shares are in turn determined by a policy rule of the form

$$(12) \quad \omega_t^m = \omega^m(P_t/P_{t-1}, Y_t; \tilde{\xi}_t),$$

where the vector-valued function $\omega^m(\cdot)$ has the property that its components sum to 1 for all possible values of its arguments. The fact that $\omega^m(\cdot)$ depends on the same arguments as $\phi(\cdot)$ means that we allow for the possibility that the central bank changes its policy when the zero bound is binding (for example, buying assets that it would not hold at any other time). The fact that it depends on the same arguments as $\psi(\cdot)$ allows us to specify changes in the composition of the central bank portfolio as a function of the particular kinds of purchases associated with quantitative easing.

The payoffs on these securities in each state of the world are specified by exogenously given (state-contingent) vectors \mathbf{a}_t and \mathbf{b}_t and matrix \mathbf{F}_t . A vector of asset holdings \mathbf{z}_{t-1} at the end of period $t - 1$ results in delivery, to the owner of a quantity $\mathbf{a}'_t \mathbf{z}_{t-1}$ of money, a quantity $\mathbf{b}'_t \mathbf{z}_{t-1}$ of the consumption good and a vector $\mathbf{F}_t \mathbf{z}_{t-1}$ of securities that may be traded in the period- t asset markets, each of which may depend on the state of the world in period t . This flexible specification allows us to treat a wide range of types of assets that may differ as to maturity, degree of indexation, and so on.²⁰

The gross nominal return $R_t(j)$ on the j th asset between periods $t - 1$ and t is then given by

20. For example, security j in period $t - 1$ is a one-period riskless nominal bond if $b_t(j)$ and $\mathbf{F}_t(\cdot; j)$ are zero in all states, while $a_t(j) > 0$ is the same in all states. Security j is instead a one-period real (or indexed) bond if $a_t(j)$ and $\mathbf{F}_t(\cdot; j)$ are zero, while $b_t(j) > 0$ is the same in all states. It is a two-period riskless nominal pure discount bond if instead $a_t(j)$ and $b_t(j)$ are zero, $F_t(i, j) = 0$ for all $i \neq k$; $F_t(k, j) > 0$ is the same in all states, and security k in period t is a one-period riskless nominal bond.

$$(13) \quad R_t(j) = \frac{a_t(j) + P_t b_t(j) + \mathbf{q}'_t \mathbf{F}_t(\cdot, j)}{q_{t-1}(j)},$$

where \mathbf{q}_t is the vector of nominal asset prices in (ex-dividend) period- t trading. The absence of arbitrage opportunities implies as usual that equilibrium asset prices must satisfy

$$(14) \quad \mathbf{q}'_t = \sum_{T \geq t+1} E_t Q_{t,T} [\mathbf{a}'_T + P_T \mathbf{b}'_T] \prod_{s=t+1}^{T-1} \mathbf{F}_s,$$

where the stochastic discount factor is again given by equation 8. Under the assumption that no interest is paid on the monetary base, the nominal transfer by the central bank to the public treasury each period is equal to

$$(15) \quad T_t^{cb} = \mathbf{R}'_t \boldsymbol{\omega}_{t-1}^n M_{t-1} - M_{t-1},$$

where \mathbf{R}_t is the vector of returns defined by equation 13.

We specify fiscal policy in terms of a rule that determines the evolution of total government liabilities D_t , here defined to be inclusive of the monetary base, as well as a rule that specifies the composition of outstanding nonmonetary liabilities (debt) among different types of securities that the government might issue. We assume that the path of total government liabilities accords with a rule of the form

$$(16) \quad \frac{D_t}{P_t} = d \left(\frac{D_{t-1}}{P_{t-1}}, \frac{P_t}{P_{t-1}}, Y_t; \tilde{\boldsymbol{\xi}}_t \right),$$

which specifies the acceptable level of real government liabilities as a function of the preexisting level and various aspects of current macroeconomic conditions. This notation allows for such possibilities as an exogenously specified state-contingent target for real government liabilities as a proportion of GDP, or for the government budget deficit (inclusive of interest on the public debt) as a proportion of GDP, among others. The part of total liabilities that consists of base money is specified by the base rule in equation 11. We suppose, however, that the rest may be allocated among any of a set of different types of securities that the government may issue; for convenience, we assume that this is a subset of the set of k securities that the central bank may purchase. If $\boldsymbol{\omega}_{jt}^n$ indicates the share of government debt (nonmonetary liabilities) at the end of period t

that is of type j , then the flow government budget constraint takes the form

$$D_t = \mathbf{R}'_t \boldsymbol{\omega}_{t-1}^f B_{t-1} - T_t^{cb} - T_t^h,$$

where $B_t \equiv D_t - M_t$ is the total nominal value of end-of-period nonmonetary liabilities, and T_t^h is the nominal value of the primary budget surplus (taxes net of transfers, if we abstract from government purchases). This identity can then be inverted to obtain the net tax collections T_t^h implied by a given rule (equation 16) for aggregate public liabilities; this depends in general on the composition of the public debt as well as on total borrowing.

Finally, we assume that debt management policy (the determination of the composition of the government's nonmonetary liabilities at each point in time) is specified by the function

$$(17) \quad \boldsymbol{\omega}_t^f = \boldsymbol{\omega}^f(P_t/P_{t-1}, Y_t; \tilde{\boldsymbol{\xi}}_t),$$

which specifies the shares as a function of aggregate conditions, where the vector-valued function $\boldsymbol{\omega}^f$ also has components that sum to 1 for all possible values of its arguments. Together the two relations in equations 16 and 17 complete our specification of fiscal policy and close our model.²¹

We may now define a rational expectations equilibrium as a collection of stochastic processes $\{p_t^*, P_t, Y_t, i_t, \mathbf{q}_t, M_t, \boldsymbol{\omega}_t^m, D_t, \boldsymbol{\omega}_t^f\}$, with each endogenous variable specified as a function of the history of exogenous disturbances to that date, that satisfy each of the conditions in expressions 2 through 6 of the aggregate demand block of the model, the conditions in equations 7 and 9 of the aggregate supply block, the asset-pricing relations equation 14, the conditions in equations 10 through 12 specifying monetary policy, and the conditions in equations 16 and 17, specifying fiscal policy in each period. We then obtain the following irrelevance result for the specification of certain aspects of policy:

21. We might, of course, allow for other types of fiscal decisions from which we abstract here—government purchases, tax incentives, and so on—some of which may be quite relevant to dealing with a liquidity trap. But our concern here is solely with the question of what monetary policy can achieve; we introduce a minimal specification of fiscal policy only for the sake of closing our general-equilibrium model, and to allow discussion of the fiscal implications of possible actions by the central bank.

PROPOSITION. The set of paths for the variables $\{p_t^*, P_t, Y_t, i_t, \mathbf{q}_t, D_t\}$ that are consistent with the existence of a rational expectations equilibrium is independent of the specification of the functions ψ (equation 11), ω^m (equation 12), and ω^f (equation 17).

The reason for this is fairly simple. The set of restrictions on the processes $\{p_t^*, P_t, Y_t, i_t, \mathbf{q}_t, D_t\}$ implied by our model can be written in a form that does not involve the variables $\{M_t, \omega_t^m, \omega_t^f\}$, and hence that does not involve the functions ψ , ω^m , or ω^f . To show this, we first note that, for all $m \geq \bar{m}(C; \xi)$,

$$u(C, m; \xi) = u[C, \bar{m}(C; \xi); \xi],$$

because additional money balances beyond the satiation level provide no further liquidity services. By differentiating this relation, we see further that $u_c(C, m; \xi)$ does not depend on the exact value of m either, as long as m exceeds the satiation level. It follows that, in our equilibrium relations, we can replace the expression $u_c(Y_t, M_t/P_t; \xi_t)$ with

$$\lambda(Y_t, P_t/P_{t-1}; \xi_t) \equiv u_c\{Y_t, L[Y_t, \phi(P_t/P_{t-1}, Y_t; \xi_t); \xi_t]; \xi_t\},$$

using the fact that expression 3 holds with equality at all levels of real balances at which u_c depends on the level of real balances. Hence we can write u_c as a function of variables other than M_t/P_t , without using the relation in equation 11, and so in a way that is independent of the function ψ . We can similarly replace the expression $u_m(Y_t, M_t/P_t; \xi_t)(M_t/P_t)$ in expression 5 with

$$\begin{aligned} \mu(Y_t, P_t/P_{t-1}; \xi_t) \equiv \\ u_m\{Y_t L[Y_t, \phi(P_t/P_{t-1}, Y_t; \xi_t); \xi_t]; \xi_t\} L[Y_t, \phi(P_t/P_{t-1}, Y_t; \xi_t); \xi_t], \end{aligned}$$

since M_t/P_t must equal $L[Y_t, \phi(P_t/P_{t-1}, Y_t; \xi_t); \xi_t]$ when real balances do not exceed the satiation level, whereas $u_m = 0$ when they do. Finally, we can express nominal profits in period t as a function:

$$\tilde{\Pi}\left[p_t(i), p_t^j, P_t; Y_t, P_t/P_{t-1}, \tilde{\xi}_t\right],$$

after substituting $\lambda(Y_t, P_t/P_{t-1}; \xi_t)$ for the marginal utility of real income in the wage demand function that is used in deriving the profit function

Π.²² Using these substitutions, we can write each of the equilibrium relations in expressions 2, 5, 6, 7, and 14 in a way that no longer makes reference to the money supply.

It then follows that in a rational expectations equilibrium the variables $\{p_t^*, P_t, Y_t, i_t, \mathbf{q}_t, D_t\}$ must satisfy in each period the following relations:

$$(18) \quad \lambda(Y_t, P_t/P_{t-1}; \xi_t) = \beta E_t \left[\lambda(Y_{t+1}, P_{t+1}/P_t; \xi_{t+1})(1+i_t) \frac{P_t}{P_{t+1}} \right]$$

$$(19) \quad \sum_{T=t}^{\infty} \beta^T E_t [\lambda(Y_T, P_T/P_{T-1}; \xi_T) Y_T + \mu(Y_T, P_T/P_{T-1}; \xi_T)] < \infty$$

$$(20) \quad \lim_{T \rightarrow \infty} \beta^T E_t [\lambda(Y_T, P_T/P_{T-1}; \xi_T) D_T/P_T] = 0$$

$$(21) \quad \mathbf{q}'_t = \frac{P_t}{\lambda(Y_t, P_t/P_{t-1}; \xi_t)} \sum_{T=t+1}^{\infty} \beta^{T-t} E_t \lambda(Y_T, P_T/P_{T-1}; \xi_T) [P_T^{-1} \mathbf{a}'_T + \mathbf{b}'_T] \prod_{s=t+1}^{T-1} \mathbf{F}_s$$

$$(22) \quad E_t \left\{ \sum_{T=t}^{\infty} (\alpha\beta)^{T-t} \lambda(Y_T, P_T/P_{T-1}; \xi_T) P_T^{-1} \tilde{\Pi}_1(p_t^*, p_t^*, P_t, Y_t, P_t/P_{T-1}, \tilde{\xi}_T) \right\} = 0,$$

along with equations 9, 10, and 16 as before. None of these involve the variables $\{M_t, \omega_t^m, \omega_t^f\}$, nor do they involve the functions ψ , ω^m , or ω^f .

Furthermore, this is the complete set of restrictions on these variables that are required in order for them to be consistent with a rational expectations equilibrium. For any given processes $\{p_t^*, P_t, Y_t, i_t, \mathbf{q}_t, D_t\}$ that satisfy the equations just listed in each period, the implied path of the money supply is given by equation 11, which clearly has a solution, and this path for the money supply necessarily satisfies expression 3 and the complementary slackness condition, as a result of our assumptions about the form of the function ψ . Similarly, the implied compositions of the central bank portfolio and of the public debt at each point in time are given by equations 12 and 17. We then have a set of processes that satisfy

22. See Woodford (forthcoming, chapter 3).

all of the requirements for a rational expectations equilibrium, and the result is established.

Discussion

The above proposition implies that neither the extent to which quantitative easing is employed when the zero bound binds, nor the nature of the assets that the central bank may purchase through open-market operations, has any effect on whether a deflationary price-level path represents a rational expectations equilibrium. Hence our general-equilibrium analysis of inflation and output determination does not support the notion that expansions of the monetary base represent an additional tool of policy, independent of the specification of the rule for adjusting short-term nominal interest rates. If the commitments of policymakers regarding the rule by which interest rates will be set, on the one hand, and the rule by which total private sector claims on the government will be allowed to grow, on the other, are fully credible, then it is only the choice of those commitments that matters. Other aspects of policy should matter in practice only insofar as they help to signal the nature of these policy commitments.

Of course, the validity of our result depends on the reasonableness of our assumptions, and these deserve further discussion. Like any economic model, ours abstracts from the complexity of actual economies in many respects. Have we abstracted from features of actual economies that are crucial for a correct understanding of the issues under discussion?

It might be suspected that an important omission is our neglect of portfolio-balance effects, which play an important role in much recent discussion of the policy options that would remain available to the Federal Reserve should the federal funds rate reach zero.²³ The idea is that a central bank should be able to lower longer-term interest rates even when overnight rates are already at zero, through purchases of longer-maturity government bonds. This would shift the composition of the public debt in the hands of the public in a way that affects the term structure of interest rates. (Because it is generally admitted in such discussions that base money and very short term Treasury securities have become near-perfect substitutes once short-term interest rates have fallen to zero, the desired effect should be achieved equally well by a shift in the maturity structure

23. See, for example, Clouse and others (2003) and Orphanides (2003).

of Treasury securities held by the central bank, without any change in the monetary base, as by an open-market purchase of long-term bonds with newly created base money.)

No such effects arise in our model, whether from central bank securities purchases or debt management by the public treasury. But this is not, as some might expect, because we have simply assumed that bonds of different maturities (or for that matter, other kinds of assets that the central bank might choose to purchase instead of the shortest-maturity Treasury bills) are perfect substitutes. Our framework allows for central bank purchases of different assets having different risk characteristics (different state-contingent returns), and our model of asset market equilibrium incorporates those term premiums and risk premiums that are consistent with the absence of arbitrage opportunities.

Our conclusion differs from that of the literature on portfolio balance effects for a different reason. The classic theoretical analysis of portfolio balance effects assumes a representative investor with mean-variance preferences. This has the implication that if the supply of assets that pay off disproportionately in certain states of the world is increased (so that the extent to which the representative investor's portfolio pays off in those states must also increase), the relative marginal valuation of income in those particular states is reduced, resulting in a lower relative price for the assets that pay off in those states. But in our general-equilibrium asset pricing model, there is no such effect. The marginal utility to the representative household of additional income in a given state of the world depends on the household's consumption in that state, not on the aggregate payoff of its asset portfolio in that state. And changes in the composition of the securities in the hands of the public *do not* change the state-contingent consumption of the representative household—this depends on equilibrium output, and although output is endogenous, we have shown that the equilibrium relations that determine it do not involve the functions ψ , ω^m , or ω^f .²⁴

Our assumption of complete financial markets and no limits on borrowing against future income may also appear extreme. However, the

24. Our general-equilibrium analysis is in the spirit of the irrelevance proposition for open-market operations of Wallace (1981). Wallace's analysis is often supposed to be of little practical relevance for actual monetary policy because his model is one in which money serves only as a store of value, so that an equilibrium in which short-term Treasury securities dominate money in terms of rate of return is not possible, although this is

assumption of complete financial markets is only a convenience, allowing us to write the budget constraint of the representative household in a simple way. Even in the case of incomplete markets, each of the assets that is traded will be priced according to equation 14, where the stochastic discount factor is given by equation 8, and once again there will be a set of relations to determine output, goods prices, and asset prices that do not involve ψ , ω^m , or ω^f . The absence of borrowing limits is also innocuous, at least in the case of a representative-household model, because in equilibrium the representative household must hold the entire net supply of financial claims on the government. As long as the fiscal rule (equation 16) implies positive government liabilities at each date, any borrowing limits that might be assumed can never bind in equilibrium. Borrowing limits can matter more in the case of a model with heterogeneous households. But in this case the effects of open-market operations should depend not merely on which sorts of assets are purchased and which sorts of liabilities are issued to finance those purchases, but also on how the central bank's trading profits are eventually rebated to the private sector (that is, with what delay and how distributed across the heterogeneous households), as a result of the specification of fiscal policy. The effects will not be mechanical consequences of the change in composition of assets in the hands of the public, but instead will result from the fiscal transfers to which the transaction gives rise; it is unclear how quantitatively significant such effects should be.

Indeed, leaving aside the question of whether a clear theoretical foundation exists for the existence of portfolio balance effects, there is not a great deal of empirical support for quantitatively significant effects. The attempt to separately target short-term and long-term interest rates under Operation Twist in the early 1960s is generally regarded as having had a modest effect at best on the term structure.²⁵ The empirical literature that has sought to estimate the effects of changes in the composition of the public debt on relative yields has also, on the whole, found effects that are

routinely observed. However, in the case of open-market operations conducted at the zero bound, the liquidity services provided by money balances at the margin have fallen to zero, so that an analysis of the kind proposed by Wallace is correct.

25. Okun (1963) and Modigliani and Sutch (1966) are important early discussions that reached this conclusion. Meulendyke (1998) summarizes the literature and finds that the predominant view is that the effect was minimal.

not large when present at all.²⁶ For example, Jonas Agell and Mats Persson summarize their findings as follows: “It turned out that these effects were rather small in magnitude, and that their numerical values were highly volatile. Thus the policy conclusion to be drawn seems to be that there is not much scope for a debt management policy aimed at systematically affecting asset yields.”²⁷ Moreover, even if one supposes that large enough changes in the composition of the portfolio of securities left in the hands of the private sector can substantially affect yields, it is not clear how relevant such an effect should be for real activity and the evolution of goods prices. For example, James Clouse and others argue that a sufficiently large reduction in the number of long-term Treasuries in the hands of the public should lower the market yield on those securities relative to short-term rates, because certain institutions will find it important to hold long-term Treasury securities even when they offer an unfavorable yield.²⁸ But even if this is true, the fact that these institutions have idiosyncratic reasons to hold long-term Treasuries—and that, in equilibrium, no one else holds any or plays any role in pricing them—means that the lower observed yield on long-term Treasuries may not correspond to any reduction in the perceived cost of long-term borrowing for other institutions. If one is able to reduce the long-term bond rate only by decoupling it from the rest of the structure of interest rates, and from the cost of financing long-term investment projects, it is unclear that such a reduction should do much to stimulate economic activity or to halt deflationary pressures.

Hence we are not inclined to suppose that our irrelevance proposition represents so poor an approximation to reality as to deprive it of practical relevance. Even if the effects of open-market operations under the conditions the proposition describes are not exactly zero, it seems unlikely that they should be large. In our view it is more important to note that our

26. Examples of studies finding either no effects or only quantitatively unimportant ones include Modigliani and Sutch (1967), Frankel (1985), Agell and Persson (1992), Wallace and Warner (1996), and Hess (1999). Roley (1982) and Friedman (1992) find somewhat larger effects.

27. Agell and Persson (1992, p. 78).

28. Clouse and others (2003). Stephen G. Cecchetti (“Central Banks Have Plenty of Ammunition,” *Financial Times*, March 17, 2003, p. 13) similarly argues that it should be possible for the Federal Reserve to independently affect long-term bond yields if it is determined to do so, given that it can print money without limit to buy additional long-term Treasuries if necessary.

irrelevance proposition depends on an assumption that interest rate policy is specified in a way that implies that these open-market operations have no consequences for interest rate policy, either immediately (which is trivial, because it would not be possible for them to *lower* current interest rates, which is the only effect that would be desired), or at any subsequent date. We have also specified fiscal policy in a way that implies that the contemplated open-market operations have no effect on the path of total government liabilities $\{D_t\}$ either, whether immediately or at any later date. Although we think these definitions make sense, as a way of isolating the pure effects of open-market purchases of assets by the central bank from either interest rate policy on the one hand or fiscal policy on the other, those who recommend monetary expansion by the central bank may intend for this to have consequences of one or both of these other sorts.

For example, when it is argued that a “helicopter drop” of money into the economy would surely stimulate nominal aggregate demand, the thought experiment that is usually contemplated is not simply a change in the function ψ in our policy rule equation 11. First of all, it is typically supposed that the expansion of the money supply will be *permanent*. If this is the case, then the function ϕ that defines interest rate policy is also being changed, in a way that will become relevant at some future date, when the money supply no longer exceeds the satiation level.²⁹ Second, the assumption that the money supply is increased through a helicopter drop rather than an open-market operation implies a change in fiscal policy as well. Such an operation would increase the value of nominal government liabilities, and it is generally at least tacitly assumed that this is a permanent increase as well. Hence the experiment that is imagined is not one that our irrelevance proposition implies should have no effect on the equilibrium path of prices.

29. This explains the apparent difference between our result and that obtained by Auerbach and Obstfeld (2003) in a similar model. These authors assume explicitly that an increase in the money supply at a time when the zero bound binds carries with it the implication of a permanently larger money supply, and that there exists a future date at which the zero bound ceases to bind, so that the larger money supply will imply a different interest rate policy at that later date. Clouse and others (2003) also stress that maintenance of the larger money supply until a date at which the zero bound would not otherwise bind represents one straightforward channel through which open-market operations while the zero bound is binding could have a stimulative effect, although they discuss other possible channels as well.

Even more important, our irrelevance result applies only *given* a correct private sector understanding of the central bank's commitments regarding future policy. Such understanding may be lacking. We have just argued that the key to lowering long-term interest rates, in a way that actually provides an incentive for increased spending, is to change expectations regarding the likely future path of short-term rates, rather than through intervention in the market for long-term Treasuries. As a matter of logic, this need not require any open-market purchases of long-term Treasuries at all. Nonetheless, the private sector may be uncertain about the nature of the central bank's policy commitment, and so it may scrutinize the bank's current actions for further clues. In practice, the management of private sector expectations is an art of considerable subtlety, and shifts in the portfolio of the central bank could be of some value in making credible to the private sector the central bank's own commitment to a particular kind of future policy, as we discuss further in the penultimate section of the paper. Signaling effects of this kind are often argued to be an important reason for the effectiveness of interventions in foreign-exchange markets, and they might well provide a justification for open-market operations when the zero bound binds.³⁰

We do not wish, then, to argue that asset purchases by the central bank are necessarily pointless under the circumstances of a binding zero lower bound on short-term nominal interest rates. However, we do think it important to observe that, insofar as such actions can have any effect, it is not because of any necessary or mechanical consequence of the shift in the portfolio of assets in the hands of the private sector itself. Instead, any effect of such actions must be due to the way in which they change expectations regarding future interest rate policy or, perhaps, the future path of total nominal government liabilities. Later we discuss reasons why open-market purchases by the central bank might plausibly have consequences for expectations of these types. But because it is only through effects on expectations regarding future policy that these actions can matter, we focus our attention on the question of what kind of commitments regarding future policy are in fact to be desired. And this question can be addressed without explicit consideration of the role of central bank open-market operations of any kind. Hence we will simplify our model—

30. Clouse and others (2003) argue that this is one important channel through which open-market operations can be effective.

abstracting from monetary frictions and the structure of government liabilities altogether—and instead consider what is the desirable conduct of interest rate policy, and what kind of commitments about this policy are desirable to make in advance.

How Severe a Constraint Is the Zero Bound?

We turn now to the question of how the existence of the zero bound restricts the degree to which a central bank's stabilization objectives, with regard to both inflation and real activity, can be achieved, even under ideal policy. The discussion in the previous section established that the zero bound does represent a genuine constraint. It is not true that equilibria that cannot be achieved through a suitable interest rate policy can somehow be achieved through other means, and the zero bound does limit the set of possible equilibrium paths for prices and output, although the quantitative importance of this constraint remains to be seen.

Nonetheless, we will see that it is not at all the case that a central bank can do nothing to mitigate the severity of the destabilizing impact of the zero bound. The reason is that inflation and output do not depend solely on the current level of short-term nominal interest rates, or even solely on the history of such rates up until the present (so that the current level of interest rates would be the only thing that could possibly change in response to an unanticipated disturbance). The expected character of *future* interest rate policy is also a critical determinant of the degree to which the central bank achieves its stabilization objectives, and this allows important scope for policy to be improved upon, even when there is little choice about the current level of short-term interest rates.

In fact, the management of expectations is the key to successful monetary policy at *all* times, not just in those relatively unusual circumstances when the zero bound is reached. The effectiveness of monetary policy has little to do with the direct effect of changing the level of overnight interest rates, since the current cost of maintaining cash balances overnight is of fairly trivial significance for most business decisions. What actually matters is the private sector's anticipation of the *future* path of short-term rates, because this determines equilibrium long-term interest rates as well as equilibrium exchange rates and other asset prices—all of which are quite relevant for many current spending decisions, and hence for optimal

pricing behavior as well. How short-term rates are managed matters because of the signals that such management gives about how the private sector can expect them to be managed in the future. But there is no reason to suppose that expectations regarding future monetary policy, and hence regarding the future paths of nominal variables more generally, should change only insofar as the current level of overnight interest rates changes. A situation in which there is no decision to be made about the current level of overnight rates (as in Japan at present) is one that gives urgency to the question of what expectations regarding future policy one should wish to create, but this is in fact the correct way to think about sound monetary policy at all times.

Of course, the question of what future policy one should wish people to expect does not arise if current constraints leave no possibility of committing oneself to a different sort of policy in the future than one would otherwise have pursued. This means that the private sector must be convinced that the central bank will not conduct policy in a way that is *purely forward looking*, that is, taking account at each point in time only of the possible paths that the economy could follow from that date onward. For example, we will show that it is undesirable for the central bank to pursue a given inflation target, once the zero bound is expected no longer to prevent that target from being achieved, even in the case that the pursuit of this target *would* be optimal if the zero bound did not exist (or would never bind under an optimal policy). The reason is that an expectation that the central bank will pursue the fixed inflation target after the zero bound ceases to bind gives people no reason to hold the kind of expectations, while the bound is binding, that would mitigate the distortions created by it. A history-dependent inflation target³¹—if the central bank’s commitment to it can be made credible—can instead yield a superior outcome.

But this, too, is an important feature of optimal policy rules more generally.³² Hence the analytical framework and institutional arrangements used in making monetary policy need not be changed in any fundamental way in order to deal with the special problems created by a liquidity trap. As we explain later in the paper, the optimal policy in the case of a binding zero bound can be implemented through a targeting procedure that

31. As we will show, it is easier to explain the nature of the optimal commitment if it is described as a history-dependent price-level target.

32. See, for example, Woodford (forthcoming, chapter 7).

represents a straightforward generalization of a policy that would be optimal even if the zero bound were expected never to bind.

Feasible Responses to Fluctuations in the Natural Rate of Interest

In order to characterize how stabilization policy is constrained by the zero bound, we make use of a log-linear approximation to the structural equations presented in the previous section, of a kind that is often employed in the literature on optimal monetary stabilization policy.³³ Specifically, we log-linearize the structural equations of our model (except for the zero bound in expression 4) around the paths of inflation, output, and interest rates associated with a zero-inflation steady state, in the absence of disturbances ($\xi_t = 0$). We choose to expand around these particular paths because the zero-inflation steady state represents optimal policy in the absence of disturbances.³⁴ In the event of small enough disturbances, optimal policy will still involve paths in which inflation, output, and interest rates are at all times close to those of the zero-inflation steady state. Hence an approximation to our equilibrium conditions that is accurate in the case of inflation, output, and interest rates near those values will allow an accurate approximation to the optimal responses to disturbances in the case that the disturbances are small enough.

In the zero-inflation steady state, it is easily seen that the real rate of interest is equal to $\bar{r} \equiv \beta^{-1} - 1 > 0$; this is also the steady-state nominal interest rate. Hence, in the case of small enough disturbances, optimal policy will involve a nominal interest rate that is always positive, and the zero bound will not be a binding constraint. (Optimal policy in this case is characterized in the references cited in the previous paragraph.) However, we are interested in the case in which disturbances are at least occasionally large enough for the zero bound to bind, that is, to prevent attainment of the outcome that would be optimal in the absence of such a bound. It is

33. See, for example, Clarida, Galí, and Gertler (1999); Woodford (forthcoming).

34. See Woodford (forthcoming, chapter 7) for more detailed discussion of this point. The fact that zero inflation, rather than mild deflation, is optimal depends on our abstracting from transactions frictions, as discussed further in footnote 40 below. As Woodford shows, a long-run inflation target of zero is optimal in this model, even when the steady-state output level associated with zero inflation is suboptimal, owing to market power. The reason is that a commitment to inflation in some period t results both in increased output in period t and in reduced output in period $t - 1$ (owing to the effect of expected inflation on the aggregate supply relation, equation 25 below); because of discounting, the second effect on welfare fully offsets the benefit of the first effect.

possible to consider this problem rigorously using only a log-linear approximation to the structural equations in the case where the lower bound on nominal interest is assumed to be not much below \bar{r} . We can arrange for this gap to be as small as we may wish, without changing other crucial parameters of the model such as the assumed rate of time preference, by supposing that interest is paid on the monetary base at a rate $i^m \geq 0$ that cannot (for some institutional reason) be reduced. Then the lower bound on interest rates actually becomes

$$(23) \quad i_t \geq i^m.$$

We will characterize optimal policy subject to a constraint of the form of expression 23, in the case that *both* a bound on the amplitude of disturbances $\|\xi\|$ and the steady-state opportunity cost of holding money $\bar{\delta} \equiv (\bar{r} - i^m)/(1 + \bar{r}) > 0$ are small enough. Specifically, both our structural equations and our characterization of the optimal responses of inflation, output, and interest rates to disturbances will be required to be exact only up to a residual of order $O(\|\xi; \bar{\delta}\|^2)$: We then hope (without here seeking to verify) that our characterization of optimal policy in the case of a small opportunity cost of holding money and small disturbances is not too inaccurate in the case of an opportunity cost of several percentage points (the case in which $i^m = 0$) and disturbances large enough to cause the natural rate of interest to vary by several percentage points (as will be required in order for the zero bound to bind).

As Woodford has shown elsewhere,³⁵ the log-linear approximate equilibrium relations may be summarized by two equations each period: a forward-looking “IS relation”

$$(24) \quad x_t = E_t x_{t+1} - \sigma(i_t - E_t \pi_{t+1} - r_t^n),$$

and a forward-looking “AS relation” (or New Keynesian Phillips curve)

$$(25) \quad \pi_t = \kappa x_t + \beta E_t \pi_{t+1} + u_t.$$

Here $\pi_t \equiv \log(P_t/P_{t-1})$ is the inflation rate, x_t is a welfare-relevant output gap, and i_t is now the continuously compounded nominal interest rate, corresponding to $\log(1 + i_t)$ in the notation used in the previous section.

35. Woodford (forthcoming).

The terms u_t and r_t^n are composite exogenous disturbance terms that shift the two equations; the former is commonly referred to as a cost-push disturbance, whereas the latter indicates exogenous variation in the Wicksellian natural rate of interest, that is, the equilibrium real rate of interest in the case that output growth is at all times equal to its natural rate. The coefficients σ and κ are both positive, and $0 < \beta < 1$ is again the utility discount factor of the representative household.

Equation 24 is a log-linear approximation to equation 2, whereas equation 25 is derived by log-linearizing equations 7 through 9 and then eliminating $\log(p_t^*/P_t)$. We omit the log-linear version of the money demand relation in expression 3, because here we are interested solely in characterizing the possible equilibrium paths of inflation, output, and interest rates, and we may abstract from the question of what might be the required path for the monetary base that is associated with any such equilibrium. (It suffices that there exist a monetary base that will satisfy the money demand relation in each case, and this will be true as long as the interest rate bound is satisfied.) The other equilibrium requirements of the earlier discussion can be ignored in the case that we are interested only in possible equilibria that remain forever near the zero-inflation steady state, because they are automatically satisfied in that case. Equations 24 and 25 represent a pair of equations each period to determine inflation and the output gap, given the central bank's interest rate policy. We will seek to compare alternative possible paths for inflation, the output gap, and the nominal interest rate that satisfy these two log-linear equations together with expression 23. Note that our conclusions will be identical (up to a scale factor) in the event that we multiply the amplitude of the disturbances and the steady-state opportunity cost $\bar{\delta}$ by any common factor; alternatively, if we measure the amplitude of disturbances in units of $\bar{\delta}$, our results will be independent of the value of $\bar{\delta}$ (to the extent that our log-linear approximation remains valid). Hence we choose the normalization $\bar{\delta} = 1 - \beta$, corresponding to $i^m = 0$, to simplify the presentation. In that case the lower bound for the nominal interest rate is again given by expression 4.

Deflation under Forward-Looking Policy

We begin by considering the degree to which the zero bound impedes the achievement of the central bank's stabilization objectives in the case

that the bank pursues a strict inflation target. We interpret this as a commitment to adjust the nominal interest rate so that

$$(26) \quad \pi_t = \pi^*$$

each period, insofar as it is possible to achieve this with some nonnegative interest rate. It is easy to verify, by the IS and AS equations above, that a necessary condition for this target to be satisfied is

$$(27) \quad i_t = r_t^n + \pi^*.$$

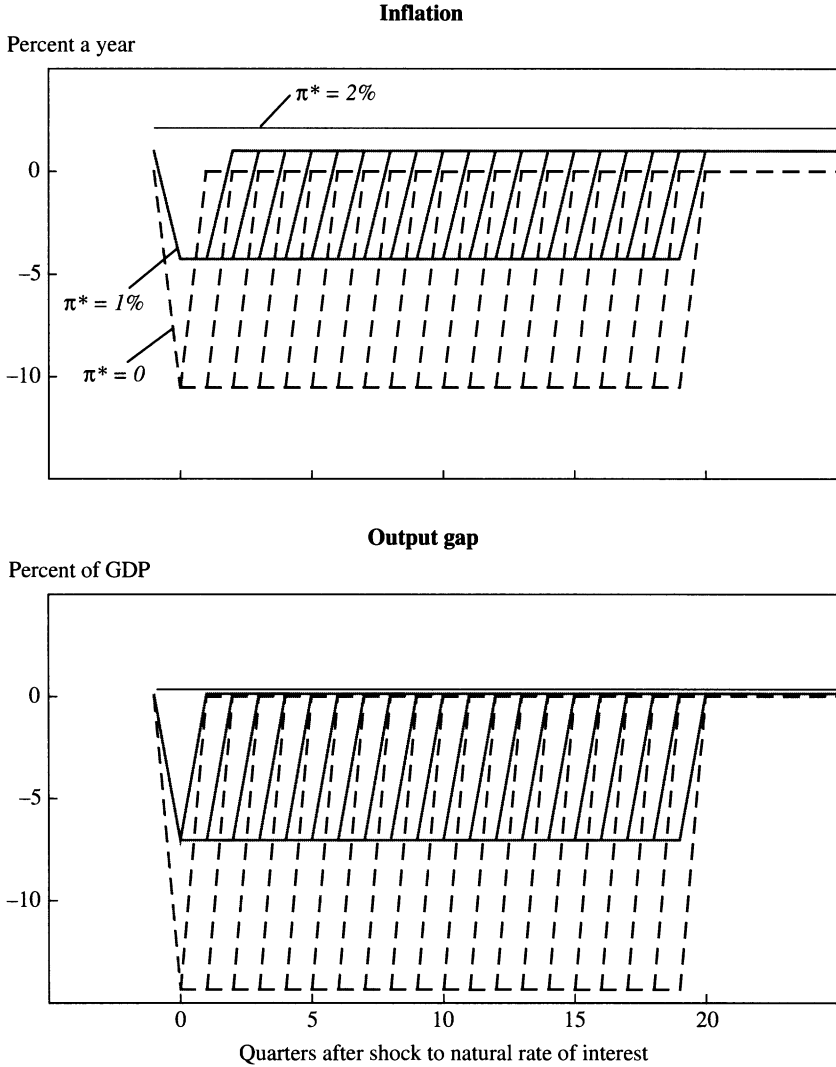
When inflation is on target, the real interest rate is equal to the natural real rate at all times, and the output gap is at its long-run level. The zero bound, however, prevents equation 27 from holding if $r_t^n < -\pi^*$. Thus, if the natural rate of interest is low, the zero bound frustrates the central bank's ability to implement an inflation target. Suppose the inflation target is zero, so that $\pi^* = 0$. Then the zero bound is binding if the natural rate of interest is negative, and the central bank is unable to achieve its inflation target.

To illustrate this, consider the following experiment. Suppose the natural rate of interest is unexpectedly negative in period 0 and reverts back to its steady-state value $\bar{r} > 0$ with a fixed probability in every period. Figure 2 shows the state-contingent paths of the output gap and inflation under these circumstances for each of three different possible inflation targets π^* . We assume in period 0 that the natural rate of interest becomes -2 percent a year and then reverts back to the steady-state value of $+4$ percent a year with a probability of 0.1 each quarter. Thus the natural rate of interest is expected to be negative for ten quarters on average at the time the shock occurs.

The dashed lines in figure 2 show the state-contingent paths of the output gap and inflation if the central bank targets zero inflation.³⁶ Starting

36. In our numerical analysis, we interpret periods as quarters, and we assume coefficient values of $\sigma = 0.5$, $\kappa = 0.02$, and $\beta = 0.99$. The assumed value of the discount factor implies a long-run real rate of interest \bar{r} equal to 4 percent a year. The assumed value of κ is consistent with the empirical estimate of Rotemberg and Woodford (1997). The assumed value of σ represents a relatively low degree of interest sensitivity of aggregate expenditure. We prefer to bias our assumptions in the direction of only a modest effect of interest rates on the timing of expenditure, so as not to exaggerate the size of the output contraction that is predicted to result from an inability to lower interest rates when the zero bound

Figure 2. State-Contingent Responses of Inflation and the Output Gap to a Shock to the Natural Rate of Interest under Strict Inflation Targeting^a



Source: Authors' calculations.

a. The targeted rate of inflation is designated by π^* . Each upward-sloping line represents the response of inflation or the output gap if the natural rate of interest returns to its steady-state value in that period.

from the left, the first dashed line shows the equilibrium that prevails if the natural rate of interest returns to the steady state in period 1, the next line if it returns in period 2, and so on. The inability of the central bank to set a negative nominal interest rate results in a 14 percent output gap and 10 percent annual deflation. The fact that in each quarter there is a 90 percent chance of the natural rate of interest remaining negative for the next quarter creates the expectation of future deflation and a continued negative output gap, which creates even further deflation. Even if the central bank lowers the short-term nominal interest rate to zero, the real rate of return is positive, because the private sector expects deflation.

The shaded lines in figure 2 show the equilibrium that prevails if the central bank instead sets a 1 percent annual inflation target. In this case the private sector expects 1 percent inflation once the economy is out of the liquidity trap. This, however, is not enough to offset the -2 percent natural rate of interest, so that in equilibrium the private sector expects deflation instead of inflation. The result of this and a negative natural rate of interest is 4 percent annual deflation (when the natural rate of interest is negative) and an output gap of 7 percent.

Finally, the solid horizontal line shows the evolution of output and inflation in the case where the central bank targets 2 percent annual inflation. In this case the central bank can satisfy equation 4 even when the natural rate of interest is negative. When the natural rate of interest is -2 percent, the central bank lowers the nominal interest rate to zero. Since the inflation target is 2 percent, the real rate is -2 percent, which is enough to close the output gap and keep inflation on target. If the inflation target is high enough, therefore, the central bank is able to accommodate a negative natural rate of interest. This is the argument given by Edmund Phelps, Lawrence Summers, and Stanley Fischer for a positive inflation target.³⁷ Krugman makes a similar argument and suggests more concretely that, in 1998, Japan needed a positive inflation target of 4 percent under its then-current circumstances to achieve negative real rates and curb deflation.³⁸

Although it is clear that commitment to a higher inflation target will indeed guard against the need for an output gap in periods when the nat-

binds. As figure 2 shows, even for this value of σ , the output contraction that results from a slightly negative value of the natural rate of interest is quite substantial.

37. Phelps (1972); Summers (1991); Fischer (1996).

38. Krugman (1998).

ural rate of interest falls, the price of this solution is the distortions created by the inflation, both when the natural rate of interest is negative and under more normal circumstances as well. Hence the optimal inflation target (from among the strict inflation targeting policies just considered) will be some value that is at least slightly positive, in order to mitigate the distortions created by the zero bound when the natural rate of interest is negative, but not so high as to keep the zero bound from ever binding (see the table in the next section). An intermediate inflation target, in contrast (like the 1 percent target considered in the figure), leads to a substantial recession when the natural rate of interest becomes negative, and chronic inflation at all other times. Hence no such policy allows a complete solution of the problem posed by the zero bound in the case that the natural rate of interest is sometimes negative.

Nor can one do better through commitment to any policy rule that is purely forward looking in the sense discussed elsewhere by Woodford.³⁹ A purely forward-looking policy is one under which the central bank's action at any time depends only on an evaluation of the possible paths for the central bank's target variables (here, inflation and the output gap) that are possible from the current date forward, neglecting past conditions except insofar as they constrain the economy's possible future path. In the log-linear model presented above, the possible paths for inflation and the output gap from period t onward depend only on the expected evolution of the natural rate of interest from period t onward. If one assumes a Markovian process for the natural rate, as in the numerical analysis above, then any purely forward-looking policy will result in an inflation rate, output gap, and nominal interest rate in period t that depend only on the natural rate in period t —in our numerical example, on whether the natural rate is still negative or has already returned to its long-run steady-state value. It is easily shown in the case of our two-state example that the optimal state-contingent path for inflation and output from among those with this property will be one in which the zero bound binds if and only if the natural rate is in the low state; hence it will correspond to a strict inflation target of the kind just considered, for some π^* between zero and 2 percent.

But one can actually do considerably better, through commitment to a *history-dependent* policy, in which the central bank's actions will depend

39. Woodford (2000).

on past conditions even though these are irrelevant to the degree to which its stabilization goals could in principle be achieved from then on. In the next section we characterize the optimal form of history-dependent policy and determine the degree to which it improves upon the stabilization of both output and inflation.

The Optimal Policy Commitment

We now characterize optimal monetary policy, by optimizing over the set of all possible state-contingent paths for inflation, output, and the short-term nominal interest rate consistent with the log-linearized structural relations in equations 24 and 25. It is assumed (for now) that the expectations regarding future state-contingent policy that are required for such an equilibrium can be made credible to the private sector. In considering the central bank's optimization problem under the assumption that a credible commitment is possible regarding future policy, we do not mean to minimize the subtlety of the task of actually communicating such a commitment to the public and making it credible. However, we do not believe it makes sense to recommend a policy that would systematically seek to achieve an outcome *other* than a rational expectations equilibrium. That is, we are interested in policies that will have the desired effect even when correctly understood by the public. Optimization under the assumption of credible commitment is simply a way of finding the best possible rational expectations equilibrium. Once the equilibrium that one would like to bring about has been identified, along with the interest rate policy that it requires, one can turn to the question of how best to signal these intentions to the public (an issue that we briefly address in the paper's penultimate section).

We assume that the government minimizes

$$(28) \quad \min E_0 \left\{ \sum_{t=0}^{\infty} \beta^t (\pi_t^2 + \lambda x_t^2) \right\}.$$

This loss function can be derived by a second-order Taylor expansion of the utility of the representative household.⁴⁰ The optimal program can be

40. See Woodford (forthcoming, chapter 6) for details. This approximation applies in the case that we abstract from monetary frictions as assumed in this section. If transactions frictions are instead nonnegligible, the loss function should include an additional term proportional to $(i_t - i^m)^2$. This would indicate welfare gains from keeping nominal interest

found by a Lagrangian method, extending the methods used by Richard Clarida, Jordi Galí, and Mark Gertler and by Woodford to the case in which the zero bound can sometimes bind, as shown by Taehun Jung, Yuki Teranishi, and Tsutomu Watanabe.⁴¹ We combine the zero bound and the IS equation to yield the following inequality:

$$x_t \leq E_t x_{t+1} + \sigma(r_t^n + E_t \pi_{t+1}).$$

The Lagrangian for this problem is then

$$\mathcal{L}_0 = E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{1}{2} [\pi_t^2 + \lambda x_t^2] + \varphi_{1t} [x_t - x_{t+1} - \sigma \pi_{t+1} - \sigma r_t^n] + \varphi_{2t} [\pi_t - \kappa x_t - \beta \pi_{t+1}] \right\}.$$

Jung, Teranishi, and Watanabe show that the first-order conditions for an optimal policy commitment are

$$(29) \quad \pi_t + \varphi_{2t} - \varphi_{2t-1} - \beta^{-1} \sigma \varphi_{1t-1} = 0$$

$$(30) \quad \lambda x_t + \varphi_{1t} - \beta^{-1} \varphi_{1t-1} - \kappa \varphi_{2t} = 0$$

$$(31) \quad \varphi_{1t} \geq 0, \quad i_t \geq 0, \quad \varphi_{1t} i_t = 0.$$

One cannot solve this system by applying standard solution methods for rational expectations models, because of the complications of the nonlinear constraint in equation 31. The appendix describes the numerical

rates as close as possible to the zero bound (or, more generally, the lower bound i^m). Nonetheless, because of the stickiness of prices, it would not be optimal for interest rates to be at zero at all times, as implied by the flexible-price model discussed by Uhlig (2000). The optimal inflation rate in the absence of shocks would be slightly negative, rather than zero as in the “cashless” model considered in this section; but it would not be so low that the zero bound would be reached, except in the event of temporary declines in the natural rate of interest, as in the analysis here. Note also that equation 28 implies that the optimal output gap is zero. More generally, there should be an output gap stabilization objective of the form $(x_t - x^*)^2$; the utility-based loss function involves $x^* = 0$ only if one assumes the existence of an output or employment subsidy that offsets the distortion due to the market power of firms. However, the value of x^* affects neither the optimal state-contingent paths derived in this section, and shown in figures 3 and 4, nor the formulas given in the earlier section for the optimal targeting rule.

41. Clarida, Galí, and Gertler (1999); Woodford (1999; forthcoming, chapter 7); Jung, Teranishi, and Watanabe (2001).

method we use to solve these equations instead.⁴² Here we discuss the results that we obtain for the particular numerical experiment considered in the previous section.

What is apparent from the first-order conditions is that optimal policy is history dependent, so that the optimal choice of inflation, the output gap, and the nominal interest rate depends on the past values of the endogenous variables. This can be seen by the appearance of a lagged value of the Lagrange multipliers in the first-order conditions. To get a sense of how this history dependence matters, it is useful to consider again the numerical example shown in figure 2. Suppose the natural rate of interest becomes negative in period 0 and then reverts to the steady state with a fixed probability in each period. Figure 3 shows the optimal output gap, the inflation rate, and the price level from period 0 to period 25. As in figure 2, the separate lines in each panel show the evolution of the variables in the case that the disturbances last for different lengths of time ranging from one quarter to twenty quarters.

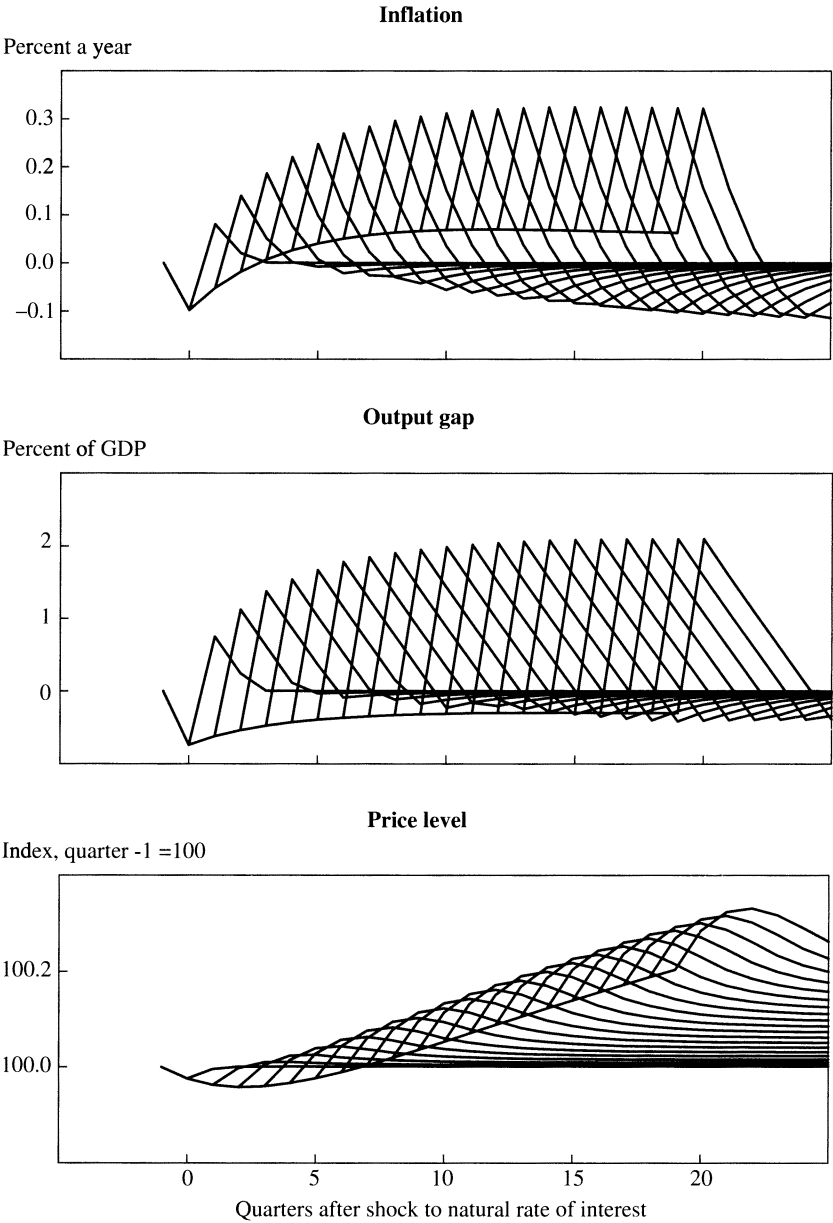
One observes that the optimal policy involves committing to the creation of an output boom once the natural rate again becomes positive, and hence to the creation of future inflation. Such a commitment stimulates aggregate demand and reduces deflationary pressure while the economy remains in the liquidity trap, through each of several channels.

As Krugman points out,⁴³ creating the expectation of future inflation can lower real interest rates, even when nominal interest rates cannot be reduced. In the context of Krugman's model, it might seem that this requires that inflation be promised quite quickly (that is, by the following "period"). Our fully intertemporal model shows how even the expectation of later inflation—which nominal interest rates are not expected to rise to offset—can stimulate current demand, because in our model current spending decisions depend on real interest rate expectations far in the future. For the same reason, the expectation that nominal interest rates will be kept low later, when the central bank might otherwise have raised them, will also stimulate spending while the zero bound still binds. And

42. Jung, Teranishi, and Watanabe (2001) discuss the solution of these equations only for the case in which the number of periods for which the natural rate of interest will be negative is known with certainty at the time that the disturbance occurs. Here we show how the system can be solved in the case of a stochastic process for the natural rate of a particular kind.

43. Krugman (1998).

Figure 3. Inflation, the Output Gap, and Prices under the Optimal Policy Commitment^a



Source: Authors' calculations.
a. See figure 2 for explanation of diagram.


finally, the expectation of higher future income should stimulate current spending, in accordance with the permanent income hypothesis. In addition, prices are less likely to fall, even given the current level of real activity, to the extent that future inflation is expected. This reduces the distortions created by deflation itself.

On the other hand, these gains from the change in expectations while the economy is in the liquidity trap can be achieved (given rational expectations on the part of the private sector) only if the central bank is expected to actually pursue the inflationary policy after the natural rate returns to its normal level. This will in turn create distortions at that later time, and this limits the extent to which this tool is used under an optimal policy. Hence some contraction of output and some deflation occur during the time that the natural rate is negative, even under the optimal policy commitment.

Also, and this is a key point, although the optimal policy involves commitment to a higher price level in the future, the price level will ultimately be stabilized. This is in sharp contrast to a constant positive-inflation target, which would imply an ever-increasing price level. Figure 4 shows the corresponding state-contingent nominal interest rate under the optimal commitment and contrasts it with the evolution of the nominal interest rate under a zero-inflation target. To increase inflation expectations in the trap, the central bank commits to keeping the nominal interest rate at zero after the natural rate of interest becomes positive again. In contrast, if the central bank targets zero inflation, it raises the nominal interest rate as soon as the natural rate of interest becomes positive again. The optimal commitment is an example of history-dependent policy, in which the central bank commits itself to raise interest rates slowly at the time the natural rate becomes positive in order to affect expectations when the zero bound is binding.

The nature of the additional history dependence of the optimal policy may perhaps be more easily seen if we consider the paths of inflation, output, and interest rates under a single possible realization of the random fundamentals. Figure 5 compares the equilibrium paths of all three variables, both under the zero-inflation target and under optimal policy, in the case where the natural rate of interest is negative for fifteen quarters ($t = 0$ through 14), but where it is not known until quarter 15 that the natural rate will return to its normal level in that quarter. Under the optimal policy the nominal interest rate is kept at zero for five more quarters ($t = 15$

Figure 4. Response of the Nominal Interest Rate under a Zero-Inflation Target and under the Optimal Policy Commitment



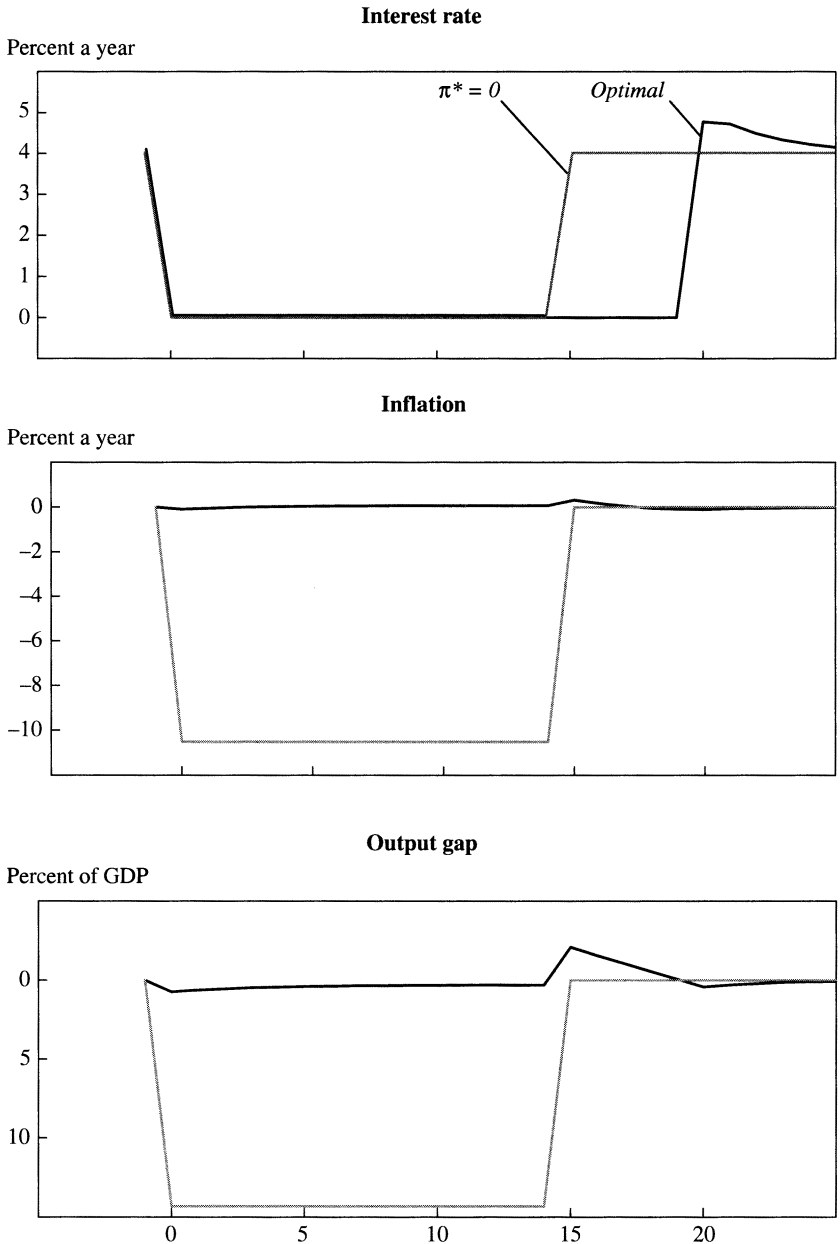
Source: Authors' calculations.

through 19), whereas it immediately returns to its long-run steady-state level in quarter 15 under the forward-looking policy. The consequence of the public anticipating policy of this kind is that both the contraction of real activity and the deflation that occur under the strict inflation target are largely avoided, as shown in the second and third panels of the figure.

Implementing Optimal Policy

We turn now to the question of how policy should be conducted in order to bring about the optimal equilibrium characterized in the previous section. The question of the implementation of optimal policy remains nontrivial, even after the optimal state-contingent paths of all variables have been identified, because in general the solution obtained for the optimal state-contingent path of the policy instrument (the short-term nominal interest rate) does *not* in itself represent a useful description of a policy

Figure 5. Response of the Nominal Interest Rate, Inflation, and the Output Gap to a Shock of Specific Duration^a



Source: Authors' calculations.
 a. Response to a fall in the natural rate of interest below zero for a period of fifteen quarters.

rule.⁴⁴ For example, in the context of the present model, a commitment to a state-contingent nominal interest rate path, even when fully credible, does not imply determinate rational expectations equilibrium paths for inflation and output; it is instead necessary for the central bank to be committed (and be understood to be committed) to a particular way of responding to deviations of inflation and the output gap from their desired paths. Another problem is that a complete description of the optimal state-contingent interest rate path is unlikely to be feasible. In the previous section we showed that one can characterize (at least numerically) the optimal state-contingent interest rate path in the case of one very particular kind of stochastic process for the natural rate of interest. But a solution of this kind, allowing for all possible states of belief about the probabilities of various future paths of the natural rate (and disturbances to the aggregate supply relation as well), would be difficult to write down, let alone explain to the public.

Here we show that optimal policy can nonetheless be implemented through commitment to a *policy rule* that specifies the central bank's short-run targets at each point in time as a (fairly simple) function of what has occurred before that date. How can this be done? One may be tempted to believe that our suggested policy is not entirely realistic or operational. Figures 3 and 4, for example, indicated that the optimal policy involves a complicated state-contingent plan for the nominal interest rate, which would be hard to communicate to the public. Furthermore, it may appear that it depends on knowledge of a special statistical process for the natural rate of interest, which is in practice hard to estimate. Our discussion of the fixed inflation target suggests that the effectiveness of increasing inflation expectations to close the output gap depends on the difference between the announced inflation target and the natural rate of interest. It may therefore seem crucial to estimate the natural rate of interest in order to implement the optimal policy. Below, however, we present the striking result that the optimal policy rule can be implemented without any estimate or knowledge of the statistical process for the natural rate of interest. This is an example of a robustly optimal direct policy rule of the kind discussed by Marc Giannoni and Woodford for the case of a general class of linear-quadratic policy problems.⁴⁵ An interesting feature of the present

44. For further discussion in a more general context, see Woodford (forthcoming, chapter 7).

45. Giannoni and Woodford (2003).

example is that we show how to construct a robustly optimal rule in the same spirit, in a case where not all of the relevant constraints are linear (owing to the fact that the zero bound binds at some times and not at others).

An Optimal Targeting Rule

To implement the rule proposed here, the central bank need only observe the price level and the output gap. The rule suggested replicates exactly the history dependence discussed in the last section. The rule is implemented as follows.

First, in each and every period there is a predetermined price-level target p_t^* . The central bank chooses the interest rate i_t to achieve the target relation

$$(32) \quad \tilde{p}_t = p_t^*,$$

if this is possible. If it is not possible, even by lowering the nominal interest rate to zero, then $i_t = 0$. Here \tilde{p}_t is an output gap-adjusted price index,⁴⁶ defined by

$$\tilde{p}_t \equiv p_t + \frac{\lambda}{\kappa} x_t.$$

The target for the next period is then determined as

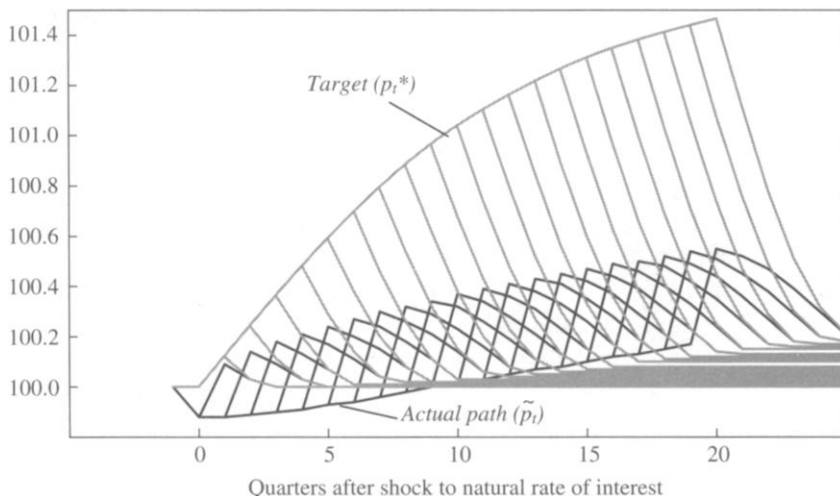
$$(33) \quad p_{t+1}^* = p_t^* + \beta^{-1}(1 + \kappa\sigma)\Delta_t - \beta^{-1}\Delta_{t-1},$$

where Δ_t is the target shortfall in period t :

$$(34) \quad \Delta_t \equiv p_t^* - \tilde{p}_t.$$

It can be verified that this rule does indeed achieve the optimal commitment solution. If the price-level target is not reached, because of the zero

46. On the desirability of a target for this index in the case that the zero bound does not bind, see Woodford (forthcoming, chapter 7). This would correspond to a nominal GDP target in the case that $\lambda = \kappa$ and that the natural rate of output follows a deterministic trend. However, the utility-based loss function derived in Woodford (forthcoming, chapter 6) involves $\lambda = \kappa/\theta$, where $\theta > 1$ is the elasticity of demand faced by the suppliers of differentiated goods, so that the optimal weight on output is considerably less than under a nominal GDP target. Furthermore, the welfare-relevant output gap is unlikely to correspond too closely to deviations of real GDP from a deterministic trend.

Figure 6. Responses of the Price-Level Target and the Gap-Adjusted Price Level to a Shock to the Natural Rate of InterestGap-adjusted index, quarter $-1 = 100$ 

Source: Authors' calculations.

bound, the central bank increases its target for the next period. This, in turn, increases inflation expectations further in the trap, which is exactly what is needed to reduce the real interest rate.

Figure 6 shows how the price-level target p_t^* would evolve over time, depending on the number of periods in which the natural rate of interest remains negative, in the same numerical experiment as in figure 3. (Here the dark lines show the evolution of the gap-adjusted price level \tilde{p}_t , and the shaded lines show the evolution of p_t^* .) One observes that the target price level is ratcheted steadily higher during the period in which the natural rate remains negative, as the actual price level continues to fall below the target by an increasing amount. Once the natural rate of interest becomes positive again, the degree to which the gap-adjusted price level undershoots the target begins to shrink, although the target often continues to be undershot (as the zero bound continues to bind) for several more quarters. (How long this is true depends on how high the target price level has risen relative to the actual index; it will be higher the longer the natural rate has been negative.) As the degree of undershooting begins to shrink, the price-level target begins to fall again, as a result

of the dynamics specified by equation 33. This hastens the date at which the target can actually be hit with a nonnegative interest rate. Once the target ceases to be undershot, it no longer changes, and the central bank targets and achieves a new constant value for the gap-adjusted price level \bar{p}_t , one slightly higher than the target in place before the disturbance occurred.

Note that this approach to implementing optimal policy answers the question of whether there is any point in announcing an inflation target (or price-level target) if one knows that it is extremely unlikely to be achieved in the short run, because the zero bound is likely to continue to bind. The answer here is yes. The central bank wishes to make the private sector aware of its commitment to the time-varying price-level target described by equations 32 through 34, because eventually it *will* be able to hit the target. The anticipation of that fact (that is, of the level that prices will eventually reach, as a result of the policies that the bank will follow after the natural rate of interest again becomes positive) while the natural rate is still negative is important in mitigating the distortions caused by the zero bound. The fact that the target is not hit immediately should not create doubts about whether central bank announcements regarding its target have any meaning, if it is explained that the bank is committed to hitting the target *if* this is possible at a nonnegative interest rate, so that, at each point in time, either the target will be attained or a zero interest rate policy will be followed. The existence of the target is relevant even when it is not being attained, because it allows the private sector to judge how close the central bank is to a situation in which it would feel justified in abandoning the zero interest rate policy; hence the current gap between the actual and the target price level should shape private sector expectations of the time when interest rates are likely to remain low.⁴⁷

Would the private sector have any reason to believe that the central bank was serious about the price-level target, if in each period all that is

47. An interesting feature of the optimal rule is that it involves a form of history dependence that cannot be summarized solely by the past history of short-term nominal interest rates; if the nominal interest rate has fallen to zero in the recent past, it matters to what extent the zero bound has prevented the central bank from pursuing as stimulative a policy as it otherwise would have. In this respect the optimal policy rule derived here is similar to the rules advocated by Reifschneider and Williams (1999), under which the interest rate operating at each point in time should depend on how low the central bank *would have* lowered interest rates in the past had the zero bound not prevented it.

observed is a zero nominal interest rate and yet another target shortfall? The best way of making a rule credible is for the central bank to conduct policy over time in a way that demonstrates its commitment. Ideally, the central bank's commitment to the price-level targeting framework would be demonstrated before the zero bound came to bind (at which time the central bank would have frequent opportunities to show that the target did determine its behavior). The rule proposed above is one that would be equally optimal both under normal circumstances and in the case of the relatively unusual kind of disturbance that causes the natural rate of interest to be substantially negative.

To understand how the rule works outside of the trap, it is useful to note that, when the nominal interest rate is positive, $\Delta_t = 0$ at all times. The central bank should therefore demonstrate a commitment to subsequently undo any over- or undershooting of the price-level target. In this case any deflation that occurs when the economy finds itself in a liquidity trap should create expectations of future inflation, as mandated by optimal policy. The additional term Δ_t implies that, when the zero bound is binding, the central bank should raise its long-run price-level target even further, thus increasing inflation expectations even more.

It may be wondered why we discuss our proposal in terms of a (gap-adjusted) price-level target rather than an inflation target. In fact, we could equivalently describe the policy in terms of a time-varying target for the gap-adjusted inflation rate $\tilde{\pi}_t \equiv \tilde{p}_t - \tilde{p}_{t-1}$. The reason we prefer to describe the rule as a price-level targeting rule is that the essence of the rule is easily described in those terms. As we show below, a *fixed* target for the gap-adjusted price level would actually represent quite a good approximation to optimal policy, whereas a fixed inflation target would not come close, because it would fail to allow for any of the history dependence of policy necessary to mitigate the distortions resulting from the zero bound.

A Simpler Proposal

One may argue that an unappealing aspect of the rule suggested above is that it involves the term Δ_t , which determines the change in the price-level target, and is nonzero only when the zero bound is binding. Suppose that the central bank's commitment to a policy rule can become credible over time only through repeated demonstrations of its commitment to act in accordance with it. In that case the part of the rule that involves the adjustment of the target in response to target shortfalls when the zero

bound binds might not come to be well understood by the private sector for a very long time, because the occasions when the zero bound binds will presumably be relatively infrequent.

Fortunately, most of the benefits that can be achieved in principle through a credible commitment to the optimal targeting rule can be achieved through commitment to a much simpler rule, which would not involve any special provisos that are invoked only in the event of a liquidity trap. Consider the following simpler rule:

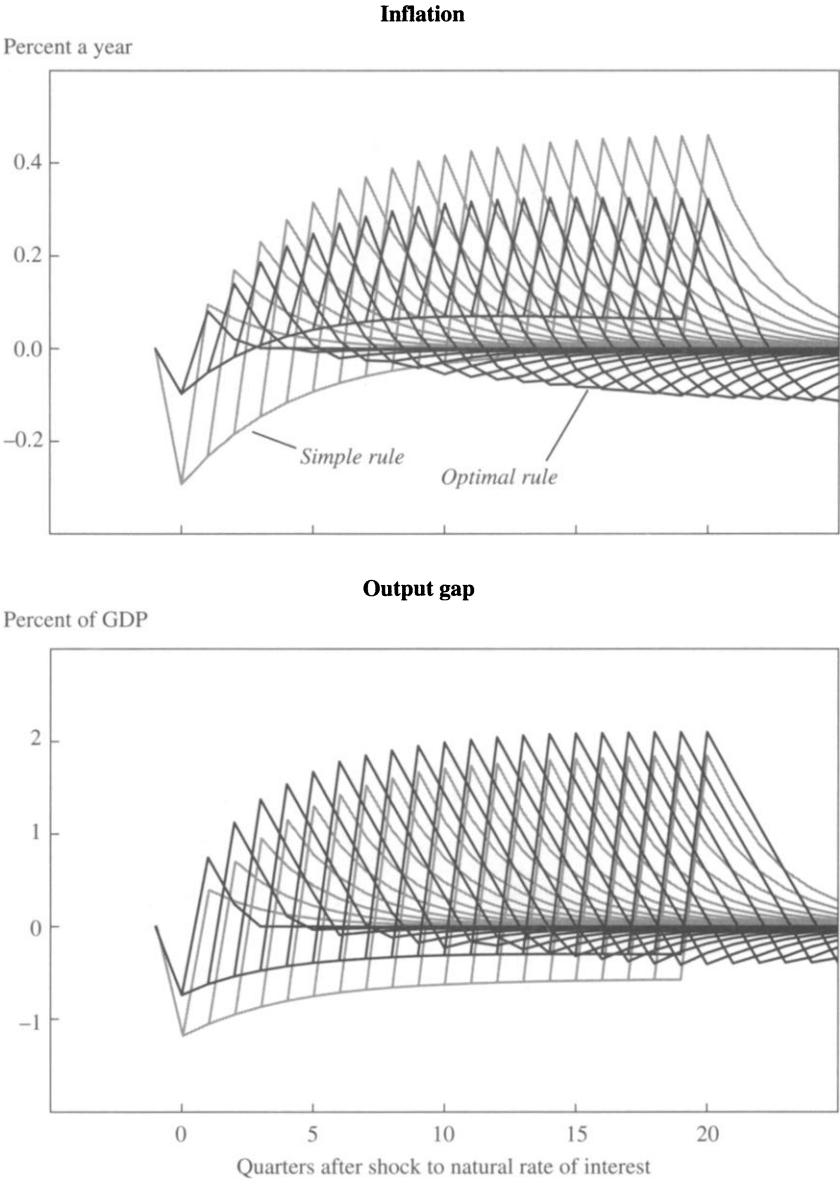
$$(35) \quad p_t + \frac{\lambda}{\kappa} x_t = p^*,$$

where now the target for the gap-adjusted price level is fixed at all times. The advantage of this rule, although it is not fully optimal when the zero bound is binding, is that it may be more easily communicated to the public. Note that the simple rule is fully optimal in the absence of the zero bound. In fact, even if the zero bound occasionally binds, this rule results in distortions only a bit more severe than those associated with the fully optimal policy.

Figures 7 and 8 compare the results for these two rules. The shaded lines show the equilibrium under the constant-price-level target rule in equation 35, whereas the dark lines show the fully optimal rule in equations 32 through 34. As the figures show, the constant-price-level targeting rule results in state-contingent responses of output and inflation that are very close to those under the optimal commitment, even if under this rule the price level falls further during the period when the zero bound binds, and only asymptotically rebounds to its level before the disturbance. The table below shows that the simple rule already achieves most of the welfare gain that the optimal policy achieves; the table reports the value of expected discounted losses, as a percentage of what could be achieved by a strict zero-inflation target (equation 28), conditional on the occurrence of the disturbance in period 0, under the various policies discussed above:

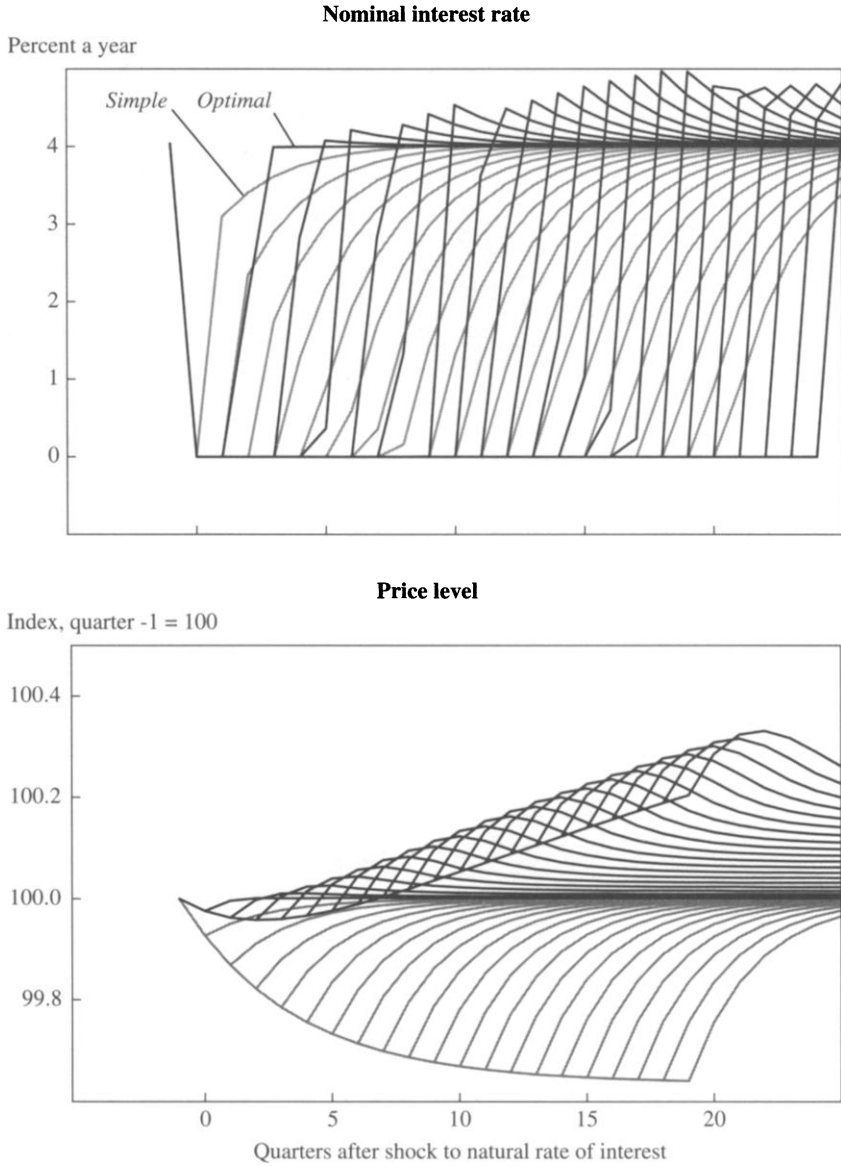
<i>Policy</i>	<i>Loss (percent)</i>
Strict inflation target, $\pi^* = 0$	100
Strict inflation target, $\pi^* = 1$	24.1
Strict inflation target, $\pi^* = 2$	32
Constant price-level target	0.0725
Optimal rule	0.036

Figure 7. Responses of Inflation and the Output Gap under the Optimal Targeting Rule and under the Simple Rule^a



Source: Authors' calculations.
a. The simple rule is described in equation 35.

Figure 8. Responses of the Nominal Interest Rate and Prices under the Optimal Targeting Rule and under the Simple Rule



Source: Authors' calculations.

Both of the latter two, history-dependent policies are vastly superior to any of the strict inflation targets. Although it is true that losses remain twice as large under the simple rule as under the optimal rule, they are nonetheless fairly small.

As with the fully optimal rule, no estimate of the natural rate of interest is needed to implement the constant-price-level targeting rule. It may seem puzzling at first that a constant-price-level targeting rule does well, because no account is taken of the size of the disturbance to the natural rate of interest. This comes about because a price-level target commits the government to undo any deflation with subsequent inflation; a larger disturbance, which creates a larger initial deflation, automatically creates greater inflation expectations in response. Thus an automatic stabilizer is built into the price-level target, which is lacking under a strict inflation targeting regime.⁴⁸

A proper strategy for the central bank to use in communicating its objectives and targets when outside the liquidity trap is of crucial importance for this policy rule to be successful. To see this, consider a rule that is equivalent to equation 35 when the zero bound is not binding. Taking the difference of equation 35, we obtain

$$(36) \quad \pi_t + \frac{\lambda}{\kappa}(x_t - x_{t-1}) = 0.$$

Although this rule results in an equilibrium identical to that under the constant-price-level targeting rule when the zero bound is not binding, the result is dramatically different when the zero bound is binding, because this rule implies that the inflation rate is proportional to the negative of the growth rate of output. Thus it mandates deflation when there is growth in the output gap. This implies that the central bank will deflate once the economy is out of a liquidity trap, because the economy will then be in a period of output growth. This is exactly the opposite of what is optimal, as we have observed above. Thus the outcome under this rule is even worse than under a strict zero-inflation target, even if this rule replicates the price-level targeting rule when out of the trap. What this underlines is that it is not enough to replicate the equilibrium behavior that corresponds to

48. Wolman (forthcoming) also stresses this advantage of rules that incorporate a price-level target over rules that only respond to the inflation rate, such as a conventional Taylor rule.

equation 35 in normal times to induce the correct set of expectations when the zero bound is binding. It is crucial to communicate to the public that the government is committed to a long-run price-level target. This commitment is exactly what creates the desired inflation expectations when the zero bound is binding.

Should the Central Bank Keep Some Powder Dry?

Thus far we have considered only alternative policies that might be followed after the natural rate of interest has unexpectedly fallen to a negative value, causing the zero bound to bind. A question of considerable current interest in countries like the United States, however, is how policy should be affected by the *anticipation* that the zero bound might well bind before long, even if it is not yet binding. Some have argued that, in such circumstances, the Federal Reserve should be cautious about lowering interest rates all the way to zero too soon, in order to save its ammunition for future emergencies. This suggests that the anticipation that the zero bound could bind in the near future should lead to tighter policy than would otherwise be justified given current conditions. Others argue, however, that policy should instead be more inflationary than one might otherwise prefer, to reduce the probability that a further negative shock will result in a binding zero bound.

Our characterization of the optimal targeting rule can shed light on this debate. Recall that the rule laid out in equations 32 through 34 describes optimal policy *regardless* of the assumed stochastic process for the natural rate of interest, and not only in the case of the particular two-state Markov process assumed in figure 3. In particular, the same rule is optimal in the case that information is received indicating the likelihood of the natural rate of interest becoming negative before this actually occurs. How should that news affect the conduct of policy? Under the optimal targeting rule, the optimal target for \bar{p}_t is unaffected by such expectations, as long as the zero bound is not yet binding, because only target shortfalls that have already occurred can justify a change in the target value p_t^* . Thus an increased assessment of the likelihood of a binding zero bound over the coming year or two would *not* be a reason for increasing the price-level target (or the implied target rate of inflation).⁴⁹

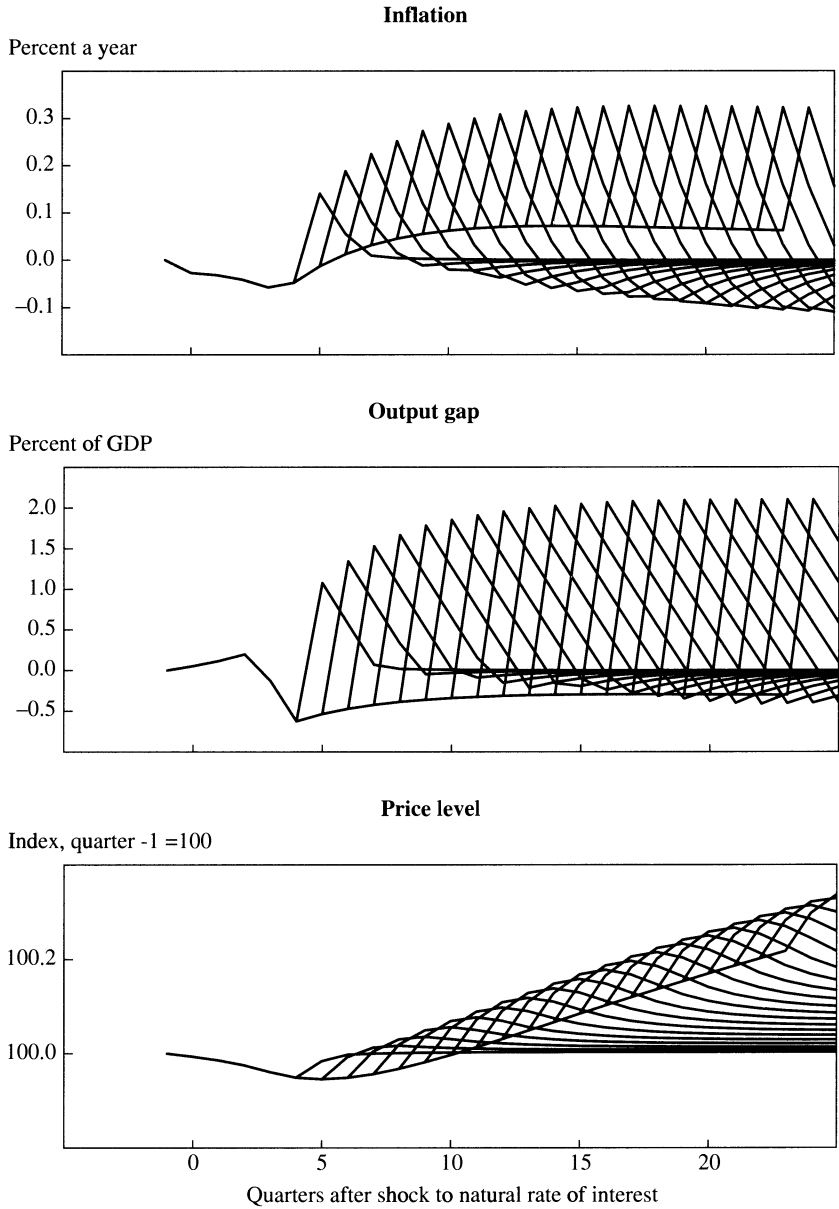
49. This conclusion, however, is likely to depend on a relatively special feature of our model, namely, the fact that our target variables (inflation and the output gap) are both

On the other hand, this news will affect the paths of inflation, output, and interest rates, even in the absence of any immediate change in the central bank's price-level target, owing to the effect on forward-looking private sector spending and pricing decisions. The anticipation of a coming state in which the natural rate of interest will be negative, and actual interest rates will not be able to fall as much, owing to the zero bound, will reduce both desired real expenditure (at unchanged short-term interest rates) and desired price increases, because of the anticipation of negative output gaps and price declines in the future. This change in the behavior of the private sector's outlook will require a change in the central bank's conduct of policy in order to hit its *unchanged* target for the modified price level, likely in the direction of a preemptive loosening. This is illustrated by the numerical experiment shown in figure 9. Here we suppose that in quarter 0 both the central bank and the private sector learn that the natural rate of interest will fall to -2 percent a year only in period 4. It is known that the natural rate will remain at its normal level of $+4$ percent a year until then; after the drop, it will return to the normal level with a probability of 0.1 each quarter, as in the case considered earlier. We now consider the character of optimal policy from period 0 onward, given this information. Figure 9 again shows the optimal state-contingent paths of inflation and output in the case that the disturbance to the natural rate, when it arrives, lasts for one quarter, two quarters, and so on.

We observe that, under the optimal policy commitment, prices begin to decline mildly as soon as the news of the coming disturbance is received. The central bank is nonetheless able to avoid undershooting its target for \bar{p}_t at first, by stimulating an increase in real activity sufficient to justify the mild deflation. (Given the private sector's shift to pessimism, this is the policy dictated by the targeting rule, given that even a mild immediate increase in real activity is insufficient to prevent a price decline, owing to the anticipated decline in real demand when the disturbance hits.) By quarter 3, however, this is no longer possible, and the

purely forward-looking variables: their equilibrium values at any point in time depend (in our simple model) only on the economy's exogenous state and the expected conduct of policy from the current period onward. There are a variety of reasons why a more realistic model may well imply that these variables are functions of lagged endogenous variables as well, and hence of past policy. In such a case, the optimal target criterion will be at least somewhat forward looking, as discussed in Giannoni and Woodford (2003).

Figure 9. Responses of Inflation, the Output Gap, and Prices to an Anticipated Shock under Optimal Policy^a



Source: Authors' calculations.

a. Response in a scenario where the shock to the natural rate of interest is anticipated four quarters in advance of its occurrence.

central bank undershoots its target for \bar{p}_t (as both prices and output decline), even though the nominal interest rate is at zero. Thus, optimal policy involves driving the nominal interest rate to zero *even before* the natural rate of interest has turned negative, when that development can already be anticipated for the near future. The fact that the zero bound binds even before the natural rate of interest becomes negative means that the price-level target is higher than it otherwise would have been when the disturbance to the natural rate arrives. As a result, deflation and the output gap during the period when the natural rate is negative are less severe than in the case where the disturbance is unanticipated. Optimal policy in this scenario is somewhat more inflationary after the disturbance occurs than in the case considered in figure 3, because in this case the optimal policy commitment takes into account the contractionary effects, in periods *before* the disturbance takes effect, of the anticipation that the disturbance will result in price-level and output declines. The fact that optimal policy after the disturbance occurs is different in this case, despite the fact that the disturbance has exactly the same effects as before from quarter 4 onward, is another illustration of the history dependence of optimal policy.

Preventing a Self-Fulfilling Deflationary Trap

Our analysis thus far has assumed that the real disturbance that results in a negative natural rate of interest does so only temporarily. We have therefore supposed that price-level stabilization will eventually be consistent with positive nominal interest rates and, accordingly, that a time will foreseeably be reached when the central bank can create inflation by keeping short-term nominal rates at a low but nonnegative level. But is it possible for the zero bound to bind *forever* in equilibrium, not because of a permanently negative natural rate, but simply because deflation continues to be (correctly) expected indefinitely? If so, the central bank's commitment to a nondecreasing price-level target might seem irrelevant; the price level would fall further and further short of the target, but because of the binding zero bound, the central bank could never do anything about it.

In the model presented in the first part of the paper, a self-fulfilling, permanent deflation is indeed consistent with both the Euler equation (equation 2) for aggregate expenditure, the money-demand relation

(expression 3), and the pricing relations (equations 7 through 9). Suppose that, from some date τ onward, all disturbances $\xi_t = 0$ with certainty, so that the natural rate of interest is expected to take the constant value $\bar{r} = \beta^{-1} - 1 > 0$, as in the scenarios considered earlier in the paper. Then the possible paths for inflation, output, and interest rates consistent with each of the relations just listed in all periods $t \geq \tau$ are given by

$$\begin{aligned} i_t &= 0 \\ P_t/P_{t-1} &= \beta < 1 \\ p_t^*/P_t &= \tilde{p}^* \equiv \left(\frac{1 - \alpha\beta^{\theta-1}}{1 - \alpha} \right)^{\frac{1}{1-\theta}} < 1 \\ Y_t &= \tilde{Y}, \end{aligned}$$

where $\tilde{Y} < \bar{Y}$ is implicitly defined by the relation

$$\Pi_1[\tilde{p}^*, \tilde{p}^*, 1; \tilde{Y}, \bar{m}(\tilde{Y}; 0), 0] = 0.$$

Note that this deflationary path is consistent with monetary policy as long as real balances satisfy $M_t/P_t \geq \bar{m}(\tilde{Y}; 0)$ each period; faster growth of the money supply does nothing to prevent consistency of this path with the requirement that money supply equal money demand in each period.

There remains, however, one further requirement for equilibrium in the earlier model, namely, the transversality condition (equation 6) or, equivalently, the requirement that households hit their intertemporal budget constraint. Whether the deflationary path is consistent with this condition as well depends, properly speaking, on the specification of fiscal policy: it is a matter of whether the government budget results in contraction of the nominal value of total government liabilities D_t at a sufficient rate asymptotically. Under some assumptions about the character of fiscal policy, such as the Ricardian fiscal policy rule assumed by Jess Benhabib, Stephanie Schmitt-Grohé, and Martin Uribe,⁵⁰ the nominal value of government liabilities will necessarily contract as the price level falls, so that equation 6 is also satisfied, and the processes described above will indeed represent a rational expectations equilibrium. In such a case, then, a commitment to the price-level targeting rule proposed in the previous section will be equally consistent with more than one equilibrium: if people

50. Benhabib, Schmitt-Grohé, and Uribe (2001).

expect the optimal price-level process characterized earlier, that will indeed be an equilibrium, but if they expect perpetual deflation, that will be an equilibrium as well.

However, this outcome can be excluded through a suitable commitment with regard to the asymptotic evolution of total government liabilities. Essentially, there needs to be a commitment to policies that ensure that the nominal value of government liabilities cannot contract at the rate required for satisfaction of the transversality condition, despite perpetual deflation. One example of a commitment that would suffice is a commitment to a balanced-budget policy of the kind analyzed by Schmitt-Grohé and Uribe.⁵¹ These authors show that self-fulfilling deflations are not possible when monetary policy is committed to a Taylor rule and the government to a balanced budget. The key to their result is that the fiscal rule includes a commitment that is as binding against running large surpluses as it is against running large deficits; then the nominal value of government liabilities cannot contract, even when the price level falls exponentially forever.

The credibility of this sort of fiscal commitment might be doubted, and so another way of maintaining a floor under the asymptotic nominal value of total government liabilities is through a commitment not to contract the monetary base, together with a commitment of the government to maintain a nonnegative asymptotic present value of the public debt. In particular, suppose that the central bank commits itself to follow a base-supply rule of the form

$$(37) \quad M_t = P_t^* \bar{m}(Y_t; \xi_t)$$

in each period when the zero bound binds (that is, when it is not possible to hit the price-level target with a positive nominal interest rate), where

$$P_t^* \equiv \exp \left\{ p_t^* - \frac{\lambda}{\kappa} x_t \right\}$$

is the current price-level target implied by the adjusted price-level target p_t^* . When the zero bound does not bind, the monetary base is whatever level is demanded at the nominal interest rate required to hit the price-level target. This is a rule in the same spirit as equation 11, specifying a

51. Schmitt-Grohé and Uribe (2000).

particular level of excess supply of base money when the zero bound binds, but letting the monetary base be endogenously determined by the central bank's other targets at other times. Equation 37 is a more complicated formula than is necessary to make our point, but it has the advantage of making the monetary base a continuous function of other aggregate state variables at the point where the zero bound just ceases to bind.

This particular form of commitment has the advantage that it may be considered less problematic for the central bank to commit itself to maintain a particular nominal value for its liabilities than for the public treasury to do so. It can also be justified as entirely consistent with the central bank's commitment to the price-level targeting rule; even when the target cannot be hit, the central bank supplies the quantity of money that would be demanded if the price level were at the target. Doing so—refusing to contract the monetary base even in a deflation—is a way of signaling to the public that the central bank is serious about its intention to see the price level restored to the target level.

If one then assumes a fiscal commitment that guarantees that

$$(38) \quad \lim_{T \rightarrow \infty} E_t Q_{t,T} B_T = 0,$$

that is, that the government will asymptotically be neither creditor nor debtor, the transversality condition reduces to

$$(39) \quad \lim_{T \rightarrow \infty} \beta^T E_t [u_c(Y_T, M_T/P_T; \xi_T) M_T/P_T] = 0.$$

In the case of the base-supply rule in equation 37, this condition is violated in the candidate equilibrium described above, since the price-level and output paths specified would imply that

$$\begin{aligned} \beta^T E_t [u_c(Y_T, M_T/P_T; \xi_T) M_T/P_T] &= \beta^t u_c[\tilde{Y}, \bar{m}(\tilde{Y}; 0); 0] \bar{m}(\tilde{Y}; 0) P_T^*/P_t \\ &\geq \beta^t u_c[\tilde{Y}, \bar{m}(\tilde{Y}; 0); 0] \bar{m}(\tilde{Y}; 0) P_t^*/P_t, \end{aligned}$$

where the last inequality makes use of the fact that, under the price-level targeting rule, $\{P_t^*\}$ is a nondecreasing series. Note that the final expression on the right-hand side is independent of T , for all dates $T \geq t$. Hence the series is bounded away from zero, and the condition in equation 39 is violated.

Thus a commitment of this kind can exclude the possibility of a self-fulfilling deflation of the sort described above as a rational expectations equilibrium. It follows that there is a possible role for quantitative easing—understood to mean the supply of base money beyond the minimum quantity required for consistency with the zero nominal interest rate—as an element of an optimal policy commitment. A commitment to supply base money in proportion to the *target* price level, and not the actual current price level, when the zero bound prevents the central bank from hitting its price-level target, can be desirable both as a way of ruling out self-fulfilling deflations and as a way of signaling the central bank's continuing commitment to the price-level target, even though it is temporarily unable to hit it.

Note that this result does not contradict our irrelevance proposition, because here we have made a different assumption about the nature of the fiscal commitment. Equation 38 implies that the evolution of total nominal government liabilities will not be independent of the central bank's target for the monetary base. As a consequence, the neutrality proposition no longer holds. The import of that proposition is that expansion of the monetary base when the economy is in a liquidity trap is necessarily pointless; rather, any effect of such action must depend either on changing expectations regarding future interest rate policy *or* on changing expectations regarding the future path of total nominal government liabilities. The present discussion has illustrated circumstances in which expansion of the monetary base—or, at any rate, a commitment not to contract it—could serve both these ends.

Nonetheless, the present discussion does *not* support the view that the central bank should be able to hit its price-level target at all times, simply by flooding the economy with as much base money as is required to prevent the price level from falling below the target at any time. Our earlier analysis still describes all possible paths for the price level consistent with rational expectations equilibrium, and we have seen that even if the central bank were able to choose the expectations that the private sector should have (as long as it were willing to act in accordance with them), the zero bound would prevent it from being able to fully stabilize inflation and the output gap. Furthermore, the degree of monetary base expansion during a liquidity trap called for by the rule in equation 37 is quite modest. The monetary base will be raised gradually, if the zero bound continues to bind, as the price-level target is ratcheted up to steadily higher

levels. But our calibrated example above indicates that this would typically involve only a very modest increase in the monetary base, even if the liquidity trap lasts for several years. There would be no obvious benefit to the kind of rapid expansion of the monetary base actually tried in Japan over the past two years. Such an expansion is evidently *not* justified by any intention regarding the future price level, and hence regarding the size of the monetary base once Japan exits from the trap. But an injection of base money that is expected to be removed once the zero bound ceases to bind should have little effect on spending or pricing behavior, as we showed in the first section of the paper.

Further Aspects of the Management of Expectations

In the first section we argued that neither expansion of the monetary base as such nor open-market purchases of particular types of assets should have any effect on either inflation or real activity, except to the extent that these actions might change expectations regarding future interest rate policy (or possibly expectations regarding the asymptotic behavior of total nominal government liabilities, and hence the question of whether the transversality condition should be satisfied). This then allowed us to characterize the optimal policy commitment without any reference to the use of such instruments of policy; a consideration of the different possible joint paths of interest rates, inflation, and output that would be consistent with rational expectations equilibrium sufficed to allow us to determine the best possible equilibrium that one could hope to arrange, and to characterize it in terms of the interest rate policy that one should wish the private sector to expect.

However, this does not mean that other aspects of policy—beyond a mere announcement of the rule to which the central bank wishes to be understood to be committed in setting future interest rate policy—cannot matter. They may matter insofar as certain kinds of present actions may help to signal the bank's intentions regarding future policy, or make it more credible that the central bank will indeed carry out those intentions. A full analysis of the ways in which policy actions may be justified as helping to steer expectations is beyond the scope of this paper, and in any event the question is one that has as much to do with psychology and effective communication as with economic analysis. Nonetheless, we

offer a few remarks here about the kinds of policies that might contribute to the creation of desirable expectations.

Demonstrating Resolve

One way in which current actions may help to create desirable expectations regarding future policy is by being seen to be consistent with the principles that the central bank wishes the private sector to understand will guide that policy. We have already mentioned one example of this: one way to convince the private sector that the central bank will follow the optimal price-level targeting rule after a period in which the zero bound has been hit is by following this rule *before* such a situation arises.

Our discussion in the previous section provides a further example. Adjustment of the supply of base money while the zero bound is binding, so as to keep the monetary base proportional to the *target* price level rather than the actual current price level, can be helpful, even though irrelevant to interest rate control, as a way of communicating to the private sector the central bank's belief about where the price level ought properly to be (and hence the quantity of base money that the economy ought to need). By putting the existence of the price-level target in greater relief, such an action can help create the expectations regarding future interest rate policy necessary to mitigate the distortions created by the binding zero bound.

As a further example, Clouse and others argue that open-market operations may be stimulative, even when the zero bound has been reached, because they "demonstrate resolve" to keep the nominal interest rate at zero for a longer time than would otherwise be expected.⁵² But an expansion of the monetary base when the zero bound is binding need not be interpreted in this way. Consider, for example, a central bank with a constant zero inflation target, as discussed previously. When the zero bound binds, such a bank is unable to hit its inflation target and should exhibit frustration with this state of affairs. If some within the bank believe it should always be possible to hit the target with sufficiently vigorous monetary expansion, one might well observe substantial growth in the monetary base at a time when the inflation target is being undershot. Nonetheless, this would not imply any commitment to looser policy

52. Clouse and others (2003)

subsequently; such a central bank would never intentionally allow the monetary base to be higher than required to hit the inflation target, if the target can be hit. The result should be the equilibrium path shown in figure 2, and there should be no effect from the quantitative easing that occurs while the zero bound binds. This shows that it matters what the private sector understands to be the principle that motivates quantitative easing; it is not simply a question of how large is the increase in the monetary base.

Similarly, open-market purchases of long-term treasury bonds when short-term rates are at zero, as advocated by Ben Bernanke and Stephen Cecchetti,⁵³ among others, may well have a stimulative effect even if portfolio balance effects are quantitatively unimportant. We argued previously that under such circumstances it is desirable for the central bank to commit itself to maintain low short-term rates even after the natural rate of interest rises again. The level of long-term rates can indicate the extent to which the markets actually believe such a commitment. If a central bank's judgment is that long-term rates remain higher than they should be under the optimal equilibrium, owing to private sector skepticism about whether the history-dependent interest rate policy will actually be followed, then a willingness to buy long-term bonds from the private sector at a price it regards as more appropriate is one way for the central bank to demonstrate publicly that it expects to carry out its commitment regarding future interest rate policy. Given that the private sector is likely to be uncertain about the nature of the central bank's commitment (in the case of imperfect credibility), and that it can reasonably assume that the central bank knows more about its own degree of resolve than others do, action by the central bank that is consistent with a belief on its own part that it will keep short-term rates low in the future is likely to shift private beliefs in the same direction. If so, open-market purchases of long-term bonds could lower long-term interest rates, stimulate the economy immediately, and bring the economy closer to the optimal rational expectations equilibrium. However, that effect follows not from the purchases themselves, but from how they are interpreted. For them to be interpreted as indicating a particular kind of commitment with regard to future policy, it is important that the central bank have itself formulated such an intention, and that it so inform the public, so that its open-market purchases will be seen in this light.

53. Bernanke (2002); Stephen G. Cecchetti, "Central Banks Have Plenty of Ammunition," *Financial Times*, March 17, 2003, p. 13.

Similar remarks apply to the proposals by Bennett McCallum and Lars Svensson that purchases of foreign exchange be used to stimulate the economy through devaluation of the currency.⁵⁴ Under the optimal policy commitment described in an earlier section, a decline in the natural rate of interest should be accompanied by depreciation of the currency, both because nominal interest rates fall (and are expected to remain low for some time) and because the expected long-run price level (and hence the expected long-run nominal exchange rate) should increase. It follows that the extent to which the currency depreciates can provide an indicator of the extent to which the markets believe that the central bank is committed to such an optimal policy; and if the depreciation is insufficient, purchases of foreign exchange by the central bank provide one way for it to demonstrate its own confidence in its policy intentions. Again, the effect in question is not a mechanical consequence of the bank's purchases, but instead depends on their interpretation.⁵⁵

Providing Incentives to Improve Credibility

A related but somewhat distinct argument is that actions at the zero bound may help render the central bank's commitment to an optimal policy more credible, by providing the bank with a motive to behave in the future in the way that it would currently wish that people would expect it to behave. Here we briefly discuss how policy actions that are possible while the economy remains in a liquidity trap may be helpful in this regard. Our point is not so much that the central bank is in need of a "commitment technology" because it will be unable to resist the temptation to break its commitments later in the absence of such a constraint. Rather, it is that the central bank may well need a way of making its com-

54. McCallum (2000); Svensson (2001). Svensson's proposal includes a target path for the price level, which the exchange rate policy is used to (eventually) achieve, and in this respect it is similar to the policy advocated here. However, Svensson's discussion of the usefulness of intervention in the market for foreign exchange does not emphasize the role of such interventions as a signal regarding future policy.

55. The numerical analysis by Coenen and Wieland (forthcoming) finds that an exchange rate policy can be quite effective in creating stimulus when the zero bound is binding. But what is actually shown is that a rational expectations equilibrium exists in which the currency depreciates and deflation is halted; these effects could be viewed as resulting from a credible commitment to a target path for the price level, similar to the one discussed in this paper, and not requiring any intervention in the foreign exchange market at all.

mitment *visible* to the private sector. Taking actions now that imply that the central bank will be disadvantaged later if it were to deviate from the policy to which it wishes to commit itself can serve this purpose.

To consider what kind of current actions provide useful incentives, it is helpful to analyze (Markov) equilibrium under the assumption that policy is conducted by a discretionary optimizer, unable to commit to specific future actions at all.⁵⁶ Consider first what a Markov equilibrium under discretionary optimization would be like in the case that the only policy instrument available is a short-term nominal interest rate, whose value is chosen each period, and the objective of the central bank is minimization of the loss function in equation 28. As shown above, if the central bank can credibly commit itself, this problem has a solution in which the zero bound does not result in too serious a distortion, although it does bind.

Under the assumption of central bank discretion, however, the outcome will be much inferior. Note that discretionary policy (under the assumption of Markov equilibrium in the dynamic policy game) is an example of a purely forward-looking policy. It then follows from our earlier argument that the equilibrium outcome will correspond to the kind of equilibrium discussed there in the case of a strict inflation target. More specifically, it is obvious that the equilibrium is the same as under a strict inflation target $\pi^* = 0$, since this is the inflation rate that the discretionary optimizer will choose once the natural rate of interest is again at its steady-state level. (From that point onward, a policy of zero inflation clearly minimizes the remaining terms in the discounted loss function.)

As shown in figure 2, if the private sector expects that the central bank will behave in this fashion, and the natural rate of interest remains negative for several quarters, the result will be a deep and prolonged contraction of economic activity and a sustained deflation. We have also seen that these effects could largely be avoided, even in the absence of other policy instruments, if the central bank were able to credibly commit itself to a history-dependent monetary policy in later periods. Thus, in the kind of situation considered here, there is a *deflationary bias* to discretionary monetary policy, although, at its root, the problem is again the one identified in the classic analysis of Finn Kydland and Edward Prescott.⁵⁷ We now consider instead the extent to which the outcome could be improved,

56. As in Eggertsson (2003a, 2003b).

57. Kydland and Prescott (1977).

even in a Markov equilibrium with discretionary optimization, by changing the nature of the policy game.

One example of a current policy action, available even when the zero bound binds, that can help shift expectations regarding future policy in a desirable way is for the government to cut taxes and issue additional nominal debt.⁵⁸ Alternatively, the tax cut can be financed by money creation, because when the zero bound binds, there is no difference between expanding the monetary base and issuing additional short-term Treasury debt at zero interest. This is essentially the kind of policy imagined when people speak of a “helicopter drop” of additional money into the economy, but here it is the fiscal consequences of such an action with which we are concerned.

Of course, if the objective of the central bank in setting monetary policy remains as assumed above, this will make no difference to the discretionary equilibrium: the optimal policy once the natural rate of interest becomes positive again will once more appear to be the immediate pursuit of a strict zero-inflation target. However, if the central bank also cares about reducing the social costs of increased taxation—whether because of collection costs or because of other distortions—as it ought if it really takes social welfare into account, the result is different. As Eggertsson has shown elsewhere,⁵⁹ the tax cut will then increase inflation expectations, even if the government cannot commit to future policy.

It may be asked why, if such an incentive exists, Japan continues to suffer deflation, given the growth in Japanese government debt during the 1990s. One possible answer is that although the gross national debt is 140 percent of GDP in Japan today, this does not reflect the true inflation incentives of the government. The ratio of gross national debt to GDP *overestimates* the government’s inflation incentives, because a substantial portion of government debt is held by other government institutions.⁶⁰ *Net* government debt is only 67 percent of GDP, and, as a result, inflation

58. As discussed in Eggertsson (2003a).

59. Eggertsson (2003a).

60. Government institutions such as the social security system, the postal savings system, postal life insurance, and the Trust Fund Bureau hold much of this nominal debt. If the part of the public debt held by these institutions is subtracted from total gross government debt, the remainder is only 67 percent of GDP. Most of the government institutions that hold the government’s nominal debt have real liabilities. For example, the social security system (which holds roughly 25 percent of the nominal debt held by the government) pays Japanese pensions and medical expenses. Those pensions are indexed to the consumer

incentives may not be much greater in Japan than in a number of other countries.

An even more likely reason for continued low inflation expectations in Japan, despite the size of the nominal public debt, is skepticism about whether the central bank can be expected to care about reducing the burden of the public debt when determining future monetary policy. The public may believe that the Bank of Japan lacks such an objective; the expressed resistance of the Bank of Japan to suggestions that it increase its purchases of Japanese government bonds, on the ground that this could encourage a lack of fiscal discipline,⁶¹ certainly suggests that reducing the burden of government finance is not among its highest priorities. As Eggertsson has stressed elsewhere,⁶² in order for fiscal policy to be effective as a means of increasing inflationary expectations, fiscal and monetary policy must be coordinated so as to maximize social welfare. The consequences of a narrow concern with inflation stabilization on the part of the central bank, together with an inability to credibly commit future monetary policy, can be dire, even from the point of view of the bank's own stabilization objectives.

Another instrument that may be used to change expectations regarding future monetary policy is open-market purchases of real assets or foreign exchange. Purchases of real assets (say, real estate) can be thought of as another way of increasing nominal government liabilities, which should affect inflation incentives in much the same way as deficit spending.⁶³ Purchases of real assets have the advantage of not worsening the overall fiscal position of the government—a current concern in Japan, given its existing gross debt—while still increasing the fiscal incentive for inflation. A further advantage of this approach is that it need not depend on a perceived central bank interest in reducing the burden of the public debt. Since the (nominal) capital gains from inflation accrue to the central bank itself under this policy, the central bank may be perceived to have an

price index. If inflation increases, the real value of social security assets will fall, but the real value of most its liabilities will remain unchanged. Thus the Ministry of Finance would eventually have to step in to make up for any loss in the value of social security assets if the government is to keep its pension program unchanged. Therefore the gains from reducing the real value of outstanding debt are partly offset by a decrease in the real value of the assets of government institutions such as social security.

61. Asahi Shimbun, "Bank of Japan Advised to 'Print Money' to Escape Deflation," Dow Jones News, February 10, 1999.

62. Eggertsson (2003a).

63. Eggertsson (2003a).

incentive to inflate simply on the ground that it cares about its own balance sheet, for example because doing so will help ensure its independence. (One can easily argue that, under a rational scheme of cooperation between the central bank and the government, the central bank should *not* choose policy on the basis of concerns about its balance sheet. But under such an ideal regime, it should choose monetary policy with a view to reducing the burden of the public debt, among other goals.)

The incentive effects of open-market operations in foreign exchange are even simpler.⁶⁴ Open-market purchases of foreign assets give the central bank an incentive to inflate in the future in order to realize capital gains at the expense of foreigners. These will be valuable if the central bank cares either about its own balance sheet or about reducing the burden of the public debt, as in the case of real asset purchases. However, capital gains on foreign exchange that result from depreciation of the domestic currency will be valuable even if the central bank cares neither about its balance sheet (for example, because it cooperates perfectly with the public treasury) nor about the burden of the debt (for example, because nondistorting sources of revenue are available to the public treasury). Capital gains at the expense of foreigners would allow an increase in domestic spending, by either the government or the private sector, and a central bank must value this if it has the national interest at heart.

Under rational expectations, of course, no such capital gains are realized on average. Still, the purchase of foreign assets can work as a commitment device, because if the central bank reneged on its inflation commitment, it would cause capital losses if the government holds foreign assets. Purchases of foreign assets are thus a way of committing the government to looser monetary policy in the future. This creates a reason for purchasing foreign exchange in order to cause a devaluation (which will also stimulate current demand), even without any assumption of a deviation from interest rate parity of the kind relied upon by authors such as McCallum in recommending devaluation for Japan.⁶⁵ Clouse and others argue that open-market purchases of long-term Treasuries by the Federal Reserve should also change expectations in a way that results in immediate stimulus.⁶⁶ The argument is that if the central bank were not to follow through on its commitment to keep short-term rates low, it would

64. As shown by Eggertsson (2003b).

65. McCallum (2000).

66. Clouse and others (2003).

suffer a capital loss on the long-term bonds that it purchased at a price that made sense only on the assumption that it *would* keep interest rates low. Similarly, Peter Tinsley has proposed a policy that would create this kind of incentive even more directly, namely, the sale by the Federal Reserve of options to obtain federal funds at a future date at a certain price.⁶⁷ The Federal Reserve would then stand to lose money if it did not keep the funds rate at the level to which it had previously committed itself.

Although these proposals should also help reinforce the credibility of the kind of policy commitment associated with the optimal equilibrium (as characterized in the first section of the paper), they have at least one important disadvantage relative to purchases of real assets or of foreign exchange. They only provide the central bank an incentive to maintain low nominal interest rates for a certain period; they do not provide it with an incentive to ensure that the price level eventually rises to a higher level. Thus they may do little to counter private sector expectations that nominal interest rates will remain low for years—but because goods prices are going to continue to fall, not because the central bank is committed to eventual reflation, as in the self-fulfilling deflation trap discussed above. This is arguably the kind of expectation that has now taken root in Japan, where even ten-year bond yields are already well below 1 percent, even though prices continue to fall and economic activity remains anemic. Creating the perception that the central bank has an incentive to continue trying to raise the price level, and that it will not be content as long as nominal interest rates remain low, may be a more successful way of generating the sort of expectations associated with the optimal equilibrium.

Conclusion

We have argued that the key to dealing with a situation in which monetary policy is constrained by the zero lower bound on short-term nominal interest rates is the skillful management of expectations regarding the future conduct of policy. By “management of expectations” we do not mean that the central bank should imagine that, if it uses sufficient guile, it can lead the private sector to believe whatever the central bank wishes it to believe, no matter what it actually does. Instead we have assumed that

67. Tinsley (1999).

there is no point in the central bank trying to get the private sector to expect something that the central bank does not itself intend to bring about. But we do contend that it is highly desirable for a central bank to be able to commit itself in advance to a course of action that is desirable because of the benefits that flow from its being anticipated, and then to work to make that commitment credible to the private sector.

In the context of a simple optimizing model of the monetary transmission mechanism, we have shown that a purely forward-looking approach to policy—which allows for no possibility of committing future policy to respond to past conditions—can lead to quite bad outcomes in the event of a temporary decline in the natural rate of interest, regardless of the kind of policy pursued at the time of the disturbance. We have also characterized optimal policy, under the assumption that credible commitment is possible, and shown that it involves a commitment to eventually bring the general price level back up to a level even higher than would have prevailed had the disturbance never occurred. Finally, we have described a type of history-dependent price-level targeting rule with the following properties: that a commitment to base interest rate policy on this rule determines the optimal equilibrium, and that the same form of targeting rule continues to describe optimal policy regardless of which of a large number of types of disturbances may affect the economy.

Given the role of private sector anticipation of history-dependent policy in realizing a desirable outcome, it is important for central banks to develop effective methods of signaling their policy commitments to the private sector. An essential precondition for this, certainly, is for the central bank itself to clearly understand the kind of history-dependent behavior to which it should be seen to be committed. It can then communicate its thinking on the matter and act consistently with the principles that it wishes the private sector to understand. Simply conducting policy in accordance with a rule may not suffice to bring about an optimal, or nearly optimal, equilibrium, but it is the place to start.

APPENDIX A

The Numerical Solution Method

HERE WE ILLUSTRATE a solution method for the optimal commitment solution discussed in the first section of the text. This same method can

also be applied, with appropriate modification of each of the steps, to finding the solution in the case where the central bank commits to a constant price-level target rule or to a constant inflation target. We assume that the natural rate of interest becomes unexpectedly negative in period 0 and then reverts back to normal with probability α_t in every period t . Our numerical work assumes that there is a final date S at which the natural rate becomes positive with a probability of 1, although this date may be arbitrarily far in the future.

The solution takes the form

$$\begin{aligned} i_t &= 0 & \forall t & \text{ if } 0 \leq t < \tau + k_\tau \\ i_t &> 0 & \forall t & \text{ if } t \geq \tau + k_\tau. \end{aligned}$$

It follows that

$$\begin{aligned} E_t x_{t+1} - x_t + \sigma(E_t \pi_{t+1} + r^n) &= 0 & \text{if } t < \tau + k_\tau \\ \varphi_{1t} &= 0 & \text{if } t \geq \tau + k_\tau. \end{aligned}$$

Here τ is the stochastic date at which the natural rate of interest returns to the steady state. We assume that τ can take any value between 1 and the terminal date S . The number $\tau + k_\tau$ is the period in which the zero bound ceases to bind contingent on the natural rate of interest becoming positive in period τ . Note that the value of k_τ can depend on the value of τ . We will first show the solution for the problem as if we knew the sequence $\{k_\tau\}_{\tau=1}^S$. We then describe a numerical method to find the sequence $\{k_\tau\}_{\tau=1}^S$.

The Solution for $t \geq \tau + k_\tau$

The system can be written in the following form:

$$(A1) \quad \begin{bmatrix} E_t \mathbf{Z}_{t+1} \\ \mathbf{P}_t \end{bmatrix} = \mathbf{M} \begin{bmatrix} \mathbf{Z}_t \\ \mathbf{P}_{t-1} \end{bmatrix},$$

where

$$\mathbf{Z}_t \equiv \begin{bmatrix} \pi_t \\ x_t \end{bmatrix} \text{ and } \mathbf{P}_t \equiv \begin{bmatrix} \varphi_{1t} \\ \varphi_{2t} \end{bmatrix}.$$

If there are two eigenvalues of the matrix \mathbf{M} outside the unit circle, this system has a unique bounded solution of the form

$$(A2) \quad \mathbf{P}_t = \mathbf{\Omega}^0 \mathbf{P}_{t-1}$$

$$(A3) \quad \mathbf{Z}_t = \mathbf{\Lambda}^0 \mathbf{P}_{t-1}.$$

The Solution for $\tau \leq t \leq \tau + k_\tau$

Again this is a perfect-foresight solution, but with the zero bound binding. The solution satisfies the following equations:

$$(A4) \quad \begin{aligned} \pi_t &= \kappa x_t + \beta \pi_{t+1} \\ x_t &= \sigma(r_t^n + \pi_{t+1}) + x_{t+1} \\ \pi_t + \varphi_{2t} - \varphi_{2t-1} - \beta^{-1} \sigma \varphi_{1t-1} &= 0 \\ \lambda_x x_t + \varphi_{1t} - \beta^{-1} \varphi_{1t-1} - \kappa \varphi_{2t} &= 0. \end{aligned}$$

The system can be written as

$$\begin{bmatrix} \mathbf{P}_t \\ \mathbf{Z}_t \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix} \begin{bmatrix} \mathbf{P}_{t-1} \\ \mathbf{Z}_{t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{M} \\ \mathbf{V} \end{bmatrix}.$$

This system has a solution of the form

$$(A5) \quad \mathbf{P}_{\tau+j} = \mathbf{\Omega}^{k_\tau-j} \mathbf{P}_{\tau+j-1} + \mathbf{\Phi}^{k_\tau-j}$$

$$(A6) \quad \mathbf{Z}_{\tau+j} = \mathbf{\Lambda}^{k_\tau-j} \mathbf{P}_{\tau+j-1} + \mathbf{\Theta}^{k_\tau-j},$$

where $j = 0, 1, 2, \dots, k$. Here $\mathbf{\Omega}^{k_\tau-j}$ is the coefficient in the solution when there are $k_\tau - j$ periods until the zero bound stops being binding (that is, when $k_\tau - j = 0$, the zero bound is no longer binding and the solution is equivalent to that in equations A2 and A3). We can find the numbers $\mathbf{\Lambda}^j$, $\mathbf{\Omega}^j$, $\mathbf{\Theta}^j$, $\mathbf{\Phi}^j$ for $j = 1, 2, 3, \dots, k$ by solving the equations below using the initial conditions $\mathbf{\Phi}^0 = \mathbf{\Theta}^0 = 0$ for $j = 0$ and the initial conditions for $\mathbf{\Lambda}^j$ and $\mathbf{\Omega}^j$ given in equations A2 and A3:

$$\begin{aligned} \mathbf{\Omega}^j &= [\mathbf{I} - \mathbf{B}\mathbf{\Lambda}^{j-1}]^{-1} \mathbf{A} \\ \mathbf{\Lambda}^j &= \mathbf{C} + \mathbf{D}\mathbf{\Lambda}^{j-1} \mathbf{\Omega}^j \\ \mathbf{\Phi}^j &= (\mathbf{I} - \mathbf{B}\mathbf{\Lambda}^{j-1})^{-1} [\mathbf{B}\mathbf{\Theta}^{j-1} + \mathbf{M}] \\ \mathbf{\Theta}^j &= \mathbf{D}\mathbf{\Lambda}^{j-1} \mathbf{\Phi}^j + \mathbf{D}\mathbf{\Theta}^{j-1} + \mathbf{V}. \end{aligned}$$

The Solution for $t < \tau$

The solution satisfies the following equations:

$$\begin{aligned}\tilde{\pi}_t &= \kappa \tilde{x}_t + \beta \left\{ (1 - \alpha_{t+1}) \tilde{\pi}_{t+1} + \alpha_{t+1} (\Lambda_{11}^{k_{t+1}} \tilde{\varphi}_{1t} + \Lambda_{12}^{k_{t+1}} \tilde{\varphi}_{2t} + \Theta_1^{k_{t+1}}) \right\} \\ \tilde{x}_t &= \sigma \left\{ r^{nL} + (1 - \alpha_{t+1}) \tilde{\pi}_{t+1} + \alpha_{t+1} (\Lambda_{11}^{k_{t+1}} \tilde{\varphi}_{1t} + \Lambda_{12}^{k_{t+1}} \tilde{\varphi}_{2t} + \Theta_1^{k_{t+1}}) \right\} + \\ &\quad \left\{ (1 - \alpha_{t+1}) \tilde{x}_{t+1} + \alpha_{t+1} (\Lambda_{21}^{k_{t+1}} \tilde{\varphi}_{1t} + \Lambda_{22}^{k_{t+1}} \tilde{\varphi}_{2t} + \Theta_2^{k_{t+1}}) \right\} \\ &\quad \tilde{\pi}_t + \tilde{\varphi}_{2t} - \tilde{\varphi}_{2t-1} - \beta^{-1} \sigma \tilde{\varphi}_{1t-1} = 0 \\ &\quad \lambda_x \tilde{x}_t + \tilde{\varphi}_{1t} - \beta^{-1} \tilde{\varphi}_{1t-1} - \kappa \tilde{\varphi}_{2t} = 0.\end{aligned}$$

Here a tilde on a variable denotes the value of that variable contingent on the natural rate of interest being negative. $\Lambda_{ij}^{k_{t+1}}$ is the ij th element of the matrix $\Lambda^{k_{t+1}}$. The value k_{t+1} depends on the number of additional periods that the zero bound is binding (recall that here we are solving for the equilibrium on the assumption that we know the value of the sequence $\{k_t\}_{t=1}^S$). We can write the system as

$$\begin{bmatrix} \tilde{\mathbf{P}}_t \\ \tilde{\mathbf{Z}}_t \end{bmatrix} = \begin{bmatrix} \mathbf{A}_t & \mathbf{B}_t \\ \mathbf{C}_t & \mathbf{D}_t \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{P}}_{t-1} \\ \tilde{\mathbf{Z}}_{t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{M}_t \\ \mathbf{V}_t \end{bmatrix}.$$

We can solve this backward from the date S on which the natural rate returns to normal with a probability of 1. We can then calculate the path for each variable to date 0. Note that

$$\mathbf{B}_{S-1} = \mathbf{D}_{S-1} = 0.$$

By recursive substitution we can find a solution of the form

$$(A7) \quad \tilde{\mathbf{P}}_t = \mathbf{\Omega}_t \tilde{\mathbf{P}}_{t-1} + \mathbf{\Phi}_t$$

$$(A8) \quad \tilde{\mathbf{Z}}_t = \mathbf{\Lambda}_t \tilde{\mathbf{P}}_{t-1} + \mathbf{\Theta}_t,$$

where the coefficients are time dependent. To find the numbers $\mathbf{\Lambda}_t$, $\mathbf{\Omega}_t$, $\mathbf{\Theta}_t$, and $\mathbf{\Phi}_t$, consider the solution of the system in period $S-1$ when $\mathbf{B}_{S-1} = \mathbf{D}_{S-1} = 0$. We have

$$\begin{aligned}
\mathbf{\Omega}_{S-1} &= \mathbf{A}_{S-1} \\
\mathbf{\Phi}_{S-1} &= \mathbf{M}_{S-1} \\
\mathbf{\Lambda}_{S-1} &= \mathbf{C}_{S-1} \\
\mathbf{\Theta}_{S-1} &= \mathbf{V}_{S-1}.
\end{aligned}$$

We can find the numbers $\mathbf{\Lambda}_t$, $\mathbf{\Omega}_t$, $\mathbf{\Theta}_t$, and $\mathbf{\Phi}_t$ for periods 0 to $S-2$ by solving the system below (using the initial conditions shown above for $S-1$):

$$\begin{aligned}
\mathbf{\Omega}_t &= [\mathbf{I} - \mathbf{B}_t \mathbf{\Lambda}_{t+1}]^{-1} \mathbf{A}_t \\
\mathbf{\Lambda}_t &= \mathbf{C}_t + \mathbf{D}_t \mathbf{\Lambda}_{t+1} \mathbf{\Omega}_t \\
\mathbf{\Phi}_t &= (\mathbf{I} - \mathbf{B}_t \mathbf{\Lambda}_{t+1})^{-1} [\mathbf{B}_t \mathbf{\Theta}_{t+1} + \mathbf{M}_t] \\
\mathbf{\Theta}_t &= \mathbf{D}_t \mathbf{\Lambda}_{t+1} \mathbf{\Phi}_t + \mathbf{D}_t \mathbf{\Theta}_{t+1} + \mathbf{V}_t.
\end{aligned}$$

Using the initial condition $\tilde{\mathbf{P}}_{-1} = 0$, we can solve for each of the endogenous variables under the contingency that the liquidity trap lasts until period S , using equations A7 and A8. We then use the solution from equations A2 to A6 to solve for each of the variables when the natural rate reverts back to the steady state.

Solving for $\{k_\tau\}_{\tau=1}^\infty$

A simple way to find the value for $\{k_\tau\}_{\tau=1}^S$ is to first assume that k_τ is the same for all τ and find the lowest $k_\tau = \bar{k}$ so that the zero bound is never violated. Using this initial guess for $\{k_\tau\}_{\tau=1}^S$, one then finds the lowest value of k_S so that the zero bound is never violated. Using this value for k_S , and \bar{k} for all other k_τ , one then finds the lowest value of k_{S-1} so that the zero bound is never violated, and so on until the lowest possible value for k_1 is found. The value thus found for the sequence $\{k_\tau\}_{\tau=1}^S$ can be used as a new initial guess for $\{k_\tau\}_{\tau=1}^{S-1}$, and the procedure just described can be repeated until the solution converges. For this paper we wrote a routine in MATLAB that applies this method. The solution converged, and we verified that the result satisfied all the necessary conditions.