

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 581

May 1997

INFLATION REGIMES AND INFLATION EXPECTATIONS

Joseph E. Gagnon

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the world wide web at www.bog.frb.fed.us.

Inflation Regimes and Inflation Expectations

Joseph E. Gagnon*

Abstract

There has been much talk in the popular press about the difficulty of attaining credibility in the bond markets for the low-inflation policies that have been adopted by a number of central banks in recent years. This credibility problem is particularly severe for those countries that have a history of high inflation. Gaining credibility is often viewed in the context of learning by the public about the central bank's true intentions. However, this paper argues that a more important aspect of credibility--at least for long-term inflation expectations--may be public views about how future changes in personnel, electoral results, or economic shocks may affect central bank behavior. In other words, there is always a positive probability that the current regime will end. Views about the nature of possible future regimes are likely to be influenced by observed past regimes.

Keywords: central bank, credibility, inflation target, monetary policy

*Acting Director of the Office of Industrial Nations and Global Analyses, Department of the Treasury. This paper was written while I was employed by the Board of Governors of the Federal Reserve System and while I visited the Reserve Bank of Australia. I thank those institutions for their support. I also thank David Bowman, Gordon DeBrouwer, Jeff Dominitz, Jon Faust, David Gruen, Dale Henderson, Andrew Levin, and Jenny Wilkinson for their helpful comments. This paper represents the views of the author and should not be interpreted as reflecting those of the Department of the Treasury, the Board of Governors of the Federal Reserve System, the Reserve Bank of Australia, or other members of their staffs.

I. Introduction and Summary

Average inflation rates in industrial countries have fallen substantially since the 1980s. In several cases, countries that experienced higher inflation than the industrial-country average during the 1970s and 1980s have achieved lower than average inflation in the 1990s. In most cases, these recent low-inflation countries have not experienced a commensurate reduction in their long-term interest rates relative to the industrial-country average, suggesting that long-term inflation expectations have not moved in proportion to recent inflation outcomes.¹

To provide concrete examples, examine two bilateral comparisons: the United States vs. Canada and Australia vs. New Zealand. According to the OECD *Economic Outlook* (December 1996) consumer price inflation (measured by the consumption deflator) was lower in Canada than in the United States for every one of the past 5 years. The average inflation rate over this period was 1.3 percent in Canada and 2.6 percent in the United States. Despite this consistent record of lower inflation, the yield on 10-year government bonds was identical across the two countries at 6.7 percent at yearend 1996. Similarly, the inflation rate was lower in New Zealand than in Australia for every one of the past 5 years, and the average over this period was 1.7 percent in New Zealand and 2.0 percent in Australia.² Nevertheless, at yearend 1996, 10-year interest rates were identical in the two countries at 7.7 percent.

One explanation for these findings is that long-term inflation expectations depend on a

¹See Ammer and Freeman (1995) and Freeman and Willis (1995).

²Although the difference in recent inflation outcomes between these countries has been relatively small, the Reserve Bank of New Zealand has a target band for inflation of 0 to 3 percent (formerly 0 to 2 percent) whereas the midpoint of the Reserve Bank of Australia's objective is 2-1/2 percent. Thus, one might expect to find lower inflation expectations in New Zealand since the center of its target band is 1 percentage point lower than Australia's. For more on the specification of inflation targets in different countries see Reserve Bank of New Zealand (1996) and pp. 108-110 of International Monetary Fund (1996).

long history of past inflation--more than just the past 5 years. Indeed, during the 1980s, inflation averaged 6.2 percent in Canada versus 5.3 percent in the United States, and inflation averaged 11.9 percent in New Zealand versus 8.2 percent in Australia. In each case the country with lower recent inflation experienced higher inflation over a long period in the past. The effect of past inflation over a long horizon can also explain the higher 10-year interest rates in Australia and New Zealand versus the United States and Canada.³

More generally, there is evidence documented by Gagnon (1996) that nominal long-term interest rates are strongly correlated with both recent inflation and past inflation over a long horizon. This correlation holds both across countries and within countries over time. One explanation for this correlation is that long-term inflation expectations are influenced by a long history of past inflation. Gagnon (1996) also presents direct evidence for this hypothesis from the spread between nominal and indexed bond yields.

This paper develops a theoretical framework to explain these empirical findings. The basic idea is that since the collapse of the gold standard earlier this century, central banks in most countries have been characterized by periodic changes in policy regime. At the most basic level, regime changes are associated with changes in the central bank governor or political party in power, depending on the institutional independence of the central bank. Other factors may give rise to regime changes: Evolving theories about economic behavior

³Another explanation for different nominal long-term interest rates is that real long-term interest rates may differ across countries. Real rates may differ due to different risk premia or to different levels of real exchange rates relative to their long-run equilibria. Both of these factors may be contributing to the interest rate differentials observed here, especially the differential between Australia and New Zealand on the one hand and the United States and Canada on the other hand.

may lead to new views on the optimal conduct of monetary policy. Or, extreme social or economic shocks may necessitate a persistent shift in monetary policy. However, in general, it is not useful to think of the regime changing with every shock. Instead, regimes are viewed as implicit or explicit rules governing the behavior of monetary policy in response to ordinary shocks.

One important outcome of different monetary regimes is different average inflation rates across regimes. When agents are considering expected inflation over a long future horizon, they must factor in the possibility that the current regime will not survive over the horizon in question. Recent inflation rates may provide a good forecast of future inflation rates if the current regime survives, but they may not provide a good forecast if the current regime is replaced. To factor in the effect of a potential new regime, agents may base their forecasts on their experience of past monetary regimes over a long horizon.

To take the example of New Zealand once again, the Reserve Bank of New Zealand's previous central target of 1 percent inflation was lower than the inflation rate in every year but one since World War II. Thus, if agents were considering the possibility of a new inflation regime in the future, it seems likely that they would expect any new regime to have average inflation greater than 1 percent. Even if agents believed that the Reserve Bank would achieve its target of 1 percent inflation in the current regime, they would have to factor the possibility of a change to a higher-inflation regime into their expectations, thereby raising expected future inflation above 1 percent. The importance of regime changes for expectations of future inflation was highlighted in the context of the recent New Zealand election. Some of the parties argued for a higher-inflation regime and no party argued for lower inflation. In

the event, the central inflation target has been raised slightly, from 1 to 1.5 percent. Moreover, there was a possibility that an even greater increase in the inflation target might have resulted after the election.

In addition to explaining long-run inflation expectations in bond markets, a model with regime changes can explain the peculiar time-series properties of actual inflation over the postwar period. For most industrial countries, it is difficult to reject a unit root in the inflation rate. Yet, recent studies have found some evidence of weak mean reversion of inflation rates over long horizons. It is well known that structural breaks in an otherwise stationary series can induce apparent unit roots into the series. If inflation has undergone a small number of regime shifts in the postwar period, it would be difficult to reject a unit root. However, if the regime shifts themselves were around a constant average inflation rate, one would expect to find some evidence of mean reversion in inflation. Moreover, within relatively long-lasting regimes it should be possible to reject a unit root, which may explain the apparent stationarity of inflation over certain subsamples.

Finally, the possibility of regime shifts leads to highly asymmetric distributions of future inflation rates. The asymmetric distribution of future inflation may explain the asymmetric distribution of survey responses on future inflation expectations. Moreover, the asymmetric distribution of future inflation may explain the frequently large discrepancies between surveys of inflation expectations and implied inflation expectations in bond yields. If survey respondents report the most likely outcome, and bondholders care about the average outcome, then the discrepancy between different measures of inflation expectations would be resolved.

II. Literature Review

1. Models of Inflation

The literature on models of inflation is too voluminous to review in depth here. For the purposes of this paper, we are less interested in the dynamic interactions of inflation and other variables over the business cycle and more interested in the determination of the rate of inflation in the long run. Driffill, Mizon, and Ulph (1990) and Woodford (1990) provide surveys of the theoretical and empirical literature on the costs and benefits of inflation. Unfortunately, the only conclusion that comes close to achieving a consensus is that inflation variability *per se* is harmful and that central banks should stabilize the inflation rate to the extent that they can without inducing costly variability in other economic variables. No consensus exists on the optimal steady-state rate of inflation.

Fischer (1990) surveys the literature on the institutional framework of monetary policy and the determination of the long-run inflation rate. The treatment is purely theoretical and focuses on the issue of "rules versus discretion." A basic conclusion is that a pure rule-based policy has not existed since the Gold Standard, and many would argue that even under the Gold Standard there was a substantial discretionary aspect to monetary policy. One drawback of discretionary policy setting is that no one has designed an institutional framework that indisputably avoids the potential inflationary bias created by the time inconsistency problem.⁴

More recently, attention has focused on the adoption of explicit inflation targets by a number of central banks. Walsh (1995) discusses the circumstances under which explicit

⁴The time inconsistency problem refers to the temptation for a central bank to induce extra output by creating more inflation than the public expects, even though it knows that this extra output cannot be sustained in the long run.

inflation targets and enforcement clauses in the central bank governor's contract are optimal. For a brief review of the international policy debate, see pp. 108-110 of International Monetary Fund (1996). At this stage it appears to be too soon to conclude much about the desirability and durability of inflation targetting.

Empirical analyses of the long-run properties of inflation rates have generally occurred in the context of the real interest rate literature. See, for example, Rose (1988) and Mishkin (1992). Using data from the entire postwar period, one cannot reject a unit root in inflation for most industrial countries using standard augmented Dickey-Fuller tests. However, for many countries one can reject nonstationarity of the inflation rate in certain subsamples.

Hassler and Wolters (1995) and Baillie, Chung, and Tieslau (1996) use the Phillips-Perron test and the KPSS test on postwar monthly inflation rates and reject both a unit root and stationarity for several countries. To reconcile these conflicting findings they turn to models with "fractional integration" and find that they are strongly supported by the data. Fractional integration allows for slow mean reversion that does not decay as rapidly as the asymptotically exponential pattern associated with standard autoregressive-moving average models. This slow mean reversion is termed "long memory."

Other researchers have sought to explain the apparent nonstationarity of inflation as the result of regime shifts in the mean and variability of the inflation rate. Chapman and Ogaki (1993), Bai and Perron (1995), and Hostland (1995) find significant evidence of regime shifts in U.S., U.K., and Canadian inflation. Evans and Lewis (1995), Ricketts and Rose (1995), and Simon (1996) estimate Markov switching models for inflation in the G-7 countries and Australia. At least two regimes are significant in all countries except Germany.

Occasional shifts in the inflation regime are more economically interpretable than fractional integration. Moreover, if there are only a small number of regimes that cycle back and forth, or if the regime-generating process is stationary, inflation rates will appear to have long memory, which is consistent with the fractional integration literature.

2. Evidence from Bond Markets

Instead of modeling the inflation process, a more direct way to learn about long-run inflation expectations is to examine the inflation premia in long-term bond markets. Fuhrer (1996) shows that the pure expectations theory of the term structure fits better when one allows structural breaks in the Fed reaction function, especially the implicit inflation target. Gagnon (1996) shows that the inflation premium in long-term interest rates is more closely correlated with a long backward average of inflation than a short backward average, implying that there is long memory in long-run inflation expectations and/or the inflation risk premium.

Focusing directly on countries that have announced explicit inflation targets, Ammer and Freeman (1995) and Freeman and Willis (1995) provide evidence that announced inflation targets have not been fully credible in terms of lowering long-term inflation expectations implicit in bond yields down to the official target range for inflation.

III. Models of Inflation Regimes

1. Complete Information

We begin with a model in which agents are fully informed. They know when a regime change has occurred. They know the inflation target of the current regime. They know the probability with which the current regime will end in the next period. And they

know the probability distribution of the inflation target across future regimes. We will relax some of these assumptions later.

π	Inflation rate
Π	Inflation target
ε	Temporary shock, $N(0, \sigma^2)$
θ	Regime shift, Bernoulli(q)
η	Target draw, $N(\mu, \kappa^2)$

$$\pi_t = \Pi_{t-1} + \varepsilon_t \quad (1)$$

$$\Pi_t = (1 - \theta_t)\Pi_{t-1} + \theta_t\eta_t \quad (2)$$

The inflation rate in each period is given by the inflation target effective in the previous period plus a random error. This lag reflects the conventional monetary transmission lag of roughly one year. The word "target" is used loosely to mean the expected inflation rate within a given regime. It does not necessarily imply that the central bank is officially or unofficially aiming for this inflation rate, only that this inflation rate is the expected average outcome of its policies. More generally, one might expect the variability and persistence of the temporary shock, ε , to be different across regimes. However, such an empirically realistic extension would add complexity to the model without altering the basic theoretical conclusion.

A regime shift ($\theta=1$) occurs with probability q . With probability $1-q$ there is no regime shift ($\theta=0$). The probability of a regime shift in each period determines the average length of regimes. The expected length of a regime is $1/q$ periods. An empirically reasonable range for inflation regimes is between 2 and 20 years, implying a value of q between .05 and .5. New inflation targets are drawn from a normal distribution with mean μ and standard deviation κ .

This specification of the regime-shifting mechanism is silent on the forces that end existing regimes and give rise to new regimes. One interpretation is that different central bank governors have different objectives with regard to the level and variability of inflation and other economic variables. These differences are not fully observable prior to the appointment of a new governor. The term of each governor is random and depends on both personal factors and the struggle of partisan politics. Alternatively, inflation regimes may be seen as the outcome of broader social and political forces that are manifested in opinion polls, public debates, and election results. Still another possibility is that regime shifts are triggered by certain large and persistent shocks, such as energy supply shocks.

One important feature of the models developed in this paper is that the process generating regimes is stationary. In the broad global and historical context this assumption is reasonable, as inflation rates tend to be bounded between a small negative and a large positive number. Hyperinflations are at most sporadic and not persistent, while hyperdeflations are unheard of. However, within these bounds it is conceivable that the process generating regimes has drifted over time. Such a drift may be the result of demographic or technological forces that operate on a time scale much larger than that of monetary policy regimes. Or, one may view the switch to fiat money standards earlier this century as the beginning of a new era in which central banks have had to learn about society's inflation preferences by trial and error. In such a world one would expect the mean of inflation regimes to drift as central bank learning proceeds. In either case, inflation regimes would appear stationary over a sufficiently long time span, but may appear nonstationary in certain finite samples.

We begin our analysis by considering the formation of inflationary expectations in this

model. Expected inflation over the next period is simply given by the current inflation target as shown in equation (3). Expected inflation in subsequent periods is a weighted average of the current inflation target and the expected value of future inflation targets, as shown by equations (4-5). The farther ahead one looks, the more likely there will be at least one regime shift, and the greater the weight attached to the expected value of future inflation targets, μ .

$$E_t \pi_{t+1} = \Pi_t \quad (3)$$

$$E_t \pi_{t+2} = (1-q)\Pi_t + q\mu \quad (4)$$

$$E_t \pi_{t+1+i} = (1-q)^i \Pi_t + q \left(\sum_{j=1}^i (1-q)^{j-1} \right) \mu \quad i = 1, \dots, \infty \quad (5)$$

One important property of this model is that the probability density of future inflation is not symmetric if there is a possibility that a regime change may affect the inflation rate in the period in question. The probability density of inflation one period ahead is symmetric because any regime shift that may occur next period will not affect inflation until the following period. For one-period ahead inflation, the probability density is simply the normal density with mean equal to the current inflation target and variance equal to that of the temporary shock (equation (6)). The notation $f_\varepsilon(x)$ refers to the probability density of the variable ε evaluated at the value x . For example, if $\sigma = 1$, $f_\varepsilon(0) = 0.4$ because ε has a standard normal distribution and the standard normal density at zero is 0.4.

$$f_{\pi_{t+1}}(x) = f_{\boldsymbol{\epsilon}}(x - \Pi_t) \quad (6)$$

$$f_{\pi_{t+2}}(x) = (1-q)f_{\boldsymbol{\epsilon}}(x - \Pi_t) + q \int_{-\infty}^{\infty} f_{\eta}(x - \boldsymbol{\epsilon}) f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}) d\boldsymbol{\epsilon} \quad (7)$$

$$f_{\pi_{t+1+i}}(x) = (1-q)^i f_{\boldsymbol{\epsilon}}(x - \Pi_t) + q \left(\sum_{j=1}^i (1-q)^{j-1} \right) \int_{-\infty}^{\infty} f_{\eta}(x - \boldsymbol{\epsilon}) f_{\boldsymbol{\epsilon}}(\boldsymbol{\epsilon}) d\boldsymbol{\epsilon} \quad i = 1, \dots, \infty \quad (8)$$

Looking two periods ahead, the probability density of inflation takes on a two-part form. The first term in equation (7) states that if the current regime survives next period (with probability $1-q$) the probability density of inflation two periods ahead is the same as for one period ahead. The second term in equation (7) states that if the current regime is replaced next period (with probability q) the probability density of inflation in subsequent periods is a convolution of two densities. The first density under the integral is the density of inflation targets across regimes and the second density is the density of inflation rates within a regime.

Looking ahead $1+i$ periods, the probability of remaining in the current regime declines to $(1-q)^i$ and the probability of moving to a new regime increases accordingly. It is possible that there may be one or more regime shifts over this horizon, but the probability density of

future inflation is independent of the number of regime shifts that may occur.⁵

We now consider an example to illustrate the properties of this model. The parameters are adapted from the three-state Markov process estimates of Ricketts and Rose (1995) (RR) for Canada over the past 40 years. RR assume that inflation cycles between three different regimes, with mean inflation rates of 1.5, 4.5, and 9 percent.⁶ Translating these estimates into the model of this section implies a mean inflation rate across regimes of 5 percent ($\mu=5$) with a standard deviation of 3 percent ($\kappa=3$). (The average inflation rate over this sample is also 5 percent.) The probability of entering a new regime is 30 percent per year ($q=0.3$).⁷ At present, Canada is in the low-inflation regime, ($\Pi=1.5$). The standard deviation of inflation in the low-inflation regime is 1 percent, and this estimate is adopted for every regime in the model ($\sigma=1$).

Figure 1 displays the probability densities of inflation under the current regime and under the assumption of a regime shift without any information on the inflation target in the new regime. Figure 2 displays the probability densities for inflation at different periods in the future. These densities are weighted averages of the two densities in Figure 1, with the weight on the regime-shift density increasing with the distance into the future. Clearly, the weighting of these two densities--each of which is symmetric--leads to a highly asymmetric density for future inflation over certain horizons.

⁵This property would not hold true if there were dependence across regimes.

⁶RR allow different serial correlation and variance of inflation in different regimes. In the high inflation regime they impose a unit root on the inflation rate which is not rejected by the data. With a unit root the population mean is undefined, but the sample mean is 9 percent.

⁷RR allow different probabilities of a regime shift depending on the regime. For Canada, the probability of exiting a regime is close to 0.3 for each of the three regimes.

Table 1 displays the mean, median, and mode of inflation from 1 to 10 periods ahead. The asymmetry, as measured by the difference between the mean and the median, grows quickly and peaks in period t+3 before declining slowly over farther horizons. In period t+10 the density is quite close to the future regime density in Figure 1, which is symmetric. The density becomes bimodal in periods t+4 through t+9, with the second peak overtaking the first peak in period t+8. Over the entire 10-year period, the average of the mean inflation rates is 3.9 percent, the average of the medians is 3.5 percent, and the average of the modes is 2.7 percent.

Table 1. Asymmetric Distribution of Future Inflation Rates: Model 1			
$\Pi_t=1.5, \sigma=1, \mu=5, \kappa=3, q=0.3$			
Date	Mean	Median	Mode
t+1	1.5	1.5	1.5
t+2	2.5	1.9	1.5
t+3	3.3	2.4	1.6
t+4	3.8	3.1	1.6
t+5	4.2	3.6	1.7
t+6	4.4	4.1	1.8
t+7	4.6	4.5	2.0
t+8	4.7	4.6	5.0
t+9	4.8	4.7	5.0
t+10	4.9	4.8	5.0
Ave.	3.9	3.5	2.7

2. Learning about the Current Regime

Of the four informational assumptions described at the beginning of the previous subsection, the most realistic are that agents know when there has been a regime change and that they know the probability of a regime change in any given period. Regime changes are likely to be associated with observable events such as a change in the party or individual in control of the central bank, an announcement by the central bank indicating that a new policy has been adopted, or a major economic or social shock such as a war. The probability of a regime change is given by the institutional structure of government and the randomness of individual career decisions

and lifespans. It does not seem unreasonable to assume that agents understand this process well, or at least that their beliefs about it are not changing over time.

The first assumption that we relax is the assumption that agents know any new target inflation rate immediately. Instead, we assume that agents learn about the current regime by observing the inflation rates that occur. During the period in which a regime shift occurs, the best any agent can do is to expect future inflation to be equal to the mean across regimes, μ . (See equation (9).) The probability density is given by the convolution of the target density and the density of deviations from target, shown in equation (10).

$$E_t \pi_{t+i} = \mu \quad i = 1, \dots, \infty \quad (9)$$

$$f_{t, \pi_{t+i}}(x) = \int_{-\infty}^{\infty} f_{\eta}(x - \boldsymbol{\varepsilon}) f_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}) d\boldsymbol{\varepsilon} \quad i = 1, \dots, \infty \quad (10)$$

In the following period, an inflation rate is observed. Assuming that there is no regime change, the optimal learning procedure is to use Bayes' rule to update the probability density of future inflation under the assumption that the current regime continues. The prior density is given by equation (10). Equation (12) displays Bayes' rule, which uses the prior density combined with observed information on inflation in the current regime to determine the conditional probability density of future inflation under the assumption that the current regime survives. Since inflation one period ahead is not affected by any future regime change, its expected value is given by the standard formula for the expectation of a continuously distributed random variable displayed in equation (11) using the density defined by equation (12).

$$E_{t+1} \pi_{t+2} = \int_{-\infty}^{\infty} x f_{t+1, \pi_{t+2}}(x) dx \quad (11)$$

$$f_{t+1, \pi_{t+2}}(x) = \frac{\int_{-\infty}^{\infty} f_{\eta}(x-\boldsymbol{\varepsilon}) f_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}) f_{\boldsymbol{\varepsilon}}(\pi_{t+1} - x + \boldsymbol{\varepsilon}) d\boldsymbol{\varepsilon}}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{\eta}(y-\boldsymbol{\varepsilon}) f_{\boldsymbol{\varepsilon}}(\boldsymbol{\varepsilon}) f_{\boldsymbol{\varepsilon}}(\pi_{t+1} - y + \boldsymbol{\varepsilon}) d\boldsymbol{\varepsilon} dy} \quad (12)$$

Once the current regime ends, the distribution of future inflation reverts to its prior distribution (equation (10)). Thus, the probability density of inflation more than one period ahead takes the compound form presented in equation (14). In all periods beyond t+2, the probability density of inflation is equal to the probability of no regime change times the density for period t+2 plus the probability of a regime change times the prior density of inflation under an unknown regime. The expected value of inflation in periods beyond t+2 (equation (13)) takes a compound form parallel to equation (14). Note that the expected value of inflation under the prior density is μ , the average inflation target across regimes.

$$E_{t+1} \pi_{t+2+i} = (1-q)^i E_{t+1} \pi_{t+2} + q \sum_{j=1}^i (1-q)^{j-1} \mu \quad i = 1, \dots, \infty \quad (13)$$

$$f_{t+1, \pi_{t+2+i}}(x) = (1-q)^i f_{t+1, \pi_{t+2}}(x) + q \sum_{j=1}^i (1-q)^{j-1} f_{t, \pi_{t+1}}(x) \quad i = 1, \dots, \infty \quad (14)$$

After observing inflation in period t+2, the conditional density of inflation in period

$t+3$ is given by Bayes' rule using the prior density (equation (10)) and two pieces of information, π_{t+1} and π_{t+2} . (This density is not shown.) As more periods of inflation are observed without a regime shift, the influence of the prior density diminishes and the conditional density of inflation approaches that of the case in which the current regime target inflation rate is known.

We now consider an example to illustrate the properties of this model. As in the previous section, suppose that the mean inflation target across regimes is 5 percent with a standard deviation of 3 percent, and that the probability of a new regime is 0.3 per period. Suppose that within a regime, the standard deviation of inflation around its target is 1 percent, and that the current inflation target is 1.5 percent. If a regime shift occurs in the current period, the conditional density of future inflation in every period is just given by the density under the assumption of a future regime shift as shown in Figure 1.

If a regime shift occurred last period and the regime survives in the current period, the distribution of next period's inflation depends on the current observation of inflation. Assuming that current inflation is 1.5 (the current inflation target) Figure 3 displays the probability density of inflation next period. For comparison, the density assuming complete knowledge of the current regime is also plotted. Note that the density with incomplete knowledge is more diffuse than that assuming complete knowledge. Figure 4 displays the probability densities for inflation at various periods in the future. These densities are weighted averages of the density in Figure 3 and the density assuming an unknown regime shift (shown in Figure 1). Once again, the weighting of these two densities--each of which is symmetric--leads to an asymmetric density for future inflation.

Table 2 displays the mean, median, and mode of inflation from 1 to 10 periods ahead, conditional on observing inflation in period $t+1$ after a regime shift in period t . The growing asymmetry is readily apparent, but not as extreme as in the case of complete knowledge of the current regime.

Table 2. Asymmetric Distribution of Future Inflation Rates: Model 2			
$\pi_{t+1}=1.5, \sigma=1, \mu=5, \kappa=3, q=0.3$			
Date	Mean	Median	Mode
t+2	1.8	1.8	1.8
t+3	2.8	2.3	1.9
t+4	3.5	2.8	2.0
t+5	3.9	3.3	2.2
t+6	4.2	3.8	2.4
t+7	4.5	4.2	2.7
t+8	4.6	4.4	3.6
t+9	4.7	4.6	4.6
t+10	4.8	4.7	4.8
t+11	4.9	4.8	4.9
Ave.	4.0	3.7	3.1

3. Learning about Future Regimes

The other assumption that we relax is the assumption that agents know the distribution of target inflation rates across future regimes. In order to simplify the analysis we will return to the assumption that agents know the current and past inflation targets.

In this model, agents must estimate the mean and standard deviation of inflation targets across regimes using data on past regimes. Each time a new regime occurs, agents update their

estimates of the mean and standard deviation. Expected inflation in the next period is simply the current inflation target, shown in equation (15). Expected inflation more than one period ahead is a weighted average of the current inflation target and the average of current and past inflation targets. (See equation (16).) Since inflation regimes typically last for more than one period, the second term in equation (16) is an average computed using the first year of each regime, denoted by the set $\{R_N\}$, which contains N elements where N is the number of

regimes. By the Law of Large Numbers, when N is large the right hand side of equation (16) approaches equality with the right hand sides of equations (2) and (3). In other words, when there have been many regimes in the past, agents can estimate the true mean of future inflation targets quite accurately.

$$E_t \pi_{t+1} = \Pi_t \quad (15)$$

$$E_t \pi_{t+1+i} = (1-q)^i \Pi_t + q \left(\sum_{j=1}^i (1-q)^j \right) \left(\sum_{k \in R_N} \frac{\Pi_k}{N} \right) \quad i = 1, \dots, \infty \quad (16)$$

The probability density of inflation one period ahead is given by equation (17), which is identical to equation (6). The probability densities of inflation more than one period ahead are given by equation (18) under the assumption of a diffuse prior distribution on the mean and standard deviation of inflation targets across regimes. As the number of past regimes increases, this density approaches those of equations (7) and (8), and we return to the case of complete knowledge about the distribution of future inflation targets.

$$f_{\pi_{t+1}}(x) = f_{\epsilon}(x - \Pi_t) \quad (17)$$

To illustrate the properties of this model, we need to specify values of current and past inflation targets. To continue with the flavor of past examples, we choose past inflation targets of 4.5 and 9 percent, and a current inflation target of 1.5 percent. The average of

$$\begin{aligned}
f_{\pi_{t+1+i}}(x) &= (1-q)^i f_{\varepsilon}(x - \Pi_i) \\
+ q \left(\sum_{j=1}^i (1-q)^{j-1} \right) &\frac{\left(\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{\varepsilon}(\varepsilon) f_{\eta}(x - \varepsilon) f_{\eta}(\Pi_1) \cdots f_{\eta}(\Pi_N) d\mu d\kappa d\varepsilon \right)}{\left(\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f_{\eta}(\Pi_1) \cdots f_{\eta}(\Pi_N) d\mu d\kappa \right)} \quad (18) \\
i &= 1, \dots, \infty
\end{aligned}$$

these targets is 5 percent and the standard deviation is 3 percent. Thus, these outcomes are consistent with our earlier assumption of $\mu=5$ and $\kappa=3$. Figure 5 displays the density of future inflation under the assumption that there is a regime shift, i.e., the ratio of the triple integral to the double integral in equation (18). For comparison, Figure 5 also displays the density of future inflation after a regime shift under the assumption of complete knowledge of the distribution of inflation targets, which was originally displayed in Figure 1. It is not surprising that the density without knowledge is more diffuse than the density with knowledge. Both densities are symmetric around the same mean, however, due to our choice of observed inflation targets with the same mean as the true distribution.

Figure 6 displays the densities of inflation in particular future periods. Once again, the densities are asymmetric whenever there is a positive probability that a regime shift may affect inflation in the period in question. Table 3 displays the mean, median, and mode of inflation in various future periods under the assumption that agents do not know the parameters of the distribution of future inflation targets and must infer them from observed inflation targets. The means and medians are identical to those displayed in Table 1 because the average of current and past inflation is assumed to be equal to the true mean of inflation targets across regimes. The only difference between Model 3 and Model 1 is that the density

of inflation after a regime change is much more diffuse. This diffuseness affects the mode of future inflation since the secondary peak at 5 percent inflation is much lower than in Model 1. This diffuseness has no effect on the mean or median of future inflation.

Table 3. Asymmetric Distribution of Future Inflation Rates: Model 3			
$\Pi_R = \{4.5, 9, 1.5\}$, $\Pi_i = 1.5$, $\sigma = 1$, $q = 0.3$			
Date	Mean	Median	Mode
t+1	1.5	1.5	1.5
t+2	2.5	1.9	1.5
t+3	3.3	2.4	1.5
t+4	3.8	3.1	1.6
t+5	4.2	3.6	1.7
t+6	4.4	4.1	1.7
t+7	4.6	4.4	1.7
t+8	4.7	4.6	1.8
t+9	4.8	4.7	1.9
t+10	4.9	4.8	2.0
Ave.	3.9	3.5	1.6

Finally, we consider the effect of a new regime on the mean of future inflation in the case of learning about the distribution of inflation targets. Suppose that a new regime occurs in the example above with an inflation target of 1.5 percent. In other words, suppose that a new central bank governor was installed who chose to continue the previous inflation target of 1.5 percent. The effect of this new regime depends on the number of previously observed regimes. If there were only three previous regimes, the average of current and past inflation targets drops

from 5 to 4.1 percent. If there were ten previous regimes, which is the number of regimes estimated for Canada by RR, the average of current and past regimes declines by much less, from 5 to 4.7 percent.

IV. Empirical Support

Estimation and testing of the models in the previous section pose a serious econometric challenge that is beyond the scope of this paper. Instead, we show that artificial

data generated by the basic model of this paper behave in a manner similar to observed inflation, and that this model may explain certain puzzling properties of the observed data. In addition, we show that the model developed here may be able to explain puzzling features of the evidence on long-run inflation expectations.

Despite the fact that this model does not incorporate any serial correlation of inflation within a regime, nor any serial correlation across regimes, it is capable of explaining much of the observed serial correlation of inflation. Over the sample examined by RR, 1954-93, the Canadian CPI inflation rate has an estimated dominant autoregressive root of about 0.85, and augmented Dickey-Fuller (ADF) tests cannot reject a unit root at any significance level. Monte Carlo data generated by Model 1 with the parameters in Table 1 for the same number of observations yield a median dominant autoregressive root of about 0.5, and ADF tests reject a unit root at the 5 percent level only about 45 percent of the time. If the model is extended to include an autoregressive lag on inflation of 0.7 (the mean of the within-regime autoregressive parameters estimated by RR) and new Monte Carlo data are generated, the median dominant root increases to 0.82 and the power of the 5-percent ADF test drops to 15 percent. For comparison, data generated by a simple autoregression with no regime shifts and a lag coefficient of 0.7 yield a median estimated dominant root of 0.66 and the power of the 5-percent ADF test is 40 percent.

In addition to explaining the near unit-root behavior of inflation over long horizons, a model with regime shifts can also explain the apparent stationarity of inflation over certain shorter horizons. Simply put, within regimes inflation is stationary, therefore one ought to be able to reject nonstationarity in a regime that is sufficiently long-lasting. For example, ADF

tests on quarterly U.S. inflation reject a unit root between 1954 and 1966 and also between 1984 and 1996. Regimes of this length are plausible for the United States, as RR estimate only a 10 percent per year probability of a regime shift ($q=0.10$) with U.S. data.

The asymmetric distribution of future inflation in these models of regime shifts may explain the asymmetric distribution of survey responses on future inflation expectations. Carlson (1975) and Lahiri and Teigland (1987) present evidence that the distribution of 1-year-ahead inflation expectations across survey respondents is usually asymmetrically distributed. Moreover, the direction of the skewness is identical to that predicted by a regime-shift model for the true distribution of future inflation.⁸ When inflation is higher than its historical average, expectations are skewed negatively. When inflation is lower than its historical average, expectations are skewed positively.

Finally, the asymmetric distribution of future inflation may explain the frequently large discrepancies between surveys of inflation expectations and implied inflation expectations in bond yields. For example, in Canada the inflation premium between nominal and indexed bonds was 3 percent at yearend 1996, while *Consensus Forecasts'* (October 1996) survey of 10-year ahead inflation expectations was 2 percent. The regime-shifting models of inflation presented above and calibrated on Canada yield a 10-year ahead inflation mean of roughly 4 percent and a mode (average across years) of roughly 3 percent. However, if the probability of a regime shift were reduced from 30 percent to 10 percent per year--possibly reflecting increased credibility of the Bank of Canada's announced inflation target--the mean 10-year-

⁸I am unaware of any research on how the distribution of a variable affects the distribution across individuals of forecasts of that variable. Nevertheless, these results are suggestive.

ahead inflation rate would drop to around 3 percent and the mode would drop to below 2 percent. If survey respondents report the most likely outcome, and bondholders care about the average outcome, then the discrepancy between different measures of inflation expectations would be resolved.⁹

V. Interpretations and Extensions

The basic point of this paper is a stark one: monetary regimes with inflation targets that are quite different from the average inflation rate across previous regimes may never be fully credible to long-term financial markets. This lack of credibility is not necessarily due to slow learning by private agents or to a lack of resolve on the part of the central bank. Even when agents understand and believe in the central bank's target inflation rate, they must attach some probability to a change in the regime. For example, the central bank governor may die or resign, or the government may change the institutional framework of monetary policy. There is no way to guarantee that these things will not happen.

The key to credibility over the long term is the expected value of inflation under the next regime. This paper considers two approaches to modeling expectations of inflation under the next regime. In the first approach, it is assumed that agents know the constant mean

⁹The professional forecasters surveyed by *Consensus Forecasts* presumably are judged by clients on the accuracy of their forecasts. I would like to thank Jeff Dominitz for pointing out that forecasters should report the mean of future inflation if the penalty for forecast errors is proportional to the squared error. They should report something between the mean and the mode if the penalty is proportional to the absolute error. They should report the mode if the penalty is constant for all errors greater than a given magnitude and zero otherwise. In practice, forecasters communicate more to their clients than a simple point forecast. It is common to talk of the forecast being the most likely scenario with unequal upside and downside risks, which would imply a forecast that is closer to the mode than the mean.

inflation rate across regimes. If the current regime inflation target is equal to this long-run inflation mean, then policy is credible in the long run. If the current regime inflation target is far from the long-run inflation mean, then policy is not credible and policy will never become credible no matter how long the current regime lasts or how often similar regimes arise. In the second approach, agents update their expectations about inflation in the next regime based on inflation in the current and previous regimes. Under this approach, a sequence of low (or high) inflation regimes would change agents' expectations of inflation in the next regime. However, as demonstrated in a simple example, significant changes in long-run inflation expectations may still require a very long time.

One plausible extension of these models is to consider learning on the part of the central bank. For example, one may argue that the high inflation of the 1970s was a mistake, that central banks have learned their lessons, and that the public understands that this episode will not recur. Under this hypothesis, agents ought to place more weight on recent inflation rates when forming expectations about inflation in the next regime, and long-run credibility would be easier to obtain than in the basic models. Nevertheless, as long as agents place some positive weight on past inflation targets in forming expectations about future inflation targets, the credibility problem will remain.

Another extension of the model would be to consider variation over time in the probability of a regime shift. One way to increase the long-run credibility of the current inflation target is to take steps to reduce the probability of a regime shift. Recent attempts in many countries to increase the independence of the central bank may be interpreted as reducing the probability of a regime shift and thus strengthening credibility. Nevertheless, it

is not possible to guarantee that any regime will last forever.

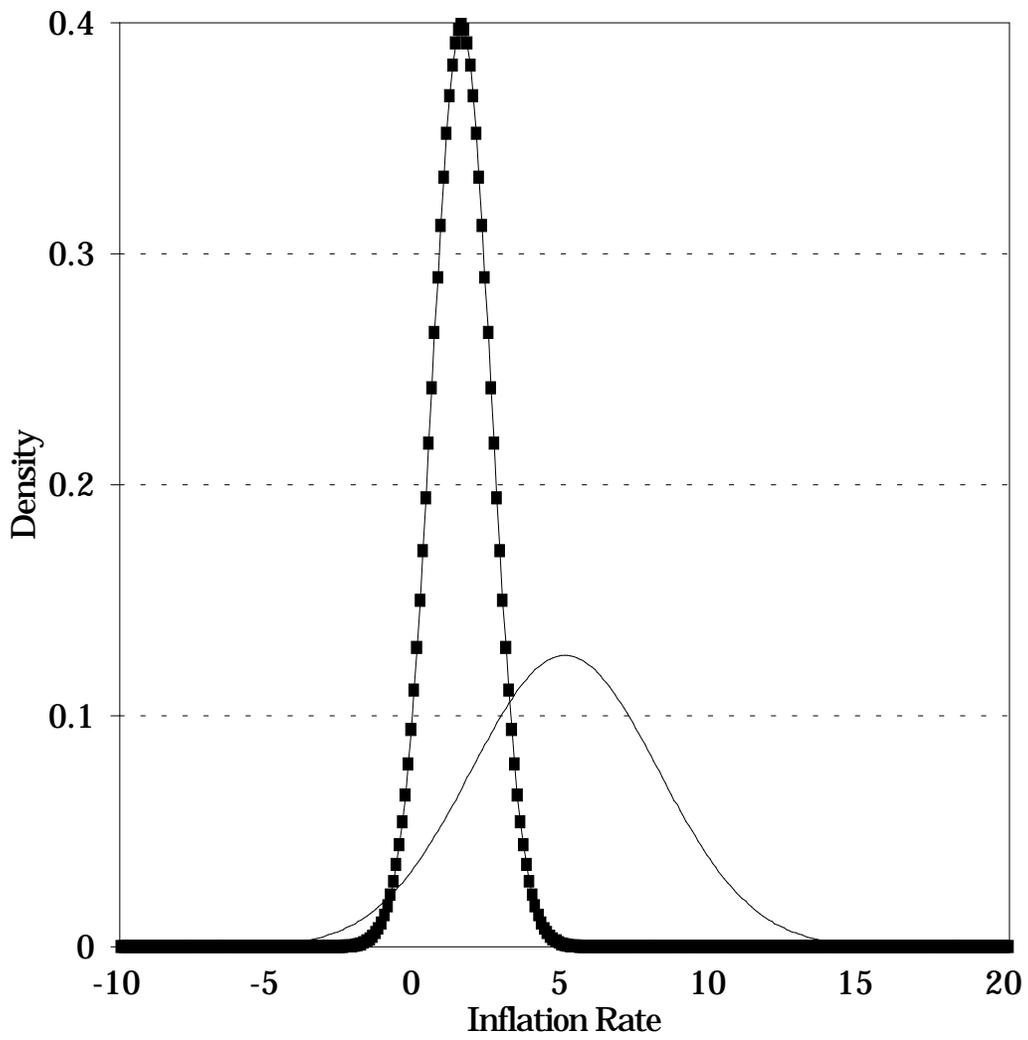
Although the hypothesis of central bank learning seems plausible and many central banks have achieved greater independence in recent years, I would like to conclude the paper by noting several caveats. First, the persistence of inflation in some countries that have a long history of inflation (particularly in less developed countries) argues for caution about the idea that a bad experience with inflation inoculates one against future inflation. At the very least, one should keep in mind that lessons learned may become lessons forgotten. Second, even if central banks have learned their lessons well and permanently, the public may be skeptical and the time needed to convince the public may be measured in decades rather than years. Third, even if one does discount the possibility of a return to double-digit inflation, it is harder to justify ignoring the possibility of a return to moderate inflation rates of around 5 percent or so. In light of the fact that no one is recommending a regime with negative inflation rates, an inflation target that is very close to zero can never be credible in the long run as long as there is some possibility of a return to positive inflation.

References

- Ammer, John, and Richard T. Freeman (1995) "Inflation Targeting in the 1990s: The Experiences of New Zealand, Canada, and the United Kingdom," *Journal of Economics and Business* 47, 165-192.
- Bai, Jushan, and Pierre Perron (1995) "Estimating and Testing Linear Models with Multiple Structural Changes," Massachusetts Institute of Technology, manuscript.
- Baillie, Richard T., Ching-Fan Chung, and Margie A. Tieslau (1996) "Analysing Inflation by the Fractionally Integrated ARFIMA-GARCH Model," *Journal of Applied Econometrics* 11, 23-40.
- Carlson, John A. (1975) "Are Price Expectations Normally Distributed?" *Journal of the American Statistical Association* 70, December, 749-754.
- Chapman, David A., and Masao Ogaki (1993) "Cotrending and the Stationarity of the Real Interest Rate," *Economics Letters* 42, 133-138.
- Driffill, John, Grayham E. Mizon, and Alistair Ulph (1990) "Costs of Inflation," in B.M. Friedman and F.H.Hahn (eds.) *Handbook of Monetary Economics* vol. 2 (Amsterdam: North-Holland).
- Evans, Martin, and Karen Lewis (1995) "Do Expected Shifts in Inflation Affect Estimates of the Long-Run Fisher Relation?" *The Journal of Finance* 50, 225-253.
- Fischer, Stanley (1990) "Rules Versus Discretion in Monetary Policy," in Friedman and Hahn.
- Freeman, Richard T., and Jonathan L. Willis (1995) "Targeting Inflation in the 1990s: Recent Challenges," International Finance Discussion Papers No. 525, Board of Governors of the Federal Reserve System, September.
- Fuhrer, Jeffrey C. (1996) "Monetary Policy Shifts and Long-Term Interest Rates," *The Quarterly Journal of Economics* 111, November, 1183-1209.
- Gagnon, Joseph E. (1996) "Long Memory in Inflation Expectations: Evidence from International Financial Markets," International Finance Discussion Papers No. 538, Board of Governors of the Federal Reserve System, February.
- Hassler, Uwe, and Jurgen Wolters (1995) "Long Memory in Inflation Rates: International Evidence," *Journal of Business and Economic Statistics* 13, January, 37-45.

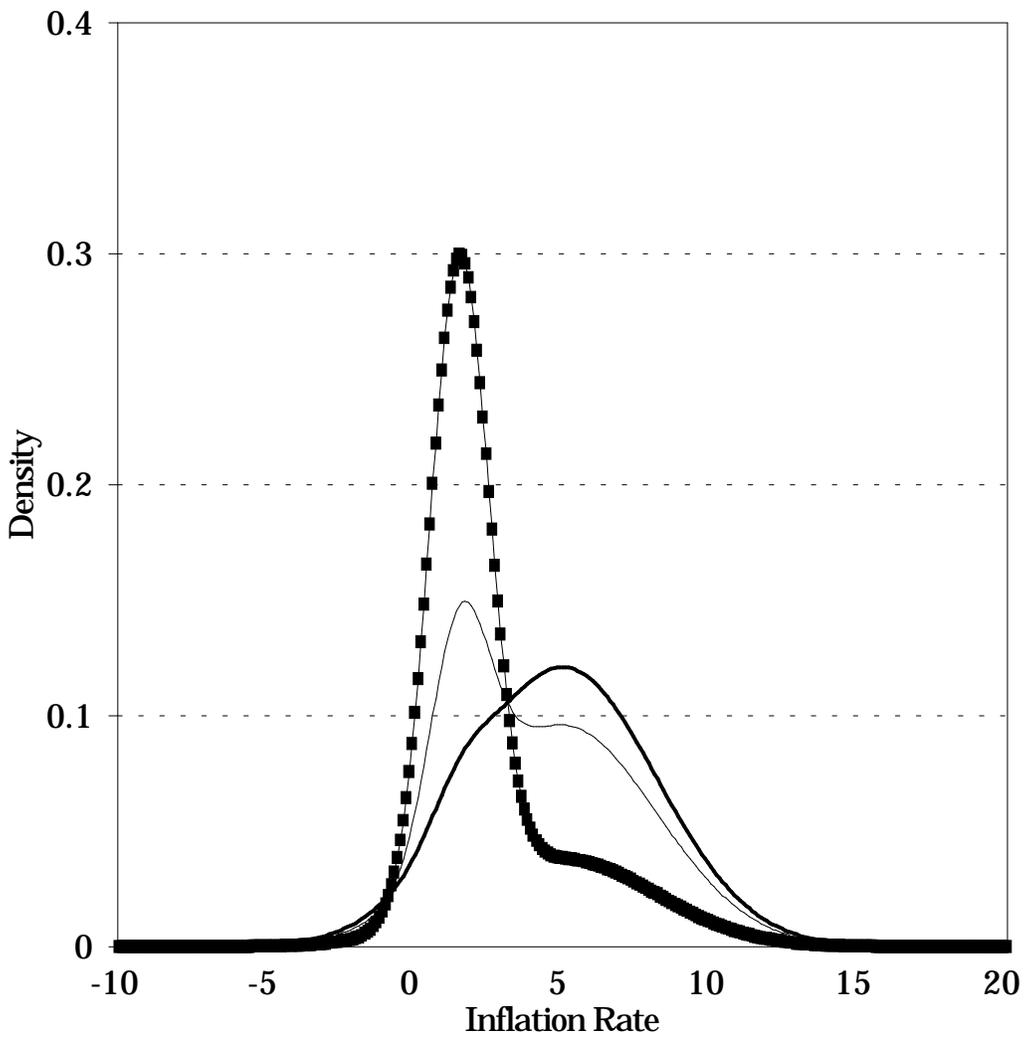
- Hostland, Doug (1995) "Changes in the Inflation Process in Canada: Evidence and Implications," Bank of Canada Working Paper No. 95-5, May.
- International Monetary Fund (1996) *World Economic Outlook*, October.
- Lahiri, Kajal, and Christie Teigland (1987) "On the Normality of Probability Distributions of Inflation and GNP Forecasts," *International Journal of Forecasting* 3, 269-279.
- Mishkin, Frederic S. (1992) "Is the Fisher Effect for Real?" *Journal of Monetary Economics* 30, November, 195-215.
- Reserve Bank of New Zealand (1996) *Policy Targets Agreement*, December.
- Ricketts, Nicholas, and David Rose (1995) "Inflation, Learning and Monetary Policy Regimes in the G-7 Economies," Bank of Canada Working Paper No. 95-6, June.
- Rose, Andrew (1988) "Is the Real Interest Rate Stable?" *Journal of Finance* 43, December, 1095-1112.
- Simon, John (1996) "A Markov-Switching Model of Inflation in Australia," Research Discussion Paper No. 9611, Reserve Bank of Australia, December.
- Walsh, Carl (1995) "Optimal Contracts for Central Bankers," *American Economic Review* 85, March, 150-167.
- Woodford, Michael (1990) "The Optimum Quantity of Money," in Friedman and Hahn.

Figure 1
Model 1



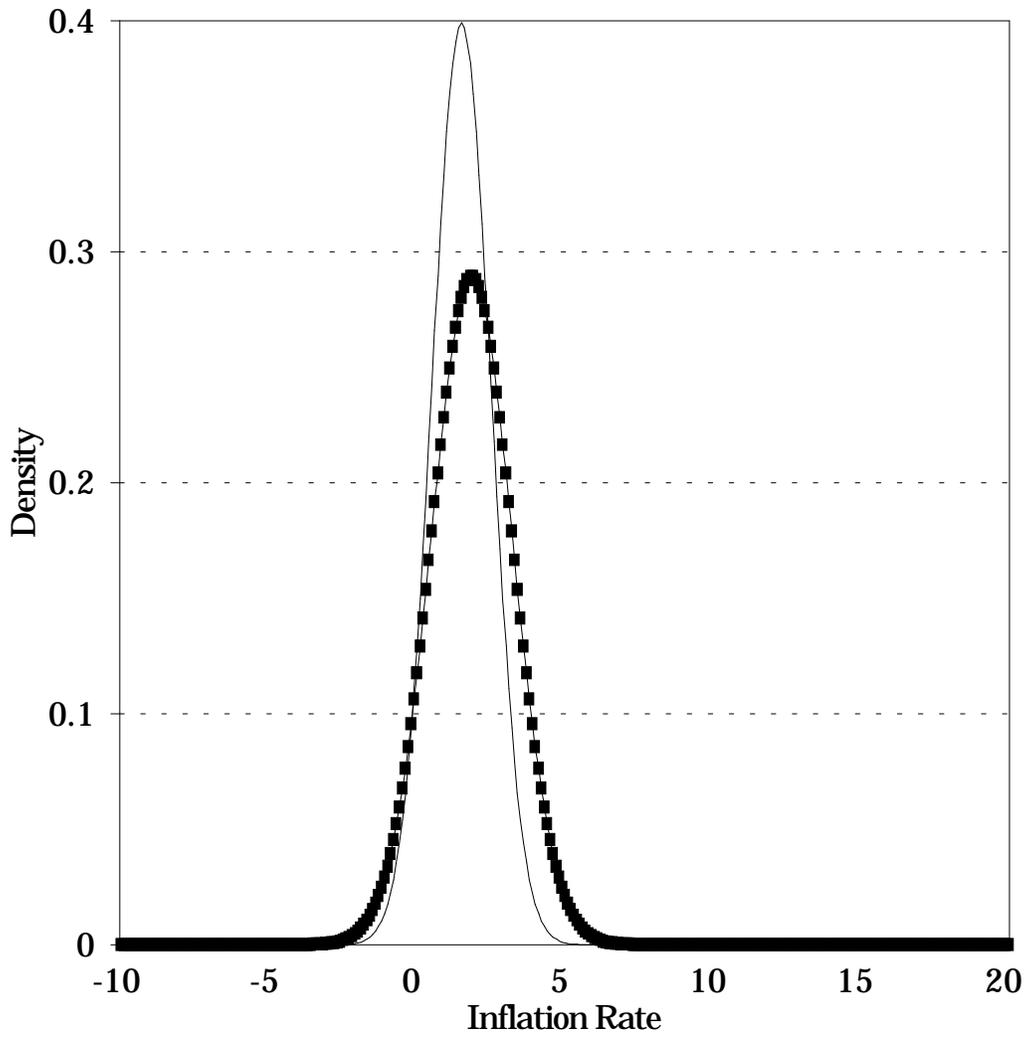
■ Current Regime — Future Regime

Figure 2
Model 1



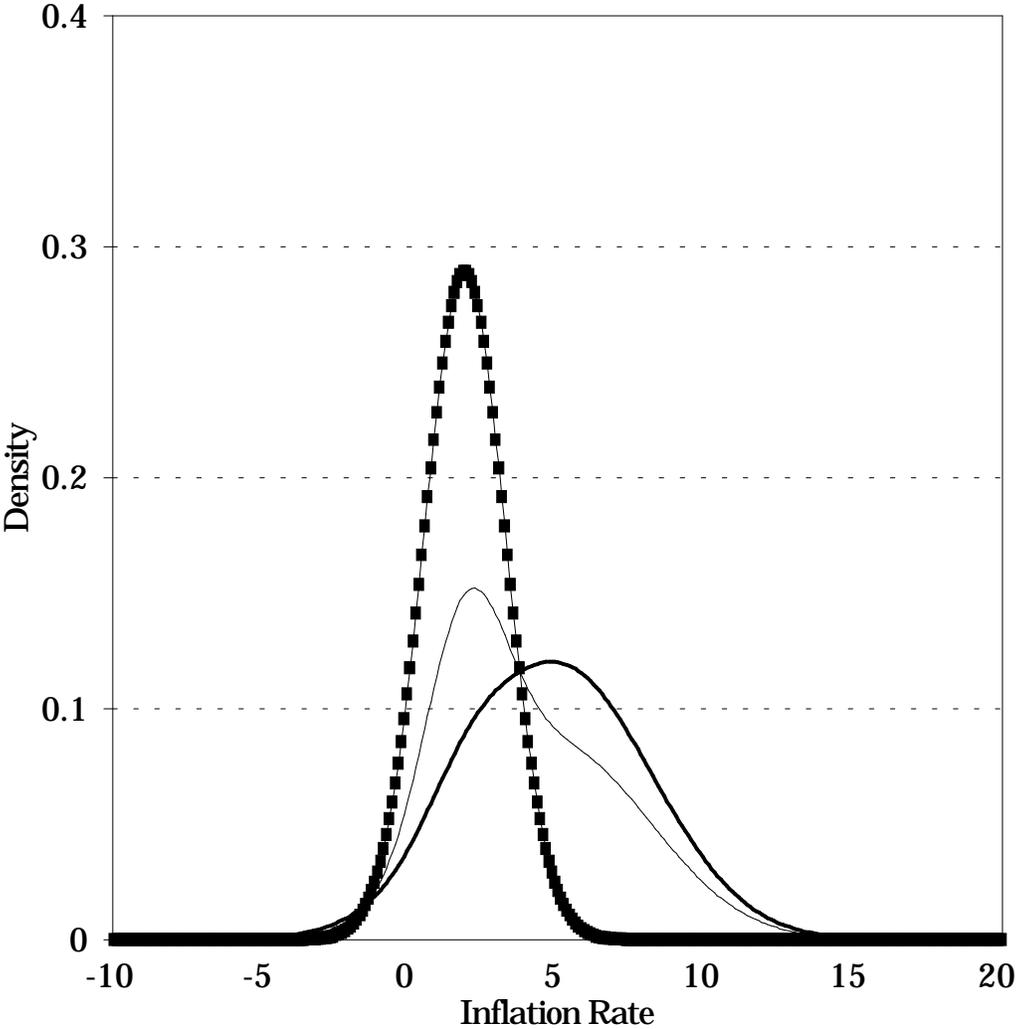
■ t+2 — t+5 — t+10

Figure 3
Model 2



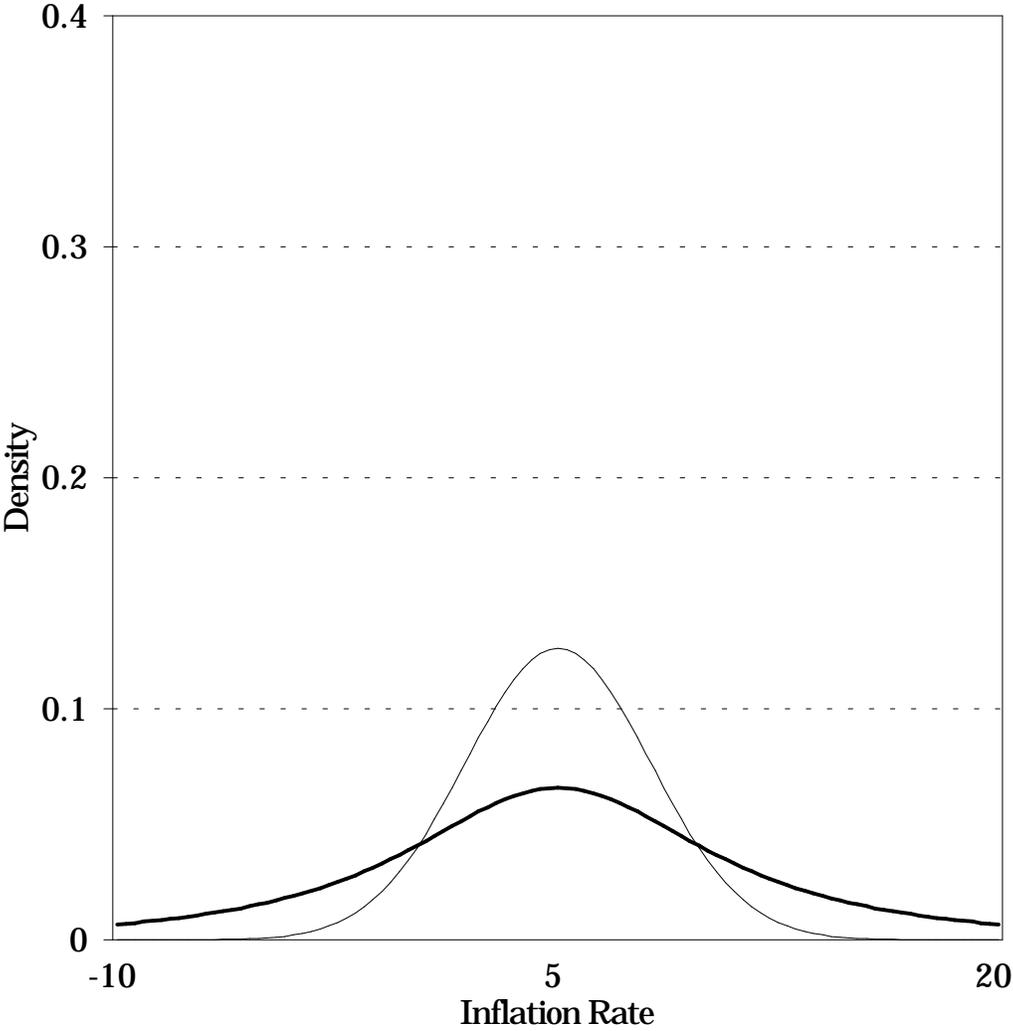
■ Current Regime Model 2 — Current Regime Model 1

Figure 4
Model 2



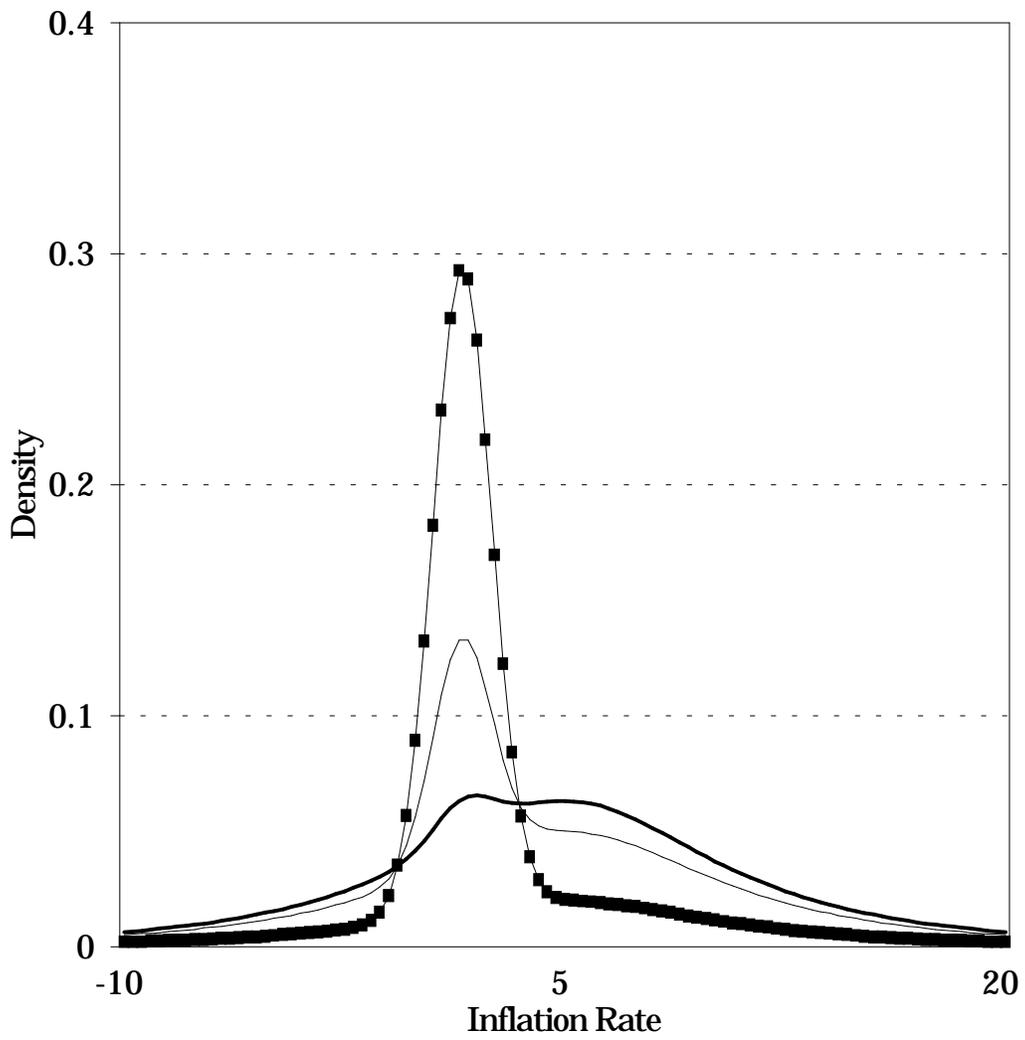
■ t+2 — t+5 — t+10

Figure 5
Model 3



— Future Regime Model 3 — Future Regime Model 1

Figure 6
Model 3



■ t+2 — t+5 — t+10