Monetary Policy in a
Low Inflation Economy with Learning

John C. Williams
Federal Reserve Bank of San Francisco

September 2006

Monetary Policy in a
Low Inflation Economy with Learning *

John C. Williams

Federal Reserve Bank of San Francisco

September 2006

Abstract: In theory, monetary policies that target the price level, as opposed to the inflation rate, should be highly effective at stabilizing the economy and avoiding deflation in the presence of the zero lower bound on nominal interest rates. With such a policy, if the short-term interest rate is constrained at zero and the inflation rate declines below its trend, the public expects that policy will eventually engineer a period of above-trend inflation that restores the price level to its target level. Expectations of future monetary accommodation stimulate output and inflation today, mitigating the effects of the zero bound. The effectiveness of such a policy strategy depends crucially on the alignment of the public's and the central bank's expectations of future policy actions. In this paper, we consider an environment where private agents have imperfect knowledge of the economy and therefore continuously reestimate the forecasting model that they use to form expectations. We find that imperfect knowledge on the part of the public, especially regarding monetary policy, can undermine the effectiveness of price-level-targeting strategies that would work well if the public had complete knowledge. For low inflation targets, the zero lower bound can cause a dramatic deterioration in macroeconomic performance with severe recessions occurring with alarming frequency. However, effective communication of the policy strategy that reduces the public's confusion about the future course of monetary policy significantly reduces the stabilization costs associated with the zero bound. Finally, the combination of learning and the zero bound implies the need for a stronger policy response to movements in the price level than would otherwise be optimal and such a rule is effective at stabilizing both inflation and output in the presence of learning and the zero bound even with a low inflation target.

---

*Prepared for the Bank of Korea International Conference 2006, June 15–17, 2006. The views expressed herein are those of the author and do not necessarily reflect those of the management of the Federal Reserve Bank of San Francisco or any other person in the Federal Reserve System.

1   Introduction

The successful reduction of inflation to low levels in many countries raises the question of how to best design monetary and fiscal policies to reduce the risk of deflation and to facilitate a rapid return to price stability if deflation occurs. The recent experience of deflation and near-zero short-term interest rates in Japan and the brief flirtation with inflation and interest rates around one percent in the United States has led to a renewal of research into the design of monetary policy that takes account of the zero lower bound on nominal interest rates. A recurring finding in this literature is that monetary policy strategies that explicitly or implicitly target the price level, as opposed to the inflation rate, should be highly effective at both mitigating the effects of the zero lower bound and at minimizing the duration and depth of deflationary episodes (see Reifschneider and Williams 2000, Svensson 2001, Eggertsson and Woodford 2003). In these models, the promise of future, indeed at times distant future, above-trend inflation aimed at restoring the price level to its target level provides a powerful pull on an economy experiencing deflation and constrained by the zero lower bound. Indeed, according to this research, a central bank can successfully target a constant price level with virtually no cost in terms of macroeconomic stabilization resulting from the zero bound.

These results rely on two crucial assumptions. The first assumption is that the central bank can credibly commit to follow such a price-level targeting policy. Eggertsson (2006) challenges the assumption the central bank can necessarily commit to future high inflation following a period of deflation associated with monetary policy being constrained by the zero lower bound. If the central bank lacks the ability to commit to future high inflation, the upward pull on inflation and output from the future is diminished as the public rightly anticipates that the central bank will choose only to bring inflation back to its target level and let the fall in the price level be bygones.

The second critical assumption is that private agents properly anticipate the implications of the monetary policy strategy on the future path of policy and the economy. Reifschneider and Roberts (2005) show that price-level-targeting monetary policy rules may lose some of their effectiveness in the presence of the zero bound when expectations are allowed to deviate from rational expectations. In this paper, we examine the role of

expectations formation on the effectiveness of monetary policy strategies in the presence of the zero bound. We follow the recent literature on learning and consider environments where agents have imperfect knowledge of the structure of the economy and monetary policy strategy and regularly update their beliefs about both based on past experience. We explore the conditions under which imperfect knowledge weakens or even disables the expectations channel that is essential to many proposed monetary policy strategies in the face of the zero lower bound. In addition, we examine the implications for monetary policy design to make it more robust to the presence of both imperfect knowledge and the zero bound.

Our framework also enables us to analyze the effects of communication strategies that help the public predict the future course of monetary policy. A number of papers that propose specific policy actions such as pegging the exchange rate, influencing larger-term bond rates, and increasing the monetary base, when the interest rate is already zero highlight the communication aspect of the such policy actions (see Meltzer (2001), Svensson (2001), McCallum (2002), Okina and Shiratsuka (2004), and McGough, Rudebusch, and Williams (2005)). But these papers typically assume that the public is fully informed about the determination of monetary policy and the behavior of the economy, so the benefits of central bank communication cannot be analyzed directly. Orphanides and Williams (2005a) show that improving the public's understanding of the policy rule reduces private expectations errors and, in so doing, improves macroeconomic performance. But, this analysis ignores the zero bound. As we show in this paper, the presence of the zero bound further complicates the public's learning problem and amplifies the costs associated with expectation errors. Therefore, the benefits of clearly communicating policy are heightened.

Our analysis reveals three main findings. First, imperfect knowledge on the part of the public, especially regarding monetary policy, can undermine the effectiveness of monetary policy strategies that would be highly effective if the public had complete knowledge. For low inflation targets, the zero lower bound can engender a dramatic deterioration in macroeconomic performance, with severe recessions occurring relatively frequently. Second, effective communication of the policy strategy that reduces the public's confusion about the future course of monetary policy significantly reduces the stabilization costs

associated with the zero bound. Third, the combination of learning and the zero bound implies the need for a stronger policy response to movements in the price level than would otherwise be optimal. Indeed, such a rule is better at stabilizing both inflation and output in the presence of learning and the zero bound, and is highly effective even if with an inflation target of only one percent.

The remainder of the paper is organized as follows. Section II describes the model and monetary policy. Section III describes the formation of expectations. Section IV outlines the model simulation methodology and describes the calibration of model parameters. Section reports the results of the monetary policy analysis. Section VI concludes.

2 The Model

This section describes the empirical macroeconomic model used for our analysis. The model is a so-called hybrid New Keynesian model (see Woodford (2003) for further details and references regarding similar models). The model contains key features of output and inflation dynamics of many recent micro-founded models used for monetary policy evaluation (Cf. Levin et al (2006)). Each period in the model corresponds to one quarter of a year.

2.1 Output and Inflation

The output gap (the deviation of output from its natural rate), denoted by y_t, is given by:

$$y_t - \eta y_{t-1} = -\phi(i_t - F_{t-1}\pi_{t+1} - r_t^*) - \phi F_{t-1}\{\sum_{j=1}^{\infty}(i_{t+j} - \pi_{t+j+1} - r_{t+j}^*)\}, \quad r_t^* \sim N(\bar{r}^n, \sigma_r^2) \quad (1)$$

where F_{t-1} refers to the agents' forecast based on information available at the end of period $t-1$, i_t is the short-term nominal interest rate, π_t is the inflation rate, and r_t^* is the stochastic natural rate of interest (around a fixed long-run value of $\bar{r}^*$), assumed to follow an i.i.d. Gaussian distribution with variance σ_r^2. The lag of the output gap in the equation captures the effects of habit in preferences. Note that because we consider deviations from rational expectations where agents have imperfect knowledge of the true structure

3

of the economy, we replace the standard mathematical expectations with private agents' forecasts. In addition, as emphasized by Preston (2005), under imperfect knowledge one cannot make the substitutions that are commonly used in the literature to rewrite this equation in terms of finite leads of the output gap. Instead, we assume that decisions are based explicitly on expectations of the fundamental determinants of the output decision.

The equation for inflation is based on a Calvo pricing model with partial indexation of prices to lagged inflation:

$$\pi_t - \rho\pi_{t-1} = \kappa(y_t - \theta y_{t-1}) + u_t + F_{t-1}\{\sum_{j=1}^{\infty}\beta^j[\kappa(y_{t+j} - \theta y_{t+j-1}) + u_{t+j}]\}, \quad u_t \sim N(0, \sigma_u^2) \quad (2)$$

where u_t is a markup shock, assumed to follow an i.i.d. Gaussian distribution with variance σ_u^2. As in the case of the output equation, we assume that pricing decisions are based on expectations of their fundamentals determinants.

2.2   Monetary Policy

We assume that the central bank's objective is to minimize the weighted sum of the unconditional variances of the inflation gap (the difference between the inflation rate and its target), the output gap, and the short-term nominal interest rate. The central bank loss, $\mathcal{L}$, is given by:

$$\mathcal{L} = \text{VAR}(\pi_t - \pi^*) + \lambda\text{VAR}(y_t) + \nu\text{VAR}(i_t), \quad (3)$$

where $\text{VAR}(x)$ denotes the unconditional variance of a variable x, λ is the relative weight on output gap variability, and ν is the relative weight on nominal interest rate variability. In the following, we assume that $\lambda = 0.5$ and $\nu = 0.1$. This choice of ν assures that the degree of interest rate variability is similar to the historical experience in the United States over the past 20 years.

Based on the findings of the theoretical literature, we assume that monetary policy follows a reaction function that reacts to the gap between the price level and a deterministic trend. We start with the "difference rule" specification of monetary policy similar to that advocated by Orphanides and Williams (2006), given by:

$$i_t = \max\{i_{t-1} + \gamma_\pi(\pi_{t-1} - \pi^*) + \gamma_{\Delta y}\Delta y_{t-1}, 0\}, \quad (4)$$

4

where Δ denotes the first difference operator, and the "max" function reflects the presence of the zero lower bound on nominal interest rates.[1] We assume that the central bank responds to data with a one-quarter lag. Note that by integrating this equation (and assuming the rule is followed without deviation), it is identical to a policy rule where the level of the interest is determined by the price level gap (the difference between the price level and a deterministic trend), the level of the output gap, and a constant. Orphanides and Williams (2006) show that rules of this form are robust to uncertainty regarding the model of agents' expectations, be it rational expectations or learning. However, that analysis abstract from the zero lower bound on interest rates.

As noted by Reifschneider and Williams (2000), the zero lower bound poses a problem for difference rules in that past deviations owing to the zero bound are carried forward into an excessively high current interest rate mechanically through the effects of the lagged interest rate. An alternative implementation that is equivalent in the absence of the zero bound but avoids this problem with the zero bound is for monetary policy to follow the integrated version of the rule:

$$i_t = \max\{\gamma_\pi (p_{t-1} - p^*_{t-1}) + \gamma_{\Delta y} y_{t-1} + \bar{i}^*, 0\}, \tag{5}$$

where p_t is the log of the price level, p^*_t is the target price level that follows $p^*_t = p^*_{t-1} + \pi^*$, and the final term:

$$\bar{i}^* = \pi^* + \bar{r}^*$$

is the long-run neutral nominal interest rate.

2.3   Fiscal Policy

Eggertsson and Woodford (2005) show that fiscal policy can be used to complement monetary policy when the zero bound is a constraint on policy. In order to explore the ability of monetary policy alone to cope with the zero bound, we do not consider the use of government spending or distortionary taxes as a complement to monetary policy. Instead, we assume that in general the fiscal authority is entirely passive. Given this

---

[1]We could impose a slightly positive lower bound of i^{LB}. In terms of our analysis, this corresponds exactly to an inflation target for $\pi^* - i^{LB}$. The experience of Japan over the past decade suggests that the lower bound is very near zero.

assumption, in periods of severe deflation, the economy can get stuck in a deflationary trap. In such cases, we assume that fiscal policy will take steps that limit the duration of such an episode to five years, at which time the economy is brought back to steady state. From then on, fiscal policy reverts to a passive role. As discussed below, this "backstop" fiscal intervention occurs very rarely when monetary policy is doing a good job of stabilizing the economy on average, and therefore is best viewed as a means of keeping the computation of model moments from being dominated by extreme outliers. Regular occurrences, on the other hand, indicate that the stipulated monetary policy rule does not stabilize the system effectively.

3   Expectations Formation

We assume that agents form expectations using a reduced-form forecasting model of the economy as opposed to using the full structural model as would be the case under model-consistent (i.e., rational) expectations. We specify the forecasting model such that it exactly corresponds to the reduced-form of the the structural model under the joint assumptions of rational expectations and the absence of the zero lower bound on nominal interest rates. We assume that agents continuously re-estimate the forecasting model based on past observations using a constant-gain least squares algorithm (see Sargent (1993) and Evans and Honkapohja (2001) for a fuller discussion of constant gain learning). Given the structure of the model and the stipulated form of the monetary policy rule, under rational expectations and ignoring the zero bound, five variables – the inflation rate, the output gap and its first lag, the interest rate, and an intercept – fully describe the state of the economy at the end of a period. We assume that agents compute forecasts using a linear forecasting model with these five explanatory variables. At the end of each period, agents re-estimate this forecasting model using the currently available data and use the resulting model to construct forecasts. We also consider alternative assumptions regarding how agents forecast interest rates within the context of their forecasting model.

Let Y_t denote the 1×3 vector consisting of the period t values of the variables to be forecast: $Y_t = (\pi_t, y_t, i_t)$. Let X_t denote the the 1×5 vector consisting of the explanatory variables: $X_t = (\pi_{t-1}, y_{t-1}, i_{t-1}, y_{t-2}, 1)$. Estimation is be described as follows: Let c_t be

the $j \times 5$ vector of coefficients of the forecasting model. Then, using data through period t, the parameters for the constant-gain least squares forecasting model can be written as:

$$c_t = c_{t-1} + \mu R_t^{-1} X_t (X_t - X_t' c_{t-1}), \tag{6}$$

$$R_t = R_{t-1} + \mu (X_t X_t' - R_{t-1}), \tag{7}$$

where $\mu > 0$ is the gain.

In the case of forecasts of the interest rate, we deviate from this simple forecasting method. First, we impose the zero lower bound on forecasts of all future nominal interest rates. Specifically, in period t we compute the forecast for $t+1$ variables. If the forecasted value of the interest rate in period $t + 1$ is negative, we set that value to zero. We then compute the $t + 2$ forecast of all variables and follow the same procedure, and so on. In this way, we enforce the zero bound both on the actual value of the interest rate and on expectations of future interest rates.[2] In principle, agents need forecasts for infinitely many periods in the future. However, to keep the problem tractable, we approximate this infinite sum with a truncated sum of k periods, replacing the terms for periods $k + 1$ and beyond with the period $k + 1$ forecast of the appropriate variables, as follows:

$$y_t - \eta y_{t-1} = -\phi(i_t - F_{t-1}\pi_{t+1} - r_t^*) - \phi F_{t-1}\{\sum_{j=1}^{k}(i_{t+j} - \pi_{t+j+1} - r_{t+j}^*)\} + F_{t-1}\{y_{t+k+1} - \eta y_{t+k}\}, \tag{8}$$

$$\pi_t - \rho\pi_{t-1} = \kappa(y_t - \theta y_{t-1}) + u_t + F_{t-1}\{\sum_{j=1}^{k}\beta^j[\kappa(y_{t+j} - \theta y_{t+j-1}) + u_{t+j}]\} + \beta^{k+1}F_{t-1}\{\pi_{t+k+1} - \rho\pi_{t+k}\}. \tag{9}$$

Given the dynamics of the system, we find that $k = 20$ is sufficient to get accurate solutions and use that value for all results reported here. The results with $k = 40$ are generally very close to those for $k = 20$.

We consider two alternative ways for agents to form forecasts of the interest rate. The first approach is simply to use the model as described above. Absent the zero bound, the interest rate equation in the forecast model is identical to that describing policy, so the fit of the forecasting equation is perfect. The presence of the zero bound, however,

---

[2]Note that this method implicitly imposes certainty equivalence by ignoring the distribution of interest rate forecasts and its effect on the expected interest rate from the zero bound. Incorporating this channel requires the use of computationally intensive nonlinear methods and is beyond the scope of this paper.

introduces positive deviations from the simple linear policy rule. The basic forecasting model implicitly treats these deviations as part of the interest rate process, and these deviations affect the forecast of future interest rates directly, through the lagged interest rate in the model, and indirectly, through the effect on estimated parameters of the interest rate equation in the forecasting model.

The second approach to modeling agents' interest rate forecasts is for agents to use the actual policy rule in forming forecasts, conditional on the forecasts of inflation and the output gap. This is accomplished by substituting the policy rule for the interest rate equation in the forecasting model. In particular, if the nominal interest rate depends on the lagged price level and output gap, then agents will not be fooled by deviations from the rules and will forecast monetary policy to eventually restore the price level to its target.

4 Model Solution and Calibration

This section describes the method used to compute model statistics and the calibration of the model parameters. Owing to the presence of the zero lower bound and learning, we cannot apply standard methods of solving and computing unconditional moments of linear rational expectations models. Instead, we use simulated moments as approximations of the unconditional moments.

4.1 Model Simulation Methodology

For a given parameterization of the model, the simulated model moments are computed based on a single stochastic simulation consisting of 101,000 periods, where the first 1000 observations are dropped in order to remove the effects of initial conditions.[3] The initial conditions for all model variables and the forecasting model matrices c and R are given by the corresponding steady-state values of the rational expectations equilibrium with

---

[3]Based on simulations under rational expectations in which we can compute the moments directly, this sample size is sufficient to yield very accurate estimates of the unconditional variances. In addition, testing indicates that 1000 periods is sufficient to remove the effects of initial conditions on simulated second moments.

no zero bound. The shocks are generated using MATLAB's Gaussian pseudo-random number generator "randn."

The presence of either the zero bound or learning introduces a nonlinearity into the model that can generate explosive behavior in a simulation of 100,000 periods, even for policy rules that are stable under rational expectations. One potential source of instability under learning is the possibility that the forecasting model itself may become unstable. To mitigate the possibility that instability in the forecasting model generates explosive behavior in the model economy, we do the following. During each period of the simulation, we compute the root of maximum modulus of the forecasting VAR excluding the constants. If the modulus of this root falls below the critical value of 1.1, the coefficients of the forecast model is updated as described above; if not, we assume that the forecast model is not updated and the matrices c_t and R_t are held at their respective previous period values. This cutoff is invoked only extremely rarely in the simulations.

However, stability of the forecasting model is not sufficient to assure stability of the full model in all situations. For this reason, we impose a second condition that restrains explosive behavior. In particular, if the absolute values of the inflation gap, output gap, or interest rate gap (the nominal interest rate less the long-run neutral rate), exceed very large values, then the offending variables are simply set to the relevant boundary value. We use a bound of 20 percentage points for the interest rate and the output gap and 10 percentage points for the inflation rate. The upper bounds are included for symmetry. Of course, this lower bound on the nominal interest rate is irrelevant given the zero lower bound that is part of the determination of the interest rate. These bounds are set wide enough that they only bind very rarely or never when policy is effective at stabilizing the economy, but bind more frequently when policy is ineffective, as discussed below.

4.2   Model Calibration

In the model simulations we consider a range of values of the constant-gain learning parameter, μ. One extreme assumption that we consider is where the public does not change their estimates at all, but rather uses the parameters associated with the rational expectations equilibrium ignoring the zero bound. Given the presence of the zero bound,

the case of $\mu = 0$ is not the same as rational expectations, but is closely related in that the parameters of the forecasting model are constant. As such, it provides a benchmark that replicates key features of outcomes under full model-consistent expectations.

For the case of learning, we take 0.02 as our benchmark value of μ, and consider alternative values of 0.01 and 0.03 as a robustness exercise. A number of researchers have estimated the value of μ within a learning framework using postwar U.S. data (see Sheridan (2003), Milani (2005, 2006), Orphanides and Williams (2005b), Branch and Evans (2005)). Although the estimates differ across specifications and samples, and are in some cases quite imprecise, the central tendency of these estimates is between 0.02 and 0.03. The value of 0.02 implies that the data from the past 10 years accounts for a little more than 1/2 of the weight in the estimation, data from the preceding 10 years accounting for 1/4 of the weight, and data more than 20 years old accounting for the remaining 20 percent. The average age of the data used in estimation is about 12.5 years, the same as would be the case if agents used standard least squares regressions with 25 years of data. This seems a plausible value given the data limitations that people face in the real world.

We cailibrate the model parameters describing the output gap and inflation dynamics using Milani's (2006) estimates of a very similar model under learning.[4] The upper part of Table 1 reports these parameter values. Note that they are fixed across the different specifications of the learning rate.

The calibration of the long-run neutral real interest rate is important in terms of interpreting our results with respect to the optimal choice of an inflation target. The neutral long-run nominal interest rate, $\bar{i}^*$, measures the average "cushion" that the central bank has to lower rates, starting from the deterministic steady state. The larger the cushion, that is, the larger is $\bar{i}_t$, the less frequently the zero lower bound constrains policy and the shorter the periods the constraint is binding. In terms of our analysis, the decomposition of the long-run neutral nominal interest rate into its real and inflation

---

[4]Milani (2006) estimates a model where the shocks to the natural rate of interest and the markup follow AR(1) processes. This model is quite similar to the one that we use in this paper, once one applies the appropriate transformation to eliminate the serial correlation to the shocks. Therefore, Milani's estimates are reasonable for the model used in this paper. Moreover, the parameter estimates are within the range of other estimates of similar models in the literature.

components is irrelevant. However, to aid in the interpretation of our results, it is useful to discuss results in terms of the inflation target as opposed to the neutral nominal rate. For this purpose, we assume that the long-run real neutral rate is 2.5 percent, about its long-run average in the postwar U.S. economy.[5] Thus, in the following, results for the case of an inflation target of x percent refer to an economy with a neutral long-run nominal interest rate of $x + 2.5$ percent.

The innovation variances are crucial for conducting analysis with the zero bound on interest rates. All else equal, the larger the variances, the more often the zero bound constrains policy and the larger are the effects of the zero bound. We therefore took pains to calibrate these variances in a manner consistent with the empirical evidence on the U.S. economy over the past 20 years. First, we computed the variances of the GDP price index inflation rate and the federal funds rate over the sample of 1985-2005. We then chose the innovation variances so that the model-generated unconditional variances assuming rational expectations and no zero bound were close to their respective empirical counterparts for the federal funds rate and the inflation rate. (We assume no covariance in the innovations.) This method yielded the values of the calibrated standard deviations of the innovations are reported in first column of the lower part of the table.

As noted by Orphanides and Williams (2005a), the presence of learning tends to raise the magnitude of fluctuations in a model economy relative to that that occurs under rational expectations. This is also true for the model analyzed in this paper. Therefore, in order to make the models with the different values of μ comparable in terms of baseline unconditional moments before introducing the zero bound, we calibrated the innovation variances separately for each value of μ, so that the model-generated unconditional variances of inflation, the output gap, and the short-term interest rate were about the same in all variants of the model.[6] The innovation variances decline slightly as the value of μ rises.

---

[5]This calculation is based on using the personal consumption deflator as the price measure. This is the same value for $\bar{r}^*$ used by Reifschneider and Williams (2000). For alternative assumptions regarding this value of $\bar{r}^*$, one can translate the results in the following section by modifying the assumed values of π^* so that the underlying values of $\bar{i}^*$ is the same.

[6]For this calibration exercise, we use a benchmark policy rule of $\gamma_1 = 1, \gamma_2 = 0.25$, and $\gamma_3 = 1$.

5 Monetary Policy Evaluation

In this section, we analyze the performance of monetary policy rules in in environments where the zero lower bound is occasionally binding under alternative assumptions regarding the formation of expectations.

5.1 Benchmark Monetary Policy Rule

We start by constructing a benchmark monetary policy rule. For this purpose, we use the methods described in Levin et al (1999) to compute the coefficient values for γ_2 and γ_3 in the monetary policy rule that minimizes the central bank loss assuming rational expectations and abstracting from the zero lower bound. The resulting coefficient values are given by: $\gamma_2 = 0.1$ and $\gamma_3 = 1$. Orphanides and Williams (2005a, 2006) show that optimal policy under learning responds more strongly to inflation than under rational expectations, so we also consider a more aggressive variant of the rule with $\gamma_2 = 0.25$. We consider two versions of the policy rule, the "difference rule" given by equation (4) and the explicit price-level targeting rule given by equation (5). As noted above, these rules are identical in the absence of the zero bound but differ in an economy where the zero bound is occasionally binding.

5.2 The Effects of the Zero Bound without Learning

We first consider the case where the public does not reestimate their forecasting model, that is, $\mu = 0$. We assume that the parameters of the forecast model are those implied under rational expectations and the absence of the zero lower bound. This might be a reasonable assumption if the zero bound had not been a constraint on policy in the past.

As expected, the "difference" specification of the policy rule fares very poorly with low inflation targets. The upper part of Table 2 shows the results under the difference rule. For these experiments, we assume that the public uses the benchmark forecasting model. For inflation targets of 1.5 percent and above, the zero bound has little effect and the economy never experiences severe recessions, as indicated by the percent of the time that the output gap is below -20 percent. But, for inflation targets of 1 percent and lower, the

zero bound causes a significant deterioration in macroeconomic performance as measured by the simulated root mean squared values of the inflation rate and the output gap. For an inflation target of zero, this policy rule no longer effectively stabilizes the economy and severe recessions are a regular occurrence.

The problem with the difference rule as specified in equation (4) is that it implicitly allows upward drift in the price level target when the zero bound is constraining policy, or is expected to constrain policy in the future. Thus, by including the lagged interest rate in the rule, this policy undermines the price-level targeting feature that is crucial for success in the face of the zero bound. For this reason, for the remainder of the paper, we focus on rules that explicitly target the price level, of the form of equation (5).

The middle panel of the table shows the results for the explicit price-level-targeting policy rule, where the public uses the benchmark forecasting model. This policy does a better job than the difference rule with low inflation targets. For inflation targets of 1 percent and above, the zero bound has little effect on macroeconomic performance. However, for inflation targets below 1 percent, the zero bound causes a marked rise in the average magnitude of fluctuations.

This deterioration in performance occurs because agents do not understand that the central bank will eventually bring the price level back to its target value. Instead, they implicitly assume that following periods when the zero bound is constraining policy, the central bank will let bygones be bygones and will act to stabilize the inflation rate, irrespective of the realized price level. For example, assume that the current interest rate is zero and policy is constrained. Agents forecast the future path of interest rates conditional on the current level of interest rates. As a result, interest rate forecasts will be higher than the monetary policy rule, which accounts for the price level, implies. As a result, the expectations channel that is so powerful and helpful when the public understands the central bank is intent on restoring the price level to its target is reduced and macroeconomic stabilization suffers.

If the public understands that the central bank is targeting the price level and incorporates this information in their forecasting model, then the zero bound has no discernable effects on macroeconomic performance even with an inflation target of zero percent. The

lower part of Table 2 reports the results. Although our framework does not encompass fully model-consistent expectations, these results where the public knows the policy rule mimic those in the literature where the zero bound is not a problem under price level targeting (see, for example Reifschneider and Williams (2000) for comparison).

5.3 The Effects of the Zero Bound with Learning

The presence of learning exacerbates the effects of the zero bound on the economy. The upper part of Table 3 reports the simulation results assuming policy follows the explicit price level targeting rule but the public uses the benchmark forecasting model with $\mu = 0.02$. The losses associated with the zero bound are much larger than in the case of no learning. Indeed, under these conditions, this policy rule does not effectively stabilize the economy for inflation targets below 2 percent. The zero bound introduces persistent deviations from agents' forecasting model, just as the case of no learning discussed above. But, with learning, there is a second channel by which the zero bound affects expectations. During a prolonged episode in which the zero bound is constraining policy, the behavior of monetary policy and the economy systematically deviates from that implied by the forecasting model. These deviations set in motion movements in the estimated parameters of the forecasting model.

Removing public uncertainty about monetary policy significantly reduces the costs associated with the zero bound under learning. The lower part of Table 3 reports the results where the public's forecasts incorporate knowledge of the monetary policy rule. However, even with full public knowledge of the policy rule, the effects of the zero bound interact with the learning involved with the other equations of the model. As a result, inflation targets below 1 percent carry significant costs in terms of stabilization. Therefore, in the face of imperfect knowledge and the zero bound, more than communication of policy intentions is needed. The parameters of the policy rule need to be modified as well, as we show next.

5.4   More Aggressive Monetary Policy

A more aggressive policy rule response to inflation is more effective at minimizing the deleterious effects of the zero lower bound. Table 4 shows the results for the economy with learning where policy follows the more aggressive version of the rule with $\gamma_2 = 0.25$. The more aggressive rule is effective because it reduces the likelihood of deflation and therefore entering a liquidity trap and it promises prompt and aggressive action once the zero bound is no longer constraining policy.

Assuming the public understands the rule, there is little cost to zero inflation under this rule. Comparing these results to those in the previous table, this rule delivers better stabilization of both inflation and output at a zero percent inflation target than does the baseline rule with a 1 percent inflation target. Figures 1 and 2 show the distributions of the inflation rate and the output gap, respectively, under the benchmark and more aggressive rules when the inflation target is zero. For these figures, we have assumed that the public forms expectations using the true monetary policy rule. For the inflation rate, we summed the observations below 5 percent into the leftmost bar (and the same for inflation rates above 5 percent). For the output gap, we summed the observations that are greater than 10 percent in absolute value. Absent learning, given the stipulated objective function, this rule stabilizes inflation too much at the cost of more variability in the output gap. However, with learning, its better containment of inflation helps anchor inflation expectations and avoids deflation and the associated severe recessions.

5.5   Robustness to Alternative Learning Rates

The qualitative results are the same for other values of the learning rate, μ, but qualitatively the losses with low inflation are much larger when the learning rate is 0.03. Tables 5 and 6 show the results for the economy with alternative learning speeds of $\mu = 0.01$ and $\mu = 0.03$, respectively, where policy follows the more aggressive version of the rule with $\gamma_2 = 0.25$. For the case of $\mu = 0.03$, if the public knows the policy rule, the costs associated with the zero bound rise for inflation targets below 1 percent.

6 Conclusion

The historical experiences of deflation with interest rates constrained at zero in the United States in the 1930s and more recently in Japan suggest that it may be prudent to avoid such situations. One solution is to target an inflation rate a few percentage points above zero. Indeed, for this reason and others, inflation-targeting central banks tend to target an inflation rate around 2 percent. Theoretical research on monetary policy yields a far more optimistic view on the ability of monetary policy to stabilize the economy even with an inflation target of zero. This paper suggests a note of caution regarding the effectiveness of monetary policy in the presence of the zero bound if one abandons the assumption that the public has perfect knowledge of the economy and the monetary policy strategy. In a world with imperfect knowledge, policies that would work well if expectations were rational can perform very poorly if the public has imperfect knowledge, especially when the public is uncertain of the policy strategy itself. Although not studied in this paper, a clear corollary of the potential difficulty in stabilizing the economy in the presence of the zero bound is the potential use of fiscal policy interventions when policy is constrained at zero, and the need for more research in this area.

The message of the paper is not, however, entirely negative. First, we show that effective communication of the monetary policy strategy can reduce the costs associated with the zero bound. In this respect, our results relate to Eggertsson's (2005) analysis of the effectiveness of the sudden regime shifts in monetary and fiscal policies in 1933 in the United States. Second, we find that a robust strategy to cope with both imperfect knowledge and the zero bound is to respond more strongly to inflation than would be optimal under rational expectations. This policy rule, assuming it is communicated effectively to the public, is highly effective at stabilizing inflation and output even with an inflation target of one percent.

References

Branch, W. A. and G. W. Evans, 2005, A Simple Recursive Forecasting Model, University of Oregon working paper.

Eggertsson, G.B., 2005, Great Expectations and the End of the Depression, NY Fed Staff Report No. 234, December.

Eggertsson, G.B., 2006, The Deflation Bias and Committing to Being Irresponsible Journal of Money, Credit, and Banking, forthcoming.

Eggertsson, G.B. and M. Woodford, 2003, The Zero Bound on Interest Rates and Optimal Monetary Policy, Brookings Papers on Economic Activity 1, 139–233.

Eggertsson, G.B. and M. Woodford, 2005, Optimal Monetary and Fiscal Policy in a Liquidity Trap, NBER-ISOM Volume, forthcoming.

Evans, G.W. and S. Honkapohja, 2001, Learning and Expectations in Macroeconomics. Princeton University Press, Princeton.

Levin, A.T., V. Wieland and J.C. Williams, 1999, Robustness of Simple Monetary Policy Rules under Model Uncertainty, in John Taylor (ed.) Monetary Policy Rules, University of Chicago Press, 1999.

Levin, Andrew T. Levin, Alexei Onatski, John C. Williams, and Noah Williams, 2006, Monetary Policy under Uncertainty in Micro-founded Macroeconomic Models, NBER Macroeconomics Annual 2005, 20.

McCallum, B.T., 2002, Inflation Targeting and the Liquidity Trap, in: N. Loayza and R. Soto, eds., Inflation Targeting: Design, Performance, Challenges (Central Bank of Chile, Santiago) 395–437.

McGough, Bruce, Glenn D. Rudebusch, and John C. Williams, 2005, Using a Long-Term Interest Rate as the Monetary Policy Instrument, Journal of Monetary Economics, July 2005, 855–879.

Meltzer, A.H., 2001, Monetary Transmission at Low Inflation: Some Clues from Japan, Monetary and Economic Studies 19 S-1, 13–34.

Milani, Fabio, 2005, Expectations, Learning, and Macroeconomic Persistence, mimeo, University of California, Irvine.

Milani, Fabio, 2006, Learning, Monetary Policy Rules, and Macroeconomic Stability, mimeo, University of California, Irvine.

Okina, K. and S. Shiratsuka, 2004, Policy Commitment and Expectation Formation: Japan's Experience under Zero Interest Rates, The North American Journal of Economics and Finance 15, 75–100.

Orphanides, Athanasios and John C. Williams, 2002, Robust Monetary Policy Rules with Unknown Natural Rates, *Brookings Papers on Economic Activity*, 2:2002, 63-118.

Orphanides, Athanasios and John C. Williams, 2005a, Imperfect Knowledge, Inflation Expectations and Monetary Policy, In *The Inflation Targeting Debate*, edited by Ben Bernanke and Michael Woodford, University of Chicago Press.

Orphanides, Athanasios and John C. Williams, 2005b, The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning, and Expectations, *Journal of Economic Dynamics and Control*, 29(11), 1927–1950, November.

Orphanides, Athanasios and John C. Williams, 2006, Monetary Policy with Imperfect Knowledge, *Journal of the European Economics Association*,forthcoming.

Preston, Bruce, 2005, Learning About Monetary Policy Rules When Long-Horizon Expectations Matter, International Journal of Central Banking, 1.

Reifschneider, David L. and John M. Roberts, 2005, Expectations Formation and the Effectiveness of Strategies for Limiting the Consequences of the Zero Bound on Interest Rates, Finance and Economics Discussion Series 2005-70, December 2005.

Reifschneider, David and John C. Williams, 2000, Three Lessons for Monetary Policy in a Low Inflation Era, Journal of Money, Credit and Banking 32(4), 936–966.

Sargent, Thomas J., 1993, *Bounded Rationality in Macroeconomics.* Oxford and New York: Oxford University Press, Clarendon Press.

Sheridan, Niamh, 2003, Forming Inflation Expectations, Johns Hopkins University, mimeo, April.

Svensson, L.E.O., 2001, The Zero Bound in an Open-Economy: A Foolproof Way of Escaping from a Liquidity Trap, Monetary and Economic Studies 19 S-1, 277–312.

Svensson, L.E.O., 2003, Escaping from a Liquidity Trap and Deflation: The Foolproof Way and Others, Journal of Economic Perspectives 17(4), 145-166.

Woodford, M., 2003, Interest and Prices: Foundations of a Theory of Monetary Policy, (Princeton University Press, Princeton).

Table 1. Model Calibration

Parameter	Calibrated Values			
ϕ			0.200	
η			0.945	
β			0.990	
κ			0.078	
ρ			0.849	
θ			0.849	
μ	0.000	0.010	0.020	0.030
σ_r	7.500	7.500	7.250	6.750
σ_u	0.550	0.539	0.528	0.507

Notes: Parameter values reported in the upper part of the table are taken from Milani (2006), Table 3.3. The calibration of the values of the long-run neutral real interest rate, $\bar{r}^*$, and the innovation standard deviations is described in the text.

Table 2. The Effects of the Zero Bound without Learning ($\mu = 0$)
Baseline Policy Rule: $\gamma_1 = 1$, $\gamma_2 = 0.1$, $\gamma_3 = 1$.

Inflation		Root Mean Square		Central	Frequency	
Target (π^*)	Inflation	Output Gap	Interest Rate	Bank Loss	$i_t = 0$	$y_t \leq -20$
Policy follows difference rule (eq. 4), and public forecasts with same						
0.0	3.7	7.5	1.8	28.0	22.8	12.0
0.5	2.2	4.5	1.8	10.1	10.2	3.4
1.0	1.2	2.6	1.8	3.4	3.7	0.4
1.5	0.9	2.0	1.8	2.2	1.5	0.0
2.0	0.9	2.0	1.8	2.1	0.7	0.0
3.0	0.9	1.9	1.8	2.0	0.1	0.0
4.0	0.9	1.9	1.8	2.0	0.0	0.0
Policy follows price level rule (eq. 5), but public forecasts with difference rule						
0.0	1.5	3.1	1.7	4.9	12.3	1.3
0.5	1.0	2.2	1.8	2.5	6.4	0.2
1.0	0.9	2.0	1.8	2.1	3.2	0.0
1.5	0.9	2.0	1.8	2.0	1.6	0.0
2.0	0.9	1.9	1.8	2.0	0.7	0.0
3.0	0.9	1.9	1.8	2.0	0.1	0.0
4.0	0.9	1.9	1.8	2.0	0.0	0.0
Policy follows price level rule (eq. 5), and public forecasts with same						
0.0	0.9	1.9	1.7	2.0	8.2	0.0
0.5	0.9	1.9	1.7	2.0	4.9	0.0
1.0	0.9	1.9	1.8	2.0	2.7	0.0
1.5	0.9	1.9	1.8	2.0	1.4	0.0
2.0	0.9	1.9	1.8	2.0	0.7	0.0
3.0	0.9	1.9	1.8	2.0	0.1	0.0
4.0	0.9	1.9	1.8	2.0	0.0	0.0

Table 3. The Effects of the Zero Bound with Learning ($\mu = 0.02$)
Baseline Policy Rule: $\gamma_1 = 1$, $\gamma_2 = 0.1$, $\gamma_3 = 1$.

Inflation	Root Mean Square			Central	Frequency	
Target (π^*)	Inflation	Output Gap	Interest Rate	Bank Loss	$i_t = 0$	$y_t \le -20$
Policy follows price level rule (eq. 5), but public forecasts with difference rule						
0.0	6.7	13.3	3.3	89.8	50.7	40.6
0.5	4.7	9.4	2.9	45.3	27.1	19.5
1.0	3.3	6.6	2.6	22.8	13.4	9.2
1.5	2.5	5.1	2.3	13.4	7.1	5.0
2.0	2.0	4.0	2.2	8.6	3.8	2.8
3.0	1.4	2.9	2.0	4.3	1.2	0.9
4.0	1.0	2.3	1.9	2.7	0.3	0.2
Policy follows price level rule (eq. 5), and public forecasts with same						
0.0	1.7	3.8	1.9	6.8	12.3	2.1
0.5	1.5	3.3	2.0	5.2	7.3	1.3
1.0	1.2	2.8	1.9	3.9	4.2	0.8
1.5	1.1	2.5	1.9	3.0	2.2	0.3
2.0	1.0	2.3	1.9	2.8	1.3	0.3
3.0	1.0	2.1	1.9	2.3	0.3	0.1
4.0	0.9	2.1	1.9	2.3	0.2	0.0

Table 4. The Effects of the Zero Bound with Learning ($\mu = 0.02$)
More Aggressive Policy Rule: $\gamma_1 = 1$, $\gamma_2 = 0.25$, $\gamma_3 = 1$.

Inflation	Root Mean Square			Central	Frequency	
Target (π^*)	Inflation	Output Gap	Interest Rate	Bank Loss	$i_t = 0$	$y_t \leq -20$
Policy follows price level rule (eq. 5), but public forecasts with difference rule						
0.0	2.0	4.4	1.9	9.2	15.9	3.0
0.5	1.5	3.4	1.9	5.3	8.4	1.3
1.0	1.2	2.8	1.9	3.7	4.5	0.6
1.5	1.0	2.5	1.9	2.9	2.3	0.3
2.0	0.9	2.3	1.9	2.4	1.0	0.1
3.0	0.9	2.2	1.9	2.3	0.2	0.0
4.0	0.8	2.2	1.9	2.3	0.1	0.0
Policy follows price level rule (eq. 5), and public forecasts with same						
0.0	0.9	2.7	1.8	3.0	10.8	0.2
0.5	0.9	2.5	1.9	2.6	6.2	0.1
1.0	0.8	2.3	1.9	2.4	3.4	0.0
1.5	0.8	2.3	1.9	2.4	1.8	0.0
2.0	0.8	2.3	1.9	2.4	0.9	0.0
3.0	0.8	2.3	1.9	2.4	0.3	0.0
4.0	0.8	2.2	1.9	2.3	0.1	0.0

Table 5. The Effects of the Zero Bound with Slower Learning ($\mu = 0.01$)
More Aggressive Policy Rule: $\gamma_1 = 1$, $\gamma_2 = 0.25$, $\gamma_3 = 1$.

Inflation	Root Mean Square			Central	Frequency	
Target (π^*)	Inflation	Output Gap	Interest Rate	Bank Loss	$i_t = 0$	$y_t \leq -20$
Policy follows price level rule (eq. 5), but public forecasts with difference rule						
0.0	1.2	2.9	1.8	3.9	13.0	0.6
0.5	1.0	2.6	1.8	3.1	7.2	0.3
1.0	0.9	2.3	1.8	2.4	3.7	0.0
1.5	0.8	2.2	1.9	2.3	1.9	0.0
2.0	0.8	2.2	1.9	2.2	0.9	0.0
3.0	0.8	2.2	1.9	2.2	0.2	0.0
4.0	0.8	2.2	1.9	2.2	0.0	0.0
Policy follows price level rule (eq. 5), and public forecasts with same						
0.0	0.8	2.3	1.7	2.3	10.2	0.0
0.5	0.8	2.2	1.8	2.2	6.0	0.0
1.0	0.8	2.2	1.8	2.2	3.2	0.0
1.5	0.8	2.2	1.8	2.2	1.7	0.0
2.0	0.8	2.2	1.9	2.2	0.9	0.0
3.0	0.8	2.2	1.9	2.2	0.2	0.0
4.0	0.8	2.2	1.9	2.2	0.0	0.0

Table 6. The Effects of the Zero Bound with Faster Learning ($\mu = 0.03$)
More Aggressive Policy Rule: $\gamma_1 = 1$, $\gamma_2 = 0.25$, $\gamma_3 = 1$.

Inflation	Root Mean Square			Central	Frequency	
Target (π^*)	Inflation	Output Gap	Interest Rate	Bank Loss	$i_t = 0$	$y_t \leq -20$
Policy follows price level rule (eq. 5), but public forecasts with difference rule						
0.0	3.8	8.1	2.5	31.3	24.7	13.6
0.5	2.4	5.3	2.2	13.1	11.4	5.0
1.0	1.7	3.9	2.0	7.0	5.5	2.2
1.5	1.4	3.3	1.9	5.1	3.1	1.3
2.0	1.2	2.9	1.9	4.0	1.7	0.8
3.0	1.0	2.6	1.9	3.1	0.7	0.4
4.0	0.9	2.4	1.9	2.6	0.3	0.2
Policy follows price level rule (eq. 5), and public forecasts with same						
0.0	1.2	3.4	2.0	4.8	11.1	0.9
0.5	1.1	3.0	2.0	3.9	6.5	0.5
1.0	1.0	2.7	1.9	3.2	3.5	0.3
1.5	0.9	2.5	1.9	2.7	1.9	0.1
2.0	0.9	2.5	1.9	2.7	1.1	0.1
3.0	0.9	2.4	1.9	2.5	0.5	0.1
4.0	0.8	2.3	1.9	2.4	0.2	0.0

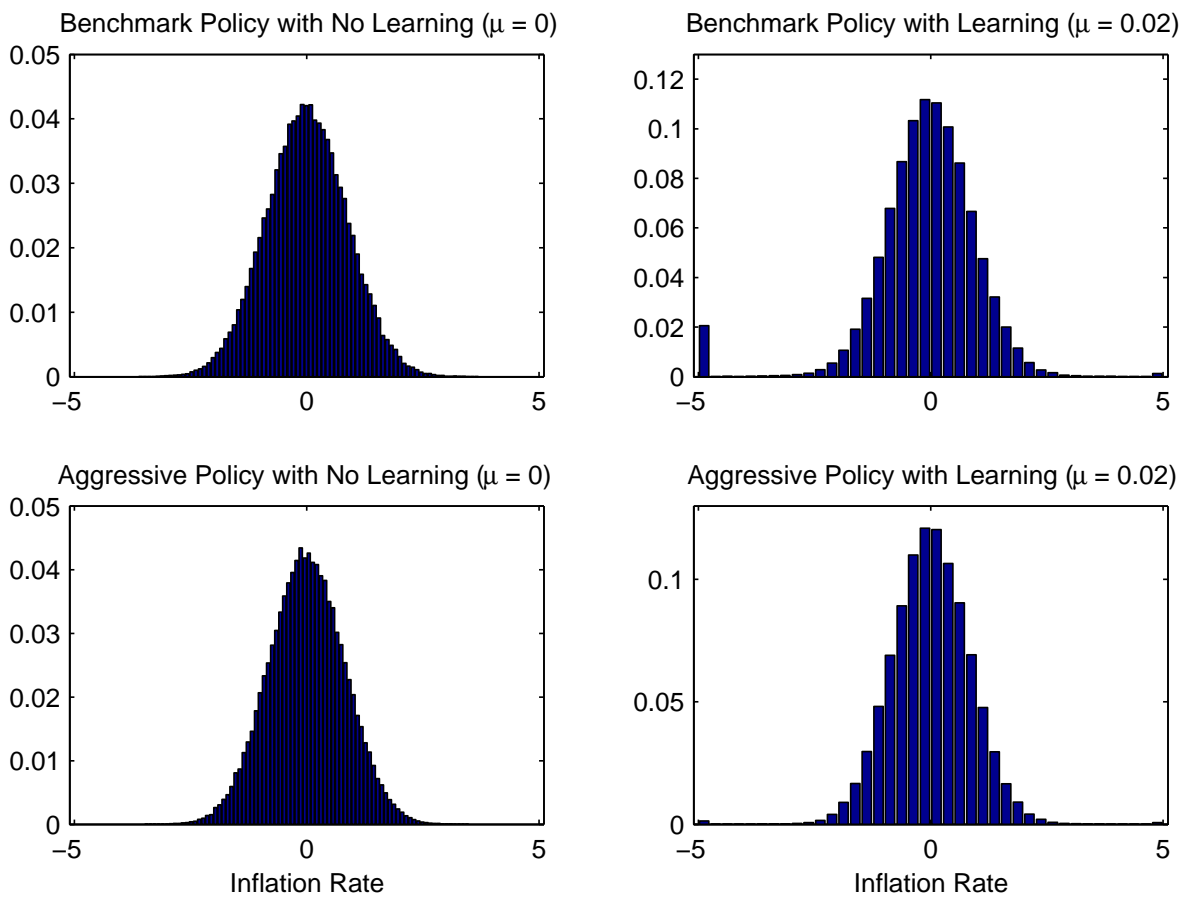Figure 1: Distributions of Inflation Rate with a Zero Inflation Target

Figure 2: Distributions of the Output Gap with a Zero Inflation Target