Working Paper 95-6 / Document de travail 95-6

Inflation, Learning and Monetary Policy Regimes in the G-7 Economies

by

Nicholas Ricketts and David Rose

Bank of Canada Banque du Canada

June 1995

INFLATION, LEARNING AND MONETARY POLICY REGIMES IN THE G-7 ECONOMIES

by

Nicholas Ricketts and David Rose

Research Department
Bank of Canada, Ottawa, K1A 0G9

Preface

We gratefully acknowledge the expert assistance of Hope Pioro in programming the routines used for the estimation reported in this paper. We thank participants in seminars at the Bank of Canada, the University of Warwick and the Université du Québec à Montréal for helpful discussion. In particular, comments by Pierre Duguay, Tiff Macklem, Louis Phaneuf, Steve Poloz and by Alain Guay and Simon van Norden helped us improve the paper. Chris Lavigne prepared the graphs. This paper was presented in draft form at the meetings of the Canadian Economics Association in June 1994 at the University of Calgary. We thank participants for their helpful comments, especially our discussant, Wai-Ming Ho, and Jeremy Rudin. Any remaining errors are our responsibility.

E-mail : nricketts@bank-banque-canada.ca
drose@bank-banque-canada.ca

Telephone: (613) 782-8378 (Ricketts)
(613) 782-8728 (Rose)

Fax: (613) 782-7163

ISSN 1192–5434

ISBN 0–662–23530–4

Printed in Canada on recycled paper

CONTENTS

Abstract

In this paper, the authors report estimates of two- and three-state Markov switching models applied to inflation, measured using consumer price indexes, in the G-7 countries. They report tests that show that two-state models are preferred to simple one-state representations of the data, and argue that three-state representations are more satisfactory than two-state representations for some countries. The preferred estimation results usually include a state that features a unit root in its dynamic structure, which concurs with results of direct tests for this property. However, the multistate representation of the data shows that for all G-7 countries these quasi-unit-root properties arise primarily from a few brief episodes of history, concentrated in the 1970s and associated with the major oil-price shocks.

For all countries there is evidence of progress towards establishing credibility of regimes with stable inflation, and in many countries there is evidence of progress in building credibility of regimes with low inflation. Credibility refers to the ex post probability assigned to the state by the Markov model, which has a large effect on how expectations of future inflation are formed. An interesting contrast arises from the results for the United States and Canada. Whereas in Canada the credibility of a regime with historically low inflation has risen sharply in the last few years, in the United States there has been convergence on a regime with a stable, but historically average, rate of inflation and not on the alternative low-inflation regime.

Résumé

Dans cette étude, les auteurs présentent des estimations de modèles de Markov à changement de régime comportant deux et trois états pour représenter l'inflation, mesurée par les indices des prix à la consommation des pays du groupe des Sept. Les tests effectués révèlent que les modèles à deux états sont supérieurs à ceux à état unique et que, pour certains pays, les modèles à trois états sont plus attrayants que ceux à deux états. Les estimations retenues incorporent généralement un état dont la structure dynamique est caractérisée par une racine unitaire compatible avec les résultats des tests directs de racine unitaire. Toutefois, les modèles à états multiples indiquent que, pour tous les pays du groupe des Sept, cette propriété, de racine plutôt quasi unitaire, résulte principalement de quelques événements de courte durée, liés aux principaux chocs pétroliers, qui sont surtout survenus dans les années 70.

Selon les résultats empiriques, tous les pays progressent vers l'établissement de régimes crédibles d'inflation stable, et de nombreux pays sont en voie d'implanter des régimes crédibles de faible inflation. La notion de crédibilité renvoie à la probabilité ex post que le modèle de Markov attribue à l'état, laquelle influe beaucoup sur la façon dont les anticipations d'inflation se forment. Un contraste intéressant apparaît lorsque les résultats obtenus pour les États-Unis à cet égard sont comparés à ceux qui s'appliquent au Canada. Alors que, au Canada, la crédibilité d'un régime de faible inflation s'est considérablement accrue au cours des dernières années, on a trouvé aux États-Unis une convergence vers un régime caractérisé par un taux d'inflation stable, mais de niveau habituellement moyen, plutôt que le niveau bas du régime de faible inflation.

1 Introduction

The data from all industrialized countries show that there have been periods of low inflation as well as periods of moderate or high inflation during the past few decades. In this paper, we ask how this history should be characterized. Should it be characterized as varying outcomes from a single regime or data-generating process, and if so, is that process stationary? Or is it better to think of this history as reflecting different policy regimes? In principle, what appears to be non-stationarity in the historical data could arise from a combination of regimes, each of which exhibits stationary properties but with different means and different dynamics, such that the overall series fails tests for stationarity in small samples. One question we ask in this paper is whether there is any empirical support for this view and, more generally, whether there is empirical support for rejecting a "single process" interpretation of inflation in the G-7 economies.

We use a Markov switching representation of a multistate environment to model the inflation process.[1] In this representation, there are alternative possible states of the world (policy regimes), each with its own average rate of inflation and dynamic properties. By assumption, agents using such a representation of the data cannot know with certainty which regime is in place. They react to the outcome for inflation and assign ex post probabilities as to which regime has been operative. Agents' forward-looking inflation expectations are formed using a weighted average of the forecasts from the models for each state, where the weights are the ex ante probabilities that the various states will occur next period, given the evaluation of the starting point. Expectations thus evolve through time, based on a combination of least-squares learning from the data about the shocks

1. Our application is similar to that in Evans and Wachtel (1992) and Laxton, Ricketts and Rose (1994). We review some other applications of the Markov switching model in Section 2.

(and their dynamic implications) and Bayesian updating of the perceived probabilities of the alternative regimes (and the implied probabilities of future transitions).

Using consumer price data for the G-7 countries, we conclude that we can generally reject, statistically, single-regime interpretations of the data against more general Markov formulations. Although we are not able to eliminate completely the need for unit-root episodes as part of a more general representation of the data, we are able to show that such properties tend to be limited to unusual and brief circumstances. Overall, the picture that emerges is that inflation can be characterized as stationary.

Another question we ask is how agents using a Markov switching framework might have perceived the evolution of monetary policy over time in the various countries. Of particular interest is the question of whether there is any systematic evidence that lower inflation regimes have become more credible in recent years. Laxton, Ricketts and Rose (1994) show that the perceptions of economic agents of what regime is in place can have dramatic consequences for the way an economy responds to shocks.

The importance of perceptions of regime shifts for expectations formation can be illustrated by considering the implied pattern of forecast errors. When a regime shift occurs in a noisy environment, agents may not realize it and incorrectly interpret their forecast errors as the result of shocks to the economy. In that case, there may be important systematic forecasting errors, as agents continue to rely on dynamic forecasting structures from the previous regime. Conversely, perception of a regime shift need not be caused by a change in the monetary authority’s objectives. Unexpected shocks may reduce the credibility of a policy regime that is truly still being followed. If agents believe that a regime shift has occurred or may occur in the near future, then their expectations will be pulled away from values consistent with continuation of the current regime. Biased forecasts of inflation and exchange rates revealed by survey data may be explained by such uncertainty about the regime.

We find that there has been a general tendency towards establishing lower and more stable inflation regimes during the 1980s and early 1990s. The probability assigned to a state with a unit root is generally very low for all G-7 countries during this period. Moreover, for many countries there is evidence that the credibility of a regime with low inflation is rising, although this is not true in all cases. Our results suggest, for example, that whereas in Canada increasing weight has been assigned during the 1990s to a regime with historically low average inflation, the same is not true for the United States. There, according to our results, high probability is being assigned to a regime with historically average inflation, but not to a regime with low inflation.

The paper is organized as follows. In the next section, we describe the Markov switching model, how the probabilities are determined and how expectations are formed. Then, in Section 3, we report statistical tests of various propositions about the data. We begin with tests for stationarity of inflation. Although the evidence is mixed, we conclude that it is not possible to reject convincingly the hypothesis that inflation is nonstationary. We then proceed to test directly, using a procedure suggested by Hansen (1992), whether a single-regime interpretation can be sustained. For most countries, the formal statistical tests favour a Markov switching structure over a single-regime interpretation. However, there is no evidence from this test that more than two states are required to explain our particular samples.

In Section 4, we discuss the results of estimations of two-state models for the G-7 countries, and we present some further diagnostic tests of the adequacy of these models. The results of the tests, and a number of unsatisfactory features of the two-state representations, lead us to conclude that further structure is needed. In Section 5, we report estimated three-state models. We argue that the three-state representations make some important contributions in explaining history, at least for some countries. Finally, in Section 6, we offer some general conclusions.

2 Markov switching model

The Markov switching model (MSM) provides a means of modelling particular forms of non-linearity in economic time series. The MSM takes all or some of the parameters of a reduced-form model as being dependent on an unobserved state variable. The state variable can take on a finite number of discrete values. The values are generated by a stochastic process, which is described by a matrix of probabilities of transition between states. The nature of the dynamics of the economic variable then depend on the outcome for the state variable. More important, however, transitions between states add an important element of non-linearity to the dynamics of the observed outcomes.

Many economic variables have been shown to exhibit different time-series properties during different periods of history. Some early applications of the MSM, including Hamilton (1989), who popularized the approach, have been in the study of business cycles, where expansions and contractions are interpreted as a process of switching between states with high and low growth rates. Another common application is to explain processes that exhibit periodic discrete shifts in regime, where agents' expectations play a major role in determining outcomes. For example, Hamilton (1988) applies the MSM to a study of the term structure of interest rates, Lewis (1989a, 1989b) examines various applications to exchange rates, and Evans and Wachtel (1992) apply the model to U.S. inflation and short-term interest rates. Laxton, Ricketts and Rose (1994) use the model to study Canadian inflation.

The basic MSM specifies a set of possible states and the dynamics for the economic variables within each state. In the most general form of the model, all the parameters can depend on the outcome of the state. Moreover, the transition probabilities between states can be time-varying, as in Filardo and Gordon (1993), Ghysels (1993), and Schaller and van Norden (1993), depending on variables not correlated with the state. Durland and McCurdy (1992) model time-varying transitions as dependent on the duration of the state.

It is common practice to restrict the number of parameters in these models to make estimation and inference less difficult. Hamilton (1989), for instance, uses an autoregressive structure that is the same in all states. Only the mean of the growth rate of GNP and the standard deviation of the random shocks differ across states in his application.

We model the inflation process in each state as a first-order autoregressive process

$$\pi_t = \alpha(S_t) + \beta(S_t) \cdot \pi_{t-1} + \varepsilon_t(S_t) \qquad \varepsilon(S_t) \sim N(0, \sigma(S_t)), \tag{1}$$

where π_t is the annual rate of inflation, S_t is the state variable and where ε is a Gaussian disturbance with a standard deviation that is state-dependent. The specification is completed with a definition of the determination of S_t

$$Pr(S_t = j | (S_{t-1} = i)) = p_{ij} \qquad i, j \in (1, 2, \ldots, n), \tag{2}$$

where p_{ij} is the probability of moving from state i to state j, and where n is the number of possible states. Together, the elements p_{ij} define a matrix of transition probabilities that describe the stochastic process for the state variable.

We allow for up to three states, in each of which inflation has a different autoregressive structure. We explore, in some detail, the implications of using different formulations of the two-state model. We consider a completely unrestricted version and two restricted versions. Recall that one question we wish to pose is whether or not there is evidence of states with nonstationary dynamics in a Markov framework. If the unrestricted two-state model indicates an unstable process for one of the states, for example, or if a freely estimated root is close to the unit circle, a natural alternative would be to assume that that state has a unit root in its dynamic structure. As part of our consideration of this issue, we estimate models with one of the states constrained to be a simple random walk.

We also estimate a version of the model with a restriction on the transition dynamics. Under the restriction that $p_{ii} = p_{ji}$, for all j, the MSM becomes a simple switching model (SSM), in which the ex ante probability of observing a state is constant and independent of the previous outcome. This is not a very interesting model, from the perspective of policy analysis; it would say, for example, that expectations would always be formed in the same way, regardless of the regime chosen by the monetary authority. We include it as part of our testing of whether or not the MSM contributes anything to an understanding of the data.

The switching models are estimated by maximizing the likelihood of the observed values of past inflation over the unobserved states. Each outcome of the unobserved state implies a particular distribution for the observed values of inflation. By specifying the distribution of inflation in each of the states and the distribution of the unobserved state variable, which in our case is based on a fixed transition matrix, we can obtain a joint conditional density function for each state. Summation over the states yields a likelihood function that can be maximized with respect to the parameters of the system.

The maximum-likelihood estimation routine[2] supplies the ex post probabilities for the states $Pr(S_T = i \mid \pi_T, ..., \pi_1)$ for i = 1, 2, ..., n. Expectations, however, are based on conditional ex ante probabilities. In the case of a two-state model, the conditional probabilities are given by:

2. The programs used in estimating the models are based on Hamilton's code with modifications by Simon van Norden, Hope Pioro and the authors.

$$Pr(S_{T+1} = 1|\pi_T, \ldots\pi_1) = Pr(S_T = 1|\pi_T, \ldots\pi_1) \cdot p_{11} + (1 - Pr(S_T = 1|\pi_T, \ldots\pi_1)) \cdot (1 - p_{22}) \tag{3}$$

$$Pr(S_{T+1} = 2|\pi_T, \ldots\pi_1) = 1 - Pr(S_{T+1} = 1|\pi_T, \ldots\pi_1)\,. \tag{4}$$

The conditional probability that the state will be 1 is given by the probability that last period's state was 1, times p_{11}, the probability of remaining in state 1, plus the probability that last period's state was 2, times $(1–p_{22})$, the probability that there will be a transition from state 2 to state 1. The conditional expectation of inflation in period T+1 is then given by the probability-weighted average of the predictions from each of the two individual dynamic models:

$$E[\pi_{T+1}|(\pi_T, \ldots\pi_1)] = Pr((S_{T+1} = 1|\pi_T, \ldots\pi_1) \cdot (\alpha_1 + \rho_1 \cdot \pi_T)) + Pr((S_{T+1} = 2|\pi_T, \ldots\pi_1) \cdot (\alpha_2 + \rho_2 \cdot \pi_T).) \tag{5}$$

Forecasts for inflation in subsequent periods can be made by iterating forward with this formula. These longer-lead expectations converge as the conditional probabilities converge on their unconditional limit values and the individual forecasts converge on the unconditional means, if they exist, or remain at the most recent value, in the case of a unit-root process in its simplest form (the random walk). The forecasts generated by the formulas will be optimal (maximum likelihood) for agents whose only information about the current state comes from observations of past inflation and who take into account the possibility of future changes in regime.

Extending the model to allow for three states, while retaining the autoregressive structure within states, requires expanding the transition matrix and introduces many more parameters to be estimated. Expectations are formed as above, but they take into account the additional possibilities for switching between states.

3 On unit roots and tests for multiple states

In this section, we investigate two important issues. First, we consider the evidence as to whether inflation in the G-7 economies, as measured by the consumer price indexes (CPIs) of these countries, can be treated as stationary, that is, whether we can reject the presence of a unit root.[3] We then offer a direct test of whether a single-regime interpretation of these data can be accepted.

Table 1 (see appendix) summarizes the results for the unit-root tests.[4] The results are ambiguous in their conclusions regarding the stationarity of the CPI data for every country. However, some patterns do emerge. The Sargan-Bhargava (SB) (1983) statistic consistently rejects the null of a unit root for all countries, while the Phillips-Perron (PP) (1988) test always fails to reject this null hypothesis. The augmented Dickey-Fuller (1979) (ADF) test rejects it only for the case of Japan, when a time trend is not included. Unlike the other tests presented, the Kwiatkowski, Phillips and Schmidt (KPS) (1991) statistic tests the null of stationarity against the alternative of a unit root. The test rejects the stationary null at the 5 per cent level only for Canada when a time trend is present.

The SB test is a relatively powerful test of the null of a unit root, but it is sensitive to departures from independently and identically distributed (IID) Gaussian errors, and diagnostic tests suggest that this problem cannot be ruled out for any of these data sets. The most likely form of non-IID behaviour is heteroscedasticity. The ADF test is not affected by this form of heterogeneity in the errors, but it may lack power for alternatives with an autocorrelation coefficient near unity. Amano and van Norden (1992) show that the KPS test also

3. We use the most aggregate measure published, the “total” or “all items” measure. The technical work for this paper was completed in the spring of 1994. The sample includes data up to 1993 for all countries except Italy (where it ends in 1992). The starting point for the sample varies by country. See Table 1 in the appendix for details.

4. We thank Simon van Norden for the program used for these tests.

lacks power when the data have a moving average element with a negative root, which is generally the case for inflation data. However, they also provide evidence that using the KPS test of the stationary null hypothesis together with the PP test of a unit-root null can reduce the number of incorrect conclusions. With our results, this provides additional support for the presence of a unit root only in the case of Canada, where these tests agree in their conclusions, when a time trend is assumed to be part of the model.

Our conclusion from these tests is that, while the evidence is not overwhelming either way, it is not possible to reject convincingly the presence of a unit root in the inflation data for any of the G-7 countries. We take this to be a standard conclusion. However, we do not accept that this means that none of the G-7 countries was able to demonstrate any degree of control over inflation. Rather, we take these results as a starting point for a further investigation and ask what type of non-linear process may have generated the quasi-unit-root properties of the data.

The Markov switching model can, in principle, describe a process where the data appear to be nonstationary over finite samples, but where the process is strictly stationary in each of its states. In this case, the apparent nonstationarity in the data arises from the additional dynamics generated by the switching process. Alternatively, there could be episodes within an essentially stationary history where the data suggest strict nonstationarity. At such times, the best model of the inflation process may indeed be a characterization with a unit root, where shocks have permanent effects on the rate of inflation and there is no true nominal anchor on the system. As long as such episodes are infrequent and brief, however, it would be inappropriate to characterize inflation as inherently nonstationary.

In terms of parameter count, the most parsimonious alternative to the simplest random-walk model would be a two-state MSM with a stable autoregressive (AR) process together with a random-walk process. A slightly more complicated two-state model is an MSM with two stable AR processes, but with different long-run means and different dynamics. In this case, the monetary authority would be seen to retain control over inflation at all times, but to have changed its target at least once during the period.[5] It is of some interest whether we can identify regimes with stable high inflation rates for any of the G-7 countries. Previous research (Evans and Wachtel 1992, for the U.S. CPI, and Laxton, Ricketts and Rose 1994, for the Canadian GDP deflator) suggests that high and stable inflation regimes are not evident in the data.

To assess the statistical significance of multistate models as opposed to single-state characterizations, conventional tests are inappropriate. Standard distributions for conventional test statistics are valid only when all parameters are identified and their score vectors have non-zero variation under the null hypothesis. Standard distribution theory requires that the likelihood surface be locally quadratic in a neighbourhood of the null that contains the global optimum. When some parameters are not identified under the null hypothesis, the likelihood will be flat in a neighbourhood of the null, or the global maximum may be quite far from the null, so that the region between them is far from quadratic. Hansen (1992) proposes a method for calculating an approximation to the distribution of a valid test statistic using an empirical distribution of an upper bound for the likelihood-ratio (*LR*) statistic. A description of this calculation follows.

5. For example, a monetary authority may prefer to target a low rate of inflation, but if unanticipated shocks drive inflation much higher than the target, it may be unwilling to incur the costs needed to bring inflation back to the original target and instead choose to set and try to maintain a new higher target level of inflation. A model of this sort is described in Ball (1992).

Let P=0 be the vector of parameter restrictions that embed the null hypothesis in the general model and let B be the vector of unidentified parameters under the null. The likelihood function under the null is unaffected by these nuisance parameters (that is, $\partial L(P=0,\ B)/\partial B\ =\ 0$). The concentrated likelihood ratio, $LR_n(P, B)$, is the difference between the maximum likelihood of the general model for a given set of parameter values (P, B) and the optimum of the likelihood under the null hypothesis. Maximizing the likelihood ratio over (P, B) is equivalent to maximizing the likelihood, since the likelihood ratio is a level shift in the likelihood surface. The likelihood ratio can be decomposed into its mean and deviations from mean,

$$LR_n(P, B) = R_n(P, B) + Q_n(P, B)\ , \tag{6}$$

where R_n reaches a maximum of zero at the true parameter values, and where Q_n is a random variable in the parameters (P, B) and in the data. Hansen assumes that Q_n has a well-defined limiting distribution,

$$\frac{1}{\sqrt{n}} Q_n(P, B) \Rightarrow Q(P, B)\ , \tag{7}$$

where $Q(P, B)$ is a mean zero, Gaussian process. When the processes are standardized to have unit variance, the primary result is that

$$P\{supLR_n^*(P, B) \geq x\} \leq P\{supQ_n^*(P, B) \geq x\} \rightarrow P\{supQ^* \geq x\}\ , \tag{8}$$

where the star denotes the standardized process and *sup* is the supremum over values of (P, B). Thus, by calculating the distribution for $supQ^*$ we obtain an envelope for the distribution of the $supLR_n^*$ statistic. The inequality in (8) implies that a test based on this envelope distribution will provide a biased evaluation of the null hypothesis. It will tend to fail to reject the restriction to the null when the alternative is true with greater than the hypothetical frequency of the test. The effect of sample size, number of parameters, and so forth on power and size will all be model- and sample-dependent, so it is not possible to provide for a general correction for this bias.

The asymptotic distribution of Q^* can be approximated using stochastic simulations of its covariance generating function. To do this one must evaluate the statistics $Q^*_n(P, B)$ for a large number of values of the parameters P and B. At each point (p, b), the distribution of Q^* can be approximated by drawing a large number of IID $N(0, 1)$ sequences of size n, which are used to generate values of Q^* according to the sample counterpart of its covariance generating function. The distribution of $supQ^*$ is then approximated by the distribution of the supremum over (p, b) of the generated Q^* values.

The computational requirements for this procedure increase with the size of the vectors P and B, and with the number of values of each parameter at which the empirical distribution is approximated. In practice, the parameter space must be restricted and a limit imposed on the division of the parameter space into grid points. The parameter space must be restricted in such a way as to ensure that the likelihood surface is not too irregular with respect to P and B. A necessary condition for the test procedure is that the distribution of $Q^*_n(P, B)$ converge to that of Q^*. This implies, for example, that adding more parameter values to the grid will help identify the envelope distribution only if the likelihood surface is smooth in that region. To ensure that this holds in our application, we attempt to restrict the parameter space and use finer grids rather than expand the parameter space with a coarse grid. When the parameter space cannot be restricted, a priori, the test statistics have reduced power against alternatives within any restricted parameter space. Andrews (1993) notes that when the parameter space is unbounded, the test statistics diverge to infinity in probability. He recommends a restriction to the 15th to 85th percentiles of the parameter space.

The null hypothesis that we wish to test is a simple random walk process. Rejection in favour of a multistate model will be evidence for a particular type of nonstationarity in the data-generating process, one which may in fact be stable at all times, but with random switching in the process parameters. However, the easiest approach to embedding the unit root in the multistate models is to make one of the Markov states a random walk. The null is then obtained by imposing restrictions on the transition matrix such that only that one state is estimated. We perform two layers of tests – one in which a two-state MSM with a random-walk state is tested against the null of a single random-walk process, and a second in which a three-state MSM with a random-walk state is tested against the two-state MSM.

For the first test, the nuisance parameters consist of the transition parameters (p_{11} and p_{22}) and the AR(1) parameters for state 1 (α, β and σ). For each country, the grid values were set as given in Table 2b. The values for p_{11} and p_{22} were chosen so that each state would have a non-trivial probability of persisting in the data. The country specific values for α were chosen so that the mean of state 1 would include the long-run average inflation rate for the country. The total number of grid points for each country was 3 750. The empirical distributions were calculated using 10 000 vectors of $N(0,1)$ random numbers. The results for the *supLR** statistic for each country are given in Table 2a. We find that the two-state models, with the exception of Germany, can reject the single random-walk specification at the 90 or 95 per cent confidence level.[6]

For the test of the three-state model against the null of the two-state MSM, the nuisance parameters consist of p_{11}, the probability of remaining in state 1, λ_3, the preference for state 1 in transitions out of state 3, and the AR(1) parameters for

6. This does not mean that we view the simple random-walk model as appropriate for Germany. We will have more to say on this in subsequent sections.

state 1. The *supLR** statistics are unanimous in providing no support for the three-state model over the two-state model. This implies that there is insufficient statistical evidence in the data for additional states, at least when a random walk is part of the alternative model. The random walk was retained in both MSM alternatives for several reasons. One is that a primary objective was to show that the quasi-unit-root properties of the data come from specific episodes of historical experience and are not pervasive. We wished to do this without adding the complication of having to test for the unit root at the same time as we tested for additional states. We were also interested in testing the hypothesis that the random-walk regime within a three-state model could serve as a transition state in moves between stable states characterized by different long-run means.[7] We do not regard the presence of a unit-root regime within a more complicated MSM as implying that monetary policy truly loses control of the inflation process from time to time. However, within the context of simple time-series characterizations of the data-generating process, this may appear to be true. This may make it reasonable for agents who are not sure about the objectives of the monetary authority to form inflation expectations as if there were no long-run anchor on the process. We return to this question later when we examine the results of estimating three-state models.

We retain the three-state model in the study, despite the above results, for a number of reasons. First, we remind the reader that the Hansen test is known to be biased against rejecting the restrictions to the null. Thus, we cannot be too confident that the test results are providing correct conclusions. Moreover, the three-state model used in the tests was a restricted version. The random-walk

---

7. This was an interpretation given by Laxton, Ricketts and Rose (1994) for the role of the unit-root state in their representation of Canadian inflation as measured by the GDP deflator.

regime was imposed to act as a transition between the two freely estimated states. This was done to reduce the computational burden of the test, which remains very high even with this simplifying assumption. However, we have estimated a more general form of the three-state model, and there is some indication that our simplifying restriction may have influenced the test results. More important, however, we think that the evidence from the two-state estimations points to several instances where a qualitative case for additional states seems to emerge. We also report some other statistical tests of the adequacy of a two-state representation in the next section, which add to this case. Finally, even if three-state representations are not required for the relatively short periods we study in these estimations, there is a consensus emerging that something has changed in the way central banks are pursuing low and stable inflation. It may be that we do not need a separate state to capture this; however, for several countries inflation is at (or appears to be headed towards) lower values than have been observed over our historical samples. We think that this necessitates special attention to the possibility that a new state will be needed to represent future data and to the question of how the transition of expectations may be proceeding.

4 Two-state models

In this section, we report characterizations of the data for the G-7 countries obtained from two-state models. We examine the parameter estimates, the ex post probabilities and one-period-ahead forecasts of inflation to determine if the model makes a contribution to understanding the inflation process in each country. We also present the results of two statistical tests designed to gauge further the adequacy of the two-state model in this respect.

In addition to unconstrained estimates, two types of restrictions to the general MSM are examined. The first is a restriction that one of the states is a simple random walk. The second is a restriction on the transition properties to give an SSM or mixture model, in which the two states have constant probabilities that are independent of the previous state. The interpretation of an SSM in terms of monetary regimes is much less clear than for the MSM. In particular, it is difficult to interpret different long-run means as a policy choice, because there is always a fixed probability of moving to another state regardless of the recent outcomes, and there is no logical role for the monetary authority to change the outcome. As noted above, we include this possibility in our tests because if we could not reject it, then the contribution of the MSM to policy analysis would be questionable.

The estimation results and test statistics for the two-state models are presented in Tables 3a–3d. The accompanying figures (Figures 1–16) display the actual and expected series for CPI inflation (top panels, labelled *a*) and the calculated ex post probabilities assigned to the two states (bottom panels, labelled *b*) for each country for selected models. In the tables, the models are labelled according to their restrictions. The general Markov switching models are designated MS, while the simple switching models are designated SS. A “U” is appended to indicate unrestricted AR(1) states, while an "R" indicates that one of the states has been restricted to be a simple random walk. Thus, for example, SSR indicates a simple switching model with the restriction that one state is a random walk.

As we report the results, we will maintain a focus on whether the particular model passes certain diagnostic checks. One test we use for this purpose is based on the properties of the score or gradient vector at the estimated parameter values. If the model is correctly specified, the score vector will be serially uncorrelated. Serial dependence in the score vector with respect to the transition probabilities is normally taken as indicative of omitted Markov effects. If the score vector shows evidence of serial correlation, this would imply that the state associated with the transition probability was not accounting for the persistence of the state in the data. The alternative in this situation might be unrestricted transition probabilities (in the case of the SSM), time-varying transitions, or additional states. Another possibility is that additional lags are needed in the dynamic model or that the dynamic model is otherwise inadequate. Under the null hypothesis of no serial correlation, the test statistic is distributed as χ^2 with one degree of freedom in the case of the simple switching model. For the model with unrestricted transition structure, the statistic for the joint test on the two switching parameters is distributed as χ^2 with two degrees of freedom.[8]

The second test statistic we consider is a Wald test of the restrictions on the transition parameters that produce the SSM structure. This test does not consider the model's adequacy in terms of its residual properties. The two tests are thus complementary. If the Wald test does not reject a restriction to SSM structure, the residual diagnostics for the MSM are less interesting. However, it can happen that

8. We have also considered, but do not report, tests for serial correlation in the residuals, autoregressive conditional heteroscedasticity (ARCH) effects and omitted Markov effects in the models for each state.

the Wald test does not reject the SSM structure, but the score vector test nevertheless shows us that the restricted model is not acceptable on other grounds. In such cases, we tend to discount the Wald test and focus on the properties of the MSM.

In what follows, "state 1" is used to denote the state with *lower* mean inflation and "state 2" is used to denote the state with higher mean inflation (in the case of an unconstrained model where this mean exists) or the random-walk state. In almost every case, the state with the higher mean in the freely estimated model also has much higher persistence; that is, it has an autoregressive root much closer to the unit circle than does the state with the lower mean.

For two of the countries, Canada and Italy, both of the MS models pass the score vector test and reject the simple switching restriction on the transition parameters (Tables 3a and 3b). For these countries we therefore show the results graphically only for the MS models (Figures 1, 2, 9 and 10).

For Canada, the estimated persistence of high inflation in the MSU model is not significantly different from one, based on the normal t-test.[9] While the ex post probabilities for this model and the model with a random-walk state are almost exactly the same, their conditional expectations have some subtle differences. The MSU model tends to overpredict when inflation is low and underpredict when inflation is high. In the late 1950s and early 1960s, anticipated inflation is persistently high as actual inflation remains below the long-run mean for state 1. The sharp fall in inflation at the end of the sample period appears difficult to assign to either state for both models, and expectations in the MSU model remain biased towards higher inflation.

9. Note that while the t-test is biased in this case, it is biased towards rejecting the unit coefficient. Thus, our result would not be affected by using the correct distribution.

For Italy, the MSU model estimates a stable high-inflation state for the period of the 1970s and early 1980s. The transition out of the high-inflation state occurs abruptly in 1984, when inflation is still above the moderate level of the late 1980s. The credibility of the moderate-inflation state rises quickly in the MSU model, as opposed to gradually in the MSR model. This is the only case we have seen in which the estimated high-inflation state is not close to being dynamically unstable.

For France, either the restriction to SS is rejected by the Wald test or there are problems indicated by the score-vector tests (Table 3a). Moreover, where the restriction to SS is rejected, the score-vector test on the MS results also indicates problems. In sum, we have no two-state result for France that stands up to the diagnostic tests. We show all four results graphically in Figures 3–6.

The estimated high-inflation state in the models with two unrestricted AR(1) states has high negative persistence, implying possibly unstable oscillations when inflation is high. Both of these models assign only the peaks in inflation to the second state. In addition, the conditional expectations from these models significantly overpredict inflation during the late 1980s, when inflation returns to a moderate level. When we impose that state 2 is a random walk, both estimates continue to show long periods of persistent errors in expectations, but not during the stable period of the late 1980s. The models assign a relatively long period, from the late 1960s to the mid 1980s, to the random-walk regime. This was the period of high and rising inflation. For the MS model, the exit from the random-walk regime is gradual, beginning only after inflation has reached the long-run mean of state 1. The SS model has trouble distinguishing between the two states in the early and later parts of the sample.

For two of the economies, the United Kingdom and the United States, the point estimates from the models with unrestricted AR(1) processes indicate unstable roots for state 2 (Tables 3c and 3d). Such results are both suspect, econometrically, and unacceptable as economic models of inflation dynamics. For these countries, then, we present the graphical results for the two restricted models (Figures 13–16).

For the United Kingdom, the restriction to SS structure is *not* rejected, but the score test indicates that the resulting model has specification problems, which is not the case for the MSR model. For the United States, no acceptable two-state model is found. The restriction to SS is rejected, but the score test on the MSR model indicates remaining specification problems.

For both of these countries, the two-state MSR models assign the random-walk regime to the high-inflation period in the 1970s. The pre- and post-1970s periods are assigned to moderate-inflation regimes with means of 5 per cent and 2.9 per cent, respectively. The results indicate that, in both countries, the credibility of the stable-inflation regime was established gradually, starting in 1981. The two-state model also seems unable to capture the strong push towards even lower inflation in the United Kingdom during the 1990s.

For Germany and Japan, some of the model estimates are suspect because of a very low estimate for inflation variability (the standard deviation) for one of the states (Tables 3b and 3c). We suspect that such extreme results may be unreliable, and in a maximum-likelihood system such doubts must extend to all estimated parameters.

For Japan, this problem arises for both models with unrestricted AR(1) dynamics. The results for the two restricted models are shown in Figures 11 and

12. The statistical tests cannot reject the SS model with a random-walk state. This seems odd. In the period between 1972 and 1981, the ex post probabilities fluctuate erratically between the random walk and the low-inflation state. There appears to be very little support in the data of this period for a regime with less than 2 per cent mean inflation, yet the switching models are unable to characterize the pre-1980 data with just one state. We are left with the impression that, despite statistical evidence to the contrary from the Hansen test, the two-state model is inadequate as a description of the Japanese data.

For Germany, the MS models appear better able to distinguish a second state in the data. The SSR model assigns most of the historical period to a moderate-inflation regime, with a mean of 3.6 per cent. The random-walk state is characterized as rather ephemeral (the probability of remaining in that state is estimated to be only 0.29), but in most periods it is nevertheless assigned between 20 and 30 per cent probability. The SSU model estimates a very low value of σ_1, which makes these estimates suspect. We therefore show only MS results in Figures 7 and 8. The MSR estimate for the mean in the low-inflation state is 2.9 per cent – higher than the Bundesbank's "target" for inflation, which is generally taken to be about 2 per cent. The MSU estimates a 2.7 per cent mean for the low-inflation state, but its estimate of the high-inflation mean is close to the peak of observed inflation with persistence close to unity.

Examining the overall results for the two-state models, we find that in several cases these models separate the data into periods of high inflation and periods of moderate or low inflation. Often the estimated moderate-inflation state is a compromise between low and somewhat higher inflation, and is unable to pick up details in either "substate" successfully. Another common feature is that the expectations series derived from the two-state representations show episodes of persistent errors in one-step-ahead forecasts. This occurs especially during periods when, according to the estimates, the regime appears to have changed and uncertainty about inflation is likely to have been high.

In several cases, the high-inflation state is itself unstable or has a persistence parameter that is close to unity, so that imposing a unit root emerges as a natural alternative. Nevertheless, the two-state characterization of inflation does make an important contribution to explaining the source of quasi-unit-root properties in the data. By and large, these properties appear to have come from the 1970s and to have virtually disappeared in the 1980s in the G-7 economies.

An interesting feature of the two-state models is that the estimated *standard deviation* of inflation in the high-inflation states is usually much higher than it is in the low-inflation states. For Canada, for example, the standard deviation of inflation is almost four times higher in the high-inflation state than in the moderate-inflation state. For the United States it is more than twice as high, confirming Evans and Wachtel's (1992) result from a slightly different model. In the context of switching models, a higher variance of inflation in one state also produces a higher conditional variance of inflation one period ahead. High inflation, whether it is described by a process with a unit root or a stable AR(1), seems to be associated with increased uncertainty about future inflation.

5 Three-state models

We now describe the results for a three-state model in which there are no restrictions on the transitions between the states.[10] The estimation results are provided in Table 4. Only MS models are considered, and we again impose that one state is a simple random walk. The results are presented in graphical form in Figures 17–23. A series of diagnostic tests of the resulting models are reported in Table 5.

In the case of Canada (Figure 17), a regime with low average inflation of about 1.6 per cent is identified and given some probability in the early 1960s. As in the two-state model, the random-walk regime is assigned to periods of high inflation, and a moderate-inflation regime is identified for the periods preceding the run-up of inflation in the early and late 1970s. A puzzling item is the weight given to the low-inflation regime, as opposed to the moderate-inflation regime, at the point where inflation falls in the early 1980s. The abruptness of the decline in inflation in 1983 may have suggested the return of low inflation, but it is the moderate-inflation regime that does become established for the remainder of the 1980s. An interesting point is that the recent further decline in inflation to levels not seen since the early 1960s has brought the credibility of the low-inflation regime back to over 40 per cent at the end of our sample.

The estimated transition probabilities imply a particular order of occurrence for the three states. The unit-root regime serves as the *only* route for transition from state 2 to state 1, but transitions from the low-inflation state are always to the moderate-inflation regime.

---

10. Recall that, for the Hansen test, we used a simplified version of this model, where the transition matrix was constrained such that the random-walk state serves as a transition state between the two stationary states in both directions. This restriction is not supported by the freely estimated transition structure, which may have contributed to our failure to find statistical support for use of a third state. We have not done the Hansen test for the general case, because the computational burden is enormous, owing to the extra parameters that must be added to the grid.

France (Figure 18) provides an interesting case where the unit-root regime is defined only by a very small number of outliers in the data. A high-inflation regime with a mean of about 8.6 per cent is estimated for the periods of high and rising inflation (it is used much like the random-walk state is for Canada during the 1970s). This regime persists well into the mid-1980s, but it is then replaced by the low-inflation regime, which has a mean of about 2.8 per cent. Credibility in the low-inflation regime is slow to build, as memory persists of the long period of high inflation that began in the late 1960s.

An interesting puzzle from the estimation for Germany (Figure 19) is that the mean of the high-inflation state, at 8.3 per cent, is beyond the range of values seen in the data. Rising inflation is expected to culminate in a level that is higher than the peaks of actual inflation. It may be that the monetary authority was able to intervene successfully to bring inflation under control before it reached the levels implied by the estimated process. However, it is interesting to speculate whether the vaunted German fear of inflation, extant based on memories of the hyperinflation during the time of the Weimar Republic, may lead agents to extrapolate small observed increases in inflation into expectations of much higher mean values of inflation. We are uncomfortable pressing such an argument too far, but we can assure the reader that we have checked this result and it seems robust.

It is important to note that the unit-root regime is almost never assigned any weight in this model for Germany. Its credibility increases only at the *troughs* in the inflation series. We think that this makes much more sense than the apparent acceptance of the single-state random-walk model from the Hansen tests and the direct tests for the presence of a unit root. This result suggests that the Bundesbank does indeed have a special sort of credibility, in that agents retain their confidence that there is a nominal anchor in the face of inflationary pressures.

The estimated model for Italy, like the one for Canada, assigns some credibility to a low-inflation state at the time when inflation falls rapidly in the early 1980s, even though the levels of inflation remain far above the mean of that state (Figure 20). The persistence estimated for the low-inflation state is such as to make the long decline appear to be a gradual adjustment to a low long-run mean. The estimated transition probabilities indicate a corner solution – moves out of state 1 are always to state 2, moves out of state 2 are always to state 3 and moves out of state 3 are always to state 1. We are not sure how to interpret such results. While we have tried to ensure that we do not have just a local peak of the likelihood function, we cannot rule out this possibility.

The results for Italy provide a good example of an important general point. If we look only at the means, it may appear that states 1 and 2 are really very similar. This is certainly the impression one gets immediately from Figure 20. But this ignores the dynamic properties of the model. These two states have radically different autoregressive structures (Table 4), with state 1 having very high persistence of deviations from the mean and state 2 much less so. This is what permits the estimator to use state 1 during the early 1980s, despite the fact that inflation remains well above even the mean for state 2.

The three-state characterization of inflation in Japan (Figure 21) is more precise than we obtained from the two-state model. The additional state is used to describe the period of moderate inflation in the late 1970s. There is a clear transition in terms of ex post probabilities from the random-walk state to the state with moderate inflation and finally to the state with low inflation. This case provides fairly convincing evidence, we submit, that despite the results of the formal tests, a three-state representation of these data is much more satisfactory.

The estimated model for the United Kingdom (Figure 22) gives credibility to only two of the states. The random walk is almost never assigned any probability, and all the details of the results suggest that we have a solution where only two states are used. Note, furthermore, that the estimated high-inflation state is stable with a persistence parameter of only 0.36. This leaves us with a puzzle. The three-state results suggest that only two states are required and that they both have stable means. Yet the two-state estimates failed to find such a point; only the model with a random-walk regime passed the residual diagnostic tests, and we did not find an acceptable model with stable high inflation. This may indicate that we have not identified the best two-state result.

Although the above discussion would suggest caution in using the results for the United Kingdom, those results indicate that the fall in inflation in the 1990s has not been interpreted as a move to a low-inflation state. Indeed, as is also the case for Italy, there is no “low-inflation” state identified; for the United Kingdom, state 1 has a mean above 5 per cent.

In contrast, the estimation for the United States does pick up a low-inflation regime, with a mean of about 2 per cent, because of the experience of the early 1960s. The other regime with a stable mean shows an average of just under 4 per cent inflation. These two states replace the composite 3 per cent regime identified from the two-state results. The use of the random-walk regime is virtually the same in the two exercises.

As in the Canadian three-state model, the decline of inflation in the early 1980s brought a rise in credibility for the low-inflation state in the United States, which was subsequently lost as moderate inflation persisted through the rest of the decade. Moreover, according to our estimates, credibility of the low-inflation regime has not returned in the United States in the 1990s as it has in Canada (Figure 23). The results suggest that the United States is likely in a regime with just under 4 per cent trend inflation.

There is not much support in the estimated models for the use of the random-walk regime solely as a transition state to be used temporarily in moves between the two stable states. The random-walk model is used primarily as a description of periods with high inflation, generally associated with the major oil-price shocks. This may mean that high inflation is inherently unstable or is not described well by the simple AR(1) specification. We have not calculated unit-root tests for an individual state within the switching models, so a direct conclusion cannot be drawn.

In the results for two-state models, we noted a clear association of higher volatility with higher levels of inflation. A similar picture emerges from the three-state models, but with some qualifications. For Canada, France and Japan there is a clear ranking of level and volatility of inflation as measured by the estimated standard deviations. For Canada, for example, the standard deviation of the state with moderate (just over 4 per cent) average inflation is 30 per cent higher than that of the state with low (1.6 per cent) average inflation; and the standard deviation of changes in inflation in the random-walk model (which is assigned to the periods of highest inflation) is five times as large as the result for the low-inflation regime. For Italy and the United States, the same major difference in volatility between the periods of high inflation (generally characterized as state 3) and the other periods also is clear. However, for these countries, state 2, with a higher average rate of inflation than state 1, has a lower standard deviation. The United Kingdom provides the only example we have seen where the estimated standard deviation for the state with a unit root is actually smaller than that obtained for one of the stationary states (state 2). However, it is important to remember that state 3 is essentially unused in the U.K. model. The basic association of high level with high volatility comes through very clearly in a comparison of the results for states 1 and 2 for the United Kingdom – the state with lower inflation has about a third the volatility. Germany, as usual, provides us with somewhat different results. state 3 has the highest standard deviation, but in this case that is not associated with episodes of high inflation. The state that is used for the periods of higher inflation (state 2) has the lowest estimated standard deviation. We are not inclined to put much weight on this result; state 3 is almost never used and state 2 only for two episodes – Germany has almost always been in state 1.[11]

This research has given us some insight into the non-linearities that make estimation of the Markov switching model difficult. As with other highly non-linear models, estimation can be complicated by the existence of many local peaks in the likelihood function. The estimates are obtained by numerical methods that solve for the roots of the gradient to the likelihood function. We have found that care is required to assure that a suggested solution is not simply a local peak. Moreover, the solution with the highest likelihood value for a given sample may not provide the asymptotic global maximum.[12] In short, we have found that the estimation of Markov systems cannot be treated as a "one-pass" process. It is very

11. Moreover, recall that for the two-state MSR model for Germany, an association between higher volatility and a higher level of inflation is suggested.

12. We have seen cases in other similar research where the identified best combination of parameters for a subsample is no longer the best point when the sample is extended. This is sometimes, though not always, a reflection of a small set of competing regions of parameter space, where we tend to get sample-dependent identification of which characterization is most likely given that particular sample.

important to try alternative starting assumptions, for example, to weed out local maxima.

We have performed a number of diagnostic tests, which are reported in Table 5, on the three-state results. These include tests for omitted Markov effects, tests for ARCH problems and tests for serial correlation problems in the residuals.[13] These tests do not give any of the estimated models a clean bill of health. For the Canadian model, for example, there is still evidence of omitted Markov effects and indication of both ARCH and serial correlation problems with the state 3 model. For the U.S. model there is also evidence of omitted Markov effects and ARCH problems in states 2 and 3. While we must take these results seriously as evidence that there may be problems remaining in the specification, it is worth noting that these tests applied to MSMs have low power and are biased towards excessive rejection of the null hypothesis in small samples (see Hamilton 1990).

To conclude this section, let us return to the issue of whether the addition of a third state to the model has improved our ability to explain the data. The Hansen test gives us no support for the need for this elaboration, but we know that this test, by its nature, has low power in small samples to identify the need for additional structure. We have presented a variety of evidence from the two-state results, including statistical tests of the adequacy of those models and some qualitative evaluation of the plausibility of the identified history of inflation, to support a case that extra structure was necessary. The results from the three-state estimations are mixed in this respect. In the case of the United Kingdom and of

13. Tests are reported for all countries. In cases, such as that of the United Kingdom, where we have not identified an acceptable three-state model, these results should be interpreted with caution. Moreover, for some countries we have some estimated parameters at the boundary of the permitted range (of probability measures in the transition matrix). For some of these results we have near-singularity in the required covariance matrix and have computed the test statistics eliminating the effect of such parameters.

Germany, the three-state results gave us some useful information but no strong case that three states are needed to describe the historical sample. For four countries, however – Canada, the United States, France and Japan – a reasonably strong case emerges that the addition of a third state has added something important to the model and our ability to describe the modern history of inflation.[14]

14. We are not quite sure what to say about the results for Italy. The third state is certainly used, but we have problems with how it is used.

6 Conclusions

This paper applies multistate, Markov switching models to CPI inflation data from the G-7 countries. A two-state Markov switching model with one state imposed to be a random walk is preferred in statistical tests to a single state (random-walk) description of the data, for all countries but Germany. The probability estimates for the Markov model show that the random-walk state is generally assigned to periods of high and rising inflation – predominantly the periods in the 1970s and early 1980s when these economies were subjected to substantial oil-price shocks. At other times, stable autoregressive processes tend to be chosen to describe the data. The evidence also indicates that the inclusion of a state with a unit root may not be necessary. For Canada, for example, we found a representation with a technically stable, albeit high variance, high-inflation state, and a stable low-inflation state that performs better than the same model with a unit root imposed on one of the states. Moreover, for the United Kingdom (and perhaps for Germany), we found evidence that multiple local maxima may have prevented the identification of two stable states.

The three-state models highlight some additional features in the inflation processes for some of the countries. For Canada, Japan and the United States, an additional state is given noteworthy probability in the explanation of episodes of history. For the North American countries, a low-inflation regime is found to have persisted for some time in the early 1960s and is given weight again in the early 1980s. Note, however, that this state is assigned rising probability in Canada in recent years, whereas the estimates for the United States show no evidence of credibility of a return to a low-inflation regime in the 1990s. Indeed, for the United States, the results show that the credibility of a low-inflation regime was building in the mid-1980s but has dissipated subsequently.

Our results indicate that Japan has been in a fully credible low-inflation regime since about 1984, while Germany, France and the United Kingdom have all been in moderate-inflation regimes since the mid-1980s. Even following the reunification, German inflation did not move out of the moderate-inflation regime that describes most of the sample. In all the countries, therefore, monetary policy in the 1980s can be characterized as having moved towards or having maintained a low or moderate rate of inflation. Some countries appear to have started later on this path than others or to have acquired (or to be acquiring) credibility in the new regime at a slower pace than others, but there is a clear general trend towards regimes with lower (and usually more stable) inflation.

There is also fairly systematic evidence of a general association of higher volatility of inflation with higher levels of inflation. This point extends to comparisons of low versus moderate inflation for some, but not all countries.

This analysis of the inflation process is directed, in part, towards establishing simple models from which to derive realistic proxies for private agents' expectations of inflation. As a proxy for inflation expectations, the conditional expectations generated from these models incorporate past information about inflation as well as a forward-looking component that takes into account the possibility of switches in monetary regimes. These expectations exhibit persistent bias at times of regime change. Laxton, Ricketts and Rose (1994) have shown that the establishment of a new regime, such as the announced low-inflation policy in Canada, can be difficult, because agents must be convinced that random disturbances that temporarily raise the inflation rate are not evidence of a return to a previous regime and because evidence of lower inflation is not sufficient, in itself, to change underlying perceptions of policy. The data from the United Kingdom provide a case in point. Inflation has fallen significantly in the

1990s, but according to our results, no low-inflation state (in the sense that we see for most other countries) is even identified as a possibility as yet for the United Kingdom.

Table 1: Tests for unit roots: augmented Dickey-Fuller (ADF), Kwiatkowski, Phillips and Schmidt (KPS), Phillips-Perron (PP) and Sargan-Bhargava (SB) tests

Country	Sample period[a]	ADF lags[b]	PP lags[b]	ADF t-ratio	KPS	PP t-ratio	SB
Canada	1954–93	1	6	-2.46	0.30	-1.77	0.31*
				-2.25	0.15*	-0.87	
France	1954–93	1	6	-2.29	0.16	-2.69	0.53*
				-2.10	0.13	-2.53	
Germany	1954–93	1	6	-2.85	0.18	-2.30	0.35*
				-2.87	0.13	-2.22	
Italy	1968–92	1	5	-1.82	0.15	-1.76	0.28*
				-2.02	0.14	-1.03	
Japan	1971–93	5	4	-3.26*	0.42	-1.80	0.63*
				-0.65	0.10	-2.55	
United Kingdom	1963–93	1	5	-2.04	0.14	-1.95	0.47*
				-1.98	0.14	-1.39	
United States	1954–93	2	6	-1.70	0.30	-2.00	0.34*
				-1.46	0.14	-1.70	

Notes: With the exception of SB, each test is conducted first without and then with a time trend included.

a. The data are first differences of logs of the consumer price indexes (annual averages, all components) for the G-7 countries taken from data bases maintained at the Bank of Canada. The data come originally from national sources.

b. The ADF tests use lagged differences of inflation to correct for additional serial dependence. The number of lags is determined using t-statistics on the lag coefficients. The PP tests use a non-parametric estimator of the variance-covariance matrix to correct for serial dependence. The truncation parameter for the adjustment factor is set to the square root of the sample size as suggested by Andrews (1991).

* indicates rejection of a unit root at the 0.05 significance level, except for the KPS test, where it indicates rejection of stationarity.

Table 2a: Specification tests for two- and three-state models

Hansen standardized *LR* test				
Country	**2-state vs. random walk**		**3-state vs. 2-state model**	
	supLR*	**p-value**	***supLR****	**p-value**
Canada	2.84	0.036	1.25	0.88
France	2.93	0.046	1.94	0.47
Germany	2.25	0.170	1.65	0.78
Italy	2.76	0.073	0.22	0.99
Japan	3.83	0.002	1.00	0.65
United Kingdom	2.74	0.068	1.74	0.83
United States	3.11	0.024	0.48	1.00

Table 2b: Grid values for the mean of state 1

Country	2-state vs. random walk	3-state vs. 2-state model
Canada	2.5, 3.0, 3.5, 4.0, 4.5	2.5, 3.0, 3.5, 4.0, 4.5
France	8.5, 9.0, 9.5, 10.0, 11.0	3.0, 4.0, 5.0, 6.0, 7.0, 8.0
Germany	2.5, 3.5, 4.5, 5.5, 6.5	2.5, 3.5, 4.5, 5.5, 6.5
Italy	4.5, 5.0, 5.5, 14.0, 15.0	4.5, 5.5, 11.5, 12.5, 13.5
Japan	2.5, 3.5, 4.5, 5.0, 5.5	3.5, 4.0, 5.0, 5.5
United Kingdom	2.5, 3.5, 4.5, 5.5, 6.5	2.5, 3.5, 4.5, 5.5, 6.5, 7.0
United States	2.5, 3.5, 4.0, 4.5, 5.0	2.5, 3.5, 4.0, 4.5, 5.0

Note: The grid values for the remaining parameters were as follows: β = (0.3, 0.4, 0.5, 0.6, 0.7, 0.8); σ = (0.7, 0.9, 1.1, 1.3, 1.5), p_{11} = (0.60, 0.675, 0.750, 0.825, 0.90), p_{22} = (0.60, 0.675, 0.750, 0.825, 0.90), λ_3 = (0.20, 0.40, 0.60, 0.80); where β is the persistence parameter in state 1, σ is the standard deviation of inflation in state 1, p_{11} is the probability of remaining in state 1, p_{22} is the probability of remaining in state 2, and λ_3 is the preference for state 1 in transitions out of state 3.

Table 3a: Maximum-likelihood estimates for two-state Markov switching (MS) and simple switching (SS) models

Country/Model[a]		Average Log L	α_1	β_1	σ_1	p_{11}	α_2	β_2	σ_2	p_{22}	Score test	Wald test
Canada												
	MSU	-1.78	0.88 (0.26)	0.74 (0.06)	0.74 (0.12)	0.85 (0.10)	1.66 (2.13)	0.80 (0.28)	2.80 (0.62)	0.69 (0.20)	2.00 (0.37)	3.33* (0.07)
	SSU	-1.81	0.90 (0.23)	0.74 (0.04)	0.69 (0.13)	0.64 (0.10)	0.83 (0.48)	0.92 (0.13)	2.70 (0.38)	0.36	4.08** (0.04)	
	MSR	-1.79	0.91 (0.23)	0.73 (0.05)	0.73 (0.12)	0.90 (0.08)	0	1	2.88 (0.40)	0.68 (0.13)	1.33 (0.51)	6.85** (0.01)
	SSR	-1.82	0.93 (0.23)	0.74 (0.04)	0.69 (0.12)	0.64 (0.10)	0	1	2.75 (0.39)	0.36	2.28 (0.13)	
France												
	MSU	-2.06	0.85 (0.43)	0.81 (0.06)	1.35 (0.20)	0.91 (0.05)	20.1 (2.34)	-0.70 (0.22)	1.99 (0.73)	0.36 (0.25)	12.4** (0.00)	1.71 (0.19)
	SSU	-2.08	0.82 (0.33)	0.82 (0.05)	1.39 (0.20)	0.89 (0.05)	20.3 (0.49)	-0.72 (0.10)	1.92 (0.42)	0.11	13.3** (0.00)	
	MSR	-2.25	2.06 (0.21)	0.28 (0.04)	0.56 (0.11)	0.80 (0.09)	0	1	3.52 (0.37)	0.87 (0.06)	6.81** (0.03)	20.0** (0.00)
	SSR	-2.37	1.72 (0.19)	0.29 (0.02)	0.27 (0.11)	0.22 (0.08)	0	1	3.16 (0.32)	0.78	2.18 (0.14)	

a. The samples used for these estimations are the same as reported in Table 1 for the unit-root tests. Standard errors are shown in parentheses for the parameter estimates. P-values are shown in parentheses for the test statistics.

* Indicates significance at the 0.1 level. ** Indicates significance at the 0.05 level.

Table 3b: Maximum-likelihood estimates for two-state Markov switching (MS) and simple switching (SS) models

Country/Model[a]		Average Log L	α_1	β_1	σ_1	p_{11}	α_2	β_2	σ_2	p_{22}	Score test	Wald test
Germany												
	MSU	-1.38	1.23 (0.54)	0.54 (0.15)	0.63 (0.21)	0.70 (0.26)	0.50 (0.62)	0.92 (0.18)	1.07 (0.20)	0.76 (0.19)	13.6** (0.00)	1.92 (0.17)
	SSU	-1.28	1.14 (0.06)	0.40 (0.02)	0.06 (0.03)	0.16 (0.06)	0.77 (0.29)	0.82 (0.08)	1.00 (0.12)	0.84	0.44 (0.51)	
	MSR	-1.39	1.34 (0.03)	0.54 (0.10)	0.59 (0.15)	0.65 (0.13)	0	1	1.13 (0.18)	0.72 (0.12)	14.4** (0.00)	4.57** (0.03)
	SSR	-1.41	0.99 (0.30)	0.73 (0.10)	0.85 (0.18)	0.71 (0.16)	0	1	1.21 (0.27)	0.29	0.44 (0.51)	
Italy												
	MSU	-2.09	2.54 (0.49)	0.54 (0.07)	0.84 (0.16)	0.84 (0.08)	12.1 (3.24)	0.28 (0.20)	2.76 (0.63	0.86 (0.11)	0.95 (0.62)	8.11** (0.00)
	SSU	-2.33	2.60 (0.35)	0.53 (0.04)	0.81 (0.15)	0.58 (0.09)	11.3 (0.63)	0.36 (0.06)	2.45 (0.44)	0.42	12.4** (0.00)	
	MSR	-2.27	2.54 (0.41)	0.54 (0.06)	0.84 (0.17)	0.90 (0.07)	0	1	4.12 (0.52)	0.90 (0.08)	0.19 (0.90)	14.6** (0.00)
	SSR	-2.39	2.77 (0.48)	0.53 (0.05)	0.67 (0.22)	0.36 (0.13)	0	1	3.48 (0.47)	0.64	3.29* (0.07)	

a. The samples used for these estimations are the same as reported in Table 1 for the unit-root tests. Standard errors are shown in parentheses for the parameter estimates. P-values are shown in parentheses for the test statistics.

* Indicates significance at the 0.1 level. ** Indicates significance at the 0.05 level.

Table 3c: Maximum-likelihood estimates for two-state Markov switching (MS) and simple switching (SS) models

Country/Model[a]		Average Log L	α_1	β_1	σ_1	p_{11}	α_2	β_2	σ_2	p_{22}	Score test	Wald test
Japan												
	MSU	-2.36	2.91 (0.08)	-0.37 (0.03)	0.03 (0.01)	0.49 (0.23)	1.80 (1.42)	0.67 (0.18)	4.10 (0.68)	0.90 (0.07)	2.99 (0.22)	2.23 (0.14)
	SSU	-2.41	2.91 (0.08)	-0.37 (0.03)	0.03 (0.01)	0.17 (0.07)	1.80 (0.61)	0.67 (0.13)	4.10 (0.48)	0.83	12.8** (0.00)	
	MSR	-2.07	0.92 (0.24)	0.47 (0.03)	0.73 (0.14)	0.78 (0.11)	0	1	5.81 (0.63)	0.42 (0.18)	5.65* (0.06)	0.80 (0.37)
	SSR	-2.09	0.93 (0.24)	0.47 (0.03)	0.73 (0.14)	0.71 (0.10)	0	1	5.67 (0.63)	0.29	0.51 (0.47)	
United Kingdom												
	MSU	-2.36	1.60 (0.78)	0.55 (0.06)	1.35 (0.32)	0.54 (0.23)	0.48 (1.70)	1.26 (0.17)	2.56 (0.59)	0.44 (0.18)	6.04* (0.05)	0.01 (0.92)
	SSU	-2.36	1.60 (0.46)	0.55 (0.05)	1.35 (0.27)	0.55 (0.11)	0.51 (0.54)	1.26 (0.09)	2.56 (0.41)	0.45	2.01 (0.91)	
	MSR	-2.47	2.45 (0.46)	0.52 (0.08)	1.68 (0.26)	0.93 (0.05)	0	1	5.62 (0.54)	0.86 (0.10)	3.49 (0.17)	15.7** (0.00)
	SSR	-2..49	2.16 (0.46)	0.53 (0.05)	1.48 (0.32	0.56 (0.13)	0	1	4.10 (0.49)	0.44	0.00 (0.03)	

a. The samples used for these estimations are the same as reported in Table 1 for the unit-root tests. Standard errors are shown in parentheses for the parameter estimates. P-values are shown in parentheses for the test statistics.

* Indicates significance at the 0.1 level. ** Indicates significance at the 0.05 level.

Table 3d: Maximum-likelihood estimates for two-state Markov switching (MS) and simple switching (SS) models

Country/Model[a]		Average Log L	α_1	β_1	σ_1	p_{11}	α_2	β_2	σ_2	p_{22}	Score test	Wald test
United States												
	MSU	-1.76	0.97 (0.37)	0.64 (0.06	1.06 (0.15)	0.85 (0.14)	0.80 (1.47)	1.21 (0.19)	1.42 (0.39)	0.54 (0.20)	3.48 (0.18)	1.74 (0.19)
	SSU	-1.79	0.87 (0.33)	0.64 (0.06)	1.05 (0.17)	0.66 (0.12)	0.46 (0.43)	1.24 (0.10)	1.35 (0.27)	0.34	4.99** (0.03)	
	MSR	-1.84	1.09 (0.28)	0.63 (0.07)	1.06 (0.14)	0.95 (0.04)	0	1	2.70 (0.41)	0.85 (0.10)	16.2** (0.00)	18.4** (0.00)
	SSR	-1.88	1.11 (0.31)	0.64 (0.06)	1.10 (0.17)	0.69 (0.12)	0	1	2.32 (0.38)	0.31	9.19** (0.00)	

a. The samples used for these estimations are the same as reported in Table 1 for the unit-root tests. Standard errors are shown in parentheses for the parameter estimates. P-values are shown in parentheses for the test statistics.

* Indicates significance at the 0.1 level. ** Indicates significance at the 0.05 level.

Table 4: Maximum-likelihood estimates for a three-state Markov switching model (State 3 is imposed to be a simple random walk.)[a]

	Canada	France	Germany	Italy	Japan	United Kingdom	United States
α_1	0.76 (0.27)	2.01 (0.22)	1.02 (0.22)	0.75 (0.57)	0.77 (0.31)	2.49 (0.46)	0.88 (0.37)
β_1	0.50 (0.06)	0.28 (0.04)	0.61 (0.07)	0.80 (0.06)	0.50 (0.11)	0.51 (0.08)	0.59 (0.09)
Mean of state 1	1.52 (0.47)	2.79 (0.21)	2.62 (0.35)	3.75 (2.24)	1.54 (0.43)	5.08 (0.70)	2.15 (0.81)
σ_1	0.47 (0.18)	0.57 (0.10)	0.64 (0.11)	1.01 (0.33)	0.75 (0.15)	1.70 (0.25)	1.10 (0.22)
p_{11}	0.64 (0.15)	0.80 (0.09)	0.81 (0.09)	0.65 (0.17)	1.00 (0.00)	0.94 (0.04)	0.86 (0.09)
λ_1[b]	1.00	1.00	0.42 (0.16)	1.00	0.45 (0.27)	1.00	1.00
α_2	1.49 (0.29)	2.33 (0.40)	2.25 (0.39)	4.76 (0.47)	3.98 (0.61)	9.54 (0.57)	0.94 (0.41)
β_2	0.66 (0.05)	0.73 (0.04)	0.73 (0.09)	0.14 (0.08)	0.34 (0.08)	0.36 (0.12)	0.75 (0.12)
Mean of state 2	4.38 (0.53)	8.63 0.93)	8.33 (1.36)	5.53 (0.22)	6.03 (0.78)	14.9 (2.79)	3.76 (1.19)
σ_2	0.61 (0.12)	1.08 (0.19)	0.45 (0.13)	0.58 (0.13)	1.80 (0.41)	4.69 (0.52)	0.93 (0.17)
p_{22}	0.77 (0.09)	0.70 (0.10)	0.72 (0.12)	0.88 (0.09)	0.84 (0.11)	0.85 (0.10)	0.93 (0.05)
λ_2	0.00	0.14 (0.10)	1.00	0.00	1.00	0.00	0.00
σ_3	2.51 (0.37)	6.98 (0.50)	1.57 (0.41)	4.53 (0.56)	7.83 (0.66)	4.66 (0.57)	2.745 (0.40)
p_{33}	0.68 (0.13)	0.00	0.00	0.83 (0.10)	0.76 (0.15)	0.00	0.89 (0.07)
λ_3	0.63 (0.16)	0.48 (0.16)	1.00	1.00	0.00	1.00	1.00
average log L	-1.71	-1.94	-1.31	-2.16	-1.90	-2.40	-1.80

a. The estimation samples are the same as reported in Table 1. Standard errors in parentheses.
b. The λs are the conditional transition probabilities; λ_1 is the preference for state 2 in transitions out of state 1, λ_2 is the preference for state 1 in transitions out of state 2 and λ_3 is the preference for state 1 in transitions out of state 3.

Table 5: Results of score vector tests on three-state models

Country	AR tests S1	AR tests S2	AR tests S3	ARCH tests S1	ARCH tests S2	ARCH tests S3	Omitted Markov effects S1	Omitted Markov effects S2	Omitted Markov effects S3	S1-S3
Canada	0.13 0.72	1.81 0.18	3.66 0.06*	0.04 0.84	0.24 0.62	3.79 0.05**	3.31 0.07*	5.66 0.02**	3.23 0.07*	11.72 0.01**
France	0.39 0.53	5.72 0.02**	24.08 0.00**	0.39 0.53	0.00 0.99	5.68 0.02**	1.05 0.31	0.34 0.56	-----	1.25 0.53
Germany	0.58 0.45	11.04 0.00**	2.83 0.09*	6.61 0.01**	5.49 0.02**	4.06 0.04**	3.28 0.07*	0.89 0.35	-----	3.56 0.17
Italy	0.33 0.56	5.08 0.02**	0.29 0.59	7.36 0.01**	1.91 0.17	1.20 0.27	12.15 0.00**	4.70 0.03**	2.04 0.15	17.17 0.00*
Japan	2.75 0.10*	0.15 0.70	53.69 0.00**	0.80 0.37	2.30 0.13	33.28 0.0**	-----	1.47 0.22	0.08 0.78	1.69 0.42
United Kingdom	1.92 0.17	16.13 0.00**	1.86 0.17	0.70 0.40	5.96 0.01**	6.78 0.01**	0.68 0.41	0.33 0.57	-----	0.79 0.67
United States	1.17 0.28	1.01 0.32	0.36 0.55	1.88 0.17	6.60 0.01**	4.41 0.04**	3.85 0.05**	4.62 0.03**	4.64 0.03**	12.88 0.00**

Notes: Significance levels (p-values) are shown in parentheses.
* Indicates significance at the 0.1 level.
** Indicates significance at the 0.05 level.

Figure 1
a) Actual and Markov expected CPI inflation
Canada - MSU

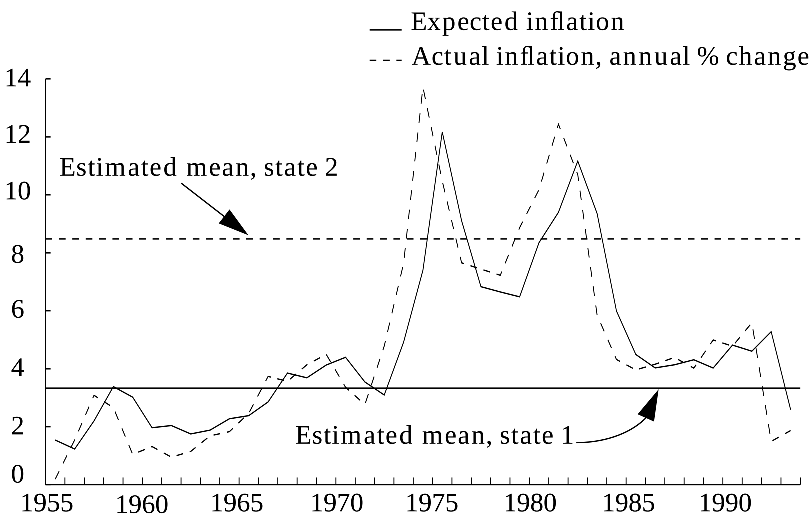

b) Probabilities of states 1 and 2

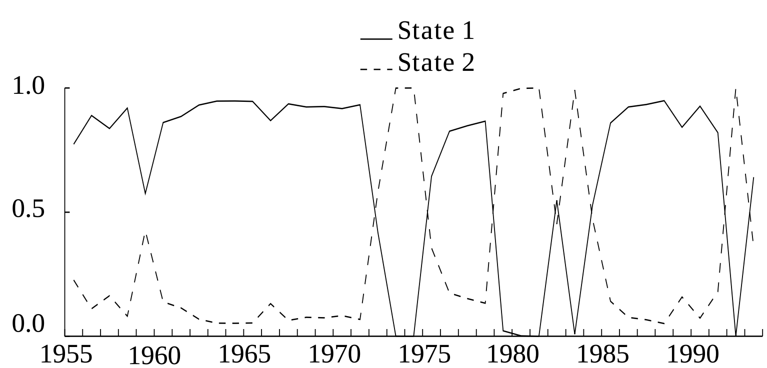

Figure 2
a) Actual and Markov expected CPI inflation
Canada - MSR

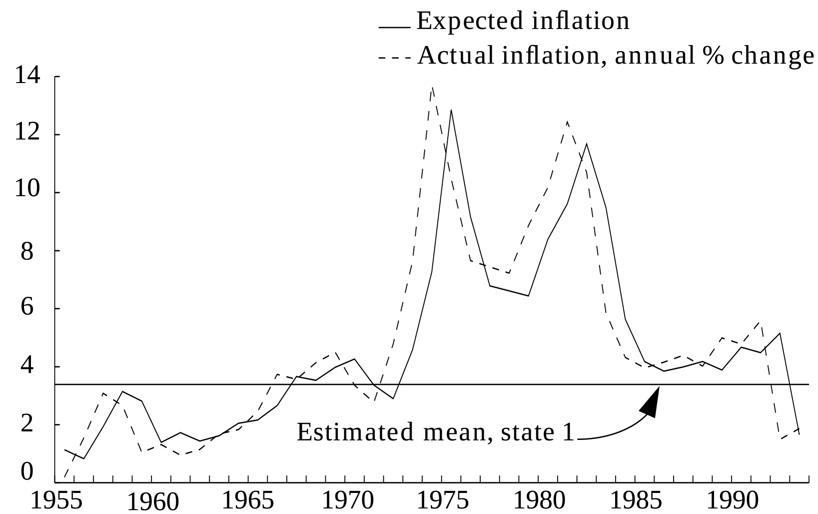

b) Probabilities of states 1 and 2

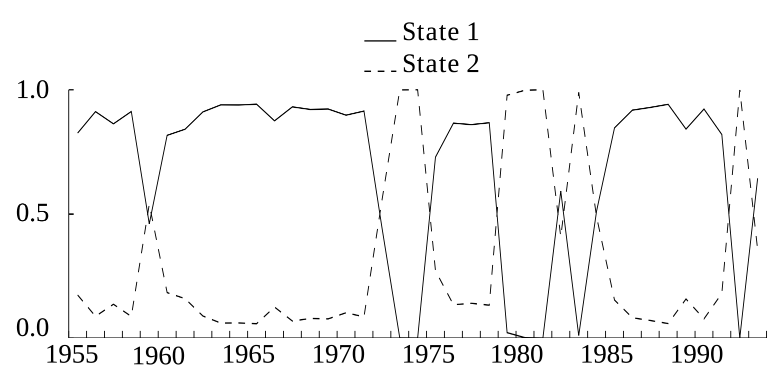

Figure 3

a) Actual and Markov expected CPI inflation France - MSU

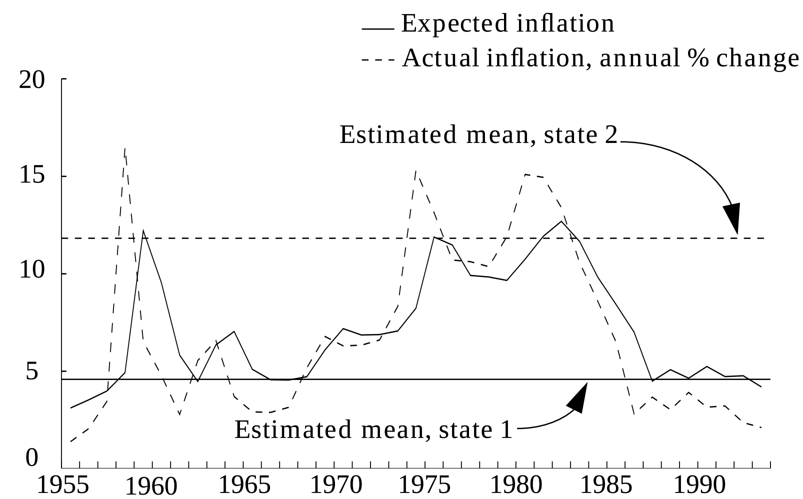

b) Probabilities of states 1 and 2

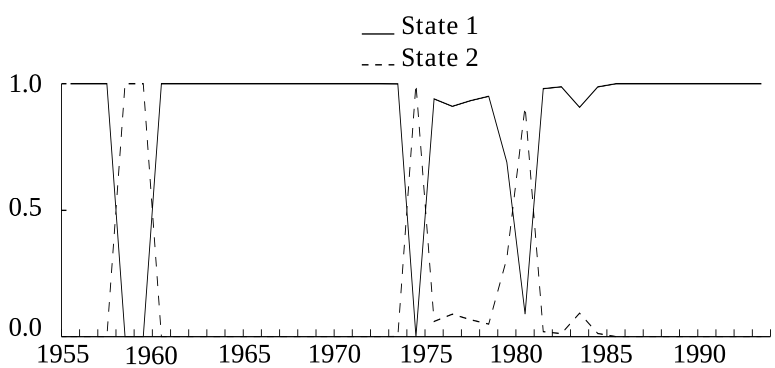

Figure 4

a) Actual and Markov expected CPI inflation France - SSU

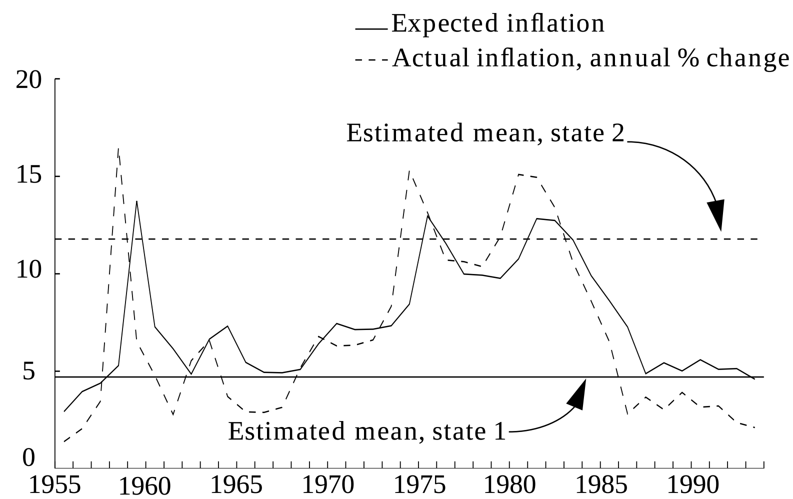

b) Probabilities of states 1 and 2

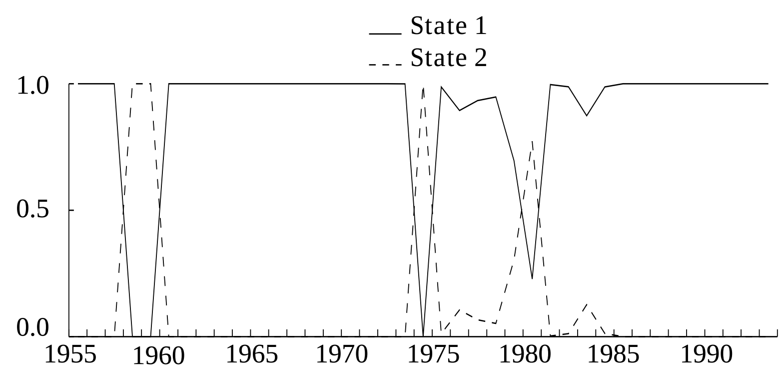

Figure 5
a) Actual and Markov expected CPI inflation France - SSU

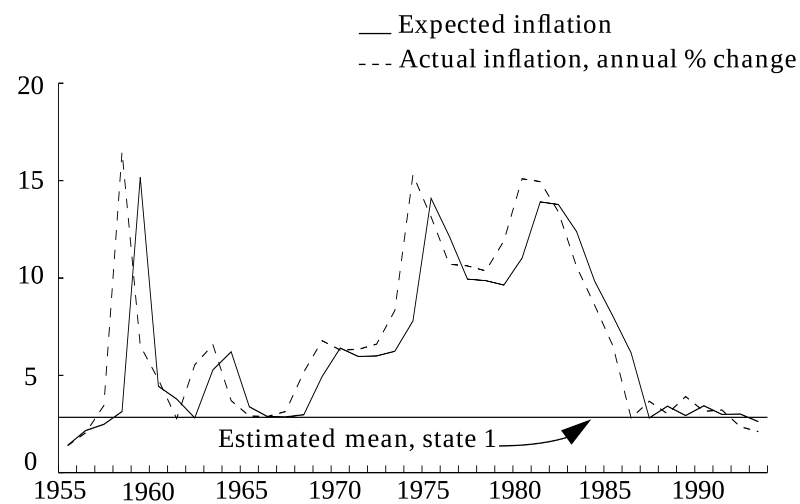

b) Probabilities of states 1 and 2

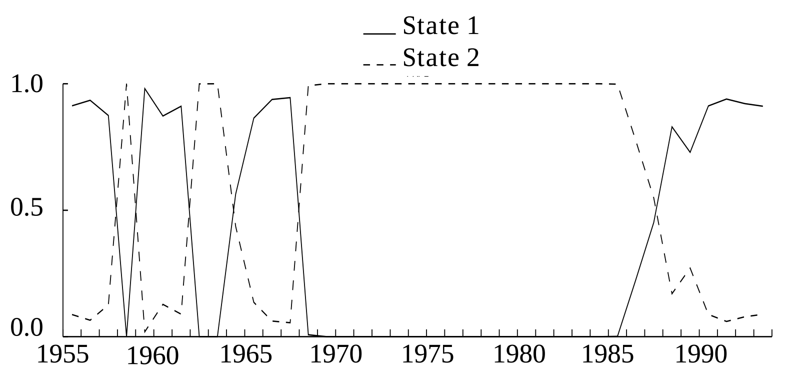

Figure 6
a) Actual and Markov expected CPI inflation France - SSR

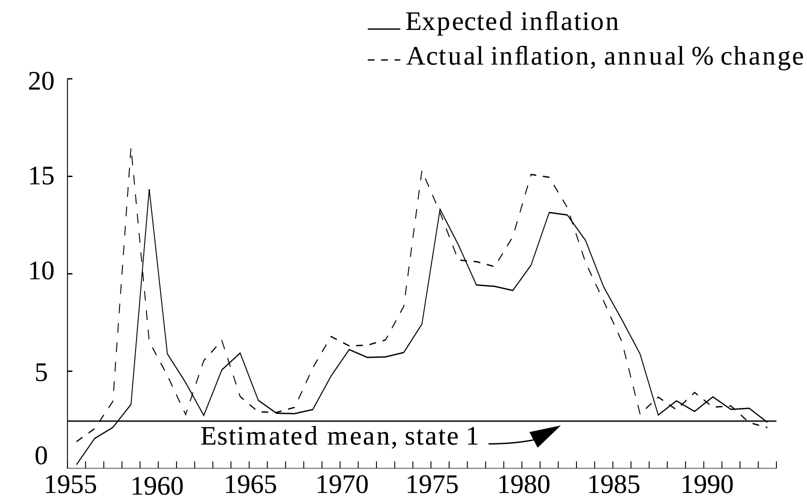

b) Probabilities of states 1 and 2

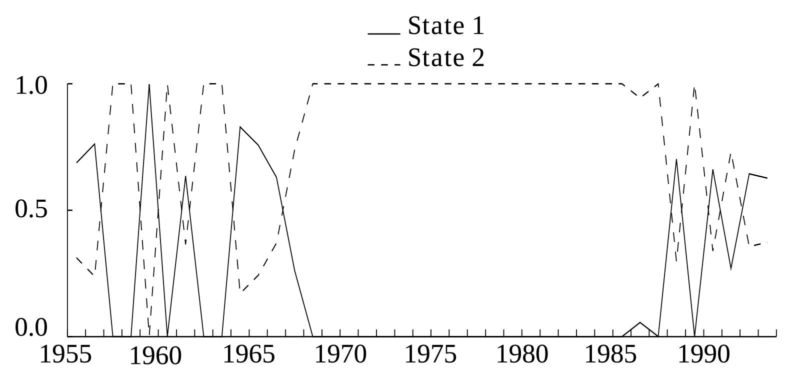

Figure 7

a) Actual and Markov expected CPI inflation Germany - MSU

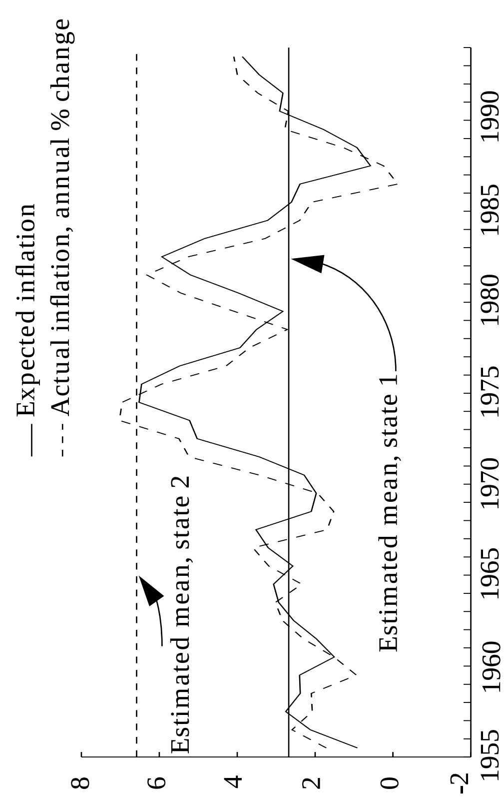

b) Probabilities of states 1 and 2

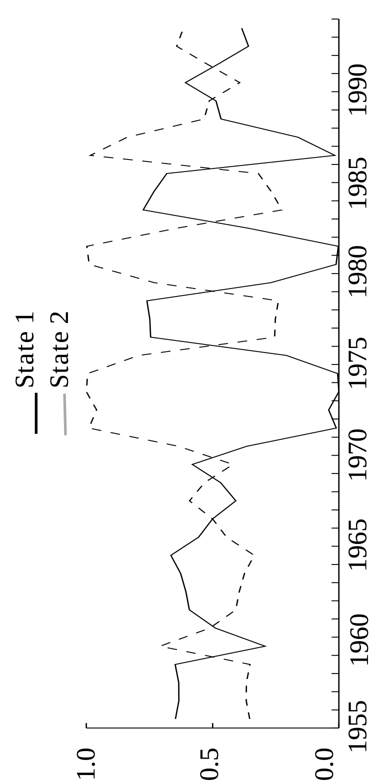

Figure 8

a) Actual and Markov expected CPI inflation Germany - MSR

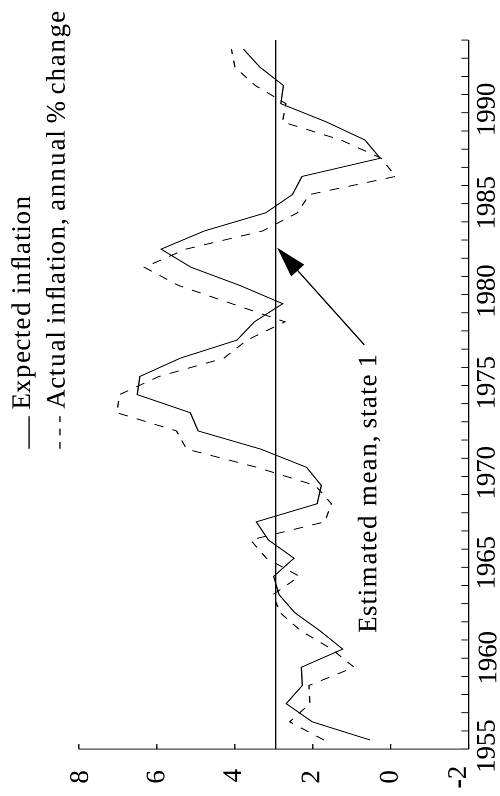

b) Probabilities of states 1 and 2

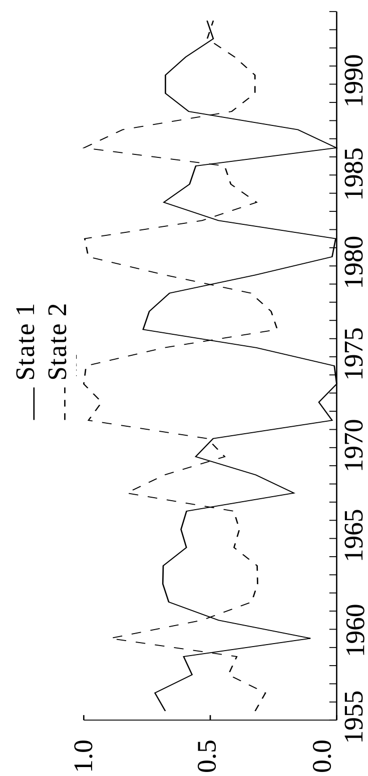

Figure 9

a) Actual and Markov expected CPI inflation Italy - MSU

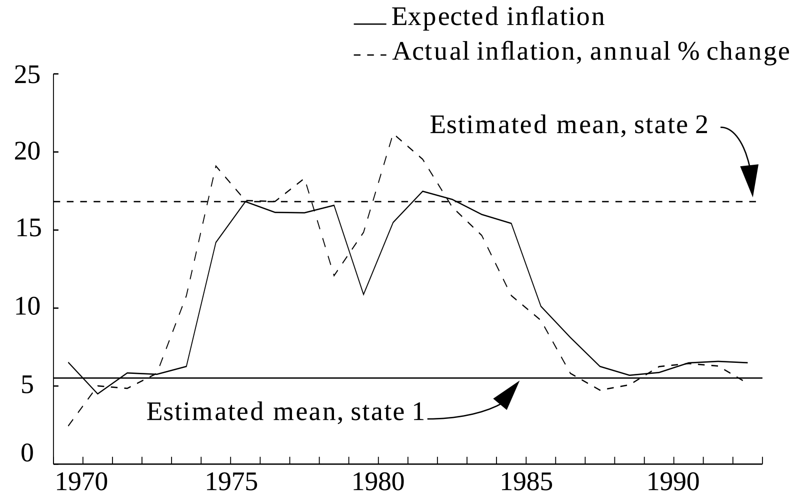

b) Probabilities of states 1 and 2

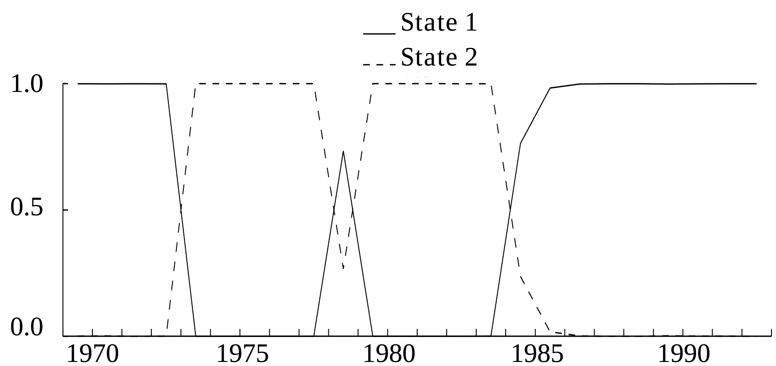

Figure 10

a) Actual and Markov expected CPI inflation Italy - MSR

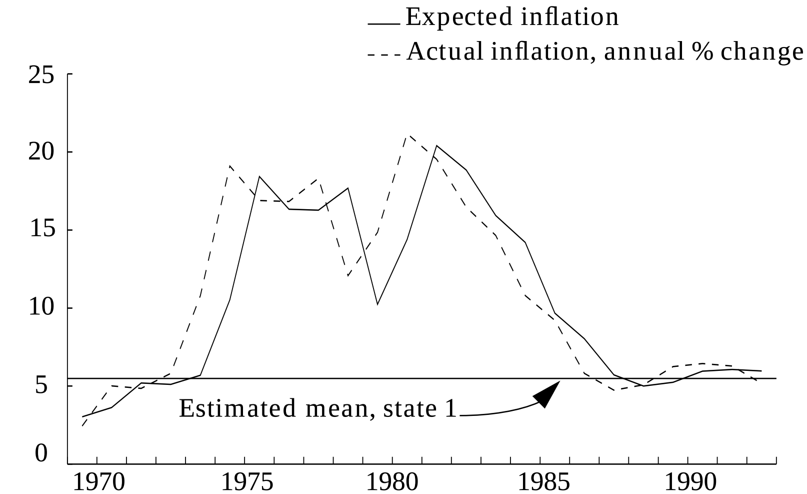

b) Probabilities of states 1 and 2

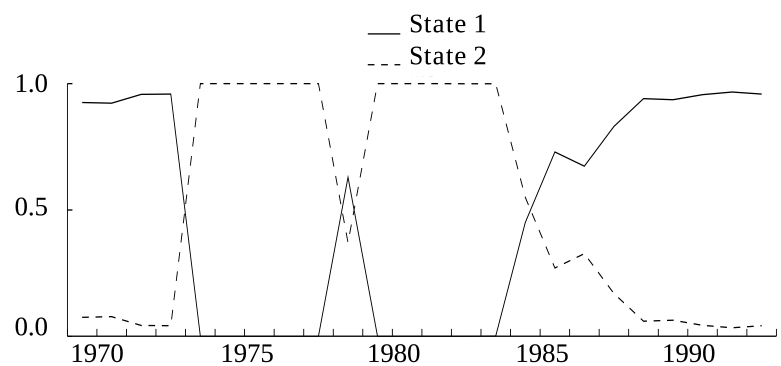

Figure 11
a) Actual and Markov expected CPI inflation
Japan - MSR

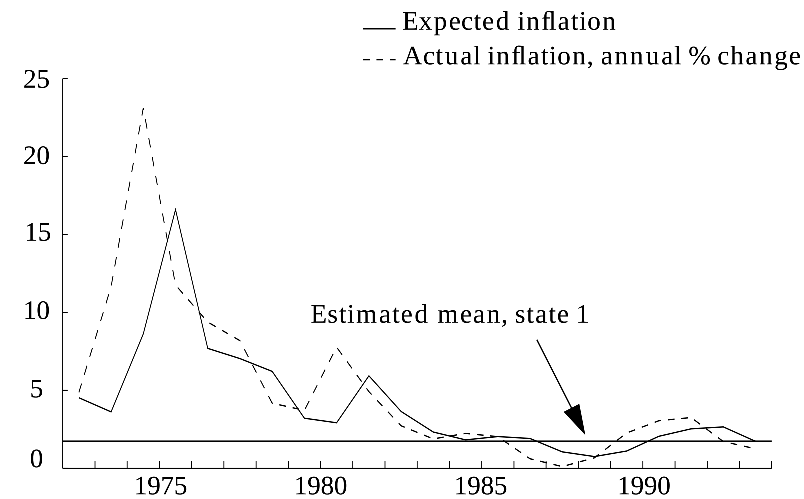

b) Probabilities of states 1 and 2

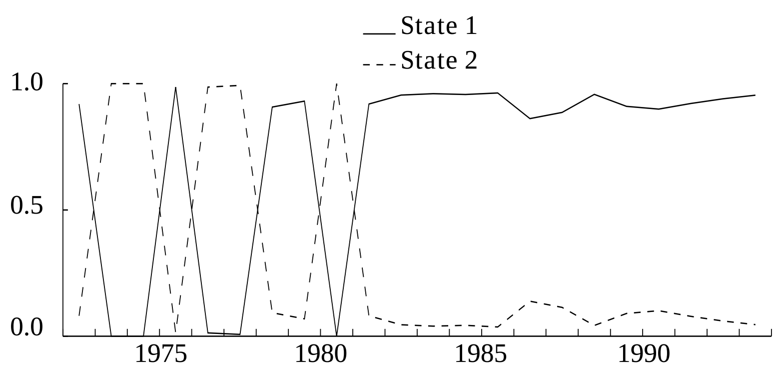

Figure 12
a) Actual and Markov expected CPI inflation
Japan - SSR

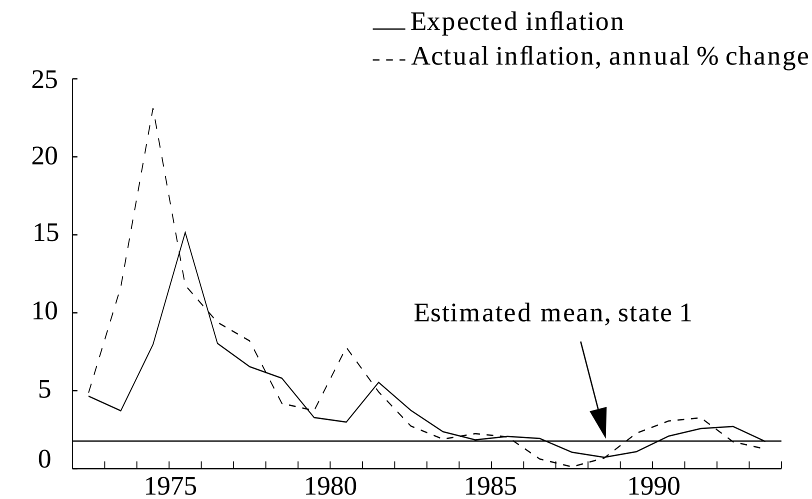

b) Probabilities of states 1 and 2

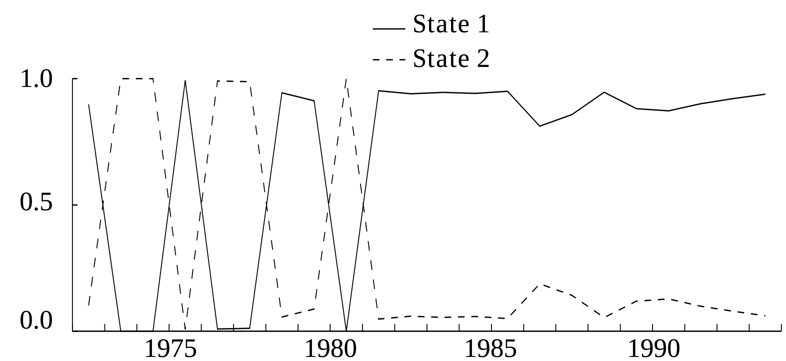

Figure 13
a) Actual and Markov expected CPI inflation United Kingdom - MSR

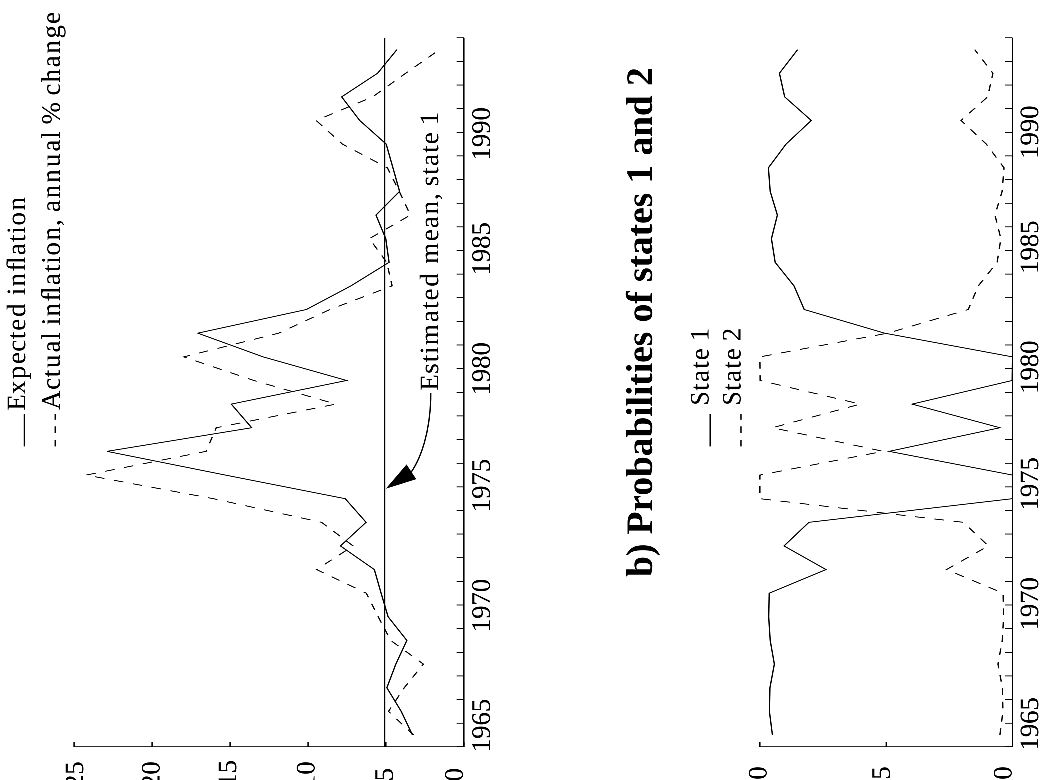

Figure 14
a) Actual and Markov expected CPI inflation United Kingdom - SSR

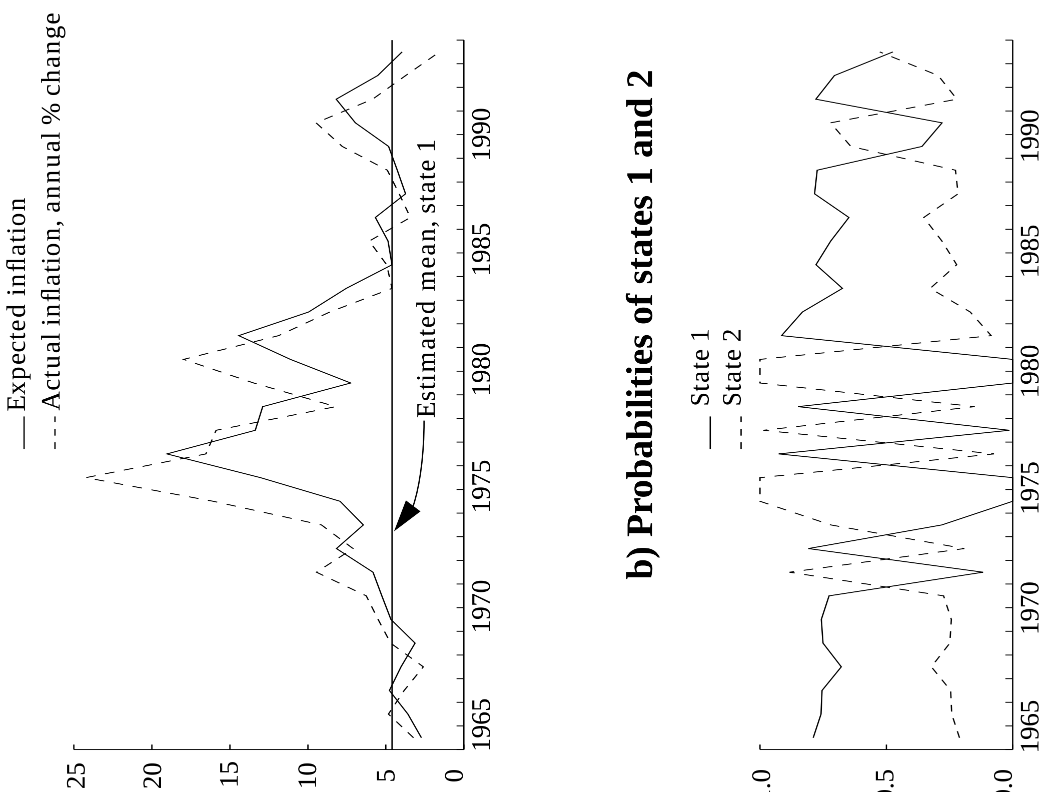

Figure 15

a) Actual and Markov expected CPI inflation United States - MSR

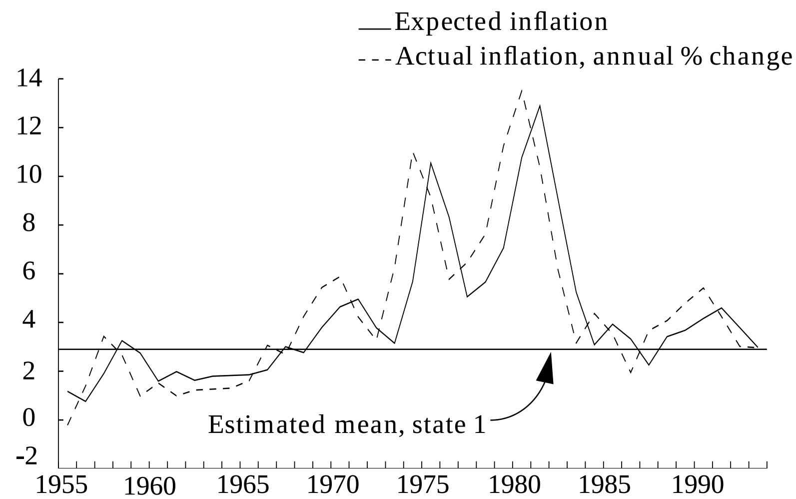

b) Probabilities of states 1 and 2

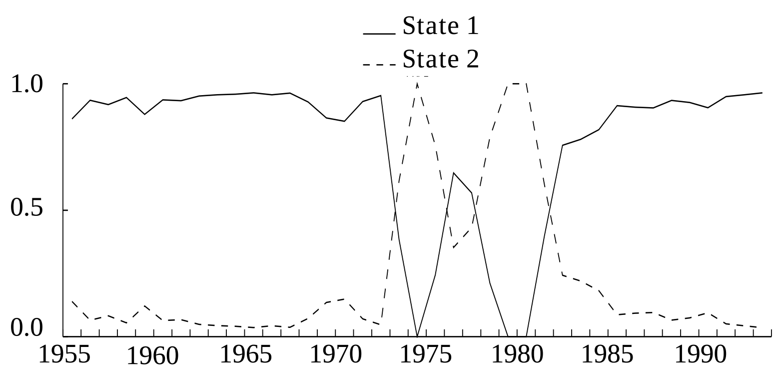

Figure 16

a) Actual and Markov expected CPI inflation United States - SSR

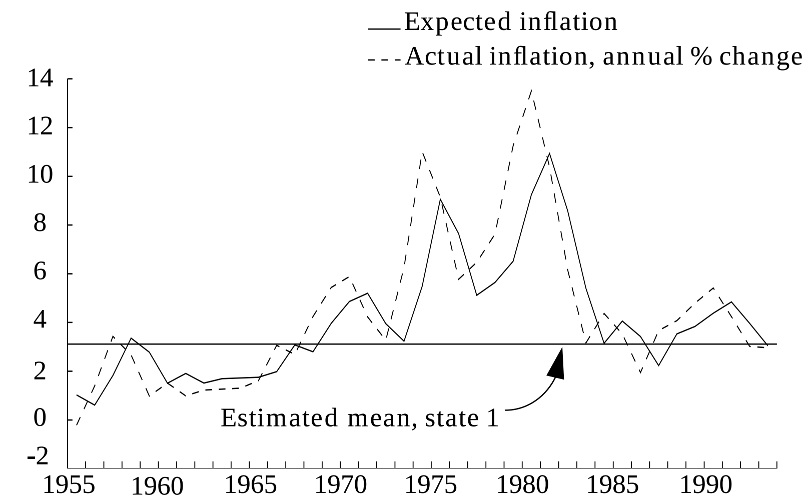

b) Probabilities of states 1 and 2

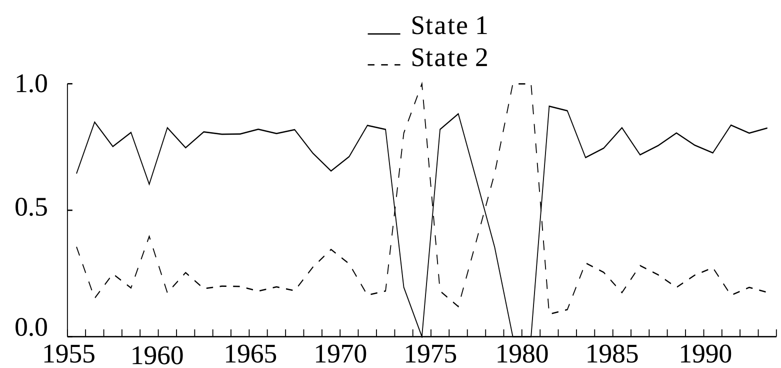

Figure 17
a) Actual and Markov expected CPI inflation
three–state model
Canada

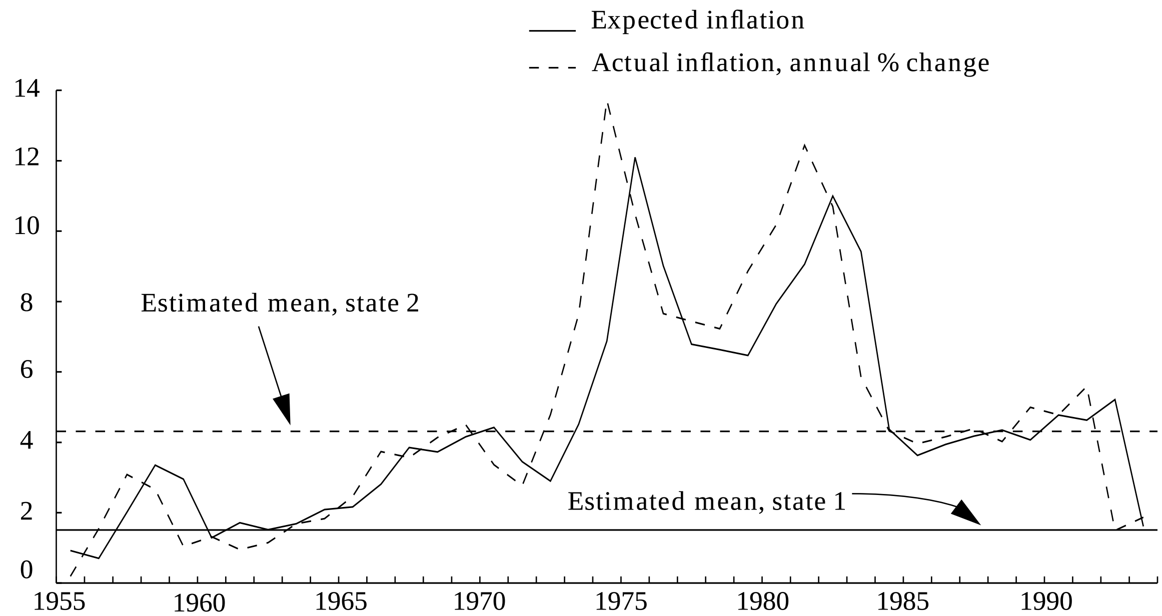

b) Probabilities of states 1, 2 and 3

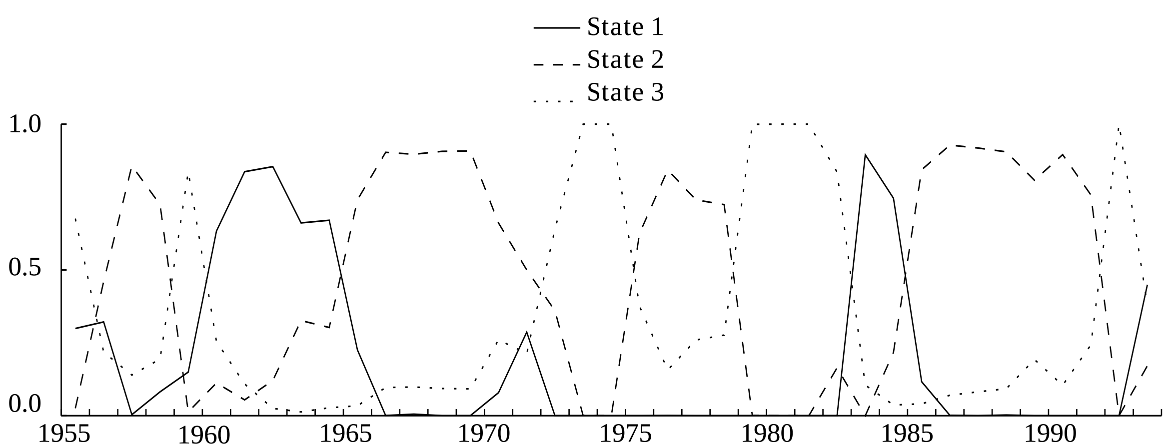

Figure 18
a) Actual and Markov expected CPI inflation
three–state model
France

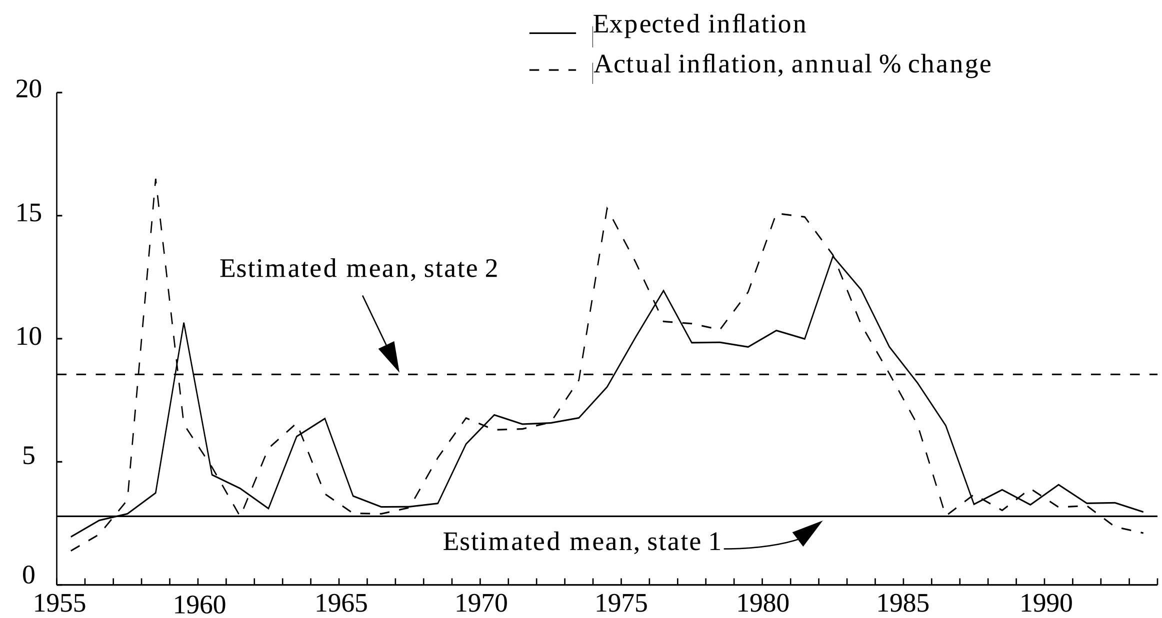

b) Probabilities of states 1, 2 and 3

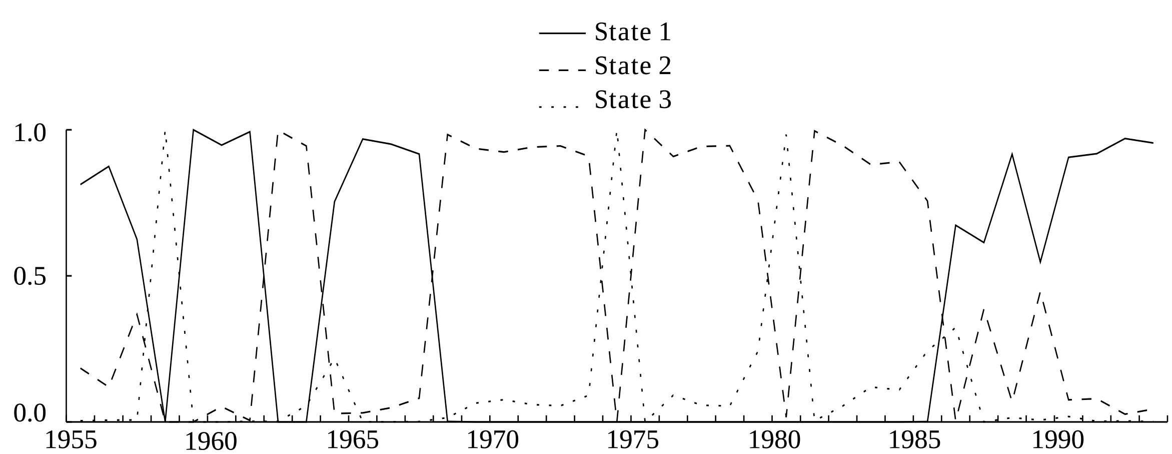

Figure 19
a) Actual and Markov expected CPI inflation
three–state model
Germany

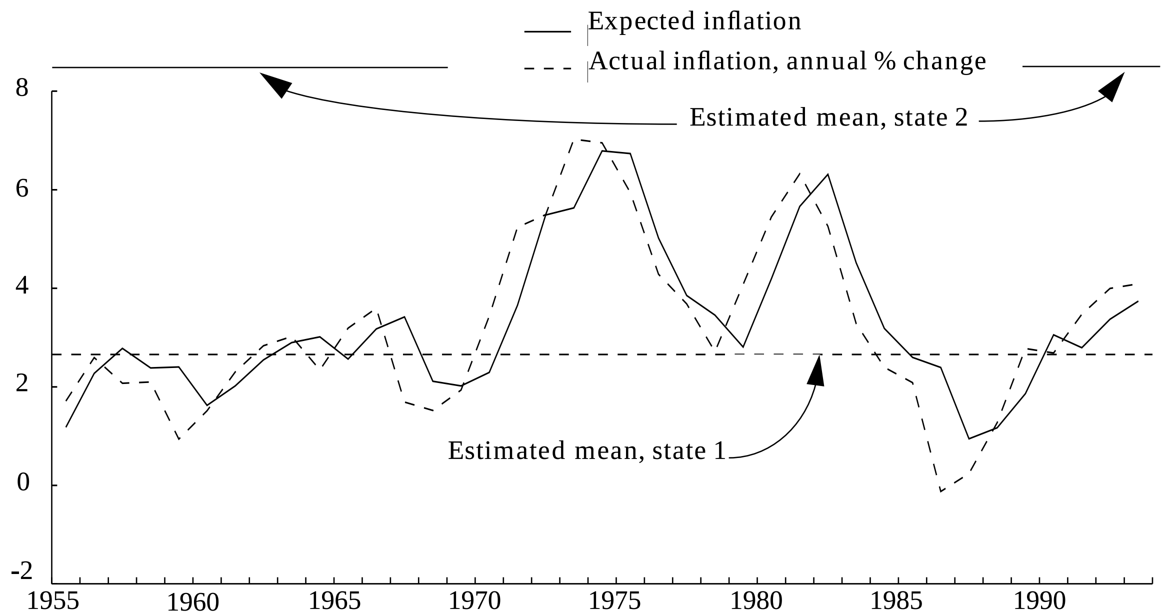

b) Probabilities of states 1, 2 and 3

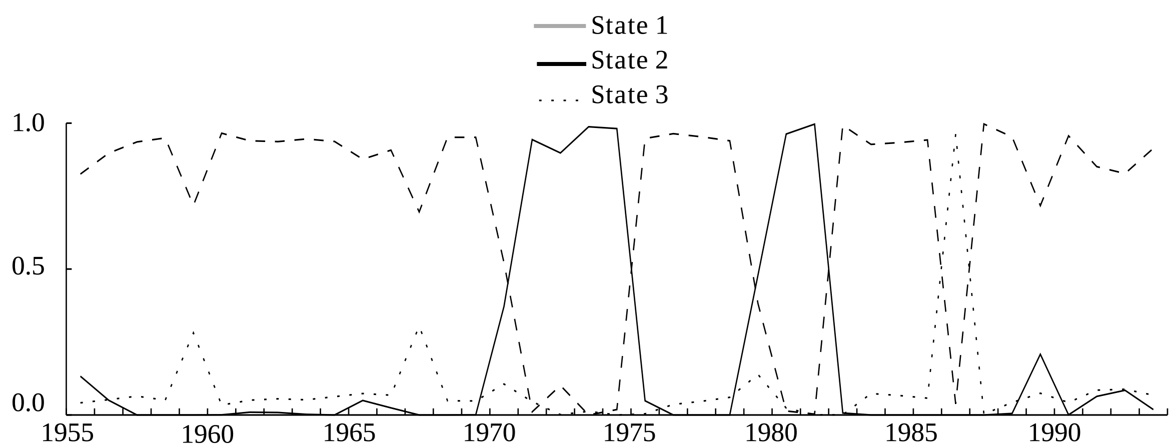

Figure 20
a) Actual and Markov expected CPI inflation three–state model Italy

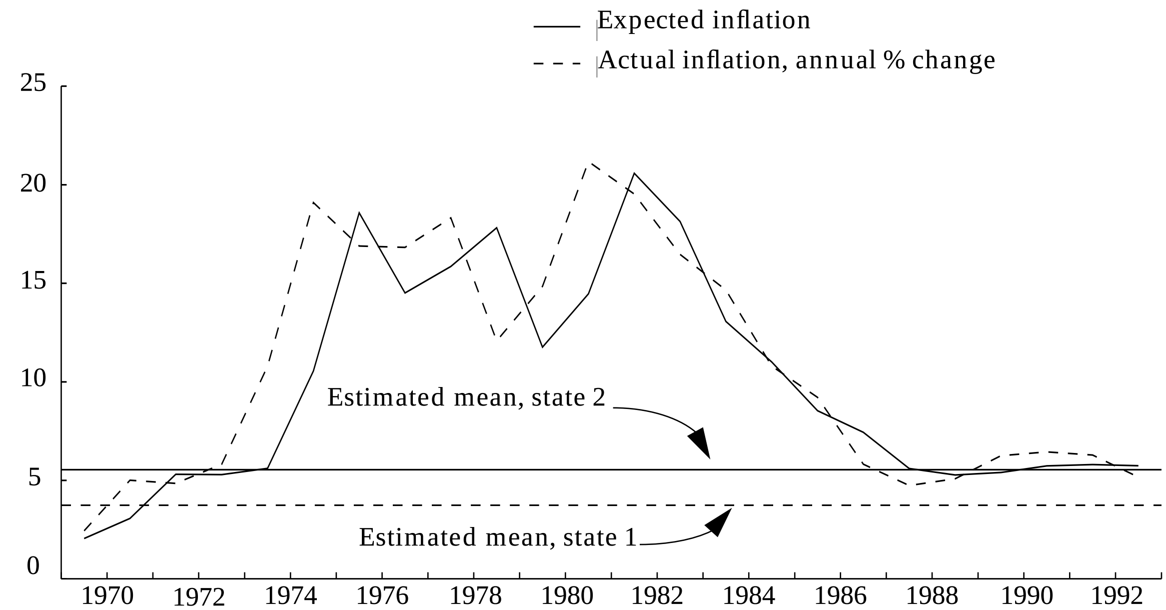

b) Probabilities of states 1, 2 and 3

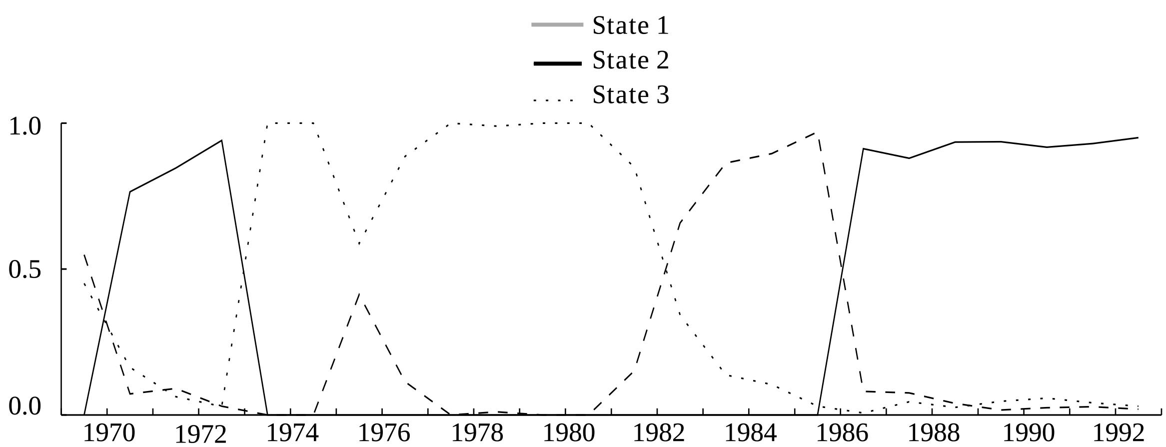

Figure 21
a) Actual and Markov expected CPI inflation
three–state model
Japan

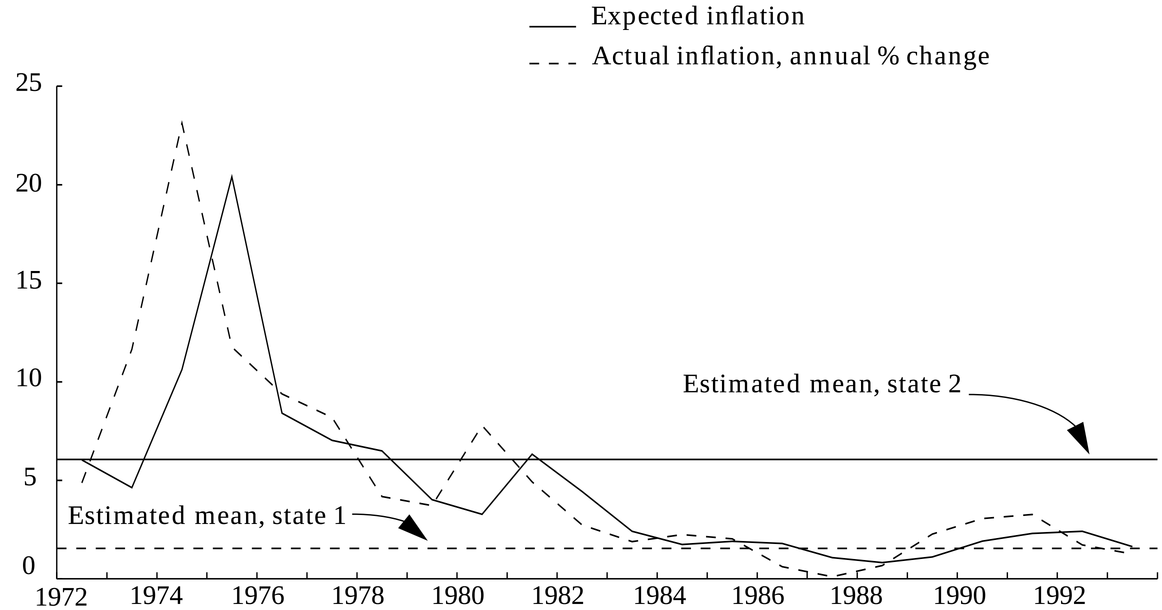

b) Probabilities of states 1, 2 and 3

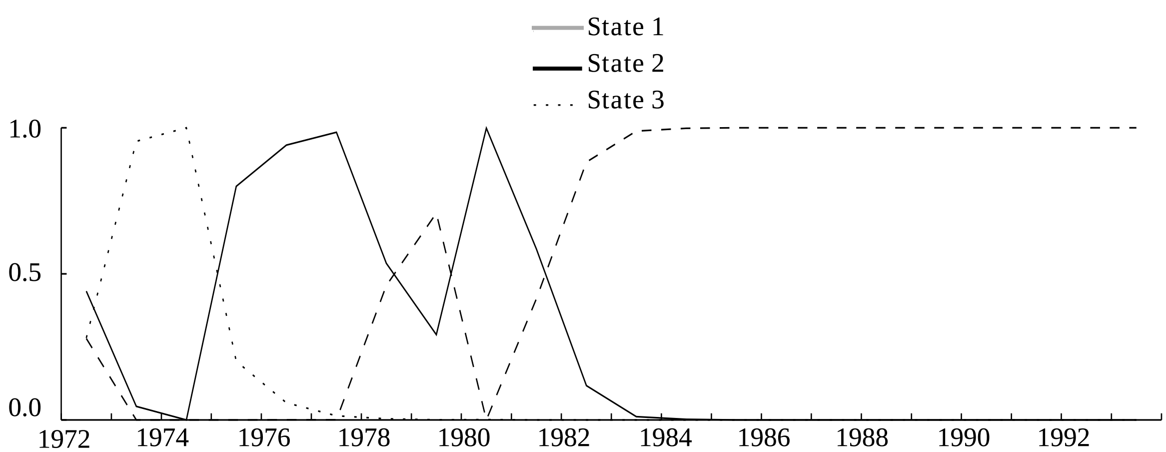

Figure 22
a) Actual and Markov expected CPI inflation three–state model)
United Kingdom

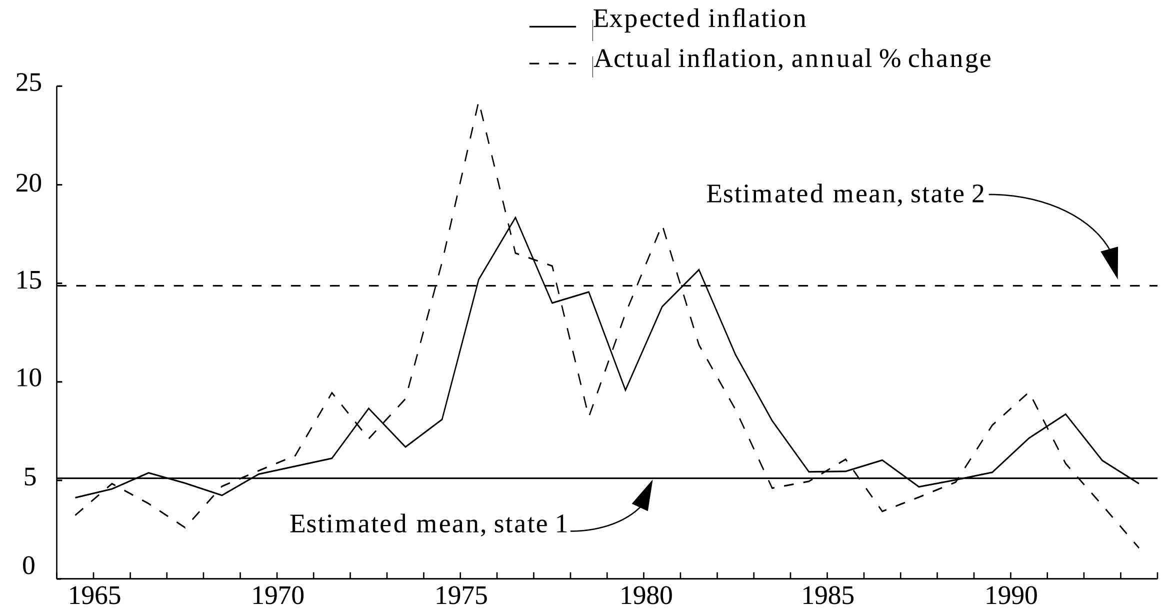

b) Probabilities of states 1, 2 and 3

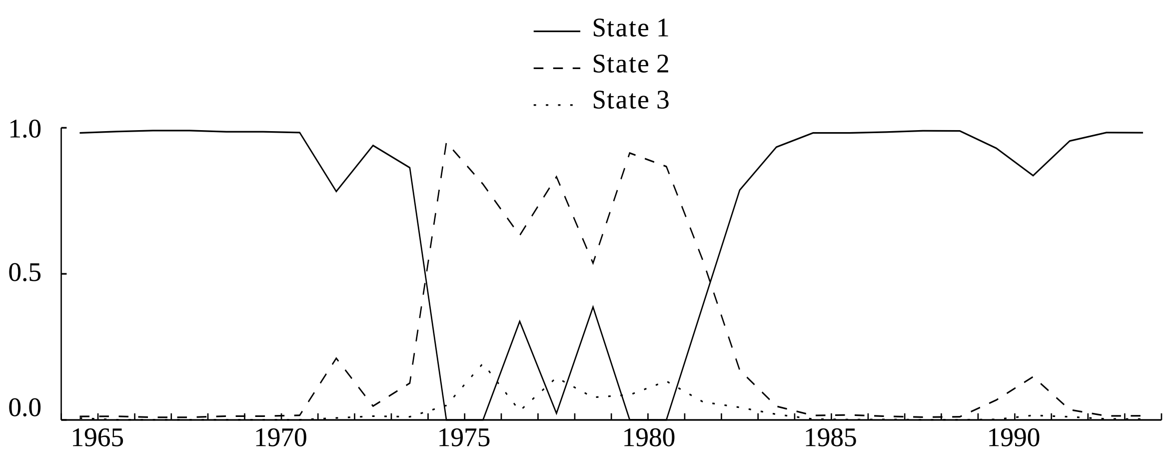

Figure 23
a) Actual and Markov expected CPI inflation
three–state model
United States

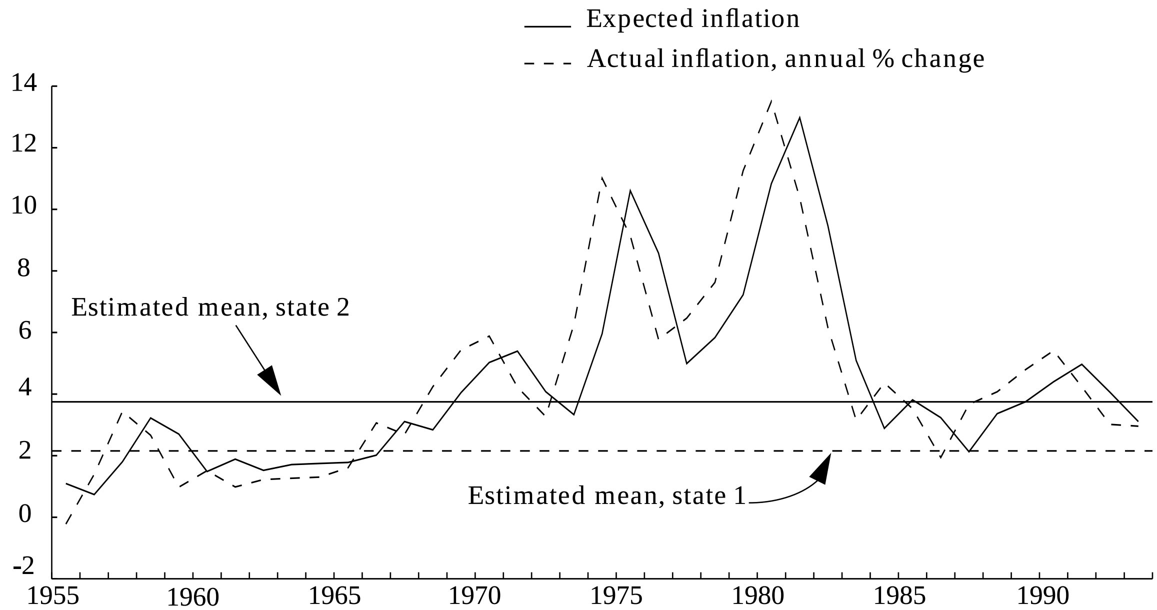

b) Probabilities of states 1, 2 and 3

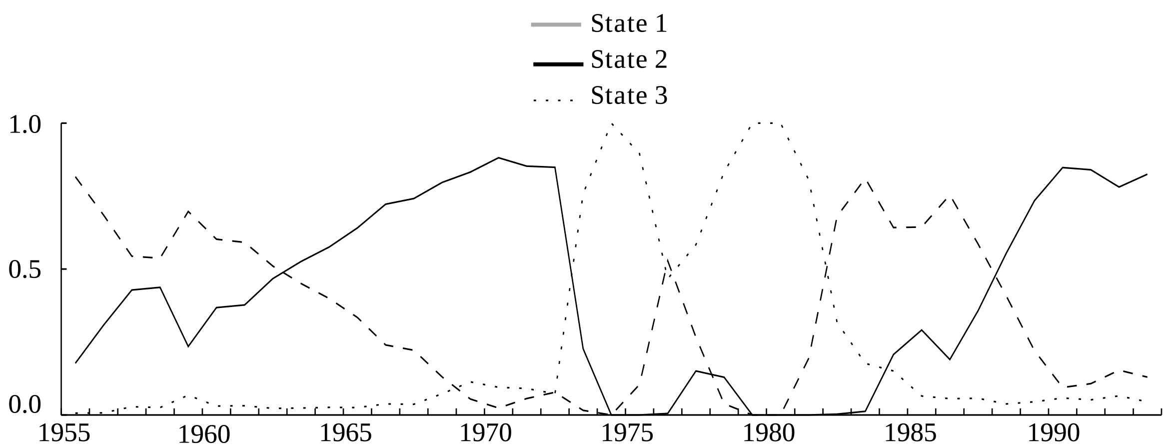

References

Amano, Robert A. and Simon van Norden. 1992. "Unit Root Tests and the Burden of Proof." Working Paper 92–7. Bank of Canada, Ottawa.

Andrews, Donald W. K. 1991. "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica* 59:817–58.

———. 1993. "Tests For Parameter Instability and Structural Change With Unknown Change Point." *Econometrica* 61:821–56.

Ball, Laurence. 1992. "Why Does High Inflation Raise Inflation Uncertainty?" *Journal of Monetary Economics* 29:371–88.

Dickey, David A. and Wayne Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Associatio*n 74:427-31.

Durland, Michael J. and Thomas H. McCurdy. 1992. "Modelling Duration Dependence in Cyclical Data Using a Restricted Semi-Markov Process." Queen's University. Unpublished paper.

Evans, Martin and Karen Lewis. 1993. "Do Expected Shifts in Inflation Affect Estimates of the Long-Run Fisher Relation?" New York University. Unpublished paper.

Evans, Martin and Paul Wachtel. 1992. "Inflation Regimes and the Source of Inflation Uncertainty." Federal Reserve Bank of Cleveland conference on inflation uncertainty.

Filardo, Andrew J. 1992. "U.S. - Canadian Business Cycles: Expansions, Contractions, and their Transitions." Federal Reserve Bank of Kansas City. Unpublished paper.

Filardo, Andrew J. and Stephen F. Gordon. 1993. "Business Cycle Durations." Federal Reserve Bank of Kansas City. Unpublished paper.

Ghysels, Eric. 1993. "A Time Series Model with Periodic Stochastic Regime Switching." University of Montreal. Unpublished paper.

Hamilton, James D. 1988. "Rational-Expectations Econometric Analysis of Changes in Regime - An Investigation of the Term Structure of Interest Rates." *Journal of Economic Dynamics and Control* 12:385–423.

———. 1989. "A new Approach to the Economic Analysis of Non-stationary Time Series and the Business Cycle." *Econometrica* 57:357–84.

Hamilton, James D. 1990. “Specification Testing in Markov-Switching Time-Series Models.” Discussion Paper 209. The Thomas Jefferson Centre for Political Economy, University of Virginia.

Hansen, Bruce E. 1992. “The Likelihood Ratio Test Under Nonstandard Conditions: Testing The Markov Switching Model of GNP.” *Journal of Applied Econometrics* 7:S61–S82.

Kwiatkowski, Denis, Peter C. B. Phillips and Peter Schmidt. 1991. “Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root: How Sure Are We That Economic Time Series Have a Unit Root?” Discussion Paper No. 979. Cowles Foundation for Research in Economics at Yale University.

Laxton, Douglas, Nicholas Ricketts and David Rose. 1994. “Uncertainty, Learning and Policy Credibility.” In *Economic Behaviour and Policy Choice Under Price Stability: Proceedings of a Conference Held at the Bank of Canada, October 1993*, 173–234. Ottawa: Bank of Canada.

Lewis, Karen K. 1989a. “Changing Beliefs and Systematic Rational Forecast Errors with Evidence from Foreign Exchange.” *American Economic Review* 79: 621–36.

———. 1989b. “Can Learning Affect Exchange Rate Behaviour? - The Case of the Dollar in the Early 1980’s.” *Journal of Monetary Economics* 23:79–100.

Marcet, Albert and Thomas Sargent. 1989. “Convergence of Least Squares Learning Mechanisms in Self-Referential Linear Stochastic Models.” *Journal of Economic Theory* 48:337–68.

Perron, Pierre. 1994. “Non-Stationarities and Non-Linearities in Canadian Inflation.” In *Economic Behaviour and Policy Choice Under Price Stability: Proceedings of a Conference Held at the Bank of Canada, October 1993*, 235–98. Ottawa: Bank of Canada.

Phillips, Peter C. B. and Pierre Perron. 1988. “Testing for a Unit Root in Time Series Regressions.” *Biometrika* 75:335-46.

Sargan, John D. and Alok Bhargava. 1983. “Testing Residuals from Least Squares Regression for Being Generated by the Gaussian Random Walk.” *Econometrica* 51(1):153-74.

Schaller, Huntley and Simon van Norden. 1993. “The Predictability of Stock Market Regime: Evidence from the Toronto Stock Exchange.” Forthcoming in *The Review of Economics and Statistics*.

Bank of Canada Working Papers

1995

95-1	Deriving Agents' Inflation Forecasts from the Term Structure of Interest Rates	C. Ragan
95-2	Estimating and Projecting Potential Output Using Structural VAR Methodology: The Case of the Mexican Economy	A. DeSerres, A. Guay and P. St-Amant
95-3	Empirical Evidence on the Cost of Adjustment and Dynamic Labour Demand	R. A. Amano
95-4	Government Debt and Deficits in Canada: A Macro Simulation Analysis	T. Macklem, D. Rose and R. Tetlow
95-5	Changes in the Inflation Process in Canada: Evidence and Implications	D. Hostland
95-6	Inflation, Learning and Monetary Policy Regimes in the G-7 Economies	N. Ricketts and D. Rose

1994

(Earlier 1994 papers, not listed here, are also available.)

94-3	The Term Structure and Real Activity in Canada	B. Cozier and G. Tkacz
94-4	An Up-to-Date and Improved BVAR Model of the Canadian Economy	D. Racette, J. Raynauld and C. Sigouin
94-5	Exchange Rate Volatility and Trade: A Survey	A. Côté
94-6	The Dynamic Behaviour of Canadian Imports and the Linear-Quadratic Model: Evidence Based on the Euler Equation	R. A. Amano and T. S. Wirjanto
94-7	L'endettement du secteur privé au Canada : un examen macroéconomique	J.-F. Fillion
94-8	An Empirical Investigation into Government Spending and Private Sector Behaviour	R. A. Amano and T. S. Wirjanto
94-9	Symétrie des chocs touchant les régions canadiennes et choix d'un régime de change	A. DeSerres and R. Lalonde
94-10	Les provinces canadiennes et la convergence : une évaluation empirique	M. Lefebvre
94-11	The Causes of Unemployment in Canada: A Review of the Evidence	S. S. Poloz
94-12	Searching for the Liquidity Effect in Canada	B. Fung and R. Gupta

Single copies of Bank of Canada papers may be obtained from

Publications Distribution
Bank of Canada
234 Wellington Street
Ottawa, Ontario K1A 0G9

The papers are also available by anonymous FTP to the following address, in the subdirectory /pub/publications:

ftp.bank-banque-canada.ca