AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
http://ageconsearch.umn.edu
aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

Reputation Effects in Dynamic Games

by

Martin Cripps[1]
Department of Economics
University of Warwick

No. 329

October 1989

This paper is circulated for discussion purposes only and its contents should be considered preliminary.

Reputation Effects in Dynamic Games

Martin Cripps[1]
Department of Economics, University of Warwick,
Coventry CV4 7AL, UK.
January 1989.

Abstract: The solution of a reputational equilibrium is given for a class of linear, quadratic, gaussian dynamic games with noisy control. Although there is imperfect monitoring, a sequential equilibrium is found where the uninformed agents always smoothly learn the type of the informed agent, there is no sudden switch in agents' strategies; a common feature of reputation models. Reputation effects are temporary in the infinite horizon case for positive discount rates, as the discount factor tends to unity there is a permanent reputation.

[1] Thanks are due to Norman Ireland, Gareth Myles, Jonathan Thomas and a seminars at Southampton University and the CentER for Economic Research, Tilburg, for helpful comments and suggestions, the remaining errors are all mine.

1. Introduction

Here we attempt to show that linear quadratic gaussian (LGQ) models and the reputation effects of game theory can be linked. We focus on the reputation models introduced by Kreps & Wilson (1982), Milgrom & Roberts (1982) and applied in many well known papers; for example Backus & Driffill (1985), Benabou & Laroque (1988). There has been work on imperfect monitoring in reputational models; Benabou & Laroque (1988), Fudenberg & Levine (1988, 1989), there has also been work on models with a continuum of types Milgrom & Roberts (1982). We show how reputation results can be established in LGQ games, which are dynamic with imperfect monitoring, a continuum of types, and contain actions sets which are a continuum. Thus it is possible to combine these various features in a tractable framework and prove a result on reputational equilibria when strategy sets are non-finite. The results presented characterize the sequential equilibrium strategy of the informed agent and establish the convergence to full information in infinite horizon games. We argue that for any positive discount rate reputation effects will be temporary in infinite horizon games when action sets are large. We show that the length of these temporary reputations will increase as discount rates tend to zero and become permanent only in the limit. We also show that the lower bound on payoffs to Nash equilibrium, derived by Fudenberg & Levine, applies in our model although there is always revelation of a player's type.

Fudenberg & Levine (1988,89) have established that in an infinite horizon reputation game, as the discount factor tends to unity, any type's payoff at a Nash equilibrium is bounded below by the payoff to masquerading as a dominant strategy type. This result holds in the model below. A necessary condition for the Fudenberg & Levine conclusion is that the dominant strategy type receives strictly positive probability in the priors of the short term agents. This is not true in our model as there is only one type (in the continuum) with a dominant strategy. There is, however, a strictly positive probability that the long term player's type lies in an arbitrarily small interval around the dominant strategy type. Thus it appears that the Fudenberg & Levine result persists in the limit. If dominant strategy types are surrounded by a neighbourhood of similar (but

non-dominant strategy types) the lower bound on equilibrium payoffs survives.

The behaviour at a sequential equilibrium of our reputation-type game differs in two ways from the process of reputation building and sudden destruction exhibited in Kreps & Wilson (1982), Milgrom & Roberts (1982,1982), Backus & Driffill (1985). First, the results below show that there is always *smooth* convergence of play to a position where the long term player's type is revealed. The models of Fudenberg & Levine and Benabou & Laroque show that in games with imperfect monitoring the process of reputation building and destruction is considerably smoother than in games without such noise. It is clear that it is the presence of imperfect information transmission that smooths the abrupt deterioration in reputation. The second distinction between the results here and those of Kreps & Wilson is the behaviour of play in the infinite time horizon case. Here the long-term player's type is revealed when play continues indefinitely. This differs from the results of Kreps & Wilson, which show that players acquire a reputation *and never lose it*, by masquerading as a dominant strategy type. One simple explanation for the absence of a long term reputation building in our model is that the dominant strategy type is of measure zero in the priors, so it is impossible for other types to credibly acquire a reputation. This explanation is suspect, because the behaviour of other types in the early periods does appear to be an attempt to imitate the dominant strategy type; there is an attempt at reputation building.

There has been some debate on how the form of the set of types may affect the type of outcomes sustainable as a sequential equilibrium in a reputation game, see for example Vickers (1986). In particular it has been noted that agents have incentives to take actions which reveal their type early on in the play; to self-screen. This incentive to screen will not operate effectively if the strategy sets of agents are very limited; different types will select the same actions if constrained in their choice. (For example, the model of Milgrom & Roberts does have a continuum of types but has two-element strategy sets.) In our model strategy sets are the continuum and there are a continuum of types, this provides much greater opportunities for self-screening by different types and the tendency to do this is much larger. Self-screening will also explain why it is impossible for types to permanently masquerade as the dominant strategy type, but it is possible to gain a short-term reputation. The imperfect monitoring of actions means that it takes time for types to self-screen, so they can enjoy a temporary reputation while this is going on. The pressure to self-screen can also provide an explanation for the

results of Benabou & Laroque; there are two types but actions are a continuum. In spite of a positive measure for both types in this case the masquerade is not permanent and in infinite time there is full revelation of types.

The introduction of noise in strategies has the effect of ensuring that every information set in the game is reached with a non-zero probability, this eliminates the need to describe off-the-equilibrium-path beliefs at the sequential equilibrium found below. There is no source of multiple equilibria from different specifications of these beliefs. One could view the introduction of noise into strategies as forcing agents to play a particular form of mixed strategy or of imposing a particular sort of tremble, which is, of course, undesirable. The presence of noise, however, seems natural for many economic environments where there is a discrepancy between planned and actual events.

The sequential equilibrium strategy of the long-term agent has a particularly simple form in this model. The action chosen by the long-lived agent in each period is a linear function of its type. The strategy is not responsive to past shocks or noise; the actions selected in any period are the same whatever the past history of shocks. So it is possible that the short-run agents in the model over estimate the long term agent's type after a particular sequence of shocks and subsequently revise this estimate downwards. Thus the beliefs, or assessments, of the short term players are not revised monotonically, but nevertheless converge to full information on the type. As the long term player's strategy converges to its limiting levels.

Section 2 below outlines the class of models we consider and examines the properties of the game when the type is common knowledge. In section 3 we characterize a sequential equilibrium for this game and show how it behaves as time tends to infinity. In sections 4 show how the structure developed here can be applied to two economic problems, one of limit pricing, the other from the optimal policy literature.

2. The Model

Player (A) is infinitely lived and aims to maximize the discounted sum of period by period payoffs. The player (B_t) has a one period life and receives its payoff at the end of period t=0,1,2,..; its period of existence. Player A plays a game against a sequence of identical players; one type B_t in each period. This structure is common to most of the literature on reputation models. The B_t's have no incentive to choose actions to affect the rate of information acquisition from A. This is because B_t's will cease to exist after the current period is over, so they do not care about better information in the future.

The type of player A is described by a parameter $z \in IR$. The prior beliefs of the B_t's on z are described by a normal distribution $N(\mu,\sigma_z^2)$, this assumption is necessary to enable a signal-extraction approach to the revision of priors. Player A observes its type before play begins and z enters as a parameter in A's preferences. The value z is unobservable by B. The period t payoffs to player A and B_t (U^A_t, U^B_t) are quadratic and are defined below

$$U^A_t = z[\alpha_1 a_t + \alpha_2 b_t + \alpha_3 x_t] - (0.5)\alpha_4 (a_t)^2, \qquad (1)$$

$$\text{where;} \quad a_t = a^p_t + n_t.$$

$$U^B_t = \beta_1 (a_t)^2 + \beta_2 a_t b_t - (0.5)\beta_3 (b_t)^2 + f(x_t, a_t), \qquad (2)$$

$$\text{where;} \quad b_t = b^p_t + m_t.$$

(We assume that; $\alpha_4, \beta_3 > 0$ for concavity.) The model is univariate to preserve the simplicity of structure, however, there is no reason why the results cannot be generalized to the multivariate case. Here a_t^p represents A's action in period t. Also b_t^p is B's action and x_t is a state variable in period t. The actions are subject to random error terms n_t, m_t each period, which are normal, independent and identically

distributed through time with distributions $N(0,\sigma_n^2)$ and $N(0,\sigma_m^2)$ respectively. The players have a noisy control over their actions, this means B_t is unable to monitor A's actions precisely as it can only observe the realization a_t. The noise in B_t's strategy is not essential to the results but is included for generality. Player A's payoff each period is almost linear, but does have a quadratic term in the control variable which has a negative sign to preserve concavity. It is not necessary to restrict A's preferences to only one quadratic term for the optimal strategy to be linear. Indeed they can be generalized considerably, but to solve for the time path of the optimal strategy some restriction is necessary. The player B_t has preferences which are quadratic in the two controls. They also have an arbitrary function of the state variable and A's control variable f(.) as an additively separable term. The state variable evolves according to the linear relationship given below:

$$x_{t+1} = \gamma_1 a_t + \gamma_2 b_t + \gamma_3 x_t, \qquad x_0 = x, \qquad -1 < \gamma_3 < +1. \tag{3}$$

Play in the game evolves in the following manner. At time t = -1 player A observes the value z and the initial value of the state variable is observed by both players. In period t=0 the players A and B_0 respectively choose a^p_0 and b^p_0 simultaneously. The the values n_0 and m_0 are determined and then players receive their payoffs. In each subsequent time period player A, and its opponent in that period B_t, observe x_t then choose a^p_t and b^p_t respectively. The information available to both players at this point in time consists of the past history of the state variable and the sequences of observed disturbed actions $(a_0,..,a_{t-1};\ b_0,b_1,...,b_{t-1})$. This game is repeated a finite or an infinite number of times, then play ends. Player A's payoffs accumulate as play progresses and the player B_t of each period collects its final payoff at the end of the period in which it acts.

The dominant strategy type is parameterized by z=0. In this case (1) becomes independent of the state variable and B_t's action. The type z=0 payoff is maximized by setting $a^p_t=0$ whatever the actions of the other agents; this is a dominant strategy. All other types will in general prefer to vary their actions in response to the state of information of their opponents.

The full information game has a simple structure, it is only the presence of uncertainty about types that complicates play. When A's type is known to the B_t's their reaction

function will be

$$bP_t = (\beta_2/\beta_3)aP_t.$$

Thus player B_t chooses bP_t as a linear function of aP_t. When we consider the full information repeated game the form of player A's optimal choice must maximize the discounted sum of its payoffs. Solving the appropriate dynamic programming problem gives

$$a_t^p = \frac{\alpha_1}{\alpha_4}z + \frac{\alpha_3\delta\gamma_1\gamma_3}{\alpha_4(1-\delta\gamma_3)}z.$$

(The term $0<\delta<1$ is player A's discount factor.) Player A's choice is linear in z and independent of bP_t. This means that every type of player A has a dominant strategy in the full information game. The first term reflects the optimal level of aP_t in the one shot game. The second term reflects the benefit from the effect aP_t has on future state variables. It is clear that in the full information game the form of the equilibrium behaviour is very simple; all types of player A have a dominant strategy and player B correctly predicts this when choosing its action.

The play in the game of asymmetric information is far from simple. All types of player A, except z=0, do *not* have an action which they prefer to play in every period. This is because their actions today affect B_t's behaviour tomorrow and hence the future payoffs of all types, other than z=0. As play evolves and B_t's beliefs alter there is a shift in the optimal action for player A. Player A's actions in each period will thus vary as the state of B_t's information varies. Also, player A's actions in any period will affect future payoffs via the information of B_t. Player A will prefer low levels of bP_t if $z\alpha_2<0$ and high levels if $z\alpha_2>0$. Low levels of bP_t can be achieved by encouraging player B_t to have respectively low (high) expectations for of aP_t if, $\beta_2>0$ ($\beta_2<0$). Thus, if player A educates player B_t's to expect low values of aP_t, then it will benefit in the future from lower values of bP_t. In choosing its action player A must trade off the benefit from future events against the present payoff. There are two possibilities we must consider; the first possibility is that player A wants to curtail bP_t, this is achieved by lower (higher) levels of aP_t in the early stages of play if $\beta_2>0$ ($\beta_2<0$). The second possibility is that player A wants higher bP_t's and to achieve this the reverse policy is followed. In the proofs below characterizing the sequential equilibrium of the reputation game these two separate possibilities will feature as the two cases of a

proposition.

3. Equilibrium

A sequential equilibrium in this model is now described. Player A's information set consists of the following pieces of information $h_t = \{z; a_0, a_1, .., a_{t-1}; b_0, b_1, ..., b_{t-1}; a^p_0, a^p_1, .., a^p_{t-1}\}$, and B's information set consists of $k_t = \{a_0, a_1, .., a_{t-1}; b_0, b_1, ..., b_{t-1}; b^p_0, b^p_1, .., b^p_{t-1}\}$. Let H_{At} denote the set of all possible h_t's and H_{Bt} denote the set of all possible k_t's. A strategy for player A consists of a sequence of functions $\Pi_A := \{ a^p_t(h_t) \mid t=1,2...\}$, which selects a value a^p_t for every possible history. Similarly, a strategy for each B_t is a sequence $\Pi_B := \{b^p_t(k_t) \mid t=1,2...\}$, which selects a value b^p_t for every possible history. Let Σ denote the set of probability measures on the real line. Then B_t's beliefs or assessments on A's type can be described by a function $\tau_t : H_{Bt} \rightarrow \Sigma$. An equilibrium consists of a pair of strategies Π_A^*, Π_B^* and assessments τ^*_t, which result from the application of Bayesian revision of B_t's priors to the strategy Π_A^* and the observable data k_t, such that

$$\Pi_A^* \text{ Maximizes } \quad E_{At}\left[(1-\delta)\sum_{s=t}^{T} \delta^s U^A_s \mid h_t \right] \qquad \text{For all } h_t, t.$$

$$\text{given} \quad x_{t+1} = \gamma_1 a_t + \gamma_2 b_t + \gamma_3 x_t, \quad \text{and} \quad \Pi_{Bt}^*.$$

$$\Pi_B^* \text{ Maximizes } \quad E_{Bt}[\, U^B_t \mid k_t \,] \qquad \text{For all } k_t \text{ and } t.$$

$$\text{given} \quad x_{t+1} = \gamma_1 a_t + \gamma_2 b_t + \gamma_3 x_t, \quad \text{and} \quad \Pi_A^*.$$

Here E_{At} represent expectations taken relative to A's information set at time t and E_{Bt} represents expectations relative to assessments τ^*_t. Also, $0 < \delta < 1$, denotes player A's subjective discount factor, which is used to weight future payoffs. This definition ensures that the strategies Π_A^*, Π_B^* form a sequential equilibrium for the game outlined above. At each possible information set the strategy employed by a player is an optimal response to the opponent's strategy.

The method of solving for a sequential equilibrium in the game outlined in Section 2 proceeds in two stages. The first step requires us to postulate a form for player A's equilibrium strategy and then deduce the nature of player B_t's knowledge about the private information z, given its knowledge of A's strategy.That is we describe how B_t determines its assessments τ^*_t. Once this is completed we will show that given the agent B's form their beliefs about z in this way, then it is indeed optimal for player A to choose a strategy of the form first postulated. To describe the strategy in greater detail we must solve for the coefficients of such a strategy. This approach to solving the model derives from the "method of undetermined coefficients", which is widely used in the macroeconomics literature. The approach applied here also has a much closer resemblance to that taken in Cukierman & Meltzer (1986).

First, assume that the B_t's believe that agent A sets $a^p_t = zC_t$; for some fixed sequence $\{C_t \mid t=0,1,..\}$ of real numbers. Given the agent B's believe that A abides by such a policy, we must show that A does indeed then prefer to use this strategy. This is how we construct an equilibrium. Note, we do not restrict player A to use a linear strategy in any of this. We show that provided the B's *believe* that A follows a linear strategy, then, it is best for A to behave in this way even if it can choose any strategy it likes. This is sufficient for the solution here to be a sequential equilibrium. That is not to say that other equilibria do not exist, indeed in general there may be a great number of sequential equilibria for any game. A general characterization of the set of all Nash equilibria in repeated games with incomplete information of this type has been provided by Hart (1985).

If player A follows the strategy suggested above we can deduce that B_t's beliefs on A's type continue to be a normal distribution. This is because the B_t's will face a signal extraction problem when they are forming their expectations of z. That is, they observe a sequence $a_0, a_1, .. a_{t-1}$,(where $a_s = zC_s + n_s$) and using this must estimate a value for z. The mean of the distribution of beliefs on z is the only moment used by B_t's in deciding how to behave and we now will present a result which shows that the mean of the distribution of beliefs is a linear function of past data. The solution to this signal extraction problem is given by a well known result.

PROPOSITION 1: Given, $a_s = zC_s + n_s$, then $E_{Bt}[\, z \mid k_t \,]$ satisfies;

$$E[\, z \mid k_t] = \left(\frac{\sum_{i=0}^{t-1} C_i a_i + r\mu}{r + \sum_{i=0}^{t-1} C_i^2} \right), \qquad r = \frac{\sigma_n^2}{\sigma_z^2}. \tag{4}$$

PROOF: This result can be proved in a number of ways, we will establish it in the appendix using a property of least squares projections.

This completes the description of player B_t's beliefs or assessments τ^*_t, if it believes that player A follows the linear strategy $a_s = zC_s + n_s$. At equilibrium player B_t maximizes its payoff given these assessments. Take the first order condition for (2) and substitute for $E[z| k_t]$ from (4)

$$0 = E_{Bt}[\ \beta_2 a_t - \beta_3(b_t) \mid k_t\ ]$$

$$\Rightarrow\ bP_t = (\beta_2/\beta_3)E[\ a_t \mid k_t] = (\beta_2/\beta_3)E[\ zC_t + n_t \mid k_t\ ]$$

$$\Rightarrow\ bP_t = W_t\left(\sum_{i=0}^{t-1} a_i C_i + r\mu \right); \qquad \text{where } W_t = \frac{\beta_2 C_t}{\beta_3}\left(r + \sum_{i=1}^{t-1} C_t^2 \right)^{-1} \tag{5}$$

B_t's behaviour is in turn a linear function of $E[\ z \mid k_t\ ]$, which is a linear function of A's actions. We now characterize A's optimization problem. First write down A's objective function at time s substituting the above expression for B's actions and solving for the state variable by repeated substitution. To write down the conditions for A's equilibrium strategy we must maximize the expression below. At any time s, given any past history h_t player A selects current and future actions $\{aP_s, aP_{s+1}, aP_{s+2}, \ldots\}$ to maximize;

$$E_{As}\left[(1-\delta)\sum_{t=s}^{T} \delta^t \left\{ z\alpha_1 a_t + z\alpha_2 W_t\left(\sum_{i=0}^{t-1} C_i a_i + r\mu \right) + z\alpha_2 m_t - \alpha_4(0.5)a_t^2 + z\alpha_3\left(\gamma_3^t x_0 + \sum_{i=0}^{t-1} \gamma_3^{t-1-i}(\gamma_1 a_i + \gamma_2 W_i \sum_{j=0}^{i-1} C_j a_j + \gamma_2 m_i) \right) \right\} \right]$$

$$\text{where } a_t = aP_t + n_t. \tag{6}$$

At every information set in the game player A will choose a^p_s to maximize the (6), given the information h_t and A's knowledge about the form of B's behaviour (5). To optimize (6) we apply dynamic programming, a solution to this problem is given by the Euler equations, for example see [Sargent (1979) Chapter XIV]. In effect we must differentiate (6) with respect to a^p_s and to set it equal to zero. Notice that a lot of the disturbance terms m_t, n_t vanish from the derivative by linearity. So current and past disturbances are irrelevant to the form of A's optimal response to (5). The Euler equation is given below.

$$0 = -\alpha_4 a^p_s + z\left[\alpha_1 + \alpha_2 C_s\left(\sum_{i=1}^{T-t} \delta^i W_{s+i}\right) + \alpha_3 \gamma_1 \sum_{i=1}^{T-t} (\gamma_3\delta)^i + \right. \tag{7}$$
$$\left. + \gamma_2\gamma_3\alpha_3\delta^2 C_s\left\{W_{s+1}\sum_{i=0}^{T-t-1} (\gamma_3\delta)^i + \delta W_{s+2}\sum_{i=0}^{T-t-2} (\gamma_3\delta)^i + \ldots + \delta^{T-t} W_T\right\}\right]$$

This equation has the same structure as that postulated for player A's strategy: it solves to give a^p_s linear in z. This indicates that we have found a sequential equilibrium, provided we equate the coefficient of z in the equation above with C_t and then solve for the sequence C_t. We will consider two possible cases below; an infinite time horizon or a finite horizon with no-state variable $\gamma_1=\gamma_2=\gamma_3=0$. First assume an infinite horizon. Now equating C_s with the coefficient of z in the Euler equation gives

$$C_s = \frac{1}{\alpha_4}\left(\alpha_1 + \frac{\alpha_3\gamma_1\gamma_3\delta}{1-\gamma_3\delta}\right) + \frac{1}{\alpha_4}\left(\alpha_2 + \frac{\alpha_3\gamma_2\gamma_3\delta}{1-\gamma_3\delta}\right) C_s \sum_{i=1}^{\infty} \delta^i W_{s+i} \tag{8}$$

The first term on the right is precisely the value C_s would take if z was common knowledge. The second term is the effect from the B_t's learning of z on A's choice of strategy. We will simplify this expression in the following way

$$C_s = \omega + \phi C_s \sum_{i=1}^{\infty} \delta^i \left(\frac{C_{s+i}}{V_{s+i}}\right); \text{ where } \omega = \frac{1}{\alpha_4}\left(\alpha_1 + \frac{\alpha_3\gamma_1\gamma_3\delta}{1-\gamma_3\delta}\right); \quad \phi = \frac{\beta_2}{\beta_3\alpha_4}\left(\alpha_2 + \frac{\alpha_3\gamma_2\gamma_3\delta}{1-\gamma_3\delta}\right).$$
$$V_t := r + \sum_{i=0}^{t-1} C_i^2 . \tag{9}$$

Now assume a finite horizon and no state variable, this gives an Euler equation

$$C_s = \frac{1}{\alpha_4}\left(\alpha_1 + \alpha_2 C_s \sum_{i=1}^{T-t} \delta W_{t+i}\right) \qquad (10)$$

The solution to (6) is not characterized solely by the Euler conditions, but also requires a transversality condition. This is

$$\lim_{t\to T} \delta^t E_t[\, z\alpha_4 C_t - \alpha_1 z \,] = 0.$$

At this juncture we have established that the form of agent A's strategy in this game is linear. Provided (9),(10) can be solved for a sequence C_t we have a sequential equilibrium. This equilibrium strategy is not a consequence of any form of restriction on A's actions, the optimization problem (6) is completely general in the scope of actions it allows A. The form of the equilibrium is a consequence of the linearity of B_t's signal extraction problem and the restricted quadratic form of A's preferences. Linearity eliminates many of the effects of noise and shocks: these simply disappear as a consequence of differentiation. The form of the strategy is linear with no effects from shocks during the course of the game. This does appear to be a very special result, which cannot be generally true. In general the past history of shocks will have effects on the equilibrium strategy in such a game.

We can go no further in characterizing the strategy of player A without considering the case of an infinite and finite time horizon separately. As the infinite time horizon case is simpler to treat we begin with this. The following result solves (9) for the optimal strategy and shows how it converges as the T tends to infinity.

PROPOSITION 2: There is a unique sequence of real numbers $\{C_t \mid t=0,1,..\}$, which satisfy (9) and the transversality condition. There are four possible forms for these sequences, but in each case C_t tends to ω as t tends to infinity;

(i) $\omega>0,\phi<0$: $0<C_t<C_{t+1}$,

(ii) $\omega<0,\phi>0$: $0>C_t>C_{t+1}$,

(iii) $\omega>0,\phi>0$: $0<C_{t+1}<C_t$,

(iv) $\omega<0,\phi<0$: $0>C_{t+1}>C_t$.

PROOF: Divide (9) through by C_t. Take the resultant equation lead it once, multiply by δ and then take the difference

$$0 = (\delta\phi/V_{t+1})C_t[C_{t+1}]^2 + C_{t+1}\{\omega - (1-\delta)C_t\} - \delta\omega C_t. \qquad (11)$$

This is a relatively simple equation determining two values of C_{t+1} for every pair C_t, V_t. Also, note that either V_t tends to a finite positive number, or to infinity. We now show that the sequence C_t cannot tend to zero and hence that the sequence V_t must tend to infinity; recall $V_t=(r+ C_o^2+..+C_{t-1}^2)$. Suppose otherwise and C_t tends to zero, then the solution to B's signal extraction problem (4) gives increasingly less weight to new observations, so future variables are less sensitive to current values of B's information. Consequently the Euler conditions (7) and the derived relation (8) converge to $aP_t = \omega z$, this is a contradiction ($\omega \neq 0$), so C_t does not converge to zero. This implies that V_t tends to infinity. Now rewrite (11) noting that $V_{t+1} := V_t + C_t^2$.

$$0 = \delta\phi C_t C_{t+1}^2 + C_{t+1}[\omega - (1-\delta)C_t](V_t + C_t^2) - \delta\omega C_t(V_t + C_t^2). \qquad (12)$$

This is a quadratic in C_{t+1} and has the following solutions;

$$C_{t+1} = \frac{-1}{2\delta\phi C_t}\left\{\left(\omega-(1-\delta)C_t\right)\left(V_t+C_t^2\right) \pm \sqrt{\left(\omega-(1-\delta)C_t\right)^2\left(V_t+C_t^2\right)^2 + 4\phi\delta^2\omega C_t^2\left(V_t+C_t^2\right)}\right\}$$

(13)

We will now use this relation between C_t and C_{t+1} to determine a solution to the equations (7) and the transversality condition. Now consider the two possible cases $\phi\omega<0$, $\phi\omega>0$.

<u>Case 1: $\phi\omega<0$.</u> (13) defines a value C_{t+1} for every C_t apart from those in a small neighbourhood about the point $\omega/(1-\delta)$. If $\omega>0$ ($\omega<0$) a value of C_t greater (less) than $\omega/(1-\delta)$ in (13) will give C_{t+1} negative (positive), because the first term in braces always dominates the second. For $C_t<0$ ($C_t>0$) the solutions are always negative (positive) for similar reasons, whilst for $0 \leq C_t \leq \omega/(1-\delta)$, ($0 \geq C_t \geq \omega/(1-\delta)$) the solutions are all positive (negative). As C_t becomes unboundedly large (negative) so does C_{t+1} either becomes unboundedly negative (positive), or it asymptotes to a negative (positive) constant. Similar properties hold as C_t tends to negative (positive) infinity. The relationship between C_t and C_{t+1} for $\phi\omega<0$, and $\omega>0$ is depicted in Figure 1. (If $\omega<0$ this figure must be reflected in the C_t axis and then reflected in the C_{t+1} axis.)

<u>Figure 1 about here.</u>

Henceforth we will only consider $\omega>0$, the proof with $\omega<0$ is equivalent *mutatis mutandis*. Any negative value for C_t will lead to the sequence of C_t's converging to zero or tending to infinity. This will violate the transversality condition above and thus provide a contradiction. If $C_t<0$, then on the higher branch of the map $C_t<C_{t+1}<0$ and C_t tends to zero which is a contradiction. Note: (13) may intersect the line $C_t=C_{t+1}$ for some negative value, such a point of intersection will occur on the lower branch AB of the map. These points cannot be a candidate for an equilibrium as the lower branch of the map AB is bounded above by $V_t(1-\delta)/\delta\phi<0$. This bound tends to minus infinity and will violate the transversality condition. Therefore, if $C_t<0$, then either C_t tends to zero, or the sequence C_t contains values tending to minus infinity. Hence no solution to the optimization problem of agent A can have negative values for C_t.

As the sequence $\{C_t\}$ has no negative values we eliminate all sequences which eventually have negative C_t's. Thus, any solution must have $0 < C_t < \omega/(1-\delta)$ for all t. We concern ourselves with the behaviour of C_t on this interval and the segment 0X in the figure. This segment intersects $C_t=C_{t+1}$ twice: at $C_t=0$ and at some intermediate value. The second intersection is from below because the slope of OX at $C_t=0$ is δ. If $V_t \geq 2$ this is a monotonic increasing function described by;

$$C_{t+1} = \frac{-1}{2\delta\phi C_t}\left\{ \left(\omega-(1-\delta)C_t\right)\left(V_t+C_t^2\right) - \sqrt{\left(\omega-(1-\delta)C_t\right)^2\left(V_t+C_t^2\right)^2+4\phi\delta^2\omega C_t^2\left(V_t+C_t^2\right)}\right\}$$

(14)

(The monotonicity of this function follows from differentiation of (14) with respect to C_t and that $2(\omega - (1-\delta)C_t - \omega V_t/C_t^2 <0$ for $0 < C_t < \omega/(1-\delta)$, if $V_t \geq 2$.) The differentiation of (14) with respect to K_t shows that as V_t increases so the segment 0X shifts downwards, because

$$-\sqrt{\left(V_t+C_t^2\right)^2\left[(1-\delta)C_t-\omega\right]^2+4\delta^2\omega\phi C_t^2\left(V_t+C_t^2\right)} > \left(V_t+C_t^2\right)\left[(1-\delta)C_t-\omega\right]+\frac{2\delta^2\omega\phi C_t^2}{\left[(1-\delta)C_t-\omega\right]}.$$

This holds provided $C_t<\omega/(1-\delta)$. Therefore let the function $C_{t+1}(C_t,K_t)$ describe (14), where defined, on $(0\ ,\omega/(1-\delta))$, this is invertible in C_t and shifts down if V_t increases.

Figure 2 about here.

Let $C^*(V_t)$ define the point where $C_{t+1}(C,V_t) = C$. As $0 < C_{t+1}(C_t,V_t) < C_t$ for all $C_t < C^*(V_t)$, any sequence which has a value C_t which satisfies this will subsequently decrease and converge to zero; this is a contradiction. The only sequences $\{C_t\}$ which do not yield a contradiction are those for which $\omega/(1-\delta) > C_t > C^*(V_t)$ for all C_t.

There is a unique sequence of points C_t satisfying (11), such that C_t tends to ω. Define $\chi_0 = C^*(V_0)$, then let χ_1 be the initial value for a sequence $\{C_t\}$ such that $C_1 = C^*(V_1)$, and χ_2 be the initial value of the sequence such that $C_2 = C^*(V_2)$. Proceed in this way to define a sequence of initial values $\{\chi_t\}$ for sequences which after t periods satisfy $C_t = C^*(V_t)$. There is a unique χ_t for all t because the function $C_{t+1}(C_t,V_t)$ is monotone and continuous in C_t and K_t over the relevant domains. The sequence $\{\chi_t\}$ is increasing and bounded above therefore it converges to a point X_0. By iterating (14) generate an entire sequence $\{X_t\}$ with an initial value X_0. This sequence converges to ω, because for all t; $C^*(V_t) < X_t < \omega$, and as V_t tends to infinity so must $C^*(V_t)$ tend to ω, from (11). This sequence is the unique sequence satisfying (11) with this property. Any other initial value for a sequence in $[0,\omega/(1-\delta)]$ must violate one of the inequalities $C^*(V_t) < C_t < \omega/(1-\delta)$. It suffices to observe that as the inverse of the function $C_{t+1}(C_t,V_t)$ has a derivative less than unity so the set of initial values which generate sequences to satisfy $C^*(V_t) < C_t < \omega/(1-\delta)$ must be an interval with a length which shrinks as V_t tends to zero. This completes the proof of case 1.

Case 2: $\omega\phi>0$. We will again concentrate on the case $\omega>0$, the proof for the case $\omega<0$ is analogous to the one given below with a change in signs. The map defined by (13) for this case is depicted in Figure 3.

Figure 3 here.

The plus and minus signs on the various branches of the map describe which of the solutions to (13) generate the arc of the map. The branches of the map which correspond to (+) for $C_t \leq 0$, and (-) for $C_t \geq 0$ are of little interest here. They are bounded below by $\phi^{-1}(1-\delta)V_t$ and above by $-\phi^{-1}(1-\delta)V_t$ respectively; any solution which lies on them must become unboundedly large or unboundedly small. The

remaining section of the map for $C_t \leq 0$ will lead to values for C_t which converge to zero and provide a contradiction. The section of the map in the positive quadrant is the only possible candidate for a stationary long run equilibrium as it contains the only values for which $C_t = C_{t+1}$ apart from the origin. (This follows from setting $C_t = C_{t+1}$ in (11) and examining possible solutions to this.) As above the section of the map shifts upwards over the range $0 \leq C_t$ as V_t increases. Hence the map shifts in the reverse manner to the case above. We can still employ a similar technique to that above to establish the existence of a decreasing sequence $\{C_t : t = 0,1,2,..\}$ which converges to ω. The details of this are identical so they are left to the reader. This completes the proof of Proposition 2.

A sequential equilibrium for the infinite horizon game outlined in section (2) has now been found. There exists a sequence of numbers $\{C_t\}$ such that, if player B_t believes that player A acts according to the rule $a^p_t = zC_t$, then it is optimal for the type z player A to follow this strategy at every information set. This has been established by first showing that the optimal response to these beliefs is linear in the type and then finding the sequence $\{C_t\}$ which is a fixed point in the map from beliefs to response.

The equilibrium strategy for all types of player A is described in this Proposition. The most obvious feature of the strategy is its convergence to a limit $a^p_t = z\omega$. This implies that the player A's type is slowly revealed to the B_t's, since the least squares projection (4) converges in probability to z if the C_t's are static. It is not surprising to learn that the limit $a^p_t = z\omega$ is the dominant strategy for player A when z is common knowledge. (The game of complete information is analyzed above.) As information is slowly revealed, so the play in the game of incomplete information settles down to the equilibrium of the game where z is common knowledge. This is shows how the full information game is also the limit of play in a game where z is not known.

The result does differ from those in other infinite horizon models where certain types permanently masquerade as others and there is never full revelation of types. For example, if $\gamma_1=\gamma_2=\gamma_3=0$, $\alpha_1=\alpha_4=\beta_2=\beta_3=1, \alpha_2=-1$ we have $\omega=1$, $\phi=-1$. The payoff for type z=0 converges to $-0.5\sigma_n^2$ as the C_t's converge to ω. The expected payoff for type z=1 converges to $-0.5(\sigma_n^2+1)$. In this example the C_t's are an increasing positive sequence converging to ω, (as $\phi<0$, $\omega>0$). This shows that there is

a temporary reputation to be gained from play in this game. Play commences with players of all types imitating z=0 by setting low values for $aP_t = C_t z$; thus they attempt to build a reputation. But as time passes C_t increases to the level ω and this period of temporary imitation dwindles away as play tends to the full information limit. Thus the process of convergence to the play in the full information game can be interpreted as a period of temporary reputation. The reasons players are able to build a temporary reputation in the game stem from the way information is processed by the B_t's in their signal extraction procedure. Early observations on a_t are very highly weighted in (4), so these will be important to the B_t's when they are forming their beliefs. All types of player A realize that in the early periods their actions reveal a lot of information on their type, so it is quite credible for them to prefer to imitate the z=0 type by reducing the absolute value of aP_t. The imitation also ensures them a higher future payoff by reducing the expectations of the B_t's on the A's type. This process of convergence to the full information steady state also reduces the rate at which the B_t's learn the type of the A's. Low values of a_t increase the relative size of the noise n_t in the B_t's signal extraction problem. It therefore takes longer for the B's priors to converge to correct information on A's type, than it would take if the A's had always set $aP_t=\omega$. Thus the route of the sequence C_t is selected to make learning of types much harder.

In answer to the question: why don't other types choose to permanently masquerade as z=0 in this example? There is the reply: it takes so long for another type to convince the B_t's that it is z=0, that any long run benefit from this is outweighed by the cost, in opportunities foregone, of acquiring such a reputation. Other types prefer instead to benefit from a short run, temporary reputation. The answer to this question also explains why the bound derived by Fudenberg & Levine continues to hold. As the discount factor tends to unity, so the equilibrium alters to make it harder to acquire a reputation as a z=0 type. Thus, the costs to such a masquerade exceed its benefits and pretending to be a z=0 type yields a payoff less than our equilibrium. In the limit the payoff from masquerading and the payoff from following our equilibrium strategy here become identical: this is explained in 3.1.

In the example above we examined the case $\omega>0,\phi<0$; the process of convergence to full information in the case $\omega<0,\phi>0$ is similar. In this class of games each type of player A prefers to masquerade as a z=0 type, but in the full information game it selects

a negative action, $a^p_t=z\omega<0$. The path to the steady state starts with the building of a temporary reputation since C_t is small in absolute value. As time passes $|C_t|$ grows and more information on A's type is revealed.

The cases where $\omega>0,\phi>0$ or $\omega<0,\phi<0$ are also similar in their path to full information, so we will examine these cases together. Consider for example $\gamma_1=\gamma_2=\gamma_3=0$, $\alpha_1=\alpha_4=\beta_2=\beta_3=\alpha_2=1$ this gives $\omega=1$, $\phi=1$. In the full information one shot game the payoff type z=0's payoff is $-0.5\sigma_n^2$, whilst type z=1 sets $a^p_t=1$ and receives a payoff of $1 - 0.5\sigma_n^2$. When $\omega>0,\phi>0$ or $\omega<0,\phi<0$ there are no benefits to masquerading as a dominant strategy type. Instead player A's payoff is increasing in the size of b^p_t, so player A prefers to encourage the B_t's to believe that their type is large. In the early periods of play in our example this is precisely what happens. The sequence C_t is positive and decreases in size until it reaches the level ω. Thus play begins with all types of player A imitating types with larger values of z and acquiring a temporary reputation as large types. As the B_t's learn the player's type play converges to the full information game, and the values of C_t decline. The large values for C_t in the early periods of play do speed the rate at which the B_t's learn from the signal extraction process (4). This is in the interest of the A's, however, as the B_t's are temporarily being convinced that player A's type is higher than it actually is. The path to the equilibrium in the game where $\phi<0,\omega<0$ is the mirror image of the example. Instead play is converging to a negative steady state.

The role of the noise is to make it more difficult for particular types to unambiguously signal their information. It is therefore interesting to examine what happens as σ_n^2 tends to zero. If $\omega>0,\phi<0$ the points $C^*(K_t)$ shift to the left, hence the values C_t will decrease in the infinite horizon case. As a result there will indeed be a longer period where player B has less information on z. Less noise therefore encourages player A to take more care in hiding its type.

We can now proceed to characterize the outcomes in a finite horizon game. The solution in this case is not entirely general because of the endpoint problems which arise when attempting to solve the relation (7) for a sequence C_t. In this case we are forced to entirely remove the dynamics from the model, to generate a soluble relation. This gives a condition (10) which can then be manipulated in the steps described in (11,12,13). In this case the relationship between C_t and C_{t+1} is identical to those

given in Figures 1 and 3. We will only consider the case where $\omega>0,\phi<0$ the results for the other three possibilities are very similar.

PROPOSITION 3 : If T is finite and if $\omega>0,\phi<0$, $\gamma_1=\gamma_2=0$, then there is a unique sequence $\{C_t \mid t=0,1,..,T\}$ of real numbers which satisfies (10) and the transversality condition $C_T = \alpha_1/\alpha_4$; such a sequence increases.

Proof : It is possible to repeat the argument used to derive (13) in Proposition 1 when T is finite. The relation (10) will govern the sequence C_t in the finite case under the assumptions above. The Euler condition in period T gives $0 = \alpha_1 z - \alpha_4 a^p{}_t$, hence we have $C_T = \alpha_1/\alpha_4 = \omega$, which ties down behaviour in the terminal period. As $C_T > 0$ we can be certain that $0 < C_t < \omega$, for $t \leq T-1$, as otherwise all subsequent values of C_t must be negative (see Figure 1). The fact that there is a unique terminal value implies that there is a unique sequence on the interval $[0, \omega/(1-\delta)]$ which satisfies the terminal condition and (10). From Proposition 2, we have shown that if $C_t < C^*(V_t)$ then subsequently the sequence $\{C_t\}$ decreases, but whilst $C_t > C^*(V_t)$ the sequence increases. But $\omega>C^*(V_t)$ for all V_t, if $C_T=\omega$ it must be the case that the sequence $\{C_t\}$ in $[0, \omega/(1-\delta)]$ satisfying (10) increases.

Player A's strategy is linear in z for both the finite and the infinite horizon models. Moreover, the stochastic events in the model have no effect on the terminal position in either case; the presence of noise or its magnitude do not alter the long run outcomes. In the infinite horizon case players B learn the parameter z, by solving the successive signal extraction problems posed in Proposition 1, so the noise plays no role here. In the finite horizon case the terminal condition is independent of the stochastic structure. There are important differences here because of the multiplicity of types, actions and the presence of noise. There is a smooth process of convergence to full knowledge, unlike the sudden loss of reputation observed in many reputational models. The equilibrium has a period of temporary reputation building but it always eventually collapses, in some cases Player A does take actions to slow B's learning of z. We now will show that the bound derived by Fudenberg & Levine does apply in this model with a continuum of types.

3.1 Discounting and the Length of Reputations

We have calculated a sequential equilibrium for the model set out above. We will now investigate how the length of the temporary reputation changes as discount factors tend to unity. We can also verify whether as discount factors become large the expected payoff to all types of player A is bounded below by the payoff to masquerading as a dominant strategy type. First take $C = C_t = C_{t+1}$ in (12) and rearrange

$$\frac{\delta\phi}{1-\delta}C^2 = (C - \omega)(V + C^2)$$

This allows us to investigate the properties of the intersections of the map (13) with the 45° line. If $\phi>0,\omega<0$; $\phi<0,\omega>0$, the intersection between the origin and ω approaches zero as δ tends to unity. Whilst, for the cases $\phi>0,\omega>0$; $\phi<0,\omega<0$ the intersection on the other side of ω becomes unboundedly large as δ tends to unity. We can use these properties of the intersections to deduce the changes in the equilibrium strategy as δ tends to unity, because our method of constructing the solution (see figure 2) relies on the properties of these intersections. In particular, as δ tends to unity the above implies that $|C_t|$ approaches zero if $\phi>0,\omega<0$; $\phi<0,\omega>0$ and $|C_t|$ approaches infinity if $\phi>0,\omega>0$; $\phi<0,\omega<0$. Thus, in both pairs of cases we have the result that as discount rates shrink, so does the sequential equilibrium strategy approach the dominant strategy type's play. As players become more far sighted the equilibrium we have found that play in all periods become closer to the masquerade.

Nevertheless, in the repeated game of incomplete information we have shown that players' strategies converge to those they employ in the full information game. Player A's expected payoff in each period will therefore tend to a value which is less than that received by the dominant strategy type for any positive rate of discount. It is not worthwhile to permanently masquerade as the dominant strategy type, because it takes so long to acquire such a reputation. As δ tends to unity the types become progressively more far sighted. To ensure that it is still not worth engaging in the permanent masquerade the benefits from following our equilibrium strategy are increased. As δ increases the terms in the sequence $\{C_t\}$ decrease, thus the period of temporary reputation grows longer. The temporary reputation achieved by the other types lasts longer and is more profitable as the discount factor tends to unity. In the

limit the temporary reputation is permanent as C_t either approaches zero, or becomes large, and all types successfully masquerade as the dominant strategy type. Thus the bound produced by Fudenberg & Levine is correct in the limit too, because in the limit the payoffs are equivalent. But for any positive discount rate it is more profitable to not masquerade permanently.

4. Possible Applications of the Results

The class of games analyzed above seem to have a number of possible applications to economics. One can think of the work on policy games and on entry as possible fields. A variant of the model proposed in Cukierman & Meltzer (1986) can be included in our framework and hence the technique described above can be applied in this context, this is done in greater depth in Cripps (1988). Agent A is the government choosing its monetary policy, and agent B is the private sector, which attempts to form correct inflationary expectations. The propensity of the government to benefit from inflationary surprises is the unknown parameter in the government's preferences. Let a^p_t be the planned rate of money supply growth and b^p_t be the private sector's expected rate of planned money supply growth, then the standard form for preferences in this game are;

$$U^A_t = z(a_t - b_t) - a_t^2,$$
$$U^B_t = -(a_t - b_t)^2.$$

These fit exactly the preferences described in Backus & Driffill (1985). Simple calculation reveals that $\phi < 0$ so that the optimal policy is a slow increase in the rate of growth of the money supply until the government's propensity to create inflation is revealed. The reputation type analysis applied in the finite horizon models of Backus & Driffill generalizes to the case where the uncertainty is generated by a normal distribution and there is noisy information transmission. The diversity of information and the size of the action sets generate a solution which is partially separating. It is not possible for certain types to screen themselves fully, because the noise prevents this. Instead there is smooth convergence to a position where the government's type is known.

Another possible application is to the field of entry deterrence. There are many results on reputation effects in this context, but these show that entry is always deterred if the time horizon is sufficiently long. Let aP_t, bP_t represent the output of incumbent and entrant respectively, then some natural restrictions on the parameters would be $\alpha_1>0,\alpha_2<0,\beta_2<0$ this gives $\omega>0,\phi>0$. If z describes market conditions which are only known to the incumbent, then the results above tell us that the entrants gradually get to learn the incumbent's private information. The incumbent's output starts at a high level and gradually declines, whilst the entrants' output starts low and gradually increases the more it discovers about the market.

7. Conclusion.

We have established that reputation effects can be found in linear quadratic gaussian (LGQ) models. Many of the results which are found to hold in the more familiar game theoretic environment will generalize to such a framework. We show that the variety of the types together with the noise prevents a large jump in agents' strategies from complete ignorance to complete knowledge. Moreover, the presence of noise considerably smooths the nature of the optimal strategy of the long lived agent, because it is now harder for agents to completely signal their types. The presence of noise also allows agents to create temporary reputations in infinite horizon games. They are eventually forced to abandon these reputations as their actions ultimately reveal their type to the short term players. This effect is due to the action sets which are a continuum. The continuum of actions does not undermine the lower bound on payoffs at Nash equilibria of repeated games derived by Fudenberg & Levine. In fact as discount factors tend to unity so does the temporary reputation become permanent. To sum up: the strong predictions obtained from earlier models are tempered into gradual adjustment processes where reputation is only ever a temporary event. The general form of payoff functions allows us to address dynamic models with asymmetric information and to investigate the reputation effects in these models.

APPENDIX

PROOF OF PROPOSITION 1. We begin by forming a linear combination of the available data $(a_o, a_1, .. a_{t-1})$ which minimizes the mean squared error, this by the properties of the normal distribution will be the true value for $E_{Bt}[\, z \mid k_t \,]$. Let $y_s = a_s/C_s$, then if $a_s = zC_s + n_s$, we have;

$$y_s = z + \psi_s, \qquad \text{where } \psi_s = n_s/C_s \text{, and } \psi_s \text{ is } N(0, \sigma_{\psi s}^2), \text{ and } \sigma_{\psi s}^2 = \sigma_n^2/C_s^2.$$

To calculate $E_{Bt}[\, z \mid k_t \,]$ we must take the linear form $\zeta_t\mu + \sum_{i=0}^{t-1} \theta_{it} y_i$, and find weights ζ, θ_{it} such that they minimize the expectation $E[(z - \zeta_t\mu - \sum_{i=0}^{t-1} \theta_{it} y_i)^2]$, these will generate the correct value for the expectation, that is $E[\, z \mid k_t \,] = \zeta_t\mu + \sum_{i=0}^{t-1} \theta_{it} y_i$.

$$E[(z - \sum_{i=0}^{t-1} \theta_{it} y_i)^2] = E[\, z(1 - \sum_{i=0}^{t-1} \theta_{it}) - \zeta_t\mu + \sum_{i=0}^{t-1} \theta_{it}\psi_i \,]^2$$

$$= (\sigma_z^2 + \mu^2)(1 - \sum_{i=0}^{t-1} \theta_{it})^2 + \sum_{i=0}^{t-1} \theta_{it}^2 (\sigma_{\psi i})^2 + (\zeta_t\mu)^2 - 2\zeta_t\mu^2 (1 - \sum_{i=0}^{t-1} \theta_{it})$$

Differentiating with respect to ζ_t, θ_{it} gives the following first order conditions for a minimum;

$$0 = (\sigma_z^2 + \mu^2)(1 - \sum_{i=0}^{t-1} \theta_{it}) - \theta_{it}\sigma_n^2/C_i^2 - \zeta_t\mu^2$$

$$0 = \zeta_t\mu^2 - \mu^2(1 - \sum_{i=0}^{t-1} \theta_{it}) \qquad \text{note } (\sigma_{\psi i})^2 = \sigma_n^2/C_i^2.$$

Substitute the from the second equation inot the first for ζ_t and rearrange indicates that $\theta_{it} = u_t C_i^2$. Substitution allows us to solve for u_t and hence for θ_{it}. Now re-write y_i in its original form and we have;

$$E[\, z \mid k_t] = \left[\frac{1}{r + \sum_{i=0}^{t-1} C_i^2}\right] \sum_{i=0}^{t-1} C_i a_i + \left[\frac{r\mu}{r + \sum_{i=0}^{t-1} C_i^2}\right]; \qquad r = \frac{\sigma_n^2}{\sigma_z^2}.$$

This completes the proof.

REFERENCES

Backus, D.K. and E.J. Driffill(1985) "Inflation and Reputation", American Economic Review 75, 530-538.

Backus, D.K. and E.J. Driffill (1985) "Rational expectations and policy credibility following a change in regime", Review of Economic Studies LII, 211-221.

Benabou R.& Laroque,G.(1988) "Using privileged information to manipulate markets: insiders, gurus, and credibility" Program for the Study of Economic Organization and Public Policy, Princeton University, Discussion Paper 19.

Cripps,M.W. (1988) "Learning rational expectations in a policy game" Warwick Economic Research Paper No. 297, 1988, University of Warwick, UK.

Cukierman, A. and A.H. Meltzer (1986) "A theory of ambiguity, credibility and inflation under discretion and asymmetric information" Econometrica 54, 1129-1160.

Dixit,A. (1980)"The role of investment in entry deterence" Economic Journal 90, 95-106.

Fudenberg,D & D.Levine (1988) "Reputation, unobserved strategies, and active supermartingales" mimeo.

Fudenberg,D & D.Levine (1989) "Reputation and equilibrium selection in games with a patient player" Econometrica 57, 759-778.

Hart,S.(1985) "Nonzero-sum two-person repeated games with incomplete information" Mathematics of Operations Research 10, 117-153.

Kreps, D. and R.Wilson (1982) "Reputation and incomplete information" Journal of Economic Theory, 27, 253-279.

Matthews,S.A. & Mirman,L.J.(1983) "Equilibrium limit pricing: The effects of private information and stochastic demand" Econometrica 51, 981-996.

Milgrom,P. and J.Roberts (1982) "Limit pricing and entry under incomplete information: An equilibrium analysis" Econometrica 50, 443-460.

Milgrom,P. and J.Roberts (1982) "Predation, reputation, and entry deterence" Journal of Economic Theory 27, 280-312.

Sargent, T.J.,Macroeconomic Theory, New York, 1979, Academic Press.

Spence,M."Entry, investment and oligopolistic pricing" Bell Journal of Economics 8, 1977,534-44.

Vickers, J. "Signalling in a model of monetary policy with incomplete information", Oxford Economic Papers, 1986.

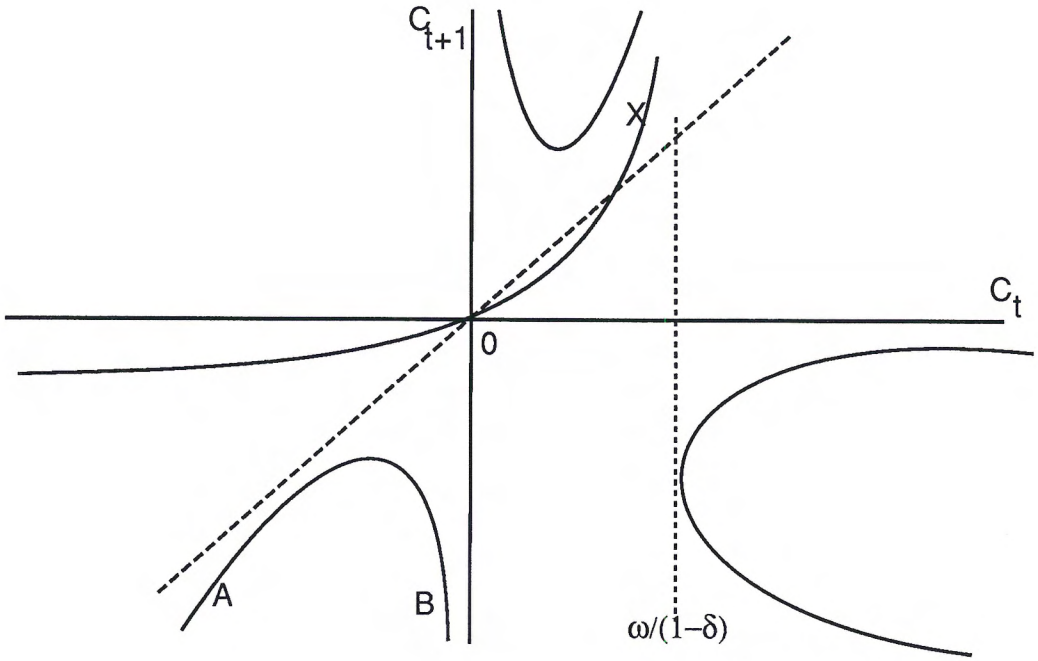

Figure 1

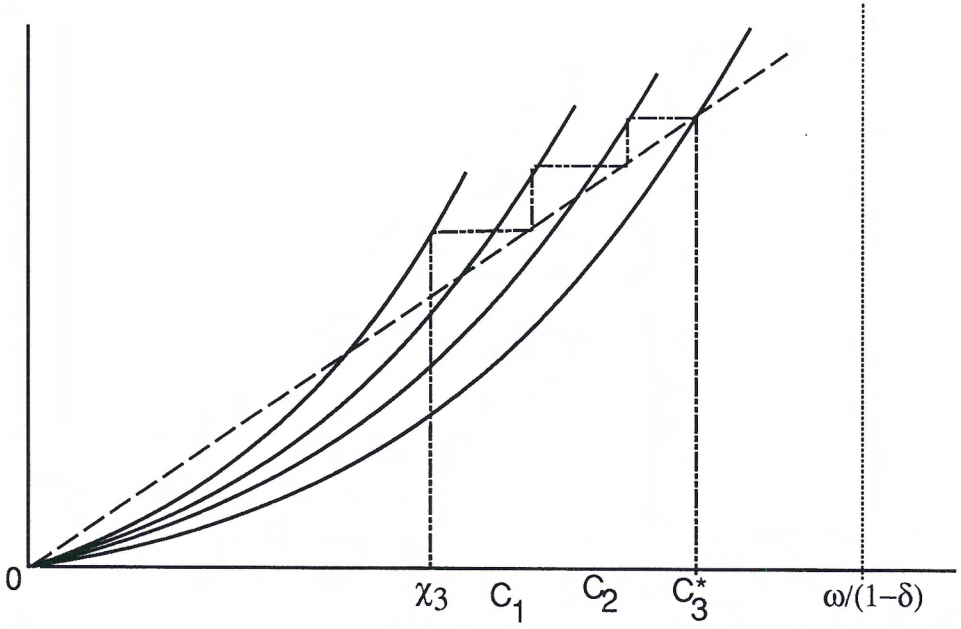

Figure 2

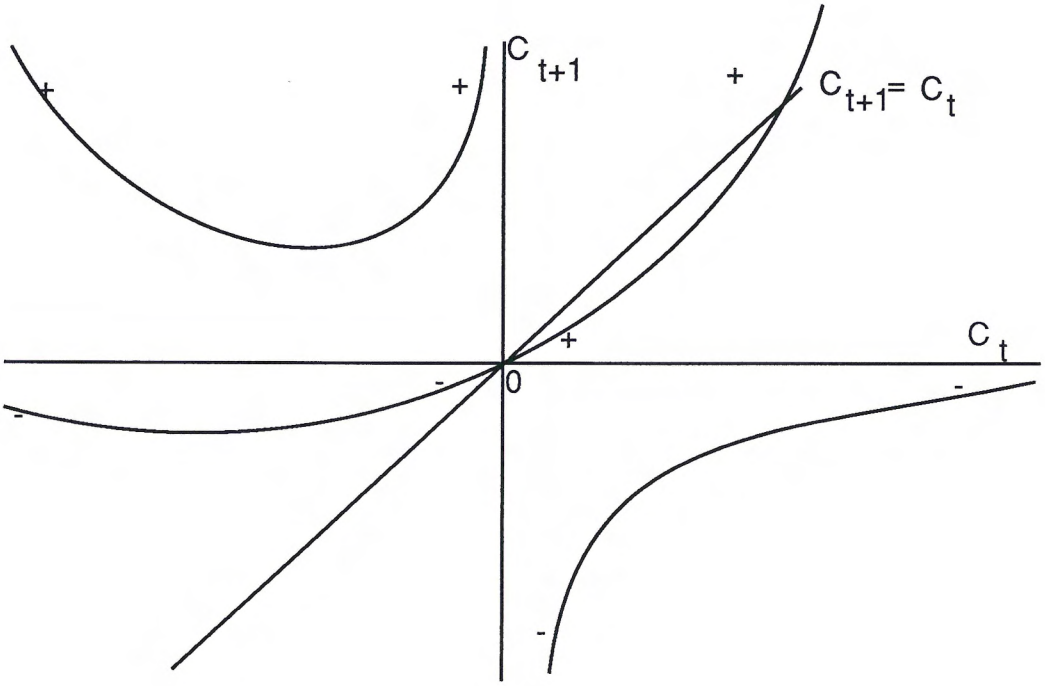

Figure 3