Expectations, Learning, and Discretionary Policymaking*

Christian Jensen
Southern Methodist University

When policy forecasts are based on the policymaker's present and past actions, current policy affects expectations of future policy, contrary to what happens when forecasters can replicate policymaking perfectly. We show that when forecasts are generated through any linear combination of present and past policy functions that produces expectations consistent with the implemented policy, the optimal discretionary policy exploiting learning converges toward the optimal commitment plan as we approach a situation where people do not discount the future. Since influencing expectations permits improving policy, successful policymakers need to know how policy expectations are formed and how they can affect these expectations.

JEL Codes: E52, E58, E61.

1. Introduction

In recent years, central banks appear to have been putting more effort into influencing people's expectations, not only by announcing explicit inflation targets, but also by incorporating predictions of future inflation and output, and in some cases even of future policy,

*Thanks for helpful comments are due to Takushi Kurozumi, Bennett McCallum, anonymous referees, and participants at the 2005 Missouri Economics Conference and the 2006 North American Econometric Society Summer Meetings. Author contact: Department of Economics, Box 750496, Southern Methodist University, Dallas, TX 75275; Tel: (214) 768-3736; e-mail: cjensen@alumni.cmu.edu.

in their policy announcements.[1] Furthermore, some of their statements indicate that they believe their actions affect the public's policy expectations and assert that they take this influence into account when deciding policy.[2] In the present paper, we assume that policy affects expectations due to the public's lack of knowledge about the policy decision process, in terms of the objectives and models policymakers use to choose policy, so that they cannot predict policy by putting themselves in the policymakers' shoes. Instead, their forecasts have to be based on present and past policy, which we assume is perfectly observable, and this way, present policy comes to affect expectations of future policy. Our objective is to show how the optimal discretionary policy can improve and become more similar to the optimal commitment plan, and even converge to it in the limiting case of no time discounting, when policymakers influence expectations through present policy. We do so in a setup where expectations are always consistent with the model and the implemented policy.

---

[1]See Bernanke and Mishkin (1997), McCallum (1996), and Svensson (1997, 1999) for discussions on inflation targeting. Some examples of announcements incorporating forecasts include the following: "Against this background, the Committee adopted a directive that was biased toward a possible firming of policy going forward. Committee members emphasized that such a directive did not signify a commitment to near-term action." (Federal Reserve System [FRS], October 5, 1999). "However, with inflation quite low and resource use slack, the Committee believes that policy accommodation can be maintained for a considerable period." (FRS, December 9, 2003). "The Committee judges that some further measured policy firming is likely to be needed to keep the risks to the attainment of both sustainable economic growth and price stability roughly in balance." (FRS, December 13, 2005). "While growth has been relatively modest so far, both external and domestic factors give reason to expect a strengthening of the recovery through 2004 and beyond." (European Central Bank [ECB], March 4, 2004). "All in all, our judgement remains that real economic growth will gradually improve over the period ahead." (ECB, June 2, 2005). "Euro area real GDP is projected to grow at rates of between 1.2% and 1.6% in 2005, and between 1.4% and 2.4% in 2006 and 2007.... The projections indicate average HICP inflation to lie between 2.1% and 2.3% in 2005, and between 1.6% and 2.6% in 2006.... For 2007 an average inflation rate of between 1.4% and 2.6% is projected." (ECB, December 1, 2005). "Looking ahead, given that our monetary policy continues to be accommodative, a progressive withdrawal of monetary accommodation will be warranted if our assumptions and baseline scenario are confirmed." (ECB, August 3, 2006). "The Executive Board's assessment is that the economic projections imply a sight deposit rate in the interval $1\frac{1}{4}$–$2\frac{1}{4}$% at the beginning of November 2004." (Norges Bank, July 1, 2004).

[2]"Furthermore, today's move by 50 basis points appeared to be the best way in which to avoid uncertainties regarding the future course of monetary policy." (ECB, November 4, 1999).

Without a credible commitment, the public expects policymakers to reconsider policy in every period and take whatever action is optimal from the perspective of that point in time. As a result, policymakers' present and past actions have no effect on the public's policy forecasts, since these are produced by going over the decision process that policymakers are expected to face in the future. However, predicting policy this way requires that forecasters be able to reproduce policymakers' decision process, which for monetary policy is difficult, if not impossible. In most countries, monetary policy is decided by a group of individuals relying on different models and beliefs about the workings of the economy and its contemporary state. In addition, there might not be agreement about the policy objectives. So how do people actually predict monetary policy? Instead of reconstructing policymakers' decision making, present and past policy is used to forecast future policy. Since the policy problem does not change much from one period to the next, recent policy is used as an indicator of the policy to be implemented in the future. An example from the literature is the Taylor rule, which is probably the most successful description of recent monetary policy in the United States.[3] While the rule is based on knowledge of monetary policymaking, it (probably) does not reproduce policymakers' decision process. Instead, it uses past policy to obtain a description of the behavior of the policy instrument. However, when individuals' policy expectations depend on policymakers' present and past actions, the discretionary solution changes, because, contrary to what is assumed under standard discretion, policy affects expectations.

Many authors question whether individuals are, in practice, able to reproduce policymaking to the extent required for this to be useful in forecasting. For instance, in the context of monetary policy, the learning literature—exemplified by Bullard and Mitra (2002), Carlstrom and Fuerst (2004), and Evans and Honkapohja (2001, 2003)—studies how equilibrium outcomes are affected when individuals learn policy, or equilibria, by applying least-squares methods to data from the past and whether these equilibria are stable and learnable. While these contributions assume that policy is exogenous and, therefore, unaffected by the learning process, later work—such

[3]See Judd and Rudebusch (1998), Orphanides (2003), and Taylor (1993, 1999).

as that by Gaspar, Smets, and Vestin (2005), Molnar and Santoro (2005), and Orphanides and Williams (2005)—not only lets policy adjust optimally to the public's expectations, but, in addition, chooses policy so as to shape these expectations in an optimal manner. A common observation in these papers is that the optimal discretionary policy exploiting adaptive learning can temporarily attain features of the optimal commitment plan. We show that when present and past policy functions are used to forecast future policy, the optimal discretionary policy exploiting learning only deviates from the optimal commitment plan due to time discounting. The reason is that the effect present policy has on expectations through these forecasts, when they are consistent with the implemented policy, is the same as the effect preannouncing policy has under commitment, apart from influencing expectations in different periods.

In the learning literature cited in the previous paragraph, it is common to assume that individuals know everything about the monetary policy problem except for having to estimate parameter values from noisy observations of past realizations. Instead, we assume that people know nothing about the policymaking process but can observe the implemented policy functions exactly. This allows people to predict policy perfectly in our model, which guarantees that our results are due to the exploitation of expectations and not because the expectations are inconsistent with actual policy. At the same time, such forecasts maximize the effect policy has on expectations, thereby bringing the discretionary solution as close as possible to the optimal commitment plan. We do this to illustrate the potential gains from exploiting expectations in a situation where these expectations are consistent with actual policy and, so, do not require that policymakers mislead the public. The actual gains will depend on exactly how people forecast policy, which is an empirical question. By focusing on policy functions instead of parameter values, our setup is closer to the Taylor-rule example than what is usual in the learning literature, in the sense that people use past policy to forecast future policy without reconstructing the policymaking process.

Since the optimal discretionary policy depends on how the public forecasts policy, it is important for policymakers to know how these forecasts are formed and how they are affected by the policymakers' actions. This knowledge can be used to influence expectations and thereby better achieve the policy objectives, much in the same way as a credible commitment to a plan can. Hence, the public's unawareness of the policymaking process, and use of present and past policy to forecast policymakers' future actions, can improve the discretionary solution. Moreover, if policymakers can affect the way in which policy expectations are formed, they can also influence the optimal discretionary, or time-consistent, policy. By determining how complicated and public the policy decision process is, policymakers can make it easier or harder for outsiders to forecast policy by reproducing this process. The less information people have about policymaking, the less likely that they can use it to forecast policy, and the more likely that their forecasts will be based on present and past policy. Thus, by affecting the knowledge the public has about the policy decision process and, thereby, how the public forecasts policy, policymakers can influence the discretionary solution.

The next section introduces a standard sticky-price model commonly used for modeling monetary policymaking. In section 3, this model is used to illustrate how the discretionary solution changes when policy affects expectations. We show that when people's next-period policy expectations are determined by the policy they see being implemented today, the discretionary solution becomes almost the same as the commitment plan, and we show that the two policies are identical when the discount factor is close enough to 1. Next, we explore how this result is modified when policy forecasts look further back in time, and when they are inconsistent with actual policy. In section 4, we argue that the optimal policy in our setup will, in general, differ from the one in a reputational model, even though policy affects expectations in both cases. We conclude that it is important for policymakers to know how the public's policy expectations are formed and, in particular, how they can influence these expectations.

2. Model

We employ a New Keynesian sticky-price model with monopolistic competition like the one used by, for example, Clarida, Galí, and Gertler (1999), McCallum and Nelson (2000), Svensson and Woodford (2003), and Woodford (1999a, 1999b, 2003). It assumes that it is costly, or impossible, to change prices frequently, so that present price setting needs to take into account expected future price movements. Assuming that the central bank can influence the rate of inflation through monetary policy, current prices will depend on both present and expected future monetary policy, so that present inflation depends on expected future inflation. The monetary policy problem is to respond to random cost-push shocks so as to keep inflation and output as close as possible to their flexible-price values, thereby minimizing the welfare loss due to price stickiness. For simplicity, and in line with current practice in the literature, we assume that the policymaker determines the rate of inflation directly and that individuals decide how much to produce, instead of the prices.

Letting π_t denote the rate of inflation and y_t the log of output, both in terms of deviations from flexible-price values, Calvo (1983), Rotemberg (1982), and Rotemberg and Woodford (1999) show that the policy problem in any period $t = 0$ is to minimize[4]

$$E \sum_{t=0}^{\infty} \beta^t \left(\pi_t^2 + \omega y_t^2\right) \tag{1}$$

subject to

$$\pi_t = \beta E_t \pi_{t+1} + \alpha y_t + u_t \tag{2}$$

$$u_t = \phi u_{t-1} + \varepsilon_t. \tag{3}$$

The shock ε_t is white noise with variance σ_ε^2 and is assumed to be observed by the policymaker before setting the inflation rate π_t. The parameters α, ω, ϕ, and σ_ε are all strictly positive, while $\beta \in (0, 1)$.[5]

As is shown by Clarida, Galí, and Gertler (1999), Currie and Levine (1993), Svensson and Woodford (2003), and Woodford

[4]We minimize the unconditional expected value instead of the conditional one to avoid optimal policies being dependent on initial conditions. The time-inconsistency problem is the same with the conditional and unconditional objective functions, so our conclusions apply for both cases.

[5]$\phi > 0$ is needed to generate persistence in the model. If $\phi = 0$, expectations of future inflation, $E_t \pi_{t+1}$, would always be zero for some of the policy equations we consider, and the dependency of current inflation on expected future inflation would in practice disappear. Persistence can alternatively be introduced by including lagged values of variables in state equation (2) or in the policy function.

(1999a), the optimal commitment policy is

$$\pi_0 = -\frac{\omega}{\alpha} y_0 \tag{4}$$

$$\pi_t = -\frac{\omega}{\alpha} y_t + \frac{\omega}{\alpha} y_{t-1}, \quad t \geq 1. \tag{5}$$

This policy minimizes the objective in (1), but, as Kydland and Prescott (1977) argue, it is not time consistent, because if we reoptimized in any later period $\tau > 0$, the policymaker would want to deviate from the commitment policy in (5) and instead commit to implementing

$$\pi_\tau = -\frac{\omega}{\alpha} y_\tau \tag{6}$$

$$\pi_t = -\frac{\omega}{\alpha} y_t + \frac{\omega}{\alpha} y_{t-1}, \quad t > \tau. \tag{7}$$

When the policymaker reoptimizes in every period, the implemented policy is always

$$\pi_t = -\frac{\omega}{\alpha} y_t, \tag{8}$$

which—as is shown by Clarida, Gali, and Gertler (1999), McCallum and Nelson (2000), Svensson and Woodford (2003), and Woodford (1999a)—is the optimal discretionary policy. Combining (8) with the state equations in (2) and (3), one can show that the discretionary solution can equivalently be written as

$$\pi_t = \frac{\omega}{\alpha^2 + \omega(1 - \beta\phi)} u_t \tag{9}$$

in terms of the exogenous cost-push shock u_t.

3. Expectations

To study the role of policy expectations in the time-consistent solution, it is useful to be more explicit about these expectations. Assume that policy is given by

$$\pi_t = h_t u_t, \tag{10}$$

where $\{h_t\}_{t=0}^{\infty}$ is a sequence of policy parameters to be determined, so that the policy problem is now to choose h_t, instead of π_t directly.[6] State equation (2) summarizes individuals' optimal behavior. The reason it has current inflation, or output, depend on expected next-period inflation is that when prices are sticky, an individual's optimal present action depends on his or her inflation expectations for the next period. Hence, it is the public's inflation expectations that matter for the present state. This implies that when policy is as in equation (10), the expectations of next-period inflation, $E_t\pi_{t+1}$, in state equation (2) satisfy

$$E_t\pi_{t+1} = E_t h_{t+1} u_{t+1} = h^e_{t,t+1}\phi u_t, \tag{11}$$

where $h^e_{t,t+1}$ denotes the policy that the public at time t expects will be implemented in period $t+1$.[7] The reduced-form solution of the model described by equations (2), (3), (10), and (11) is given by (10) and

$$y_t = \frac{h_t - \beta\phi h^e_{t,t+1} - 1}{\alpha} u_t, \tag{12}$$

where the values of the endogenous variables, π_t and y_t, depend on the policy that is implemented in the contemporaneous period, h_t, and the policy that in period t is expected to be implemented in the period after, $h^e_{t,t+1}$. Inserting into the objective function in (1), we have

$$E\sum_{t=0}^{\infty}\beta^t\left(h_t^2 + \frac{\omega}{\alpha^2}\left(h_t - \beta\phi h^e_{t,t+1} - 1\right)^2\right)u_t^2. \tag{13}$$

The value of the objective function in any period $t = 0$ depends on both the policymaker's future actions, $\{h_t\}_{t=1}^{\infty}$, and the public's policy expectations, $\{h^e_{t,t+1}\}_{t=0}^{\infty}$. We distinguish between the two, because the policymaker is able to predict policy by resolving the policy problem, while the public might not be able to do so. The

[6] Writing policy as a function of y_t instead of u_t might be more in accordance with what policymakers actually do, but it makes the computation of the optimal policy more arduous.

[7] From the perspective of period t, individuals' policy expectations for the period after, $h^e_{t,t+1}$, are nonstochastic.

optimal time-consistent policy is found by minimizing the objective in (13) with respect to h_0, assuming that all later policy actions, $\{h_t\}_{t=1}^{\infty}$, will be chosen the same way.

With standard discretion, the present policy action, h_0, has no effect on policy expectations, so

$$\frac{\partial h_{t,t+1}^e}{\partial h_0} = 0, \quad \forall t, \tag{14}$$

and the first-order condition from minimizing (13) with respect to h_0 yields

$$h_0 = \omega \frac{\beta \phi h_{0,1}^e + 1}{\alpha^2 + \omega}. \tag{15}$$

Since future policy, $\{h_t\}_{t=1}^{\infty}$, and future policy expectations, $\{h_{t,t+1}^e\}_{t=1}^{\infty}$, are independent of present policy, h_0, the optimal present policy action only depends on present expectations of next-period policy, $h_{0,1}^e$. Because future policy will be determined by resolving exactly the same optimization problem in the future, the optimal policy will be the same in every period. Under standard discretion the public is assumed to forecast policy by replicating the policymaker's decision process, so they will be able to predict perfectly, and therefore we have

$$h_{0,1}^e = h_1 = h_0 = h_t, \quad \forall t. \tag{16}$$

Inserting this into the first-order condition in (15) and solving for h_t, we find that the time-consistent solution is

$$h_t = \frac{\omega}{\alpha^2 + \omega(1 - \beta\phi)}, \tag{17}$$

which gives the standard discretionary policy in (9).

Imagine now that the present policy action, h_0, affects the public's present or future policy expectations so that condition (14) does not hold. Then, the optimal time-consistent policy will differ from the standard discretionary solution in (9). When policy affects future policy expectations, $\left\{h_{t,t+1}^e\right\}_{t=1}^{\infty}$, the optimal present policy action will not only depend on the effect it has on expectations, $\left\{\frac{\partial h_{t,t+1}^e}{\partial h_0}\right\}_{t=1}^{\infty}$, but also on the expectations themselves, $\left\{h_{t,t+1}^e\right\}_{t=1}^{\infty}$, and the policy the policymaker will implement in the future, $\{h_t\}_{t=1}^{\infty}$. While the present policy action has no direct effect on the policymaker's future actions, it can affect these actions through the public's expectations when the policymaker reoptimizes in future periods.

For a simple, but explicit, example of policy affecting expectations in a context where these are consistent with the model and actual policy, imagine that policy expectations are

$$h^e_{t,t+1} = h_t, \quad \forall t, \tag{18}$$

so that the policy that is implemented today is expected to be implemented in the next period too.[8] While these expectations are plain, they should be reasonable in a situation where people have little knowledge about the policy decision process but in which they can observe the implemented policy, as we assume here. The public's policy expectations are not static in (18); they take into account that the policymaker will reoptimize in every period. Present policy, h_0, determines present expectations of next-period policy, $h^e_{0,1}$, but has no effect on future expectations, since $h^e_{1,2} = h_1$, $h^e_{2,3} = h_2, \ldots$. Hence, the policymaker takes into account that if the present policy action determines expectations of next-period policy today, future policy actions will determine expectations in the future. With a policy function like (10), this implies, in the present model, that the policy implemented in any period t will only affect inflation and output in period t. Inserting the expectations in (18) into the objective function in (13) and minimizing with respect to h_0, we find that the optimal discretionary policy of form (10) is now

$$\pi_t = \frac{\omega(1-\beta\phi)}{\alpha^2 + \omega(1-\beta\phi)^2} u_t, \tag{19}$$

no matter what policy was implemented in the past or will be implemented in the future. As Clarida, Galí, and Gertler (1999) show, this policy is the optimal commitment rule of the form $\pi_t = hu_t$.[9] Since

[8] We are not assuming that people expect the same inflation rate, but that they expect the same policy equation.

[9] The effect of exploiting learning is in this case the same as if the weight on output stabilization ω was lowered to $\omega(1 - \beta\phi)$ in the standard discretionary

the optimal policy is the same in every period, using present policy to predict next-period policy, as in (18), should arguably be a reasonable way to forecast in this setup, especially given the assumption that individuals cannot reproduce the policymaker's decision process. In fact, these forecasts are exactly correct in terms of the policy function, and in terms of the inflation rate, they only miss due to the unpredictable white noise shock ε_t, the same as when forecasts are generated by reproducing the policymaker's decision process.

When people expect next-period policy to be the same as the present policy, the optimal discretionary policy will, in general, not be of the form in (10). For $\alpha = .05$, $\beta = .99$, $\omega = .25$, and $\phi = .8$, the policy

$$\pi_t = 3.251u_t + .414y_{t-1} \tag{20}$$

is a stationary discretionary equilibrium that minimizes the objective in (1) with a policy function of the same form as the optimal commitment policy,[10]

$$\pi_0 = 4.587u_0 \tag{21}$$

$$\pi_t = 3.247u_t + .455y_{t-1}, \quad t \geq 1. \tag{22}$$

With a policy function of the form in (20), the policy implemented in period t does not only affect inflation and output in period t, since the lagged output term, y_{t-1}, in the policy function links periods together. Consequently, the optimal present policy action depends on what policy was implemented in the past. By a stationary equilibrium, we mean that the reported policy is optimal given that the same policy was implemented in the past. In other words, the reported policy is the one we would converge toward as the effect of policy implemented prior to the initial period, $t = 0$, dissipates.

Comparing equations (20) and (22), we see that when the present policy is expected to be implemented in the next period too, the

---

solution, and due to Rogoff's (1985) conservative central banker result, we know that lowering the weight on output stabilization can increase welfare.

[10]The commitment policy in (21)–(22) is equivalent to the one in (4)–(5). The policy in (20) can be computed algebraically, but the general solution is too extensive to reproduce here.

discretionary solution becomes very similar to the optimal commitment policy.[11] The reason is that the influence policymakers have on expectations with such forecasts is the same as the influence they have on expectations under commitment. The optimal commitment plan takes into account that the policy action to be implemented in any particular period will determine people's policy expectations for that period, since $h^e_{t,t+1} = h_{t+1}$ for $t = 1, 2, \ldots$, using the notation from above. When people forecast the next-period policy action to be the same as the present one, the action to be implemented in a period will determine people's expectations for the period immediately afterward, and $h^e_{t,t+1} = h_t$ for $t = 0, 1, \ldots$. Apart from the timing, the effects on expectations are the same. When the discount factor β is close enough to 1, periods are weighted almost the same, and the difference in timing becomes irrelevant, making the optimal policies under commitment and discretion the same. In fact, one can prove that as β converges toward 1, the optimal stationary discretionary policy and the optimal commitment policy converge, no matter what the values of α, ω, and ϕ (see the appendix).

Likewise, when expectations are determined by the policy action implemented $J > 0$ periods back in time, so that

$$h^e_{t,t+1} = h_{t-J}, \tag{23}$$

discretion will match commitment only if β^{J+1} is close to 1. The effect on policy expectations is the same, no matter what the value of J (including for commitment, which corresponds to the case with $J = -1$), but the larger J is, the more these effects are discounted away and ignored when determining the optimal discretionary policy. From this it follows that when expectations are a weighted average of past policy actions,

$$h^e_{t,t+1} = \sum_{j=0}^{J} \theta_j h_{t-j}, \tag{24}$$

[11] The parameters have been chosen somewhat arbitrarily from previous work. The key variable is the discount factor β, since, as shown below, the discretionary solution converges toward the optimal commitment plan as β approaches 1. A discount factor of .99 is fairly standard in the literature for a quarterly model.

where θ_j are arbitrary weights, the discretionary solution will be close to the optimal commitment rule only if

$$\sum_{j=0}^{J} \beta^{j+1}\theta_j \tag{25}$$

is close to 1.[12] For the expectations in (24) to be consistent with actual policy,

$$\sum_{j=0}^{J} \theta_j = 1 \tag{26}$$

needs to be satisfied, so the larger J is, the less similar discretion will be to commitment, given that $\beta \in (0,1)$. Still, whatever effect policy has on expectations, it will be optimal to exploit this influence to improve policymaking.

It is much easier to match commitment if expectations are not required to be consistent with actual policy, and one can even outdo it when policy has no influence on expectations. Using state equation (2) to substitute for y_t in the standard discretionary policy in (8), we find that the optimal inflation rate under discretion when policy has no effect on expectations is

$$\pi_t = \frac{\omega}{\alpha^2 + \omega}(\beta E_t \pi_{t+1} + u_t), \tag{27}$$

while the associated output level is

$$y_t = -\frac{\alpha}{\alpha^2 + \omega}(\beta E_t \pi_{t+1} + u_t), \tag{28}$$

both in terms of the public's inflation expectations, $E_t\pi_{t+1}$.[13] By making these expectations equal $-\frac{1}{\beta}u_t$, one can make both inflation and output be zero at all times, which is much better than what can be achieved under commitment. Such expectations would, however, not be consistent with the implemented policy, defined in (27). To

[12] This can be verified numerically, or even algebraically for a small J.

[13] One can show that this is the optimal policy by substituting state equation (2) into the objective in (1) and minimizing with respect to π_t and $E_t\pi_{t+1}$ as two separate and independent variables.

guarantee that the gains from exploiting learning that we derive in the present study are not the result of inconsistent expectations or fooling the public, we assume throughout that policy is perfectly observable, which in our framework implies that policy forecasts are dead-on.

The learning literature is generally more optimistic than we have been above in terms of the knowledge the public is assumed to have about the policy decision process. In our setup, people are assumed to know nothing about this process, but, instead, we imagine that they can observe the implemented policy perfectly, which is also too optimistic. When policy is observed with error, after a lag, or when policymakers cannot exactly control what policy is implemented, policymakers' influence on expectations could be weakened, thus reducing the possible gains from its exploitation. Of course, such assumptions could also make it harder for forecasters to weed out expectations that are inconsistent with actual policy, facilitating a situation in which policymakers mislead the public.

Policy can also affect expectations when individuals are able to reproduce the policymaking process, as long as they cannot do so perfectly. An example is provided by Gaspar, Smets, and Vestin (2005), where forecasters are assumed to know everything about the model used to determine policy except the parameter values. Since the data that are used to estimate these parameters are observed with error, it can be difficult to distinguish policymakers' preferences, and parameter values in general, from strategic manipulation of the data, and policymakers should exploit this to influence the public's expectations. However, once policy forecasts are exactly correct because the true parameter values have been deduced, present policy has no effect on estimates and forecasts, and the optimal discretionary solution falls back to the standard one where expectations are independent of policy. In contrast, our forecasting scheme allows policy to affect expectations also when these are exactly correct, even if the magnitude of these effects depends on how far back forecasters look and how heavily people discount the future.

4. Discussion

We assume that there is no way for policymakers to credibly commit to implementing a policy. Barro (1986) and Rogoff (1987) argue that policymakers that have a reputation for implementing the policy they preannounce could achieve credible commitment.[14] As in our framework, policy affects expectations in their setup but in a different way. In a reputational model, the policymaker's past actions determine the credibility of the commitment and only influence policy expectations by affecting how likely the public believes it is that the announced policy will actually be implemented. If the policymaker's credibility is perfect, policy expectations will be determined solely by the announced policy. If the policymaker has no credibility, people will expect the standard discretionary policy. In our setup, there are no announcements, but policy affects expectations because people use it to learn what policy is, and will be, implemented. In one case, people learn whether they should trust the policymaker's announcements; in the other, they learn what policy is being implemented and never trust the policymaker's announcements. In a reputational model, the standard discretionary policy is not implemented because it would ruin the policymaker's good reputation and ability to credibly commit. In our setup, the standard discretionary policy is not implemented because when present policy is used to learn what policy is—and will be—employed, the optimal present action is not standard discretion.

Evans and Honkapohja (2006) and McCallum (1995) suggest that central banks should implement the optimal commitment policy even if they have no way of credibly committing to it, not even a reputation. This recommendation must be based on a belief that, sooner or later, people would learn that the standard discretionary policy is not being implemented and that they would therefore stop expecting this policy, since otherwise it would be optimal to implement ordinary discretion, not commitment. However, if this learning takes place, it must be because policy affects people's expectations of future policy, and, in this case, the optimal policy depends on the exact effect it has on expectations. If the effect is large enough, the optimal policy can be similar to the optimal commitment policy, but it is small; or if people discount the future too heavily, the optimal policy can be closer to the standard discretionary solution, and

[14] See Barro (1986), McCallum (1990), and Rogoff (1987) for discussions on why a reputation might be insufficient for the policymaker to achieve the desired commitment.

implementing the optimal commitment plan can give worse results in terms of achieving the policy objectives than standard discretion.[15]

5. Conclusions

The optimal time-consistent policy depends on the public's policy forecasts. We know how people should forecast policy assuming they are rational and have perfect and complete information; they should replicate the policymaker's decision process. When forecasts are formed this way, present and past policy is irrelevant for expected future policy, as is assumed in the standard discretionary solution. However, if forecasters do not have all the information required to reproduce the policymaker's decision process, they have to use present and past policy to learn about this process and predict policy. When they do so, time-consistent policy can reap the benefits of influencing expectations, just as in the commitment solution, to a degree determined by the exact learning/forecasting process and the level of time discounting in the economy. It is therefore important for policymakers to know how the public predicts policy and how their actions affect these predictions. This way, policymakers can influence policy expectations even if they are not able to commit and credibly preannounce their actions. If, in addition, policymakers can influence the way in which individuals forecast policy—for example, by providing more or less information about the policy decision process—they can affect the time-consistent solution and thereby how well discretionary policy achieves their objectives.

Appendix

This appendix outlines the proof of the proposition that the optimal (stationary) discretionary policy and the optimal commitment

[15] This discussion also applies to timeless-perspective commitment policies like the ones proposed by Svensson and Woodford (2003), Woodford (1999a, 2003), or Jensen and McCallum (2002), where the policymaker commits to a procedure for determining policy instead of a specific equation. The proposed timeless-perspective policies are only optimal if the policymaker can credibly commit and thereby dictate individuals' policy expectations. If policy, or commitment, is learned over time, the optimal timeless-perspective policies can differ from the ones proposed in the references above.

policy converge as β approaches 1, when expectations are such that the policy action implemented in the present period is expected to be implemented in the next period too. With a policy function

$$\pi_t = f_t u_t + g_t y_{t-1} \tag{29}$$

and expectations

$$f^e_{t,t+1} = f_t, \quad \forall t \tag{30}$$

$$g^e_{t,t+1} = g_t, \quad \forall t, \tag{31}$$

the reduced form of the model is given by (29) and

$$y_t = \frac{(1-\beta\phi)f_t - 1}{\alpha + \beta g_t} u_t + \frac{g_t}{\alpha + \beta g_t} y_{t-1}, \tag{32}$$

which we can use to write the policy objective in (1) as

$$E \sum_{t=0}^{\infty} \beta^t \left((f_t u_t + g_t y_{t-1})^2 + \omega \left(\frac{(1-\beta\phi)f_t - 1}{\alpha + \beta g_t} u_t + \frac{g_t}{\alpha + \beta g_t} y_{t-1} \right)^2 \right). \tag{33}$$

Minimizing with respect to f_0 and g_0, and then assuming stationarity ($\cdots = f_{-1} = f_0 = f_1 = \cdots$ and $\cdots = g_{-1} = g_0 = g_1 = \cdots$) and letting β converge toward 1, we find that the optimal discretionary policy converges toward

$$\pi_t = \lambda u_t + \rho y_{t-1}, \tag{34}$$

where

$$\lambda = \frac{\rho(2\omega(1-\phi^2) - \alpha^2\phi^2) + \alpha\omega(1+\phi^2)}{\rho(1+\phi)(\alpha^2\phi + (\alpha^2 + 2\omega)(\phi - 1)^2) + \alpha((3 - \phi - \phi^2 - \phi^3)\omega + \alpha^2)} \tag{35}$$

$$\rho = -\frac{1}{2}\alpha + \frac{1}{2}\sqrt{\alpha^2 + 4\omega}. \tag{36}$$

Combining equations (2), (3), and (5), one can show that when β converges toward 1, the optimal commitment policy converges toward

$$\pi_t = \frac{\omega}{\alpha^2 + \omega(1 - (r - 1 + \phi))} u_t + \frac{\omega}{\alpha}(1 - r) y_{t-1}, \tag{37}$$

where

$$r = \frac{\alpha^2 + 2\omega - \alpha\sqrt{\alpha^2 + 4\omega}}{2\omega}. \tag{38}$$

Showing that the discretionary policy given in (34)–(36) and the commitment policy given by (37)–(38) are identical is mere algebra.

References

Barro, R. J. 1986. "Reputation in a Model of Monetary Policy with Incomplete Information." *Journal of Monetary Economics* 17 (1): 3–20.

Bernanke, B. S., and F. S. Mishkin. 1997. "Inflation Targeting: A New Framework for Monetary Policy?" *Journal of Economic Perspectives* 11 (2): 97–116.

Bullard, J., and K. Mitra. 2002. "Learning about Monetary Policy Rules." *Journal of Monetary Economics* 49 (6): 1105–29.

Calvo, G. A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Carlstrom, C. T., and T. S. Fuerst. 2004. "Learning and the Central Bank." *Journal of Monetary Economics* 51 (2): 327–38.

Clarida, R., J. Galí, and M. Gertler. 1999. "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37 (4): 1661–1707.

Currie, D. A., and P. Levine. 1993. *Rules, Reputation and Macroeconomic Policy Coordination.* Cambridge University Press.

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

———. 2003. "Expectations and the Stability Problem for Optimal Monetary Policies." *Review of Economic Studies* 70 (4): 807–24.

———. 2006. "Monetary Policy, Expectations and Commitment." *Scandinavian Journal of Economics* 108 (1): 1–182.

Gaspar, V., F. Smets, and D. Vestin. 2005. "Optimal Monetary Policy under Adaptive Learning." ECB Working Paper No. 644.

Jensen, C., and B. T. McCallum. 2002. "The Non-Optimality of Proposed Monetary Policy Rules Under Timeless-Perspective Commitment." *Economics Letters* 77 (2): 163–8.

Judd, H. P., and G. D. Rudebusch. 1998. "Taylor's Rule and the Fed: 1970–1997." *Economic Review* (Federal Reserve Bank of San Francisco) 3:3–16.

Kydland, F. E., and E. C. Prescott. 1977. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85:473–92.

McCallum, B. T. 1990. "Inflation: Theory and Evidence." In *Handbook of Monetary Economics*, Vol. 2, ed. B. M. Friedman and F. H. Hahn. North Holland.

———. 1995. "Two Fallacies Concerning Central-Bank Independence." *American Economic Review* 85 (2): 207–11.

———. 1996. "Inflation Targeting in Canada, New Zealand, Sweden, the United Kingdom, and in General." NBER Working Paper No. 5579.

McCallum, B. T., and E. Nelson. 2000. "Timeless Perspective vs. Discretionary Monetary Policy in Forward-Looking Models." NBER Working Paper No. 7915.

Molnar, K., and S. Santoro. 2005. "Optimal Monetary Policy when Agents are Learning." Working Paper, Universitat Pompeu Fabra.

Orphanides, A. 2003. "Historical Monetary Policy Analysis and the Taylor Rule." *Journal of Monetary Economics* 50 (5): 983–1022.

Orphanides, A., and J. Williams. 2005. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy." In *The Inflation-Targeting Debate*, ed. B. Bernanke and M. Woodford. Chicago: University of Chicago Press.

Rogoff, K. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (4): 1169–89.

———. 1987. "Reputational Constraints on Monetary Policy." *Carnegie-Rochester Conference Series on Public Policy* 26:141–82.

Rotemberg, J. J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *Review of Economic Studies* 49 (October): 517–31.

Rotemberg, J. J., and M. Woodford. 1999. "Interest Rate Rules in an Estimated Sticky Price Model." In *Monetary Policy Rules*, ed. J. B. Taylor. Chicago: University of Chicago Press.

Svensson, L. E. O. 1997. "Optimal Inflation Targets, 'Conservative' Central Banks, and Linear Inflation Contracts." *American Economic Review* 87:98–114.

———. 1999. "Inflation Targeting as a Monetary Policy Rule." *Journal of Monetary Economics* 43 (3): 607–54.

Svensson, L. E. O., and M. Woodford. 2003. "Indicator Variables for Optimal Policy." *Journal of Monetary Economics* 50 (3): 691–720.

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

———. 1999. "A Historical Analysis of Monetary Policy Rules." In *Monetary Policy Rules*, ed. J. B. Taylor, 319–41. Chicago: University of Chicago Press.

Woodford, M. 1999a. "Commentary: How Should Monetary Policy Be Conducted in an Era of Price Stability?" In *New Challenges for Monetary Policy*, 277–316. Kansas City, MO: Federal Reserve Bank of Kansas City.

———. 1999b. "Optimal Monetary Policy Inertia." NBER Working Paper No. 7261.

———. 2003. *Interest and Prices*. Princeton, NJ: Princeton University Press.