LEARNING RATIONAL EXPECTATIONS IN A POLICY GAME

MARTIN CRIPPS

UNIVERSITY OF WARWICK

No. 297

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

UNIVERSITY OF WARWICK
COVENTRY

LEARNING RATIONAL EXPECTATIONS IN A POLICY GAME

by Martin Cripps[1]

Rational expectations is a maintained assumption in the analysis of economic policy. Here we examine how two types of learning rational expectations (rational and econometric) affect the time profile of optimal policy. In both cases the government adopts policies which delay convergence to rational expectations. There is also a reduction in the inflationary bias, in one case permanently in the other temporarily.

1. INTRODUCTION

Recent work on policy games (Barro & Gordon (1983), Backus & Driffill (1985), Cukierman & Meltzer (1986)) requires the public sector to form rational expectations in a complex strategic environment. In this paper we consider the problem of learning, or convergence to, rational expectations in such models. This process will take place in the early years of government's terms of office when the public are discovering the government's propensity to create inflationary surprises. The issue is; does the process of learning rational expectations have important effects on the time profile, or limit, of optimal policy. In particular is it in the government's interest to adopt strategies which delay the acquisition of rational expectations, with the resultant costs to the public sector in ambiguity? Also, how does the process of learning influence the credibility properties of the eventual steady state? These issues are addressed in the context of a variant of the model used in Cukierman & Meltzer (1986). The difference in the approach taken here is that we examine the process of convergence to the steady state and not the steady state itself, hence to simplify matters the government's preferences do not evolve through time.

Most models of policy games maintain the assumption of rational expectations even after the government has deviated (or cheated) on the current policy. However Lucas (1981) has convincingly argued that rational expectations will only prevail at a long run, or stable, policy. One of the costs of deviating from a current policy must therefore be the breakdown, at least temporarily, of rational expectations. To capture this effect one must explicitly model the process agents use to up-date their beliefs: their learning. We find that the credibility of policy and its relationship with learning has important implications for governmental behaviour in the early periods of a regime.

Learning can occur in two different ways; the first assumes the public is informed about, or can deduce, the strategic nature of the government's behaviour and correctly takes this into account when learning unknown information: <u>rational learning</u>. The second models the private sector as a less sophisticated learner, who does not know the true likelihood generating the observations it makes. Instead the private sector only has access to standard econometric techniques when forming expectations, this approach is exemplified by Bray & Savin (1986) and is usually termed non-rational learning.

Below we show how the government's optimal strategy under rational learning results in a steady increase in inflationary bias. Thus, it is credible for the government to initially set low levels of inflation, but as learning converges the government is forced to expand the growth rate of the money supply. So governments initially understate their propensity to create surprise inflation. This is not irrational behaviour, but a result of the optimal strategy given the public is learning rational expectations. It also suggests a different interpretation for the reputation effects of Backus & Driffill (1985); which is based upon a government's desire to slow the public's learning of its preferences. There is one piece of work, Basar and Salmon (1987), which looks at rational learning in a Stackelberg equilibrium, this results in a different type of equilibrium behaviour.

Non-rational learning is compared with the rational learning. We show that if the public adopts Ordinary Least Squares (OLS) in forming expectations, then the optimal policy is identical to the limit of rational learning. However, the other forms of learning we consider all generate consistently lower inflationary biases. Non-rational learning generates lower credible levels of inflation than rational learning. This result may also provide a rationale for the punishment periods in the model of Barro & Gordon (1983), which sustain lower credible levels of inflation.

The sections below are arranged in the following way; Section 2 contains an outline of the model used and Section 3 characterizes the rational learning solution. Section 4 characterizes a class of econometric non-rational learning processes and presents the optimal policy under non-rational learning. Section 5 contains the conclusion.

2. THE MODEL

The model of government policy used in this paper is based mainly on that developed by Cukierman & Meltzer (1986). There is one substantive difference in that the government's preferences do not evolve through time, but instead are determined at the outset by a one-off random selection of a weight on surprise inflation in the preferences. The notation will be similar to facilitate comparison.

There are a sequence of periods t=0,1,2,.. . In each period the public sets its expectations for the rate of growth of the money supply $E[m_t|I_t]$ and the government simultaneously sets the planned money supply growth $m^p{}_t$. Neither agent is able to observe their opponent's current action. Ex-post the public is not able to observe the choice of planned money supply growth $m^p{}_t$, instead they observe the actual money supply growth m_t. Actual and planned money supply growth are related in the following way;

(1) $$m_t = m^p{}_t + n_t.$$

The government's choice $m^p{}_t$ is disturbed by an independent sampling n_t from a normal distribution with mean zero and variance $\sigma_n{}^2$. Hence the government has a noisy control of the variable m_t, and the public face a signal extraction problem when trying to deduce the actual policy used by the government.

The objective function of the private sector in this model is specified below. We assume that the aim of the private sector is to correctly form their expectations of money supply growth rates $E[m_t|I_t]$. This is justified by giving the public an objective function based on minimising the mean square error of some variable m^e_t;

$$\max_{m^e_t} E[\ -(m^e_t - m_t)^2 \ |\ I_t\ ],$$

$$\text{so;} \qquad m^e_t = E[m_t|I_t].$$

The government chooses a sequence of m^p's to maximise its objective function, this is described in (2). This function reflects the government's preferences for surprise inflation $(m_t-E[m_t|I_t])$, which receives a weight of $(A+p)$ [2]. The value A is some positive constant and p is the private information of the government. The government's preferences also reflect its desire to minimise the absolute level of the money supply growth $(m^p_t)^2$. The expectations operator E_{G0} represents expectations taken relative to the government's information set, whilst the notation $E[.|I_t]$ is used to denote expectations relative to the public's information I_t, (these are described in some detail below and in every circumstance the government's information set includes that of the public). There is a discounting parameter β which weights future payoffs;

$$(2) \qquad \max_{\{m^p_t\}} E_{G0} \sum_{t=0}^{\infty} \beta^t\{\ (A+p)(\ m_t-E[m_t|I_t]\ ) - \tfrac{1}{2}(m^p_t)^2\ \},$$
$$(A>0,\ 0<\beta<1).$$

The value p is determined by an independent drawing from a normal distribution with zero mean and variance σ_p^2. The value p is observed before the start of play by the government, but not the private sector. This means that the government's propensity to create inflation is unknown to the public, hence they are unable to calculate their rational expectation of the credible level of inflation. Instead they estimate the value of p by observing the government's behaviour over time, it is this unknown parameter which the public must learn. We compare the rational learning of p with a non-rational approach in Sections 3 and 4.

3. RATIONAL LEARNING

In this section we outline the government's optimal strategy when the public engages in rational learning. We prove that it is optimal to steadily increase the rate of growth of the money supply as the private sector learns the value of p. Hence under rational learning the optimal policy slowly converges to the equilibrium in the static game where the government's propensity to create inflation is known. There is no long run reduction in the level of inflationary bias.

Under rational learning of the observation p the relevant distributions underlying the model, the government's preferences and the past realisations $m_0, m_1, m_2, ..m_{t-1}$ are all common-knowledge. Using this information the public is able to deduce the functional form of the government's optimal strategy. The public then uses its knowledge of this functional form and statistical inference to estimate the value of p. The government has private information $\{p, m^p_t\}$, which are never directly observed by the public.

The determination of the rational learning solution follows a very similar route to that used by Cukierman & Meltzer (1986). We first assume that the government adopts a linear strategy and then solve using the method of undetermined coefficients. The differences here are twofold; first there is one sampling of the government's preferences, so these do not evolve through time. Second, as we are considering the learning of the value p by the private sector, we are in effect considering the process of convergence to the steady state and not the steady state itself. This implies that the government follows a non-stationary strategy and we must therefore characterize the convergence of this strategy. We begin by assuming the government determines planned money supply growth using the rule;

$$(3) \qquad m^p_t = B_t(A+p).$$

The weight on inflationary surprises in the government's preferences determines the planned money supply growth. The constant B_t is chosen optimally by the government at each point in time and is independent of p. The sequence $\{B_t\}$ therefore determines the time profile of the optimal inflationary bias.

Rational learning implies that the public knows the government's preferences, the structure of the model and therefore the values $\{B_t\}$. Hence the public must solve a signal extraction problem to learn about p. This is because the public observes past rates of money supply growth ($m_{t-1}=B_{t-1}(A+p)+n_{t-1}$), but the noise n_{t-1} prevents direct observation of p. This signal extraction process enables the public to deduce an expected value for p and hence the current $E[m_t|I_t]$. Given the structure of the model the least squares projection result can be used to calculate $E[p_t|I_t]$.

PROPOSITION 1: *Given* $m^p_s = B_s(A+p)$ *for* $s=0,1,2,..t-1$; *the value* $E[m_t|I_t]$ *satisfies;*

$$E[m_t|I_t] = AB_t + B_tE[p|I_t]$$

(4) $$= AB_t + B_tW_{t-1}\sum_{i=1}^{t}(B_{t-1}m_{t-1} - B^2_{t-1}A);$$

$$W_{t-1} = \left(r + \sum_{s=0}^{t-1}B^2_s\right)^{-1}; \qquad r = \sigma_n^2/\sigma_p^2.$$

PROOF: Given $m^p_t = B_t(A+p)$ and normality one can deduce that the conditional expectation $E[p|I_t]$ is a linear function of the available data, so let;

(5) $$E[p|I_t] = \sum_{i=1}^{t}\theta_{it}(m_{t-1} - AB_{t-1})(B_{t-1})^{-1} = \sum_{i=1}^{t}\theta_{it}(p+\phi_{t-1});$$
$$\phi_{t-1} = (n_{t-1})/(B_{t-1}).$$

The θ's in this expectation are constants, which can be calculated by minimising the mean squared error;

(6) $$E\left[p - \sum_{i=1}^{t}\theta_{it}(p+\phi_{t-1}) \right]^2 = \sigma_p^2\left(1 - \sum_{i=1}^{t}\theta_{it}\right)^2 + \sigma_n^2\sum_{i=1}^{t}\theta^2_{it}B^{-2}_{t-1}.$$

Minimising this with respect to the θ's yields the first order conditions;

(7) $$0 = -\left(1 - \sum_{i=1}^{t}\theta_{it}\right) + r\,\theta_{it}/B^2_{t-1}; \qquad r = \sigma_n^2/\sigma_p^2.$$

This implies that the last term on the right hand side is independent of i and only depends on t. So define $\theta_{it} = k_t(B_{t-1})^2$ and substitute into (7), solving this then gives the expression;

(8) $$k_t = \left(r + \sum_{s=0}^{t-1}B^2_s \right)^{-1}.$$

Using the relation $\theta_{it} = k_t(B_{t-1})^2$ we can now calculate the value $E[p|I_t]$ from (5) and in turn calculate $E[m_t|I_t]$ using the definitions above.∎

The expectation of money supply growth is calculated by giving each of the observations $(m_{t-i}-AB_{t-i})$ a relative weighting of B_{t-i}. Although each observation appears to contain equal information on the unknown data "p". The weighting B_{t-i} takes account of the scaling each observation receives in the government's policy. The periods when planned money supply growth (B_{t-i}) is large will give more accurate information on the magnitude of p, because the observed money supply growth rates contain relatively less noise. Therefore the magnitude of B_{t-i} is also the way the government controls the public's rate of learning about its preferences. Notice that if the policy variable B_t is tending to a finite limit, then each observation will receive equal weight, and in the limit new information has zero weight.

The next step in the solution of the rational learning case is to calculate the government's optimal strategy conditional on the private sector's expectations formation (4). We begin by writing down the government's optimisation problem, this simply takes (4) and substitutes into (2);

$$\max_{\{m^p_t\}} E_0 \sum_{t=0}^{\infty} \beta^t \{ (A+p)[m^p_t+n_t-AB_t-B_tW_{t-1}\sum_{i=1}^{t-1} B_{t-i}(m_{t-i}-B_{t-i}A)] - (1/2)(m^p_t)^2. \quad (9)$$

The government chooses a sequence $\{m^p_t\}$ to maximise this expression. The solution to this problem can be found by using standard dynamic programming arguments, (for example Sargent (1981) Chapter XIV). These give the following set of first order conditions;

$$0= (A+p) - m^p_t - (A+p)\{\beta B_tW_tB_{t+1} + \beta^2B_tW_{t+1}B_{t+2}.. \}.$$

Also there is a transversality condition:

(10) $$\lim_{t\to\infty} \beta^t E_{\Omega 0}[(A+p)(1-\beta\{B_{t+1}W_tB_t+\beta B_{t+2}W_{t+1}B_t..\})-m^p_t] =0.$$

Rearranging the first of these we can solve for the government's policy rule;

$$m^p_t = (A+p)\ [\ 1-\ \beta B_tB_{t+1}W_t\ -\ \beta^2B_tB_{t+2}W_{t+1}\ -..\ \ ].$$

Employing the method of undetermined coefficients we get;

$$B_t^{-1} = 1 + \beta B_{t+1}W_t + \beta^2 B_{t+2}W_{t+1} +\dots .$$

Leading this once, multiplying through by β and differencing gives the following difference equation for the B_t's;

(11) $$B_{t+1} - \beta B_t = B_tB_{t+1}(\ 1-\beta+\beta W_tB_{t+1}).$$

We now can use this information to characterize the time path of the sequence $\{B_t\}$ and hence the time path of optimal policy. We prove below that there is a unique sequence satisfying (11) which also solves (9), under circumstances of full rational learning and credibility.

PROPOSITION 2: *There is a unique credible optimal policy under rational learning. This satisfies:*

(i) $m^p_t < m^p_{t+1} < A+p$; $\quad (B_t < B_{t+1} < 1)$;

(ii) $\lim_{t\to\infty} m^p_t = A+p$; $\quad (\lim_{t\to\infty} B_t =1)$.

PROOF: See Appendix.

This result shows that rational learning in this model must result in the repeated game converging to the Rational Expectations Equilibrium in the static game with full information. (It is easy to verify that an inflationary bias of A+p is the inflationary bias in the static game with rational expectations where p is known.)

Clearly, the rational learning solution in this game converges to complete knowledge of the value of p as information accrues and as the government's strategy approaches the steady state $B_t=1$. Hence the rational learning equilibrium can be interpreted as a process of convergence to the full information steady state. It is also worth noting that for this example the act of learning rational expectations has little implication for the eventual steady state. There has been no change in the long run level of the rate of growth in the money supply as a result of the private sector's rational learning.

However, there is some interest in the route to the steady state. Initially lower levels of inflationary bias are credible, but over time the bias increases until it approaches the full information level. There are two ways of interpreting this process. From the point of view of information transmission low levels of bias are ineffective in conveying information on p to the public. So one could view this process of convergence as the government attempting to slow down the public's rate of learning, by choosing relatively uninformative levels of inflationary bias in the early stages of play. This adds support to Cukierman & Meltzer's views on optimality of ambiguity in economic policy.

The second interpretation of the gradual increase in levels of inflationary bias stems from the credibility of the policy under rational learning. In the early stages of the learning process, there are low levels of inflationary bias. This is credible, because in the early periods the public have relatively little data to use in forming their expectations, so they place a high weight on new information. In these circumstances there would then be a significant revision in the public's beliefs when the government creates surprise inflation. This revision would impose costs on the government in the future, which outweigh the current benefits from the surprise even though there are (currently) low inflationary expectations. As the rational learning process converges it places equal weight on all observations. In the limit there will be zero weight on new information in the expectations, so the creation of surprise inflation generates no costs in the future because beliefs are not responsive to a current surprise. In the long term therefore the current level of inflation must be credible without the presence of long term costs. In summary; the government desires to delay the public's learning, because this generates lower credible levels of inflation which improves the government's utility.

There are heavy informational requirements placed on the private sector's behaviour under rational learning, it must solve an extremely complex fixed-point problem using its entire knowledge of the government's preferences and the distributions to do this. In the microeconomic literature on convergence to Rational Expectations Equilibria the informational requirements of rational learning have been convincingly criticised as an unsatisfactory model of learning. The approach also presents a highly idealised version of the learning process. Hence, there are many advocates of a less elaborate approach to the modeling of learning.

4. NON-RATIONAL LEARNING

In considering non-rational modes of learning we assume that the public is much less sophisticated than the public outlined in Section 3. The public does not know the government's preferences nor does it know the relevant distributions underlying the model. Hence instead of employing the sophisticated reasoning described above we imagine the public employing an econometrician. One could justify the use of econometrics by suggesting that the public faces information processing costs and so are forced to employ an ad-hoc technique in an attempt to solve the sophisticated inference problem. The real appeal of the non-rational learning approach is its similarity to the real learning process, where agents start off with an incorrectly specified likelihood, but through interaction with the data arrive at a correct likelihood. The rational learning approach is essentially an exercise in statistical inference where individuals do have access to the correct likelihood generating the events they observe and are simply estimating a parameter in the likelihood. This is not really in accord with general notions on the nature of learning, which is much more than a simple statistical process.

In the model outlined in Section 2 the private sector is attempting to predict the rate of growth of the money supply in an environment where it is possible for the government to pervert this process. So the set of non-rational econometric models we consider in this section must be sufficiently robust to cope with this problem. Here the public is assumed to model the process they are learning as a simple one-parameter econometric model;

$$m_t = \mu + \epsilon_t \qquad E[\epsilon_t] = 0.$$

This models the public's belief that the money supply is composed of an unknown constant term μ and an error ϵ_t. This is clearly very crude in comparison with the beliefs specified in the Section 3 and is a severely misspecified likelihood. In particular the parameter value μ may change through time as the government changes its policies, because the government's strategy could well change through time in response to the public's learning. There is also the possibility of heteroskedasticity in the process generating the errors, this again could be caused by the public's learning. For these reasons we assume that the public use an exponentially weighted estimator of μ;

$$(12) \qquad E[m_t|I_t] = \mu^t := (1-\gamma)\sum_{i=1}^{\infty} \gamma^{i-1} m_{t-i} \qquad (0< <1).$$

This is the Discounted Least Squares (DLS) estimator of the parameter μ and is recommended for circumstances where there is a belief that the process being estimated is unstable, see for example Harvey (1983) Chapter 6. It is also identical in form to the Generalised Least Squares (GLS) estimator, which is often used in cases of heteroskedasticity. One can also show that this estimator generates the optimal forecast, (based on the Kalman Filter) when the parameter μ is itself generated by a random walk: a time varying parameter model. This sort of econometric model may well capture many of the features of fully rational learning, in particular the "return to normality" model, (see Harvey Chapter 6,) has many similarities with rational learning.

The scalar gamma which determines the exponential weighting also describes a class of different learning schemes. As gamma approaches unity the relative weight on recent observations becomes smaller, whilst the converse is true as gamma approaches zero. This class of models includes OLS as a special case (gamma equals unity), hence we include the most well-known form of non-rational learning. It is also worth comment that the observations $(m_{-1}, m_{-2}, m_{-3}, \ldots)$ are used to represent the public's prior, but as the conclusions we derive here are entirely independent of these values it does not seem worth dwelling upon this assumption.

We now characterize the government's optimal strategy given that the public learns in the non-rational manner specified here. We will assume that the government is entirely aware of the nature of the econometric process used by the public and the value of gamma, although it is not necessary for the government to know the public's priors.

PROPOSITION 3: *The government's optimal strategy given the non-rational learning specified above is;*

$$m_p = (A+p)(1-\beta)(1-\gamma\beta)^{-1}$$

PROOF: The government optimises conditional on the private sector's learning, this gives the following optimisation problem:

$$\max \quad E_{\Omega 0} \sum_{t=0}^{\infty} \beta^t [(A+p)[m_t^p + n_t - (1-\gamma) \sum_{i=1}^{\infty} \gamma^{i-1}(m_{t-i}^p + n_{t-i})] - \tfrac{1}{2}(m_t^p)^2]$$

The solution to this is again characterized by the Euler conditions;

$$0 = (A+p) - m^p_t - (A+p)(1-\gamma)[\beta + \gamma\beta^2 + \gamma^2\beta^3 + \ldots].$$

Or,

$$m^p_t = (A+p)(1-\beta)(1-\gamma\beta)^{-1} < (A+p). \blacksquare$$

Under non-rational learning the optimal strategy for the government has two significant properties. First, the non-rational learning converges to full rational expectations. The estimator used by the public will converge to the value $(A+p)(1-\beta)(1-\beta\gamma)^{-1}$,

$$E[\mu^t] = (1-\gamma^t)(A+p)(1-\beta)(1-\beta\gamma)^{-1} + o(\gamma^{t+1}).$$

The fact that the government adopts a stationary strategy also implies that the specification of the public's econometric model is actually quite a close approximation to reality. The process generating the observations made by the public is;

$$m_t = (A+p)(1-\beta)(1-\gamma\beta)^{-1} + n_t.$$

Whilst the public estimates;

$$m_t = \mu + \epsilon_t.$$

Thus the public's econometric model should withstand a number of specification tests, because under non-rational learning the money supply is made up of a constant term plus noise. The only feature of the public's model specification which may be in doubt is the implicit assumption which underlies the WLS estimation procedure. Econometricians tend to use WLS in situations where there is doubt about parameter stability, or where there is heteroskedasticity present. Neither of these features are exhibited by the government's strategy under non-rational learning. Hence specification tests for parameter stability might accept this hypothesis, in which case OLS is the only legitimate estimation procedure. As OLS is a subclass of the learning procedures considered here this does not seem to be a major problem.

The second significant conclusion from Proposition 3 is that the level of monetary growth and hence inflationary bias is uniformly lower than the rational learning case; provided gamma is strictly less than unity. That is under non-rational learning it becomes credible and optimal for the government to set lower levels of monetary growth. From the point of view of information transmission the low level of bias slows down the public's learning, because the relative noise in the government's strategy is large. However, the important feature of the result is how the low level of inflationary bias becomes credible. The learning process specified in (12) does have the property that new information never receives zero weight in the private sector's expectations. This implies that cheating on the optimal policy outlined in Proposition 3 always results in a non-negligible shift in the public's expectations. Hence cheating always imposes non-negligible future costs on the government, because there is a significant shift in beliefs. As in the models of Barro & Gordon, these future costs then sustain a lower credible level of inflation. The comparative statics of the equilibrium strategy make this clear. The optimal level of monetary growth is an increasing function of gamma and a decreasing function of β. Small gammas give a high relative weight on current information in the public's expectations, so the public's reaction to a deviation becomes larger. This imposes larger expected costs from a deviation and lower levels of inflation become credible. As the government's discounting parameter becomes larger the weight on future costs in its utility function increases, so the future costs are perceived to be higher.

The results here contrast with those under rational learning, because under rational learning the learning process used converges to a situation where zero weight is put on new information. In the limit the only credible policy is one which generates the full inflationary bias. The same outcome happens under non-rational learning in two polar outcomes; as β tends to zero, or as gamma tends to unity. The first case occurs through governmental myopia, this must result in convergence to the one shot game. The second case occurs when the public's learning approaches OLS. This estimator ultimately puts zero weight on new information, hence there is no future cost to cheating when expectations converge. As a result the only credible level of inflation is that which we observe in the one-shot game with full rational expectations. This suggests that if the public engages in specification tests (like those discussed above) and eventually rejects the notion of unstable parameters, then the non-rational learning process will converge to precisely the outcome of rational learning.

The argument above seems to suggest that even under non-rational learning the public will eventually converge to the outcome of fully rational learning, because the public will find OLS to be the correct econometric technique. Another argument in favour of OLS can be constructed, if we notice that the public's payoff is a function of the parameter gamma in their learning. This parameter will determine the public's payoff in the following way;

$$\lim_{t \to \infty} E[-(m^e_t - m_t)^2] = -2\sigma_n^2(1+\delta)^{-1}.$$

The public's payoff is increasing in gamma, so if the public adjusts its learning to maximise its expected payoff, then there will also be a tendency for the public to choose gamma close or equal to unity. The intuition behind this is clear; OLS minimises the variance in the estimator, as it puts progressively smaller weights on the noisy current observation. This optimum is based on the assumed preferences of the public, in this context it seems unreasonable to rule out the possibility that the public also cares about the absolute level of inflation and not just its mistakes in prediction. If this were so the public would trade off a smaller prediction error against lower inflation.

5. CONCLUSION

The process of learning rational expectations appears to have significant consequences for observed government behaviour. If we assume rational learning there will be low levels of initial inflationary bias in government policy and the government slows the public's learning. Thus policy ambiguity in the early years of a regime will be high and credibility low. If rational learning is seen as unsatisfactory, then the results on non-rational learning may be of some interest. The results are promising; the public starts life with little information on the structure of the model and yet ends up with an unbiased predictor of money supply growth. More significantly there is a permanently reduced inflationary bias under most forms of non-rational learning. So there is a reduction in the credible sustainable level of inflation when the public is learning using WLS. The public's econometric model will not be readily controverted by the data, although it may be vulnerable to tests for non-stable parameters.

Department of Economics, University of Warwick,
Coventry, CV4 7AL, UK, February 1988.

APPENDIX

PROOF OF PROPOSITION 2: We first show that there is a unique sequence $\{B_t\}$ which satisfies (11) and solves the government's optimisation problem. From (11) we know the sequence must satisfy;

$$(A.1) \qquad 0 = B_{t+1}^2(\beta W_t) + B_{t+1}(B_t(1-\beta) - 1) + \beta B_t$$

There are multiple possible time paths for the sequence $\{B_t\}$ defined by this relation, because the quadratic defines two possible values of B_{t+1} for every possible B_t. One can use the optimality and credibility properties of the government's optimal strategy to deduce a unique increasing sequence of B_t's satisfying (A.1). We first eliminate the possibility that the sequence converges to zero. Suppose B_t tends to zero, then $E[m_t|I_t]$ also must converge to zero, this cannot be a credible equilibrium because as B_t converges the government's policy choice in period t converges to;

$$\max E_{G0}\ (A+p)(m^p{}_t + n_t - 0) - (1/2)(m^p{}_t)^2 + 0.$$

The future effects of a deviation become negligible because W_t has converged to zero, so the choice of an optimal strategy in period t becomes essentially myopic. Hence the government prefers to set $m^p{}_t=A+p$ and to deviate from the putative equilibrium $B_t=0$. This is a contradiction. The series B_t does not converge to zero and so the series $W_t=[r+\Sigma(B_s)^2]^{-1}$ will converge to zero.

Suppose $B_t\geq(1-\beta)^{-1}$, then there is no positive solution to (A.1). If $B_t\leq 0$, then there are two possible solutions to (A.1)

$$\text{(A.2)} \qquad B_{t+1} = (2\beta W_t)^{-1}\{1-(1-\beta)B_t \pm \sqrt{[(1-(1-\beta)B_t)^2-B_t\beta^2 W_t]}\}$$

This is obtained by solving the quadratic equation (A.1). One solution to (A.2) satisfies $B_t<B_{t+1}<0$. The other implies $B_{t+1}>(1-\beta)^{-1}$. These two observations establish that there are two possible time paths for $\{B_t\}$ with observations outside the interval $(0,(1-\beta)^{-1})$. Either it tends to zero from below, which is not consistent with credibility, or the sequence alternates positive and negative values with the positive values satisfying (A.2). These must tend to infinity as W_t tends to zero, this contradicts the transversality condition (10), $m^p{}_t$ will become unboundedly large. Thus both these possibilities can be eliminated.

On the interval $[1,(1-\beta)^{-1})$ the B_t values satisfy $B_{t+1}>B_t$ and tend to infinity as B_t approaches $(1-\beta)^{-1}$. Thus any sequence in this interval eventually becomes one of the two classes described above.

The behaviour on (0,1) remains to be described. There are three types of possible sequences satisfying (A.1) on this interval, given a particular starting value $0<B_0<1$ either B_t tends to zero, or B_t tends to infinity as above; there is also a unique value B_0 such that B_t tends to unity.

This last sequence will now be proved to exist. The mapping (A.2) from B_t to B_{t+1} is modified into a function $f_t(B_t)$

$$\text{(A.3)} \qquad B_{t+1} = (2\beta W_t)^{-1}\{1-(1-\beta)B_t - \sqrt{[(1-(1-\beta)B_t)^2-B_t\beta^2 W_t]}\} \\ =: f_t(B_t)$$

The other root to the quadratic will generate arbitrarily large values for B_{t+1} as W_t tends to zero. (A.3) is well defined provided B_t is in (0,1). The functions f_t intersect the origin, (B_t=B_{t+1}=0 is always a solution to (A.1)). One can also show that B_{t+1} increases in B_t for all t, provided $r \geq 1$; and as W_t tends to zero f_t shift upwards on (0,1). These functions also intersect the 45° line at $B_t=B^*_t$ where

$$\text{(A.4)} \qquad B^*_t = (1-\beta)/\{1-\beta+\beta W_t\} \qquad (0<B^*_t<1).$$

We now prove the existence of a sequence $\{B'_t|t=0,1,..\}$ which satisfies A.1 and converges to unity from below. Given the properties outlined above f_t has an inverse f_t^{-1}. Define a sequence $\{B^n_t|t=1,..\}$ as follows

$$\text{(A.4)} \qquad B^n_t := \begin{cases} f_t^{-1}(B^n_{t+1}) & t < n \\ B^*_t & t \geq n. \end{cases}$$

This sequence is takes the value B^*_t for observations greater than or equal to t, and for smaller elements of the sequence it takes a value B^n_t, such that $B^t_t=B^*_t$. It is obvious that the sequence $\{B^n_t|t=1,2..\}$ converges and is increasing in t. Now define $B'_t= \lim_{n\to\infty} B^n_t$ this we claim is the unique stable sequence of policy weights. The sequence $\{B^n_t|n=1,2,..\}$ converges, as it is increasing on [0,1]; each B'_t is well defined. To establish that the sequence $\{B'_t\}$ thus defined converges we can also use the fact that it is increasing and bounded above, and since B^*_t tends to unity as t becomes large we must deduce that B'_t tends to unity. The uniqueness of this convergent sequence stems from the observation that any sequence $\{C_t\}$ on [0,1] with the property $C_t<B^n_t$ for some n and t must tend to zero. A similar upper bound can be constructed.

REFERENCES

BACKUS,D., AND J.DRIFFILL (1985)"Rational expectations and policy credibility following a change in regime." *Review of Economic Studies,* 55, 211-221.

BARRO,R.J., AND D.B.GORDON (1983) "Rules, discretion and reputation in a model of monetary policy." *Journal of Monetary Economics,* 12, 101-122.

BASAR,T. AND M.SALMON (1988) "Credibility and the value of information transmission in a model of monetary policy and inflation." mimeo 1987.

BRAY,M.M., AND N.E.SAVIN (1986) "Rational expectations equilibria, learning and model specification." *Econometrica,* 54, 1129-1160.

CUKIERMAN,A., AND A.H.MELTZER (1986)"A theory of ambiguity, credibility and inflation under discretion and asymmetric information." *Econometrica,* 54, 1099-1128.

HARVEY,A.C.(1983): *Time Series Models.* Oxford: Phillip Allen.

LUCAS,R.E.(1981) "Rules discretion and the role of the economic advisor." in: *The Phillips Curve and Labor Markets,* ed by K.Brunner and A.M.Meltzer, 19-46.

SARGENT,T.J.(1979): *Macroeconomic Theory.* New York: Academic Press.

FOOTNOTES

1. I am grateful for the comments and suggestions of Douglas Gale, Norman Ireland, John Moore, Gareth Myles and Mark Salmon.

2. The incentive to create surprise inflation derives from the increase in output that results from such a surprise, see Barro & Gordon (1983).