A MARKOV-SWITCHING MODEL OF INFLATION IN AUSTRALIA

John Simon

Research Discussion Paper
9611

December 1996

Economic Group

Reserve Bank of Australia

The opinions expressed in this paper are those of the author and should not be attributed to the Reserve Bank of Australia. The paper has benefited from comments by Alison Tarditi.

Abstract

This paper applies the methodology of Markov-switching models to describe the inflation process in Australia in the period since the early 1960s. In contrast to conventional modelling, the approach makes explicit allowance for the possibility of structural change: inflation is modelled within a framework that allows endogenous switching between simple inflation equations. The approach may be relevant to understanding shifts in inflation expectations if the public also uses relatively simple forecasting rules in formulating expectations. The results suggest that inflation is reasonably well represented by relatively simple functions of past inflation and an output gap term, with major regime changes occurring in the early 1970s and early 1990s.

JEL Classification Numbers C22, E31, O56.

Table of Contents

A MARKOV-SWITCHING MODEL OF INFLATION IN AUSTRALIA

John Simon

1. Introduction

Models of the inflation process typically specify inflation as a function of a wide set of macroeconomic and policy-related variables including wages, import prices, commodity prices and business cycle conditions, often involving complicated dynamic structures. These models can be highly successful in tracking actual inflation, given the behaviour of the explanatory variables. Recent work by de Brouwer and Ericsson (1995), for example, shows that an error-correction model including the variables listed above has good explanatory power for inflation in Australia since 1977. An issue not addressed by this kind of modelling, however, is that structural changes may have occured in the underlying processes generating inflation, with possible implications for inflation expectations.

To address these issues, this paper applies an alternative modelling approach based on some recent studies that use Markov-switching models to describe the inflation process.[1] The distinctive feature of this approach is the use of very simple equations for inflation, within a framework that allows for discrete 'regime shifts' – ie shifts among a set of alternative equations that can govern the inflation process at different points in time. Specifically, Markov-switching models allow for two (or more) processes to exist with a series of shifts between the states occurring in a probabilistic fashion, so that shifts occur endogenously rather than being imposed by the researcher. The modelling strategy thus imposes a simpler-than-conventional structure on the inflation process within any given regime, but gains power to fit the historical data by allowing regimes to change.

By removing many of the standard explanatory variables this approach clearly ignores information contained in more conventional models. Nonetheless, the approach may be relevant to understanding shifts in inflation expectations if it is true that members of the public also use simple forecasting rules to formulate their

---

1 See Hamilton (1989, 1990) and Ricketts and Rose (1995).

expectations, changing the rule when experience deviates significantly from an established pattern. The analysis allows a number of relevant issues to be addressed. These include the forms of simple rules that best fit the data in this kind of framework, the frequency of regime changes and the issue of what constitutes statistical evidence of a regime change from the point of view of an observer using simple forecasting rules.

The results in this paper suggest that the 1970s and 1980s can be characterised by a high-inflation process with relatively persistent deviations from the mean, although the process is ultimately mean reverting. In contrast, the 1960s and 1990s can be characterised as a process with a low mean and less persistent deviation from that mean. The data choose this model in preference to one where the high inflation 1970s and 1980s are characterised by a random walk (and hence do not revert to any particular long-run mean).

Section 2 introduces Markov-switching models and the particular model used in this paper is specified in Section 3. Empirical results are reported in Section 4 and Section 5 concludes.

2. Markov-Switching Models

There has been some debate in the literature about the correct characterisation of inflation dynamics. A framework emphasising the integrated nature of inflation has been popular for some time. An integrated process is one which is non-stationary: shocks to the level of the series are permanent rather than temporary. This paper makes use of an alternative time-series characterisation for inflation that allows for distinct and differing periods of inflationary behaviour, each characterised by its own time-series properties. This alternative approach has both intuitive and empirical support. It describes the inflation process as being governed by two different regimes where switches between them are based on a probabilistic process. This approach is intuitively appealing, as the behaviour of economic time series often seems to go through distinct phases. It is also consistent with the fact that inflation is often found to be integrated of order 1 (ie non-stationary) – breaks in the mean of a series could lead to that series appearing to be non-stationary.[2]

The methodology employed is a 'Markov-switching model'. A Markov process is one where the probability of being in a particular state is only dependent upon what the state was in the previous period. Transitions between differing regimes are governed by fixed probabilities. Similar analysis in the literature has commonly been univariate – no independent variables have been included in modelling the series of interest. Initial work was done by Hamilton (1989,1990) with applications to business cycles. Recent work by Evans and Wachtel (1993) and Ricketts and Rose (1995) has applied the technique to inflation. This technique has several advantages, including endogenising structural breaks and encompassing ARCH models, each of which is discussed in more detail below. The technical details of the 'Hamilton filter' estimation are discussed in the Appendix, and the particular Markov-model specification is discussed in Section 3 below.

2.1 Structural Breaks

The Markov-switching model posits that two (or more) regimes could have prevailed over the course of history. However, it differs from models with imposed breaks in that the timing of breaks is entirely endogenous. Indeed, breaks are not explicitly imposed, but inferences are drawn on the basis of probabilistic estimates of the most likely state prevailing at each point in history.

Estimates of parameters for the two most likely regimes are generated using maximum likelihood techniques. With the parameters identified, it is then possible to estimate the probability that the variable of interest (in this case inflation) is following one of the alternative regimes. This involves identifying where in the probability distribution of each regime the observation falls at each point in time.

---

2 The point has been made by Perron (1989) and reiterated frequently within this literature. Clearly, given enough breaks, any I(1) series could be indistinguishable from an I(0) series around a broken trend. However, by allowing a finite number of processes, the Markov methodology is not directly open to this sort of criticism. Also, it is still modelling the time-series properties within each regime and obtaining estimates which are statistically distinguishable from one I(1) series. If a series were truly I(1), it should be found to be I(1) within each sub-period; it is only statistical error that would prevent definitive findings.

That is, the likelihood is calculated for each possible state. The probability that a particular state is prevailing is obtained by dividing the likelihood of that state by the total likelihood for both states. Thus, the sum of all the probabilities will equal one. With this estimate of the probabilities it is common to infer that a state is prevailing when the probability estimate for that state is greater than 50 per cent. In the models considered in this paper, values close to zero or one tend to occur, making identification of the prevailing state relatively easy.

2.2 Inflation Uncertainty

In practice, univariate models of inflation have commonly been characterised as having non-constant variance (commonly modelled as autoregressive heteroskedastic errors (ARCH)).[3] It is possible that these findings are related to the common observation that inflation is more variable during periods of high inflation. Ball (1992) posits a model where high inflation is associated with regime uncertainty whereas low inflation is not. Consequently, he suggests that the observation that high inflation has higher variance than low inflation reflects regime uncertainty. Another angle on the correlation between inflation and inflation variability is presented by Taylor (1981). His paper suggests that countries which place a high weight on output and employment stabilisation will have correspondingly higher and more variable inflation and vice versa. Thus, high inflation may be associated with high variability because it is correlated with choices by policy makers to focus on output and employment rather than inflation.

Whichever is the true explanation, regime switches, posited as the cause of higher volatility in high-inflation periods, could result in the identification of ARCH errors in single regime models. That is, ARCH processes suggest that volatility in one period is related to volatility in previous periods and that this volatility changes over time. If we allow for two possible states, one with a higher volatility than the other, then single regime models (ie. traditional models) might incorrectly identify the changes in volatility as symptomatic of ARCH. Instead, by correctly modelling the regime shifts, one possible reason for the finding of ARCH errors is eliminated. The estimation procedure thus allows us to investigate separately two questions with respect to inflation uncertainty or variability. First, are high-inflation states

---

[3] For example, Mishkin and Simon (1995) find significant ARCH in their investigation of inflation in Australia.

associated with high variability of inflation? And second, are ARCH effects present within regimes, implying that volatility within regimes tends to persist following a shock? These issues are discussed in the light of the results obtained in Section 4.

3. Model Specification

The starting point for the model used in this paper is a simple autoregressive model of inflation as used in earlier studies:[4]

$$p_t = a + b p_{t-1} + e_t \tag{1}$$

where p_t is the inflation rate. Apart from the simplicity of the specification the framework is also appealing as it can capture inflation expectations effects. The long-run mean value of inflation in equation (1) is simply:

$$p = \frac{a}{1-b} \, . \tag{2}$$

This might also be thought of as a central inflation expectation. Then, if we set $g = 1 - b$, a simple rearrangement of equation (1) into a partial adjustment framework highlights the importance of inflation expectations:

$$p_t = p_{t-1} + g(p^e - p_{t-1}) + e_t \tag{3}$$

where p^e is the expectation of the average inflation rate as defined in Equation (2). If people expect inflation to return quickly to its expected value then g will be high (putting price setting frictions to one side). If, however, people do not have firmly anchored expectations g will tend to be low. The extreme of this situation is when expectations have no long-run anchor, and hence $g = 0$ and inflation behaves as a random walk.

---

4 For example, Evans and Wachtel (1993) and Laxton, Ricketts and Rose (1994).

One justification for this interpretation is that it captures the effects of wage demands on costs. If wage increases lead to price increases through a simple markup model such that:

$$\Delta p_t = f\Delta w_t + h \tag{4}$$

and wage demands are based upon expected inflation as given in equation (5):

$$\Delta w_t = (1-d)p_{t-1} + dp^e + k \tag{5}$$

we can see how inflation expectations should affect inflation. In equations (4) and (5) h and k are constants that capture the influence of other variables on prices and wages; for example, k should capture productivity effects. If people have a firmly fixed inflation expectation then d should be close to 1. If, however, they do not have a firmly fixed expectation, then wage claims are likely to have a stronger autoregressive element with d close to 0. Reducing equations (4) and (5) yields $\Delta p_t = (f dp^e + fk + h) + f(1-d)p_{t-1}$ which is exactly the same form as Equation (1). Within this setup d, which reflects expectations, has a direct analogue in Equation (1) as the autoregressive coefficient.

The foregoing captures the central conceptual reason for regime changes within this paper – that inflation expectations change, possibly reflecting changes in policy objectives or in the nature of shocks hitting the system. As expectations change, the autoregressive parameter in the above model will change. This paper identifies periods with differing autoregressive parameters as coming from differing regimes. Following on from this discussion we specify the general model used in this paper in Equation (6) based on the general form of Equation (1):

$$p_t = c(S_t) + (1 - g(S_t)) \cdot p_{t-1} + w(S_t) \cdot Y_{t-1}^{GAP} + e(S_t)$$

$$e_t(S_t) \sim N(0, s(S_t)[h + f \cdot e_{t-1}]) \tag{6}$$

where p_t is the quarterly underlying inflation rate, S_t is the state variable (either 0 or 1), Y_t^{GAP} is the output gap and e_t is an ARCH process with a state-dependent scaling term, $s(S_t)$. The state variable is assumed to evolve following a standard Markov process as described in equations (7) and (8):

$$\Pr(S_t = 1 \mid S_{t-1} = 1) = p \tag{7}$$

$$\Pr(S_t = 0 \mid S_{t-1} = 0) = q \tag{8}$$

3.1 Why This Form?

The form chosen is a slight modification of standard univariate models of inflation used in previous applications of the Markov-switching methodology. An output gap is included to improve the model by including a significant exogenous explanator, and the ARCH process is included to capture information on the nature of inflation uncertainty.

The inclusion of both a scaling term for the error and an ARCH process allows for the separate identification of the reasons that errors would vary over time. A significant scaling term would suggest that high inflation periods were associated with higher volatility and, thus, that variability about the mean of the current regime was proportional to the level of that mean. A significant ARCH term would seem more indicative of regime uncertainty. That is, people would be unable to identify if a particular shock was an example of random fluctuation or a change in regime.

3.2 Learning

When regime switches are considered, one question that can be asked is whether people recognise the changes when they occur and modify their behaviour accordingly. That is, how quickly do people learn about regime changes? To add to the structure of the model and allow for a more flexible fitting of the data, a specific learning process is introduced. This is implemented by imposing:

$$g(S_t) = P_0 b_0 + (1 - P_0) b_1 + b(S_t) \tag{9}$$

where P_0 is the estimate of the probability of being in state zero at time t and b_i is the value of b in state 0 or 1. In this context b can be interpreted as the underlying nature of the regime and the probability weighted terms as the expectational effect. Thus, as the transition from one state to another occurs, more and more people learn about the true regime and adjust their price setting accordingly. Thus, if a regime change occurs (and $b(S_t)$ changes), yet no one recognises it, the autoregressive parameter will move less than when a change occurs that everyone recognises. Thus, the probability estimate generated by the model is also used as an estimate of the proportion of people who recognise the regime change. This allows the autoregressive parameter on the inflation process to adjust smoothly between the two states, rather than in a discrete fashion. This should also allow the model to accommodate some intermediate inflation state (where the probability of being in either state is around 50 per cent). This is most useful for forecasting if we believe that periods of intermediate inflation and regime uncertainty are possible.

One technical identification problem is introduced by the learning regime. The model could find it hard to distinguish between an inflation process with a low autoregressive parameter that has a low probability of occurrence and a regime with a higher autoregressive parameter that has a high probability. Nonetheless, this is only a problem at an instant in time. Given a longer history of observations and the fact that maximum likelihood techniques are used, the model identifies the most likely regimes which will, consequently, have a high estimated probability of occurrence.

4. Empirical Results

In developing the final model reported in this paper, a number of different specifications were estimated. The initial model (Model 1) imposes that one state is a random walk. This is based on the work of Ricketts and Rose (1995) who found that this was a good description of inflation for a number of G7 countries. Importantly, it is based on their observation in Laxton, Ricketts and Rose (1994) that 'We tried to find a formulation that would estimate a 'high and stable' inflation regime as one of the alternative states. We were not successful.' It also addresses the question of whether inflation, while stationary overall, has periods of non-stationary behaviour.

Model 2 relaxes the random walk assumption in specifying the high inflation state. Instead, inflation is allowed to follow an autoregressive process in both regimes and the data accept this as the preferred specification. The first regime has high and *persistent*, yet stable, inflation; the second regime has low and less persistent inflation. Even in periods of high inflation, people retain a longer-term anchor for their inflation expectations. The work by Ricketts and Rose (1995) only identified one G7 country with this kind of process – Germany. They comment that ‘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.’ Rickets and Rose estimate their models using annual data; it may be that Germany is the only country that has sufficient control over inflation in high inflation periods to be statistically distinguishable in the annual data. Our estimates are on quarterly data, and are not strictly comparable with those in Ricketts and Rose. Finally, unlike previous studies, we introduce the output gap as an explanatory variable. In this way Model 3 addresses, at least in part, the criticism that univariate models of inflation are inadequate. We also allow for learning in the autoregressive parameter over regime shifts, which gives the model additional flexibility in dealing with transitional periods.

The learning regime is only implemented in Model 3. This allows for the comparison of results in Models 1 and 2 with previous international work and the nesting of the models for hypothesis testing. All estimations are conducted over a sample period from December 1959 to September 1995 and use an underlying measure of inflation. The Treasury underlying measure (only available from March 1971) was spliced on to the consumption deflator to provide the longer run of data. This was necessary because of the relatively few low inflation observations over the period for which the Treasury underlying series is available.[5] The model parameter estimates are set out below; standard errors are in brackets. The notation follows the standard for ARCH and s_t denotes the state (either 0 or 1) in period t (state 1 is the low inflation regime). The figures show actual inflation (grey line) and the calculated probability that the process is in the random walk state (black line).

[5] Qualitatively similar results were also generated using headline inflation.

Model 1: Random Walk and AR(1)

State 0: Random walk

$$p_t^0 = p_{t-1} + e_t^0$$

$$e_t^0 = z \cdot \underset{(0.17)}{1.15}\sqrt{s_t^2}$$

$$p(s_t = 0 \mid s_{t-1} = 0) = 0.989$$

State 1: Stable low inflation

$$p_t^1 = \underset{(0.12)}{0.52} + \underset{(0.15)}{0.25}p_{t-1} + e_t^1$$

$$e_t^1 = z\sqrt{s_t^2}$$

$$p(s_t = 1 \mid s_{t-1} = 1) = 0.983$$

$$z \sim N(0,1) \qquad s_t^2 = \underset{(0.03)}{0.13} + \underset{(0.14)}{0.38}e_{t-1}^2$$

Log likelihood = 23.92

Figure 1: Random Walk Imposed

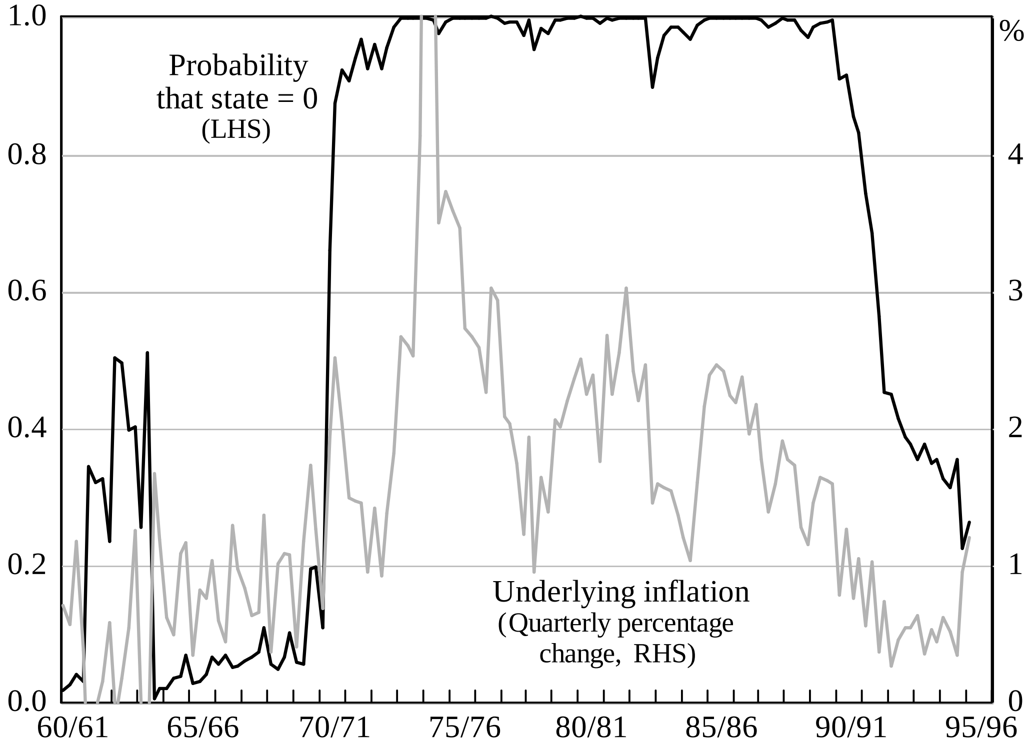

The autoregressive process of State 1 implies a steady-state rate of inflation of around 2.8 per cent.[6] The probability measure is quite volatile. It is possible to smooth the series by using hindsight to estimate the most likely state in past periods; however, this process can induce phase shifts in the probability measure. For this reason it is not reported here. The increase in the probability of being in a random-walk state early in the 1960s reflects the fact that inflation became too low to be consistent with the estimated autoregressive model. That is, a random walk could emerge at any level of inflation, not necessarily always at high levels.

Diagnostics

As the results from this model can encompass simple autoregressive models (typically implying that inflation is integrated, or at least close to it) we test the parameter restrictions involved.[7] Unfortunately, the distribution of these statistics is unlikely to be standard.[8] Nonetheless, following the approach in other papers, we use the likelihood ratio test statistic.[9] The log likelihood of the one-state process is

---

6 That is, $p^{ss} = \left[\frac{0.52}{1-0.25}\right] * 4 \approx 2.8$.

7 The simple alternative model is estimated to be $p_t = 0.2 + 0.85 p_{t-1} + e_t$ where e_t is an ARCH process as defined for the Markov model. This has an equilibrium inflation rate of around 5.5 per cent, the average of inflation over the estimation period.

8 The discussion of this issue revolves around the fact that some parameters are unidentified under H_0. If this is the case, it is claimed, the usual regularity conditions justifying the c^2 approximation do not apply. Hansen (1992) has proposed some asymptotically valid statistics although they impose a large computational burden. Hamilton and Susmel (1994), noting the results of Hansen, have suggested, of the c^2 statistics, that 'we regard these as a useful descriptive summary of the fit of alternative models'. Nonetheless, it seems that the statistics, provided the number of restrictions are counted properly, should be better than suggested by Hansen. Hansen suggests that a c_1^2 test (or t-test) on the autoregressive parameter is invalid. This is clearly true, but what is not discussed, in the context of this paper, is whether a c_3^2 test on the joint hypothesis that *p=1, q=0* and ratio=1 is invalid. With reference to Hansen (1992), the difference between the one-state and two-state models is not the restriction $m_d = 0$, but the addition of three new parameters; m_d, *p* and *q* to the one state model. Based upon the information provided in Tables 1 and 2 of Hansen (1992), a c_3^2 test would reject the Markov model in favour of the one-state model (with a p-value of 0.19); the same result obtained from the Hansen test and the opposite of the result from a simple t-test on m_d.

9 See, for example, Hamilton and Susmel (1994).

17.76; thus the likelihood ratio is 12.3 and should be compared with a c_3^2 distribution (with a 5 per cent critical value of 7.8 assuming standard distributions).[10] This test suggests that the Markov-switching model is a statistically superior model of underlying inflation.

The model does, however, have difficulty choosing which state is prevailing in recent periods (evidenced by intermediate probabilities rather than extreme zero or one probabilities). Therefore, we relax Ricketts and Rose's assumption that the high-inflation regime follows a random walk. Instead, we freely estimate the AR specification for both states (and test whether the coefficient on p_{t-1} in state 0 is significantly different from 1). The results for this are presented below.

Model 2: Two Autoregressive Regimes

State 0: Stable high inflation

$$p_t^0 = \underset{(0.15)}{0.32} + \underset{(0.08)}{0.83} p_{t-1} + e_t^0$$

$$e_t^0 = z \cdot \underset{(0.15)}{1.16} \sqrt{s_t^2}$$

$$p(s_t = 0 \mid s_{t-1} = 0) = 0.987$$

State 1: Stable low inflation

$$p_t^1 = \underset{(0.11)}{0.51} + \underset{(0.14)}{0.26} p_{t-1} + e_t^1$$

$$e_t^1 = z \sqrt{s_t^2}$$

$$p(s_t = 1 \mid s_{t-1} = 1) = 0.978$$

$$z \sim N(0,1) \qquad s_t^2 = \underset{(0.03)}{0.12} + \underset{(0.14)}{0.38} e_{t-1}^2$$

Log likelihood = 27.09

---

[10] The three variables which are allowed to vary in the switching model are the error variance ratio, p and q. The error variance ratio is implicitly set at one in the single process model, p is set to one (thus, once state 1 is entered it is never left) and q is implicitly zero (if state 0 is ever entered it is immediately left).

Figure 2: Two AR(1) Processes

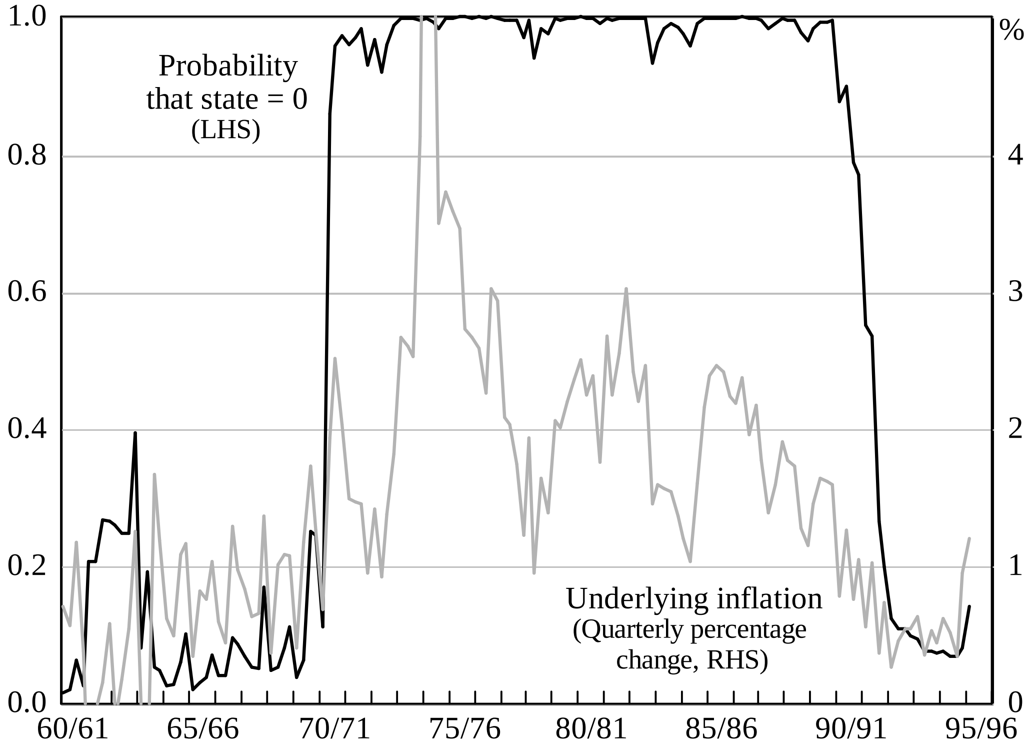

The model indicates a steady-state inflation rate in the 'low inflation' regime (State 1) of 2.8 per cent; mean underlying inflation is estimated to be 7.7 per cent in the high inflation state. The implied probabilities of being in each state are very similar to the previous model.

Diagnostics

By relaxing the restriction that the first state is a random walk the likelihood function is improved. The results should be compared with a c_2^2 distribution with a 5 per cent critical value of 6.0; the test statistic is 6.3, indicating a statistically significant improvement. As can be seen from the standard errors, the autoregressive coefficient is approximately two standard deviations away from one. This explains the nearness of the LM test to the 5 per cent critical value. However, it would seem unreasonable to expect a stronger rejection of the random walk given the expected high persistence of the high inflation state. The model is also a significant improvement over a simple one-regime autoregressive model. The appropriate test for this comparison is to add the likelihood ratios from the previous two tests and compare them with a c_5^2 distribution. This gives a test statistic of 18.6, compared to the 5 per cent critical value of 11.1.

Model 3: Output Gap Included in Both AR(1) Models

One criticism of the models developed above is that they are univariate – they do not take account of other independent explanators. To address this shortcoming, we estimate a third model that includes the output gap as an explanator. Import prices were also included but these did not significantly improve the results and so were omitted from the reported results. The output gap is measured with a Hodrick-Prescott filter on GDP(A). In the results below, *GAP* is the output gap expressed as a per cent deviation from trend. Another innovation is allowing for evolution in the autoregressive parameter based upon learning. In the results reported below, the autoregressive parameter is quoted as if it was certain that the relevant state was prevailing, that is the coefficient is $1-2b$ where b is the underlying nature of the regime as identified earlier.

State 0: Stable high inflation

$$p_t^0 = \underset{(0.17)}{0.40} + \underset{(0.07)}{0.81}p_{t-1} + \underset{(0.03)}{0.09}GAP_{t-1} + e_t^0$$

$$e_t^0 = z \cdot \underset{(0.14)}{1.04}\sqrt{s_t^2}$$

$$p(s_t = 0 \mid s_{t-1} = 0) = 0.989$$

State 1: Stable low inflation

$$p_t^1 = \underset{(0.12)}{0.54} + \underset{(0.14)}{0.34}p_{t-1} + \underset{(0.03)}{0.11}GAP_{t-1} + e_t^1$$

$$e_t^1 = z\sqrt{s_t^2}$$

$$p(s_t = 1 \mid s_{t-1} = 1) = 0.980$$

$$z \sim N(0,1) \quad s_t^2 = \underset{(0.02)}{0.10} + \underset{(0.19)}{0.56}e_{t-1}^2$$

Log likelihood = 37.56

Figure 3: Output Gap Included

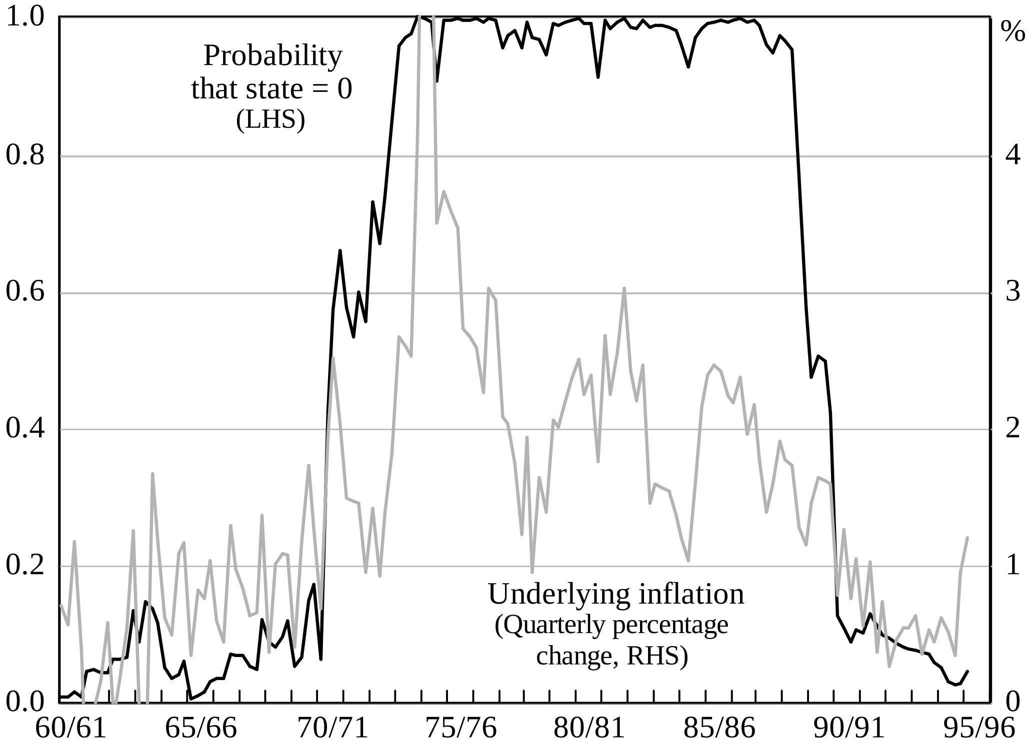

Diagnostics

Model 3 is a generalisation of Model 2 as we are allowing the parameters on GAP to be non-zero. The introduction of learning is a change to the specification but does not introduce any more parameters. Thus, we can compare the likelihood ratios for these models to a c_2^2 distribution. The ratio is 20.95 – representing a statistically significant improvement in the model. The introduction of the *GAP* parameter alone, without the change in the specification to include learning, also leads to a significant improvement in the fit of the model. To test the robustness of the parameters the sample was stopped in 1989, before the major recent falls in inflation. The estimates obtained were very similar to those estimated over the full sample period; the only major difference was in the estimated probability of transition from a high inflation state to a low inflation state. This transition probability was estimated as practically zero, as no such transition occurs over the reduced sample period. This is hardly surprising. Nonetheless, the other parameter estimates do not change much over the past five years, which suggests that the results are quite robust.

4.1 Discussion of Final Results

The results are very similar to those obtained from the univariate specifications of Model 1 and Model 2, although an important difference is that the change to a low inflation regime occurs about a year earlier. This provides an interesting comparison with the Westpac-Melbourne Institute survey of inflation expectations (Figure 4) – the only long-run series of directly measured inflation expectations available.

Figure 4: Comparison of Expectations

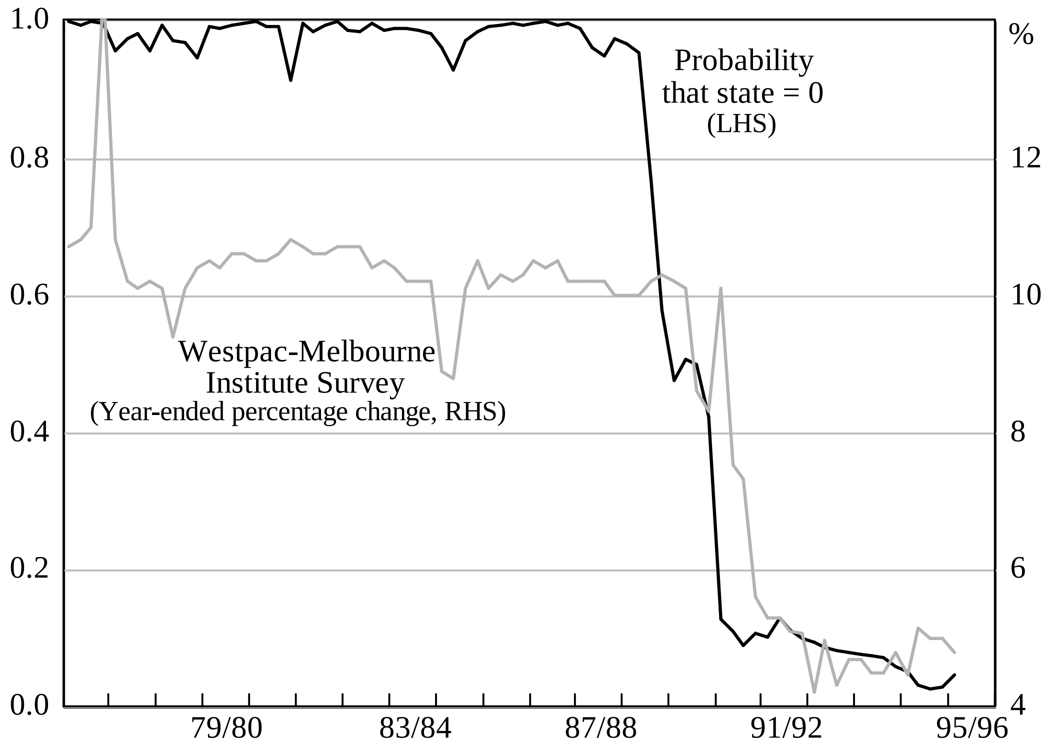

It is clear that this estimated probability series precedes the change in survey respondents' expectations (which occurred at the same time as headline inflation fell) by around a year. It is also around a year ahead of previous models' estimates. Thus, this more sophisticated model picks the trend to lower inflation before the survey respondents. While the probability estimate falls below 20 per cent in the same period that underlying inflation falls below 1 per cent (on a quarterly basis), this is not the sole reason for the change. Inflation had fallen to these levels before without causing much change in the estimated probability of being in the high inflation regime. One reason for the earlier estimated transition is the inclusion of the output gap; the inference is that rational observers, seeing the strong growth in output in the late 1980s without a corresponding increase in inflation, could have anticipated lower inflation rates if they based their forecasts on this particular model. There is also a noticeable change in the dynamics of inflation around this time, with shocks becoming less persistent. The survey results suggest that people only adjusted their expectations when headline inflation fell sharply – a change which suggests a strong backward-looking element in expectations. This model, by including the output gap and modelling the inflation process, identifies the signs of a changing regime much earlier than simple backward-looking expectations do.

Another interesting point about the results is that ARCH is clearly identified. Indeed, the scaling factor on the errors is not significantly different from one. This may suggest that the greatest cause of volatility in inflation is uncertainty about the regime rather than uncertainty about the mean level of the current regime. This highlights the potential that announced inflation targets have to reduce the volatility of inflation, in that they are associated with less uncertainty about the true regime. If people are more certain about the true state, then shocks should be recognised, rather than interpreted as signals that inflation is moving to a new regime.

Another important point about the inflation variability results is that conditional volatility is no higher in the high-inflation period than the low-inflation period, but inflation is more variable. The reason for this is the explicit allowance for differing regimes with differing shock persistence. That is, periods of high inflation are associated with longer shock persistence, which implies that the measured variance will be higher in these periods.[11]

[11] The variance of a simple autoregressive process is $s^2/(1-b^2)$ where s^2 is the variance of the errors and b is the autoregressive parameter. As shock persistence rises the measured variance will as well. Even if an autoregressive model is estimated for inflation, high inflation periods would still be identified with higher variance. This occurs because the high inflation period has greater shock persistence than the low inflation period (as estimated in this paper) which means that any adjustment to the variance would be biased against the high inflation period.

The final test of this model is to check its forecasting performance. To do this, the probability that inflation is in state 0 is projected into the future using the estimated transition probabilities. That is:

$$P_{t+1}^0 = P_t^0 \cdot P(s_t = 0 \mid s_{t-1} = 0) + P_t^1 \cdot P(s_t = 0 \mid s_{t-1} = 1) \qquad (10)$$

This is then used in conjunction with the inflation equation and actual output gap to generate forecasts. That is, we are assuming knowledge of the path of the output gap into the future. For the forecasts out to September 1996 it is assumed that the output gap linearly closes to zero. The results are shown in Figure 5 below.

Figure 5: Inflation Forecasts

Quarterly percentage change

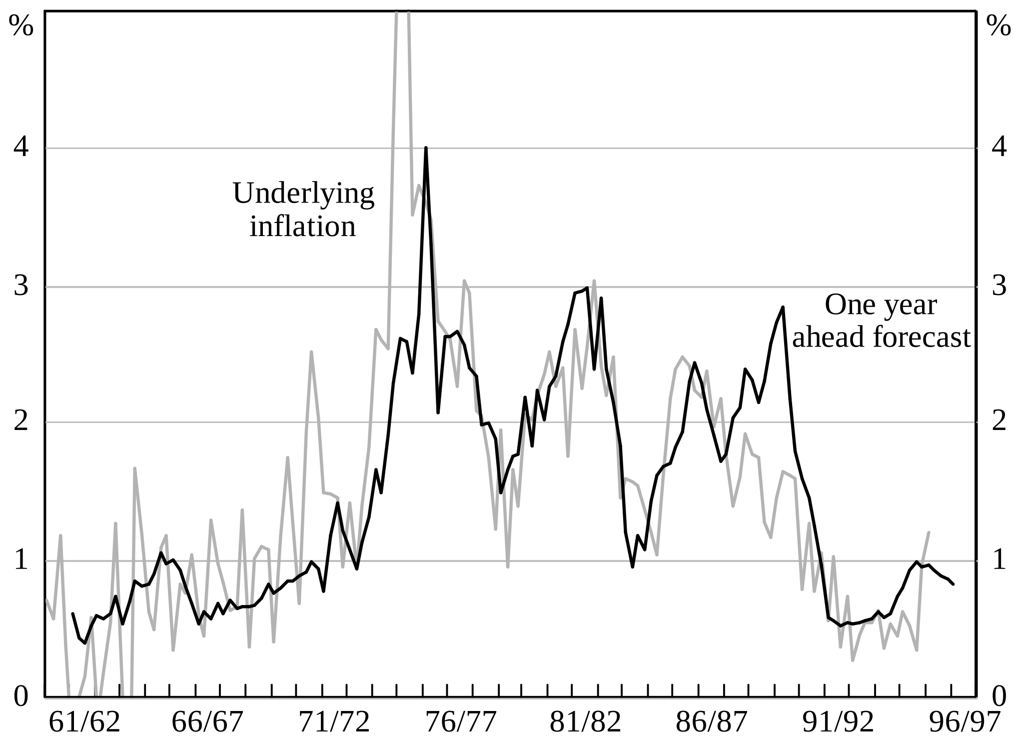

The performance would seem to be relatively good. The only period where it misses significantly is the transition to low inflation in the late 1980s. The reason for this is that the late 1980s were also associated with a high output gap. Thus, when making the forecasts, evidence had not yet arrived suggesting that the regime was changing. On the basis of the high output gap (output above trend), inflation would have tended to rise, hence the high forecasts. Indeed it is this divergence that leads to the change in the estimate of the most likely state.

4.2 Possible Extensions

While this model has achieved good results within the simple framework used, a number of extensions are possible. The inclusion of more exogenous variables would be one obvious extension. However, it should be noted that wages are not an exogenous variable in the framework used in this paper – wage demands are based upon inflation expectations. Another possible extension is to look at modelling the transition probabilities in different ways. This would be most useful for the purposes of forecasting. The estimates of the transition probabilities in this paper are quite small, since there are only two transitions over the entire sample period. Other papers have made use of time-varying transition probabilities and such an extension of this model may yield better results and improve the forecasting ability of the model.[12] That said, the forecasting properties of the model seem relatively good, as they are not solely based upon realised inflation rates, they incorporate the flexibility of the learning specification and they can deal with regime shifts.

5. Conclusion

By focusing on relatively simple equations in a framework that allows regimes to change, Markov-switching models would appear to provide a useful supplement to conventional modelling strategies for inflation. The results suggest that, within this framework, inflation in Australia since the early 1960s is reasonably well modelled by a two-regime specification, with regime changes occurring in the early 1970s and early 1990s. Within each regime, inflation in the preferred model is characterised by a simple autoregressive process supplemented by information about the output gap. The analysis may provide some insight into the behaviour of expectations, suggesting that within this limited-information framework it may be rational for imperfectly informed observers to change the forecasting rule only infrequently.

[12] See, for example, Durland and McCurdy (1994).

Appendix A: The Hamilton Filter and Maximum Likelihood Estimation

The estimation method used in this paper makes use of the algorithm described by Hamilton (1989).[13] The section below provides a brief description of the procedure. For technical details, please refer to the paper by Hamilton.

The Hamilton filter is an iterative procedure which provides estimates of the probability that a given state is prevailing at each point in time given its previous history. These estimates are dependent upon the parameter values given to the filter. A by-product of this process is the likelihood function for the given parameter estimates. Running the filter through the entire data, provides a log likelihood value for the particular set of estimates used. This filter is then repeated to optimise the log likelihood to obtain the MLE estimates of the parameters. With the maximum likelihood parameters, the probability of state 0 at each point in time is calculated and these are the probabilities that are reported in the paper.

The Hamilton Filter

The Hamilton filter starts with a vector of estimates of the probability that a particular sequence of states has led to period *t*-1. That is, it starts with a $2^r \times 1$ vector $P[S_{t-1} = s_{t-1}, S_{t-2} = s_{t-2}, \ldots, S_{t-r} = s_{t-r} \mid y_{t-1}, y_{t-2}, \ldots]$, where r is the length of the path. The length that it is necessary to keep track of is dependent upon the specification of the model dynamics. For example, in Hamilton's paper it is necessary to keep track of the past four states as he specifies a fourth order moving average process. In this paper it is only necessary to keep track of the past two states as a first order ARCH process is being used. From this vector, the probability that $S_t = s_t$ is calculated by making use of the

[13] The procedure was coded in GAUSS based on a program obtained from Thomas Goodwin, Claremont Graduate School.

Markov-switching probability estimates. Thus, if $P[S_{t-1} = 0 \quad S_{t-1} = 1] = [x \quad y]$ then:

$$P\begin{bmatrix} S_t = 0, S_{t-1} = 0 \\ S_t = 0, S_{t-1} = 1 \\ S_t = 1, S_{t-1} = 0 \\ S_t = 1, S_{t-1} = 1 \end{bmatrix} = \begin{bmatrix} q \cdot x \\ (1-p) \cdot y \\ (1-q) \cdot x \\ p \cdot y \end{bmatrix} . \qquad \text{(A.1)}$$

The next step is to calculate the likelihood that y_t occurs given the previous path of states and variables. That is, evaluate the value of the normal distribution at the point given by the residual. Consider, for example, the simple autoregressive model:

$$f(y_t \mid S_t = s_t, \ldots, y_{t-1}, \ldots) = \begin{bmatrix} \dfrac{1}{\sqrt{2p}\,s} \exp\left[-\dfrac{1}{2s^2}(y_t - a_0 y_{t-1}) \right] \\ \dfrac{1}{\sqrt{2p}\,s} \exp\left[-\dfrac{1}{2s^2}(y_t - a_1 y_{t-1}) \right] \end{bmatrix} . \qquad \text{(A.2)}$$

Here the top element of the matrix is the likelihood in state 0 and the bottom element is the likelihood in state 1 (the difference is in the alpha parameter). This is then multiplied by the probability estimate (A.1) to give the joint conditional density distribution of y_t and $(S_t, S_{t-1}, \ldots)$. The overall likelihood of y_t, $f(y_t \mid y_{t-1}, \ldots)$, is just the sum over all possible state paths of the joint conditional density function. That is, the probability-weighted likelihood for all possible paths. This is saved and used in calculating the likelihood of a particular set of estimates. This is:

$$\log f(y_T, y_{T-1}, \ldots, y_1) = \sum_{t=1}^{T} \log f(y_t \mid y_{t-1}, \ldots) \ . \qquad \text{(A.3)}$$

To generate the estimate of the probability that various paths lead to period t, $P[S_t = s_t, S_{t-1} = s_{t-1}, \ldots, S_{t-r} = s_{t-r} \mid y_t, y_{t-1}, \ldots]$, divide the conditional likelihood for each path by the total likelihood for all the paths. To obtain the input for the next iteration of the filter, collapse the probability matrix by summing over the possible states at time *t-r*. Thus, for example:

$$P\begin{bmatrix} S_t = 0, S_{t-1} = 0 \\ S_t = 0, S_{t-1} = 1 \\ S_t = 1, S_{t-1} = 0 \\ S_t = 1, S_{t-1} = 1 \end{bmatrix} \Rightarrow P\begin{bmatrix} S_t = 0 \\ S_{t-1} = 1 \end{bmatrix}. \tag{A.4}$$

This can then be used to run through the filter to get estimates for t+1 and so on. At each point in time the estimate of the probability that the current state is 0, given information available up to that time, is obtained by summing the probability vector in the same way as is illustrated in (A.4).

References

Ball, L. (1992), 'Why Does High Inflation Raise Inflation Uncertainty?', *Journal of Monetary Economics*, 29(3), pp. 371-388.

Durland, M. J. and T.H. McCurdy (1994), 'Duration-dependent Transitions in a Markov Model of U.S. GNP Growth', *Journal of Business and Economic Statistics*, 12(3), pp. 279-288.

de Brouwer, G. and N.R. Ericsson (1995), 'Modelling Inflation in Australia', Reserve Bank of Australia Research Discussion Paper No. 9510.

Evans, M. and P. Wachtel (1993), 'Inflation Regimes and the Sources of Inflation Uncertainty', *Journal of Money, Credit and Banking*, 25(3) Part 2, pp. 475-511.

Hamilton, J.D. (1989), 'A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle', *Econometrica*, 57(2), pp. 357-384.

Hamilton, J.D. (1990), 'Analysis of Time Series Subject to Changes in Regime', *Journal of Econometrics*, 45, pp. 39-70.

Hamilton, J.D. and R. Susmel (1994), 'Autoregressive Conditional Heteroskedasticity and Changes in Regime', *Journal of Econometrics*, 64, pp. 307-333.

Hansen, B.E. (1992), 'The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model of GNP', *Journal of Applied Econometrics*, Vol. 7 Supplement, S61-S82.

Laxton, D., N. Ricketts and D. Rose (1994), 'Uncertainty, Learning and Policy Credibility', *Economic Behaviour and Policy Choice under Price Stability*, Proceedings of a Conference, Bank of Canada.

Mishkin, F.S. and J. Simon (1995), 'An Empirical Examination of the Fisher Effect in Australia', *Economic Record*, Vol. 71, pp. 217-229.

Perron, P. (1989), ‘Testing for a Unit Root in a Time Series with a Changing Mean’, Princeton University, Econometric Research Program Research Memorandum No. 347.

Ricketts, N. and D. Rose (1995), ‘Inflation, Learning and Monetary Policy Regimes in the G-7 Economies’, Bank of Canada Working Paper No. 95/6.

Taylor, J.B. (1981), ‘On the Relation Between the Variability of Inflation and the Average Inflation Rate’, *Carnegie-Rochester Series on Public Policy*, 15, pp. 57-86.