



In Noise We Trust? Optimal Monetary Policy with Random Targets

Ethan Cohen-Cole and Bogdan Cosmaciuc

Abstract:

We show that a monetary policy in which the central bank commits to a randomized inflation target allows for potentially faster expectations convergence than with a fixed target. The randomized target achieves faster convergence in particular in transition environments: those demonstrating either particularly high or low inflation.

Keywords: monetary policy, asymmetric information, Bayesian rational expectations, commitment

JEL Codes: E52, E61, E42

Ethan Cohen-Cole is a financial economist at the Federal Reserve Bank of Boston. His email address is ethan.cohen-cole@bos.frb.org. Bogdan Cosmaciuc is unaffiliated.

We are grateful for suggestions and comments from Patrick de Fontnouvelle, John Morgan, Giovanni Maggi, Rodolfo Manuelli, Jonathan Parker, Helene Rey, Eric Rosengren, Ananth Seshadri, and Kenneth West. We are also grateful for research assistance provided by Nick Kraninger and Jon Larson. All errors remain our own.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bos.frb.org/economic/wp/index.htm>.

The views expressed in this paper are solely those of the authors and do not reflect official positions of the Federal Reserve Bank of Boston or the Federal Reserve System.

This version: October 2006 (10.11.6a)

I. Introduction

This paper addresses the intersection of two lines of debate in monetary policy. The first of these is the degree to which central banks have the ability to influence the real economy in the short run. Perspectives on this topic are well-known. For discursive and illustrative purposes, and acknowledging recent advances, we use a particularly extreme position, complete rational expectations. Since our conclusions can be extended logically into less strict formulations, the use of the monetarist framework is appropriate. Under such conditions, Taylor's (1975) well known conclusion is that the monetary authority can impact the real economy only during an expectations formation period at the onset of a new policy regime. At all other times, monetary policy has well-understood limited effects on the real economy.

The second debate is the appropriate monetary policy framework or mechanism to use in transition environments. Such transitions are clearly important in both persistently high and persistently low inflation environments. Case examples of these include Latin America in the 1980s, Japan in the late 1990s, and Zimbabwe today. Current conclusions are that an inflation targeting regime can contribute to a relatively rapid inflation decline in conditions of hyperinflation (see, for example, Bernanke et al. 1999). The practice involves ever bolder declarations of inflation targets, and more highly skilled central bankers to implement the plans. In each of these cases, credibility plays a crucial role in determining the efficacy of the monetary policy regime—and arguably the appropriate method of monetary policy action. In high inflation environments, the population will be unlikely to believe announcements of low-inflation policy. Even under the assumption of fully rational expectations, if the central bank's promise is not credible, it will take time until the population's expectations converge to the actual central bank target. Similarly, given the Japanese populace's expectations of the central bank's low-inflation prejudice, they are unlikely to believe the bank's desire to raise inflation (see, for example, Krugman 1998). In practice, this

“transitional expectations” period is a common occurrence in countries undergoing regime changes and/or in countries with “sticky” expectations (see Sachs 1990).

In our model, we find that randomization of monetary policy in transition cases provides a better response mechanism to the population’s behavior, because it eliminates a portion of the information asymmetry problem. Through such a policy, the monetary authority can “encourage” rapid convergence without an increase in credibility or reputation. Consider the following example. Suppose a new government takes over monetary policy in a developing country with a “traditional” inflation rate of 40-50 percent a year. The new government “knows” that the optimal inflation level is 5 percent. However, since a 5 percent target would not be credible, the government would potentially have to announce a higher target or to implement a relatively slow convergence plan in an inflation targeting framework.¹ We show that if the government commits to a so-called “random target,” the situation can change. The government would enact a policy whereby it commits to implement the targeting decisions of a computer run by some independent third party.² The computer in each period would generate a random number based on a target around 5 percent, with some small variance.³ Notice that the government’s commitment to a randomized policy is credible, because there is no *a priori* reason why the government is better off under the fixed target. Notice as well the distinction here from *inflation targeting* in that the stated commitment of the central bank is to a target that is itself moving, not to a band around a fixed target. Thus, the remaining credibility question is the ability of the bank to implement the computer-generated policy. We will show that this mechanism allows for faster expectations convergence than a fixed target.

Monetary policy randomization has potential application in at least three cases. Two of these cover the

¹ We argue below that with non-zero discounting, the optimal inflation level is greater in an economy without government credibility.

² The third party is used to illustrate the feasibility of a trustworthy random draw. The computer could even be placed within the offices of a “trusted” institution outside the country.

³ In fact, the model will specify that the computer generate a random walk target, though for the sake of explanation, random is sufficient at this stage.

“high” inflation situation, hyperinflation and persistent inflation. In situations of hyperinflation, a faster convergence rate is clearly desirable, especially since most “losses” occur in the short run. The Japanese deflation-cum-liquidity trap case is also an example in which expectations are too sticky (this time downwards) and in which, for structural and credibility reasons a change in inflationary expectations is needed where traditional monetary policy has failed.

We have focused our discussion on such “extreme” cases because we believe lack of credibility to be an important limitation in the availability of monetary policy options for such countries. While credibility can be built over time, these countries may find themselves too often in a “transitional” period, in which credibility building is impractical. Our solution prescribing inflation targeting with a random target rate seems to provide at least some improvement in the efficacy of monetary policy in such cases.

The existing literature on gradual vs. “shock” approaches to transition draws either on credibility, or on trend or targeting approaches. Guillermo Calvo’s (1983) model of stochastic price adjustments suggests that immediate disinflation can be costless, and Laurence Ball (1994) suggests that a credible and rapid disinflation can even prompt a boom. Thomas Sargent (1994) has suggested that four historic cases of hyperinflation were halted through *credible* institutional change. While these go some distance in resolving new Keynesian concerns of wage and price friction in economic adjustment, they don’t capture the credibility constraints at hand in transition environments. The practical constraints implied in targeting approaches are briefly discussed above.

At the intersection of the two debates above, we discuss the interaction between central banks and the population under the framework of economic conditions of very high or very low inflation. As alluded to above, these cases are reasonably considered to have a set of characteristics dissimilar to “normal” times, and consequently a set of monetary policy actions that may reasonably be different as well. Of these, we point to two in particular. First, use of credibility as a monetary-policy enhancement tool can potentially

result in the two cases mentioned above. In addition to being a successful mechanism for the direction of expectations, it can also fail (as in Latin America in the 1980s), or be too successful (Japan in the late 1990s). In each of these cases, the actual actions of the central bank take on the central role in the formation of expectations; thus, a monetarist style adjustment through intent announcement may be sub-optimal. Similarly, gradualist approaches that rely on perfect credibility will be constrained. Second, the use by the monetary authorities of a loss function that ignores the real cost of inflation can be inappropriate in some contexts.⁴ See Woodford (1999) for a derivation of monetary policy loss functions from microeconomic foundations. See also Walsh (2003) and Clarida et al. (1999), where using the squared-deviation from optimal inflation as a determinant of social losses may not capture the importance of the expected inflation's deviation from the socially optimal level. This particular issue is not a new debate; see Rogoff (1985), for example, for discussion of anticipated and unanticipated inflation. This can be seen easily in the limiting case of hyperinflation, where the real costs to the economy of high inflation cannot be measured solely by the deviation from expected or target inflation. There is some real cost to the inflation itself.

With such a real cost, the government must optimize some social loss function that includes not only the cost of deviation from normal inflation, but also the cost of “sticky” inflationary expectations, that is, the cost of its own lack of credibility. It must do so in a manner that is consistent with the population's continued attempts to predict inflation. In what follows, we will construct a governmental loss function through which the government targets the population's inflationary expectations.⁵ Such a social loss function seems more appropriate to situations of hyperinflation or chronic deflation than to “normal”

4 The most common generalization of this is, $\Lambda_t = (n_t - n_t^*)^2 + \chi(\pi_t - \pi^*)^2$, where n is some macro-variable (real GDP, unemployment, etc.), n_t^* is the target level of the variable at time t , π is a measure of inflation, and π^* is the target level of inflation (Svensson, 1999).

5 By targeting public inflationary expectations, we simply mean that these expectations become part of the government's loss function in a way to be described in more detail below. We do not mean, as do Bernanke and Woodford (1997), that the government uses public expectations as a signal for some disturbance in the economy that the public alone can observe; in our framework the government has full information.

times. We then solve for the government's optimal inflation target in light of a population with rational expectations, but with incomplete information.

Two situations are discussed: a "classic" one in which the government commits to a particular inflation target; and our proposal for a "randomized" target, in which the government pre-commits to a particular randomization of its inflation target. In either case, the target inflation rate is unknown to the public. We cover in Section II the rationale behind the ability of the government to obtain real gains through monetary policy, even under rational expectations, and we address the effect on expectations formation of increased uncertainty in inflation policy. In order to achieve this increased uncertainty, we use a "random target." Specifically, this refers to a periodic change in the target inflation rate of the country's monetary authority. In each time period the target is moved up or down a "notch." These "notches" are normally distributed. The result is a random walk in the target inflation rate. Section III describes our model of monetary policy equilibrium in the case of a fixed monetary target. Subsequently, Section IV explains optimal monetary policy when the government chooses a random walk inflation target. The benefits of the randomized policy are discussed. We show simulated results in Section V and conclude in Section VI.

II. A transitional expectations framework

In a world of full information and fully rational expectations, monetary policy has limited effect. Indeed, economic agents, given full information, can rationally determine the inflation that the government wants to achieve, and adjust their prices and wages accordingly. This is clearly not the case in what Taylor (1975) describes as a transitional period towards rational expectations, a period during which the central bank's target inflation rate is either not known or not credible. Under such conditions of asymmetric information, Taylor shows that monetary policy has a significant impact in the short run, that is, until expectations fully adjust to reflect the actual chosen rate.

Taylor uses an output-inflation targeting loss function (that is, the loss is measured in squared deviations from some target output level and some target inflation rate, both of which are held constant across time). In particular, Taylor is concerned with the potential gains from generating some extra “unexpected” inflation. Indeed, suppose that for some reason (fiscal policy distortions, etc.) the current rate of inflation (which is fully known by the public and thus also expected by the public) is below the optimal rate. If the central bank unexpectedly moves to target the optimal rate, it will obtain some output gains in the transitional period until the population learns the new rate. The central bank can capitalize on creating “confusion” about the true inflation rate.

At the outset of this project, our hypothesis about the benefits of a random target was the opposite of our current conclusion. Consider the following logic: if an unexpected change in the inflation target generates some social gains (or reduce losses due to transition), then increasing the transition period should increase those gains (decrease losses). Indeed, if the central bank could “confuse” the public for a longer period, it would be better off than under the Taylor model. One might expect, as we did, that randomization would create this additional confusion. In fact, further randomization of the monetary rule led to an *increased* ability of the population to discover the target rate, and thus decreased the ability of the central bank to generate output gains during the period of transitional expectations.⁶ Before illustrating this point, let us walk through Taylor’s construction.

First, we explain, as Taylor has, the logic behind the inefficacy of monetary policy under rational expectations. Suppose the relationship between inflation and unemployment can be shown as⁷

$$\pi(t) = \phi(u(t)) + x(t), \quad \phi'(\cdot) < 0, \quad \phi(u^*) = 0, \quad u^* > 0. \quad (1)$$

6 The surprising result led to the formulation of the remainder of this paper: increased randomization can facilitate convergence to the government target rate—a particularly useful property in some situations.

7 For notational convenience and comparability purposes, we have generally adopted Taylor’s notation.

Here $\pi(t)$ and $x(t)$ are the actual and expected rates of inflation, respectively, at each time t . Taylor further assumes that the monetary authority has direct control over the inflation rate through control of the money supply,⁸ so that $\pi(t)$ is effectively a choice variable. Long-run effects on employment are not possible in (1), but short-run effects are. This is a well-known argument made by Phelps, who considers that inflationary expectations adapt according to the linear rule:

$$dx(t) = \beta[\pi(t) - x(t)]dt. \quad (2)$$

The government chooses an inflation target that maximizes:

$$\int_0^{\infty} e^{-\rho t} W(x, u) dt, \quad (3)$$

where $W(x, u)$ is the instantaneous social utility rate, u is the unemployment rate, and β and ρ are discount rates. The short-run effects mentioned above are obtained when $\pi(t)$ is systematically greater than $x(t)$. Rational expectations theory suggests that over a period of time, agents understand that their expectations based on (2) are biased and modify accordingly. Taylor explains that under rational expectations the formulation appears as:

$$x(t|t_0) = E[\pi(t)|I(t_0)]. \quad (4)$$

Here $I(t_0)$ represents information available at time t_0 . This explains that under a deterministic policy with no uncertainty, $\phi[u(t)] = 0$ and $u(t) = u^*$ for all t (Taylor 1975).

⁸ We later relax this assumption in order to add reality and as the mechanism for further uncertainty.

III. Monetary policy commitment to a fixed target

From the above argument, it follows that some information asymmetry is necessary in order for the short-run Phillips curve to have an effect on output. We will thus model the monetary policy effects of central bank policies as an incomplete information game between the central bank and the public. The game is set in continuous time, so that stochastic random variables are used to model random disturbances.⁹ The public does not have knowledge of the government's methods or desires, but is able to learn from the government's prior actions. This is an abstraction that seems appropriate in a situation of low government credibility. Essentially it assumes zero credibility – any announcement is equivalent to none at all. The public derives all new information from the actions themselves. Taylor illustrates such a system: during the time period when old beliefs mix with new information, the rational expectations hypothesis no longer holds, and monetary policy has real and significant effects.

THE MODEL

A. The Government (and Central Bank)

The government commits itself at time $t = 0$ to a target inflation rate for the economy, which we denote by μ . Once the target inflation rate is set, the economy generates an actual inflation rate $\pi(t)$. We assume that the government does not have perfect control over the actual value of the inflation rate that is generated. More precisely, $\pi(t)$ is modeled as a normally distributed random variable, where $E[\pi(t)] = \mu$ (thus, the government can control the average inflation rate it generates, but not the actual value).

⁹ Taylor models the information asymmetry by having the government cum central bank follow an inflation target that is *unknown* to the public. We consider this construction to be a particularly realistic one.

Furthermore, we assume that $\text{var}[\pi(t)] = \sigma$, where σ is a known constant depending on the characteristics of the interaction between the central bank and the economy.¹⁰

To translate these ideas into stochastic calculus equations that can be solved, let us define price level at each time t as $p(t)$, and then let $z(t) = \log[p(t)]$. When $\sigma = 0$, the government has complete control over inflation, and by definition $\pi(t) = \mu(t)$. With $\sigma > 0$, inflation is generated through the diffusion process:

$$dz(t) = \mu dt + \sigma dv, \tag{5}$$

where $v(t)$ is a standard Wiener process with a zero mean and unit variance (that is, at each time t , $dv(t)$ is randomly chosen from a normal distribution with mean zero and unit variance).

Once inflation is generated, the public observes it and (as discussed below) forms some inflationary expectations $x(t) = E_{\text{public}}[\pi(t)]$. We assume that the government knows the process through which these expectations are formed, so the government can infer $x(t)$.

The government also has a negative payoff function or loss function $L(t)$ (see Svensson 1999).¹¹ At this point, we will not specify L ; suffice it to say that L depends on two arguments, actual inflation $\pi(t)$ and some other target variable (unemployment, or real GDP, etc.) which depends again on inflation and on public inflationary expectations, $x(t)$. Since by choosing $\mu(t)$ the government determines both $\pi(t)$ and $x(t)$, the government's goal is to choose a target μ that minimizes the expected value of the loss function. Indeed, both $\pi(t)$ and $x(t)$ are stochastic, and the government can control only their distribution, not the

¹⁰ A lower σ implies a more precise monetary policy. In particular, we assume that σ is the minimum possible variance; that is, the government cannot improve the precision with which it generates inflation beyond this point.

¹¹ As discussed above, Taylor uses a social welfare function instead of a loss function. We use the loss function because the terminology is compatible with more recent papers on monetary policy targeting.

actual realized values. In a full commitment regime, the government chooses a value μ that minimizes the

(discounted) expected future losses, $\int_0^{\infty} e^{-\delta t} E[L(t)]dt$.

B. The Public

At each time t , the public observes the inflation rate $\pi(t)$ generated in the economy. The public, however, does not have a full knowledge of the government's target inflation rate μ . More precisely, consider that at time zero the public expects the government's inflation target to be $\mu(0) = \mu_0$, different from μ . In other words, $E_{\text{public}}[\mu(0)] = x(0) = \mu_0$. Also, the public belief about the accuracy of this initial guess is such that $\text{var}_{\text{public}}(\mu(0)) = \sigma_0$. We assume throughout that the government knows both μ_0 and σ_0 . The fact that the initial public expectations are different from the actual government target may result, as Taylor emphasizes, from the fact that the initial basis for expectations of inflation setting is a combination of expectations from a new regime and information from an old regime.¹² Moreover, the same situation can result from a lack of credibility of the “new” government. Such a construction captures the reality that an emerging market government’s open commitment to a particularly “low” inflation level, μ , may not be entirely credible.

The public is assumed (as in Bernanke and Woodford 1997) to have a quadratic loss function, depending on the square of the difference between its inflationary expectations $x(t)$ and the true inflation $\pi(t)$. It follows that the public's goal is to generate at each time $t > 0$ the most appropriate inflationary expectations. The public observes the actual inflation generated at each time t , and based on it and on its current inflationary expectations "infers" the most likely inflation target chosen by the government and

¹² Under complete rational expectations, the population would be able to calculate the government’s initial action perfectly.

adjusts its own expectations. Note that this method assumes that the public knows σ , but not μ .¹³ Under these conditions, it can be shown that the population's inflationary expectations are given by the following stochastic differential equation:

$$dx(t) = \frac{1}{t + \sigma^2 / \sigma_0^2} [dz(t) - x(t)dt]. \quad (6)$$

Note that since the above is a stochastic differential equation, the solution will also be stochastic. This implies that even under perfect information, the government will not be able to have deterministic knowledge of the expectations path $x(t)$. However, as it will become apparent shortly, the government only needs to know the mean and the variance of the populations' expectations.¹⁴

If we define:

$$\begin{aligned} \bar{x}(t) &= E_{government}[x(t)] \\ \tilde{x}(t) &= \text{var}_{government}[x(t)] = E_{government}[x(t)^2] - \bar{x}(t)^2, \end{aligned} \quad (7)$$

then, from the theory of linear stochastic differential equations (Arnold 1974) it follows that these deterministic variables are given by the following deterministic differential equations:

$$d\bar{x}(t) = \left[-\frac{1}{t + \sigma^2 / \sigma_0^2} \bar{x}(t) + \frac{\mu}{t + \sigma^2 / \sigma_0^2} \right] dt, \text{ with } \bar{x}(0) = \mu_0, \text{ and} \quad (8)$$

$$d\tilde{x}(t) = \left[-\frac{2}{t + \sigma^2 / \sigma_0^2} \tilde{x}(t) + \frac{\sigma^2}{(t + \sigma^2 / \sigma_0^2)^2} \right] dt, \text{ with } \tilde{x}(0) = \sigma_0^2. \quad (9)$$

Solving the above equations for the mean and variance of the inflation expectations we get:

¹³ In particular, it follows that $\sigma_0 \geq \sigma$, because σ imposes a natural limit on the precision of the public's "guess."

¹⁴ This is the case because the loss functions that are currently used are quadratic forms.

$$\bar{x}(t) = \mu + \frac{\mu_0 - \mu}{t + \sigma^2 / \sigma_0^2}, \text{ and} \quad (10)$$

$$\tilde{x}(t) = \frac{\sigma^2}{t + \sigma^2 / \sigma_0^2}. \quad (11)$$

From this knowledge of the expected path of people's expectations, the government can set μ appropriately to minimize some given loss function (which we outline below). The population, as implied above, simply attempts to have full understanding of the inflation rate in order to set prices and wages to match.

C. The loss function and the resulting Bayesian equilibrium

The loss function suggested by Taylor in order to derive the government's optimal policy is a traditional one (See Kydland and Prescott (1977) and Barro and Gordon (1983)), associated with a target output level (as well as inflation):

$$L_t = (\pi_t - \pi^*)^2 + \lambda(y_t - y^*)^2. \quad (12)$$

However, such a function does not necessarily suggest that rapid transition is optimal; on the contrary, the government can benefit if the transition is delayed. To the extent that unanticipated inflation impacts unemployment, and thus output, the benefits of a rapid convergence to π^* depend on the magnitude of output loss and λ .

Such logic, however, seems to reflect poorly the conditions of emerging markets with systemically high inflation or the case of Japanese deflation. Clearly the output losses sustained by continued deflation in

