Learning and Hysteresis in a Dynamic Coordination Game

Rungcharoenkitkul, Phurichai

University of Oxford

1 January 2006

Learning and Hysteresis in a Dynamic Coordination Game

Phurichai Rungcharoenkitkul*

January 01, 2006
This version: August 17, 2011

Abstract

This paper introduces a dynamic coordination game with incomplete information defined by a state variable that evolves stochastically. Incomplete information enables us to use iterated dominance argument in order to resolve the indeterminacy issues. The key endogenous variable is the belief that each agent holds about the state of the world. We show that as agents update their heterogeneous beliefs through learning sequentially, they adjust their beliefs to justify the status quo. This effect induces equilibrium actions that support the status quo, a property we call hysteresis.

Keywords: dynamic coordination game, global games, hysteresis

JEL Classification: D83, C73

---

*Asia Pacific Department, International Monetary Fund. Please send comments to prungcharoenkitkul@imf.org and (after March 2012) phurichr@bot.or.th. This paper is based on a chapter of my D.Phil thesis. I am grateful to Kevin Roberts, Godfrey Keller, Nicolas Melissas and Tom Norman for their comments. All errors are mine.

1 Introduction

The challenge of explaining fluctuations or diversity using the notion of coordination failure is in striking a good balance between an intuitively appealing story, and a rigorous theory with clear predictions. Models with multiple equilibria are heavily biased towards the former objective, prompting several game theorists in recent years to propose various selection mechanisms. Unfortunately, by removing the multiplicity, some of these models have lost the very engine to explain the phenomenon that they set out to address.

This paper attempts to resolve this dilemma by considering an alternative mechanism for fluctuations, drawing on other properties of a coordination game aside from multiplicity. We take a game theoretic viewpoint that multiplicity is a problem that must be dealt with. In a typical coordination game, there may exist multiple Nash equilibria which are robust to small changes in the payoffs. In other words, there exists an interval of payoff parameter, such that more than one equilibrium may be supported. To resolve such indeterminacy, one may employ an equilibrium selection mechanism. This paper adopts the global game framework, and presents a coordination game with a perturbed information set where an equilibrium may be uniquely selected.

Because the game we study is dynamic in nature, the application of global game technique is time-dependent. In particular, agents are able to observe the outcome of the game in each period in the past and learn over time about the payoff structure of the game. Global game analysis helps select an equilibrium from the initially indeterminate set, but because the underlying belief is time-dependent, the resulting equilibrium also exhibits a hysteresis property, namely the past matters. Hysteresis is in fact merely a refined offspring of the multiplicity; the system may be locked in an equilibrium, but the past helps uniquely identify that equilibrium.

Hysteresis is an important property of any fluctuations phenomenon, with examples ranging from heat reaction of magnetisation in physics to business cycles in economics. Over the business cycles, the booms and recessions are defined by NBER by the measure of persistence; the economy needs to be doing well for a successive number of quarters to qualify as an expansion. Stock market fluctuations are often described in terms of bull and bear markets, suggesting that persistence or hysteresis is the underlying property. Explaining hysteresis is an important, if not obligatory, step towards explaining fluctuations.

Technically this paper introduces one way of taking global game analysis to a dynamic level. We choose to focus on the interaction between the dynamics and the posteriors which are central to global games, thus the payoff itself will be kept as simple as possible. The selection procedure is based on the original insight of Carlsson and van Damme (1993), but there will be some reservations regarding the effectiveness of this technique in a dynamic model. These issues aside, the dynamic pattern of selected equilibrium is a substitute for multiplicity in explaining the fluctuations.

There are two especially related works to ours. Abreu and Brunnermeier (2003) consider a model of bubble formation, where traders choose when to exit the market given that they are randomly and privately informed of the bubble existence. Their modelling technique of letting traders be sequentially informed of the bubble existence inspires our choice of correlation structure between agents' cost functions. However their focus is on the tension between the incentives to sell assets early to secure lower but safer profits, versus waiting to yield higher returns at the risk of the bubble bursting. Therefore, in terms of payoff functions, their model is closer in spirit to a war of attrition than a coordination games as considered here, and hence their persistence result is generated by a distinct mechanism. As a result, there is no role for global game analysis in their model, and the issue of uniqueness is bypassed by restricting attention to a particular class of equilibria[1].

Chamley (1999) on the other hand does use global game as a tool in his analysis of regime switching, and is closer in spirit to our work. In his model there is an unobserved exogenous stochastic variable defining a coordination game, and agents are forming posteriors about that moving target. Initially, beliefs are becoming more diffuse over time, as there is virtually no learning about the moving target, except that prescribed by the exogenous process. In this sense, agents in Chamley (1999) lose rather than gain information over time, and there is significant learning only late on, notably when there is an information bang at a regime switch. In our model, the opposite is true; agents learn and gain more information slowly. One implication is that the iterated dominance argument applies relatively trivially in Chamley (1999), as more information diffusion always ensures uniqueness. Also in our model, the stochastic variable is endogenously determined by agents' actions, a feature reminiscent of stochastic games (Shapley (1953)). In particular, our model has a built-in notion of stability, in the sense that a coordinated action that leads to a particular regime today, tends to shift fundamentals in favour of a regime switch in the future. Our objective is to ask if learning can generate a hysteresis effect that is strong enough to override such stability.

The paper is organised as follows. A dynamic coordination game is presented, which still retains the multiple equilibria as in traditional models. The model is then subjected to perturbation introduced to agents' information sets as an equilibrium selection device. Standard static global game exercise and sequential learning are then considered in comparison. A solved example is presented as an illustration, and the last section concludes.

[1] Incidentally they focus on the class of switching equilibria, which are guaranteed to be unique in global games as shown in Carlsson and van Damme (1993).

2 The Model

2.1 General structure and state variable dynamics

Time t is discrete. There is a large population of size normalised to n, and we denote the population set by $[z, z+n]$. There is a single consumption good, the stock of which is measured in discrete units $y_t \in Y = \{0, 1, ..., N\}$. At the beginning of period t, agent $i \in [z, z+n]$ decides whether to contribute to the production of y_t, by choosing $x_{it} \in \{0, 1\}$ denoting inactivity and investment respectively. The evolution of y_t is determined by joint investment efforts of all agents, $\int_{i\in[z,z+n]} x_{it} di$, and is modelled as a discrete-time birth-and-death process on Y.[2] Specifically, the transition rates are given by

$$\Pr\left(y_{t+1} = y \middle| y_t, \int_{i\in[z,z+n]} x_{it} di\right) = \begin{cases} b_t & \text{for } y = y_t + 1, \\ \delta & \text{for } y = y_t, \\ 1 - b_t - \delta & \text{for } y = y_t - 1, \end{cases}$$

where $\delta > 0$ is a constant, and b_t is an increasing step function

$$b_t = \begin{cases} b_h & \text{for } \int_{i\in[z,z+n]} x_{it} di > n - \kappa, \\ b_l & \text{for } \int_{i\in[z,z+n]} x_{it} di < n - \kappa, \end{cases} \tag{2.1}$$

where $b_h > b_l$.[3] Any tie-breaking rule for the case when $\int_{i\in[z,z+n]} x_{it} di = n - \kappa$ can be chosen without affecting the key results, and is omitted here for simplicity. Assume that $b_h > \frac{1-\delta}{2}$ and $b_l < \frac{1-\delta}{2}$, so that if the birth rate b_t equals b_h then the birth rate is higher than the death rate (and conversely for $b_t = b_l$). Thus, aggregate investment in period t determines the likelihood of an extra unit of y being produced or destroyed (by depreciation) in period $t+1$. If more than $n - \kappa$ agents are investing, then y_t is more likely to rise than to fall in the next period. Accordingly, it is natural to think of b_h and b_l as representing the boom (high) and recession (low) phases or regimes respectively.

2.2 Objective function

The good y_t is a public good providing equal utility $u(y_t)$ to all agents for free in each period t. On the other hand, contributing to the production of y_t is costly. The

[2] A birth-and-death process is a Markov process where the state y_t transits to $y_t + 1$ and $y_t - 1$ with some transition probabilities called birth and death probabilities respectively. Transitions to other states (within one period) occur with probability zero.

[3] Exogenous reflection applies at the boundaries, i.e.

$$\begin{aligned} \Pr\left(y_{t+1} = N - 1 \middle| y_t = N, \int_{i\in I} x_{it} di\right) &= 1 - \delta \\ \Pr\left(y_{t+1} = 1 \middle| y_t = 0, \int_{i\in I} x_{it} di\right) &= 1 - \delta. \end{aligned}$$

contribution or investment costs differ across agents, and agent $i \in [z, z+n]$, defined explicitly as the agent with cost function $c_i(y_t)$, incurs an investment cost $c_i(y_t)$, if she decides to invest in period t. However each agent who invests at time t will earn an extra private lump-sum gain f at time $t+1$ if $y_{t+1} = y_t + 1$ (i.e. if the production at time t is successful and yields an extra unit of goods next period). Therefore in this set-up, choosing to invest amounts to buying a lottery that pays if and when y_t goes up next period. The extra return f provides private incentive to the production of y_t.

The objective function of agent i at time t is then given by the lifetime discounted payoff

$$U_{it} = E_t \sum_{s=t}^{\infty} \beta^{s-t} \left[u(y_s) - x_{is} c_i(y_s) + f_{is} \right]$$

where β is the discount factor and

$$f_{it} = \begin{cases} f & \text{if } y_t - y_{t-1} = 1, \text{ and } x_{it-1} = 1 \\ 0 & \text{otherwise.} \end{cases}$$

The cost functions $c_i(y_t)$ are assumed to take the following linear form

$$c_i(y_t) = \frac{\beta f (b_h - b_l)}{\theta} (y_t - i) + \beta f b_l, \qquad i \in [z, z+n] \tag{2.2}$$

where $\theta > 0$. By defining the cost as an explicit function of i, the identity index i represents the cost ranking of agents ($i = z$ being the highest-cost agent). The parameter z determines the average investment cost of the population, and is therefore a measure of fundamentals. Because the cost functions only depend on the difference $y_t - i$, they satisfy the quasilinearity property, i.e. any two functions are a horizontal parallel displacement of each other. The fact that $i \in [z, z+n]$ also implies a specific correlation structure between the agents' costs. Otherwise, there is no loss of generality in our parametric choice, as free parameters θ and z capture the slope and intercept respectively. The following assumption is needed for what follows, however.

Assumption 1. *(Dominance regions)* $\kappa > \theta + 1$ *and* $n - \kappa > \theta + 1$.

The assumption means that the size of a critical mass needed for a regime switch is large relative to the size of a 'coordination window', θ. It is needed for establishing a dominance region in lemma 1 below, which is an important prerequisite for the iterative dominance argument used subsequently. The assumption will also help simplify the form of beliefs held by agents (see section 7).

Figure 1 plots the cost functions $c_i(y)$ against $\mathbb{R} \supset Y$.

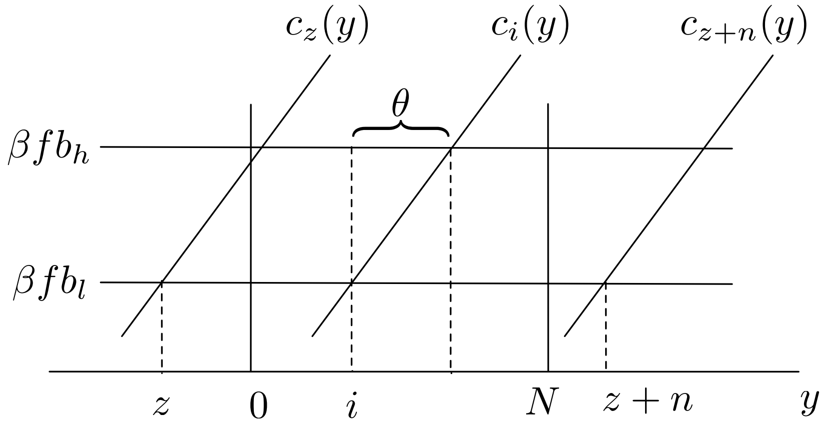

Figure 1: Cost functions $c_i(y)$

2.3 Information structures

The information set of agent i at time t is given by

$$\mathcal{I}_t^i = \left\{c_i, \{y_s\}_{s\leq t}, \{b_s\}_{s<t}\right\}. \tag{2.3}$$

That is, agents are informed of their cost functions, the current state, and past history of state and regime. The observability of the past regimes is essentially a shorthand for two joint assumptions. Firstly, agents are assumed to understand the dynamic structure of the economy, i.e. they know the economy's transition probabilities conditional on aggregate investment. Secondly, agents are able to observe the aggregate investment only in the past periods, and only imperfectly. Specifically, they can only discern whether aggregate investment was greater or lower than the threshold $n-\kappa$ in the preceding periods.

The information set specified in 2.3 is the case of primary interest, but in leading up to it, it would be useful to consider the truncated version

$$\mathcal{I}_t^{\prime i} = \{c_i, y_t\}, \tag{2.4}$$

analysed in section 6. In this case, the agents are simply assumed to be memoryless, so that any learning is forbidden.

The parameter z is assumed to be fixed, but unobservable. The correlation structure is, however, common knowledge. As apparent in Figure 1, the correlation structure is such that any set $\{c_i(y)\}_{i\in[z,z+n]}$ is uniquely determined by the mean cost parameter z. Therefore the information about the payoff structure of all agents amounts to the information about z. If z were observable, then any agent's information about her opponents' costs is complete. When z is not observed, all agents hold the same prior that z follows a distribution $H(z)$ (with density $h(z)$) on $[\underline{z}, \overline{z}]$. Let $|z| \equiv \overline{z} - \underline{z}$.

Assumption 2. *(Diffuse prior: support)* $\underline{z} \leq -\theta - n$, *and* $\overline{z} \geq N$ *(implying* $|z| \geq$

$N + n + \theta$).

Assumption 3. *(Diffuse prior: distribution) The density $h(z)$ is atomless,*

$$\min_{z \in [\underline{z}, \overline{z}]} h(z) \geq \frac{\alpha}{|z|}, \qquad 0 < \alpha < 1, \tag{2.5}$$

and is uniformly bounded,

$$\max_{z, z' \in [\underline{z}, \overline{z}]} |h(z) - h(z')| \leq \epsilon_h, \qquad \epsilon_h > 0. \tag{2.6}$$

The two conditions in assumption 3 are independent for $\alpha < 1$ (if $\alpha = 1$, condition 2.6 is redundant). In the incomplete information case, each agent uses the prior $H(z)$ and the information set $\mathcal{I}_t^i$ to compute an estimate for z, employing Bayes' rule whenever possible.

3 Interpretation and Applications

Let us discuss the idea underlying the model. An important feature of our model is that the investment decision is cast as a coordination game. The fact that we do not incorporate the investment decision directly into the utility function $u(y)$ stems from our intention to construct a coordination game version of an optimal growth problem. In the standard optimal growth, Robinson Crusoe decides whether to plant coconuts (costly investment), and he derives utility solely from the stock of coconut trees at his disposal (y_t). In our model, planting coconuts require joint efforts among many people, and the key private incentive lies in the short-run gain in coordinating successfully. There are several applications for this structure, and we outline two.

3.1 Business cycle

Firms decide whether to invest in new plants, which are only worthwhile if the aggregate demand remains strong. Aggregate demand depends on the aggregate level of investment by all firms. Firms derive a fixed amount of profit in every period that they remain invested and the demand stays high. As the number of plants rises, the land price rises and product prices fall, and it becomes more costly to invest. A firm may remain invested despite the cost rise, if it believes that other firms are investing and thereby sustaining the aggregate demand. Firms may or may not be able to observe the aggregate demand, before it commits to an investment decision in each period. If each firm's investment decision is common knowledge, there may exist multiple equilibria in investment strategies.

3.2 Asset price bubbles

A group of speculative investors with varying (privately observed) liquidity constraints choose in the morning of each day whether to buy or sell some units of a given stock. The price of the stock only move when there is enough buying or selling pressure from the investors. Each investor's payoff depends on the capital gain at the end of the day, so that it pays to buy additional units of the stock if the price is expected to rise. When the price is too low, many investors with excess liquidity may buy the stock regardless of their expectations about the price movement, whilst for a very high price, the cost of acquiring extra units may be too high for any liquidity position. When the price is in an intermediate range, investors' beliefs about the decisions of others are crucial in determining their optimal strategies. It is possible that both buying and selling can be supported as an equilibrium.

The indeterminacy, roughly capturing Keynes' concept of 'animal spirits', is often put forward as causing asset price volatility. In the context of our model however, the volatility is explained as a bubbles phenomenon. Large swings in asset prices are caused by rational investors trying to guess others' liquidity positions by relying on past stock prices. The stock price is inflated during a bubble because investors infer from observing a high stock price that others have high liquidity.

4 Best Response Functions

In this section, we first show that the lifetime discounted payoff defined by the stock of goods y_t is subject to strategic complementarity, and hence our model is essentially a series of coordination games with learning over time. Multiple equilibria exist in the complete information case, leading to indeterminacy in the dynamics of y_t in some range of initial conditions. We then discuss the implications of incomplete information in this setting. We investigate in particular the case where the only information available is the private signal about own payoffs, without record of history. This particular case is susceptible to the usual global games analysis (i.e. without sequential learning), and we carry out the equilibrium selection procedure to highlight the role of incomplete information abstracting from learning. Indeed, the case serves as a natural initial condition for a full model of learning.

4.1 Best response and strategic complementarity

Agent i uses all information available at time t to form a belief about the value of b_t, expressed as a perceived probability over $\{b_l, b_h\}$. In view of this, let us define

$$\pi_t^i = \Pr\left(b_t = b_h | \mathcal{I}_t^i\right) \tag{4.1}$$

and

$$E_t^i(b_t) = \pi_t^i b_h + (1 - \pi_t^i) b_l.$$

The value of being agent i at time t, $V(\mathcal{I}_t^i)$, is given by

$$V(\mathcal{I}_t^i) = \max_{x \in \{0,1\}} \left\{ u(y_t) - xc_i(y_t) + \beta \left[E_t^i(b_t) xf + \tilde{V} \right] \right\}. \tag{4.2}$$

To understand equation 4.2, note that $u(y_t) - xc_i(y_t)$ is simply the period-t payoff conditional on investment decision x. It is convenient to use $x \in \{0,1\}$ as an indicator function, so that the private return in period $t+1$, conditional on $y_{t+1} = y_t + 1$ and x, is simply given by xf. Thus, the expected private return in period $t+1$ as of time t conditional on x is given by $E_t^i(b_t) xf$. $\tilde{V}$ is a summation of terms involving $\pi_t^i, b_h, b_l, \delta$ and the corresponding $V(\mathcal{I}_{t+1}^i)$. When optimising, each agent takes $\tilde{V}$ as a constant, as no single agent can unilaterally affect the evolution of the system.

It follows from the functional representation 4.2 then that investment ($x = 1$) is optimal for agent i at time t if

$$\beta f E_t^i(b_t) > c_i(y_t) \tag{4.3}$$

or, in other words,

$$\begin{aligned} \pi_t^i &> \frac{c_i(y_t) - \beta f b_l}{\beta f (b_h - b_l)} \\ &= \frac{y_t - i}{\theta} \end{aligned} \tag{4.4}$$

That is, agent i will invest at time t if he believes the expected value of b_t to be sufficiently high, i.e. the probability of a high regime is high. But the expected b_t is high if and only if the strategy profile underlying the belief π_t^i assigns a sufficiently large number of agents to the investment strategy, as agents know the rule determining the regime (equation 2.1). The payoff in each period is therefore subject to strategic complementarity; the return to investment is increasing in the number of opponents who are investing.

However strategic complementarity does not affect all agents' decisions, as for any agent i, a dominant strategy may exist for sufficiently extreme values of y_t. In the optimality condition 4.3, since $E_t^i(b_t) \in [b_l, b_h]$, $x = 1$ is a dominant strategy for agent i if $c_i(y_t) < \beta f b_l$, i.e. if $y_t - i < 0$. Similarly $x = 0$ is a dominant strategy if $c_i(y_t) > \beta f b_h$, i.e. if $y_t - i > \theta$. Therefore any agent $i \in [z, z+n]$ does not have a dominant strategy if and only if $y_t \in [i, i+\theta]$. Equivalently, for any fixed y_t, the set of agents without a dominant strategy is given by $[y_t - \theta, y_t]$.

The large population assumption abstracts the individual's problem from complex intertemporal considerations leading to the simple optimality condition 4.3. The full dynamic flavour of the model, however, comes from the evolution of the belief π_t^i. This depends on the structures of agents' payoffs and the availability of information about them.

4.2 Multiplicity under complete information

Suppose that z were common knowledge, so that agents know the costs of others (and know that others know and so on), and hence are playing a coordination game of complete information. To determine the equilibrium set, it is simplest to start by eliminating dominated strategies. Let the κ^{th}-highest cost agent be denoted by

$$z^* \equiv z + \kappa.$$

Consider Figure 2. Note that for $y_t < z^*$, at least $n - \kappa$ agents have $x = 1$ as their dominant strategy, and hence according to equation 2.1 $b_t = b_h$ with probability one. Inferring that the regime is guaranteed to be high for all states $y_t < z^*$, it is optimal for all agents $i \in [y_t - \theta, y_t]$ to choose $x = 1$. On the other hand, for $y_t > z^* + \theta$, more than κ agents find $x = 0$ their dominant strategy, and $b_t = b_l$ with probability one. Over this range of states, agents $i \in [z, y_t]$ optimally choose $x = 0$, and the rest have a dominant strategy to choose $x = 1$. Thus if iterated elimination of dominated strategies can be applied, not only is a regime selected with certainty, but so are actions of all agents.

When $y_t \in [z^*, z^* + \theta]$, agents $i \in [y_t - \theta, y_t]$ do not have a dominant strategy, and the regime b_t cannot be established by the iterated dominance argument. Effectively, agents $i \in [y_t - \theta, y_t]$ are playing a coordination game, and pure-strategy Nash equilibria in which all agents $i \in [y_t - \theta, y_t]$ choose $x = 0$ and $x = 1$ (with corresponding phases b_l and b_h) can both be established. Expectations are self-fulfilling and there is a multiplicity of equilibria over this range as shown in Figure 2. The dynamics of y_t is therefore indeterminate.

The coordination problem arises whenever the agent z^* does not have a dominant strategy. It will turn out that agent z^* plays a similarly pivotal role in the incomplete information set-up. For these reasons, agent z^* will also be referred to as the decisive (or pivotal) agent.

The position of the indeterminacy range of y_t under complete information is determined by the the fundamental parameter z. However, as long as $0 \leq z^* \leq N - \theta$, the length of the indeterminacy range is invariably given by θ. More generally provided that z (and hence z^*) is not too extreme relative to n (one possible measure of the degree of heterogeneity), the indeterminacy issue arises.

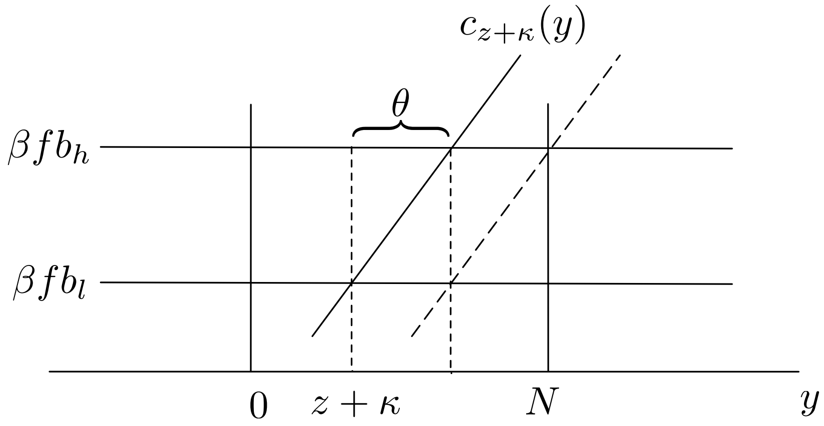

Figure 2: Multiplicity range

5 Incomplete Information and Iterated Dominance

Consider an arbitrary case of incomplete information where the belief about z that each agent i holds is given by a conditional probability distribution $\Pr(z > z' | \mathcal{I}_t^i)$. As long as $\Pr(z > z' | \mathcal{I}_t^i)$ is a well-defined function of $\mathcal{I}_t^i$, the information set $\mathcal{I}_t^i$ is arbitrary in the present section, and needs not equal that given in equations 2.3 or 2.4. An equilibrium is defined as a perfect Bayesian equilibrium, with each subgame being defined by the date t. Equivalently, the equilibrium can also be defined as a Markov perfect equilibrium, where the state space Y is extended to include the beliefs that agents hold at each date t. With this generalised state space, the game is a stochastic game and the dynamical rule governing the state is specified by the learning process.

An equilibrium in our game can be characterised much more sharply than normally possible for a general stochastic game. As mentioned earlier, the game is effectively a sequence of static coordination games, each of which is susceptible to global game selection mechanism proposed by Carlsson and van Damme (1993). This section establishes conditions under which a unique equilibrium in switching (Markovian) strategies is obtained.

Global game method relies on iterated elimination of dominated strategies in selecting an equilibrium, and requires the existence of dominance regions from which to start off iteration. In our application, this requires that, at any date t, agents without a dominant strategy $i \in [y_t - \theta, y_t]$ must believe that the decisive agent z^* has a dominant strategy with a positive probability. The next lemma establishes the condition for the existence of a dominance region.

Lemma 1. *At any date t, if every agent in the set $[y_t - \theta, y_t]$ believes that she is agent z $(z + n)$ with some probability, then she believes that the decisive agent z^* has a dominant strategy to invest (be inactive) with a positive probability.*

Proof. Let agent $i \in [y_t - \theta, y_t]$ assigns a positive probability to herself being the highest-cost type z, so that she assigns a positive probability to the decisive agent being $z^* \equiv z + \kappa = i + \kappa$. Thus, agent $y_t - \theta$ believes with a positive probability that

$$\begin{aligned} z^* &= y_t - \theta + \kappa \\ &> y_t, \end{aligned}$$

where the inequality follows from assumption 1 that $\kappa > \theta + 1$. In other words, agent $y_t - \theta$ assigns a positive probability to z^* having a dominant strategy to invest at time t. Since all agents' costs are monotone in their indices, all agents $i > y_t - \theta$ assign a positive probability to $z^* > y_t - \theta + \kappa > y_t$, so that any $i \in [y_t - \theta, y_t]$ believes agent z^* invests as a dominant strategy with a positive probability. An analogous argument follows for the other dominance region, except that $n - \kappa > \theta + 1$ from assumption 1 is used instead. $\square$

The existence of at least one dominance region is essential for the game to be solvable by iterated dominance. The following proposition is then a direct application of the iterated dominance argument used in Carlsson and van Damme (1993).

Proposition 1 (Carlsson and van Damme (1993))**.** *Let the belief that each agent i holds about z at date t be given by a distribution* $\Pr\left(z > z' | \mathcal{I}_t^i\right)$*. Then at each date t, there exists a unique Bayesian Nash equilibrium surviving iterated elimination of dominated strategies if and only if there exists a unique solution i^* to the equation*

$$\Pr\left(z^* > i^* | \mathcal{I}_t^{i^*}\right) = \frac{y_t - i^*}{\theta}. \tag{5.1}$$

Under such equilibrium, every agent $i < i^*$ chooses $x = 0$, and every $i > i^*$ chooses $x = 1$.

Proof. Only the decisions of agents $i \in [y_t - \theta, y_t]$ who do not have a dominant strategy need to be considered. For an agent $i \in [y_t - \theta, y_t]$, the probability of the regime being low, $1 - \pi_t^i$, cannot be lower than her conditional probability of agent z^* (and hence more than κ agents) having inactivity as a dominant strategy, thus

$$\begin{aligned} 1 - \pi_t^i &\geq \Pr\left(z^* < y_t - \theta | \mathcal{I}_t^i\right) \\ \pi_t^i &\leq \Pr\left(z^* > y_t - \theta | \mathcal{I}_t^i\right). \end{aligned} \tag{5.2}$$

Suppose that $\Pr\left(z^* > y_t - \theta | \mathcal{I}_t^{y_t - \theta}\right) < 1$, i.e. the agent $y_t - \theta$ at state y_t believes that he has a lower cost than the decisive agent z^* with some positive probability. It follows from inequality 5.2 that $\pi_t^{y_t - \theta} < 1$. However, for $i = y_t - \theta$, the R.H.S of equation 4.4 is 1, thus the best response rule dictates that agent $y_t - \theta$ must optimally choose to be inactive. Iterated dominance arguments can be applied for successively higher i who take as given that agents with lower signals will choose inactivity with probability one. Let i_1 be the smallest i such that inactivity cannot be established by iterated dominance arguments.

Agent $i_1 \in [y_t - \theta, y_t]$ knows that all $i < i_1$ do not invest, thus $1 - \pi_t^{i_1}$ cannot be lower than $\Pr\left(z^* < i_1 \left| \mathcal{I}_t^{i_1}\right.\right)$, implying

$$\pi_t^{i_1} \leq \Pr\left(z^* > i_1 \left| \mathcal{I}_t^{i_1}\right.\right).$$

By definition, agent i_1 must at least weakly prefer investment, thus i_1 is the smallest solution to the inequality

$$\Pr\left(z^* > i_1 \left| \mathcal{I}_t^{i_1}\right.\right) \geq \pi_t^{i_1} \geq \frac{y_t - i_1}{\theta}. \tag{5.3}$$

Similar iterated dominance arguments can be applied starting from where investment is a dominant strategy. Let i_2 be the largest i such that investment cannot be established by iterated dominance. Because the regime is known to be high for sure when $z^* > i_2$, $\pi_t^{i_2}$ is bounded from below by the (conditional) probability of this happening, i.e.

$$\pi_t^{i_2} \geq \Pr\left(z^* > i_2 \left| \mathcal{I}_t^{i_2}\right.\right),$$

combining with the definition of i_2, we have i_2 being the largest solution to the inequality

$$\Pr\left(z^* > i_2 \left| \mathcal{I}_t^{i_2}\right.\right) \leq \pi_t^{i_2} \leq \frac{y_t - i_2}{\theta}. \tag{5.4}$$

If there exists a unique solution i^* to equation 5.1, it must follow that $i^* = i_1 = i_2$, and iterated dominance leads to a unique equilibrium in period t, where each agent adopts a trigger strategy and invest only if her private signal exceeds i^*. This proves the 'if' part. If there are more than one solution to equation 5.1, then $i_1 < i_2$ and the iterated dominance argument cannot determine the equilibrium strategy played by agents $i \in (i_1, i_2)$. This establishes the 'only if' part. □

When there exists a unique solution to equation 5.1, proposition 1 states not only that there exists a unique switching equilibrium, but the equilibrium strategy of agent i is necessarily Markovian with respect to the states y_t and $\mathcal{I}_t^i$. The effect of history on the current strategy is summarised by the current state y_t and current information set $\mathcal{I}_t^i$. Most importantly, proposition 1 simplifies the equilibrium characterisation problem, suggesting that the key variable determining the equilibrium properties is $\Pr\left(z^* > i \middle| \mathcal{I}_t^i\right)$, which is readily computable given the information set $\mathcal{I}_t^i$.

The proof to proposition 1 also implies that it is not necessary for there to be two dominance regions, only that at least one exists. With only one dominance region, iterated dominance can still lead to a unique equilibrium, but in that equilibrium all agents in the set $[y_t - \theta, y_t]$ must adopt the same action as dictated by that dominance region. For instance, if only a dominance region for the decisive agent z^* investing exists, then the unique equilibrium selected by iterated dominance would prescribe all agents in the set $[y_t - \theta, y_t]$ to investing.

6 Static Game of Incomplete Information

We begin by considering the case where the dynamics of y_t is effectively suppressed from the problem, either because agents do not observe the evolution or the state is assumed fixed. In this case, the information set is limited to

$$\mathcal{I}_t'^i = \{c_i, y_t\}$$

a strict subset of the set shown in equation 2.3. One immediate consequence of this restriction is that $\Pr\left(z^* > i \middle| \mathcal{I}_t'^i\right)$ does not depend on t, since y_t does not provide any useful information about z. Accordingly, we may write

$$\Pr\left(z^* > i \middle| \mathcal{I}_t'^i\right) = \Pr\left(z^* > i \middle| i\right).$$

In this section only, we dispense with the dynamic issues, and focus on equilibrium play at each state $y \in \{0, 1, ..., N\}$ which defines a particular static game of incomplete information. This exercise serves as a natural initial condition for the full problem where no history is available, but can also be interpreted as a conventional static global game analysis as considered in Carlsson and van Damme (1993).

Because the correlation structure is common knowledge, agent $i \in [z, z+n]$ learns from her private signal that z must take values on the interval $[i-n, i]$. Combining this information with the prior about the range of z, agent i believes that $z \in [i-n, i] \cap [\underline{z}, \overline{z}]$. Let the state be fixed at y, and consider the beliefs of agents $i \in [y-\theta, y]$, i.e. those without a dominant strategy at state y. Assumption 2 implies that

$$\begin{aligned} \min_{y \in \{0,1,\ldots,N\}} \min_{i \in [y-\theta, y]} (i-n) &= -\theta - n \geq \underline{z} \\ \max_{y \in \{0,1,\ldots,N\}} \max_{i \in [y-\theta, y]} i &= N \leq \overline{z} \end{aligned}$$

so that for any state y, agents $i \in [y-\theta, y]$ believe that $z \in [i-n, i] \cap [\underline{z}, \overline{z}] = [i-n, i]$. Thus every agent updates her belief about z by truncating both sides of the publicly available prior distribution of z using private signals. The length of the posterior's support is invariant to private signals, being equal to the population size n.

Given this belief, the posterior of z conditional on any state y and private signal $i \in [y-\theta, y]$ is given by the truncated distribution on the interval $[i-n, i]$,

$$\Pr\left(z < z' \middle| i\right) = \frac{H(z') - H(i-n)}{H(i) - H(i-n)}. \tag{6.1}$$

6.1 Iterated dominance

Note that since each agents i holds an updated belief that z lies on the interval $[i-n,i]$, at any date each agent must assign a positive probability to herself being z or $z+n$, so that lemma 1 applies and iterated dominance may be employed.

Let us define

$$\begin{aligned}\mathbf{H}(i) &\equiv \Pr(z^* > i\,|i) \\ &= \frac{H(i)-H(i-\kappa)}{H(i)-H(i-n)}.\end{aligned} \tag{6.2}$$

Because of assumption 3, that $H(i)$ is atomless, $\mathbf{H}(i) \in (0,1)$. Using equation 6.2, a direct application of proposition 1 results in the following corollary.

Corollary 1. *For any fixed y, there exists a unique Bayesian Nash equilibrium surviving the iterated elimination of dominated strategies if and only if there is a unique solution i^* to the equation*

$$\mathbf{H}(i^*) = \frac{y-i^*}{\theta}. \tag{6.3}$$

Under such equilibrium, every $i < i^*$ chooses $x=0$ and every $i > i^*$ chooses $x=1$.

Existence of an equilibrium can be easily established in the spirit of Brouwer's fixed point theorem. Since over the compact interval $[y-\theta,y]$, the R.H.S of equation 6.3 is linear with extrema zero and one, and $\mathbf{H}(i)$ always lies between zero and one, at least one 'fixed point' i^* is guaranteed if $\mathbf{H}(i)$ is continuous, which is obviously the case.

Lemma 2. *There exists at least one solution i^* to equation 6.3.*

Some restrictions are required if uniqueness were to be guaranteed.

6.2 Uniqueness condition

By a contraction-like argument, it is easy to see that a sufficient condition for a unique solution i^* to equation 6.3 is that the infimum of the derivative of $\mathbf{H}(i)$ is greater than $-1/\theta$, or equivalently

$$\sup_{i\in[y-\theta,y]}\left[-\frac{\partial\mathbf{H}(i)}{\partial i}\right] < \frac{1}{\theta}. \tag{6.4}$$

Differentiating $\mathbf{H}(i)$, one gets

$$-\frac{\partial\mathbf{H}(i)}{\partial i} = \frac{h(i-\kappa)-h(i)}{H(i)-H(i-n)} + \mathbf{H}(i)\left[\frac{h(i)-h(i-n)}{H(i)-H(i-n)}\right].$$

We obtain a bound on $\sup[-\partial\mathbf{H}(i)/\partial i]$ in two steps. First, we look for the upper bound for the numerator of $-\partial\mathbf{H}(i)/\partial i$. Using the bound on $h(z)$ in assumption 3, this is given by

$$\sup_{i\in[y-\theta,y]}\left\{h\left(i-\kappa\right)-h\left(i\right)+\mathbf{H}\left(i\right)\left[h\left(i\right)-h\left(i-n\right)\right]\right\}\leq\epsilon_h$$

To see this step, note that $h(i) - h(i-n)$ can only be made large at the cost of making $h(i-\kappa) - h(i)$ small, due to the uniform bound. Given that $\mathbf{H}(i) < 1$, setting $h(i) = h(i-n)$ and $h(i-\kappa) - h(i) = \epsilon_h$ maximises the expression.

Next, we obtain the lower bound for the denominator of $-\partial \mathbf{H}(i)/\partial i$. Imagine a benchmark case of discontinuous density $h(z) = \overline{h} - \epsilon_h$ for all $z \in [i-n, i]$, and $h(z) = \overline{h}$ otherwise, where $\overline{h} = (1 + n\epsilon_h)/|z|$. It is easy to verify using assumption 3 that the lower bound for $H(i) - H(i-n)$ is $n\left(\overline{h} - \epsilon_h\right)$ if $\overline{h} - \epsilon_h > \alpha/|z|$, and $n\alpha/|z|$ otherwise (i.e. depending on which of conditions 2.5 and 2.6 in assumption 3 is binding). Thus

$$\begin{aligned}\inf_{i\in[y-\theta,y]}H\left(i\right)-H\left(i-n\right) &> \max\left\{n\left(\overline{h}-\epsilon_h\right),\frac{n\alpha}{|z|}\right\}\\ &= \max\left\{\frac{n\left(1-\epsilon_h\left(|z|-n\right)\right)}{|z|},\frac{n\alpha}{|z|}\right\}.\end{aligned}$$

It can be shown that $\overline{h} - \epsilon_h > \alpha/|z|$ if and only if $\epsilon_h < (1-\alpha)/(|z|-n)$. Hence, combining the two bounds, we have

$$-\frac{\partial\mathbf{H}\left(i\right)}{\partial i}<\begin{cases}\frac{\epsilon_h|z|}{n(1-\epsilon_h(|z|-n))} & \text{if } \epsilon_h<\frac{1-\alpha}{|z|-n}\\ \frac{\epsilon_h|z|}{n\alpha} & \text{otherwise.}\end{cases}.$$

Therefore the sufficient condition 6.4 is given by

$$\frac{1}{\theta}>\begin{cases}\frac{\epsilon_h|z|}{n(1-\epsilon_h(|z|-n))} & \text{if } \epsilon_h<\frac{1-\alpha}{|z|-n}\\ \frac{\epsilon_h|z|}{n\alpha} & \text{otherwise.}\end{cases}. \tag{6.5}$$

The parameter configurations determine whether or not the sufficient condition 6.5 is satisfied. Key roles are played by the parameters θ and ϵ_h. Clearly the condition is more likely to be satisfied for a small θ, and as θ is only bounded from below by zero, a unique equilibrium can always be ensured by a sufficiently small θ, so long as other parameters are finite. As for ϵ_h, it is also easy to see that uniqueness always obtains in the limit as $\epsilon_h \to 0$, since the R.H.S of condition 6.5 tends to zero. This limiting case corresponds to $H(z)$ being a uniform distribution. More generally, because the R.H.S is strictly increasing in ϵ_h, one can always find ϵ_h low enough that the sufficient condition holds. These key findings can be summarised as follows.

Proposition 2. *There exists a sufficiently small positive ϵ_h (and θ) for which a unique equilibrium is guaranteed for any given finite values of other parameters.*

Consider the effects of n and $|z|$. As n becomes large, so that the first line of the R.H.S of condition 6.5 applies, the R.H.S is smaller and uniqueness is more likely.

The R.H.S is increasing in $|z|$, so that smaller $|z|$ is more conducive for uniqueness. The extent to which a unique equilibrium may result from a high n and low $|z|$ is however limited by the restriction imposed by assumption 2 that $|z| \geq N + n + \theta$. Supposing that the restriction is just binding so that $|z| = N + n + \theta$, the R.H.S of the sufficient condition becomes "$\epsilon_h (N + n + \theta) / n (1 - \epsilon_h (N + \theta))$ if $\epsilon_h < (1 - \alpha) / (N + \theta)$ and $\epsilon_h (N + n + \theta) / n\alpha$ otherwise". Clearly as $n \to 0$ (and $|z| \to N + \theta$), the R.H.S explodes and the condition fails. On the other hand, as $n \to \infty$ (and hence $|z| \to \infty$), the R.H.S becomes "$\epsilon_h / (1 - \epsilon_h (N + \theta))$ if $\epsilon_h < (1 - \alpha) / (N + \theta)$, and ϵ_h / α otherwise", therefore the sufficient condition for uniqueness may or may not hold depending on other parameters. Lastly a larger α would make the R.H.S of 6.5 smaller, but α is itself bounded by 1.

The parameter θ is intuitively a measure of the number of coordination games each agent will potentially play. A smaller θ reduces the coordination problem, so that a unique equilibrium becomes more likely. A smaller ϵ_h means agents' priors become more diffuse, allowing the iterated dominance argument to select an equilibrium.

When a unique equilibrium obtains, let us denote the unique solution to equation 6.3 by $i^* (y)$.

6.3 Switching strategy and switching state

The following lemma states that when there exists a unique equilibrium in switching strategies i^*, there exists a unique switching state y^*.

Lemma 3. *Given that the condition 6.5 is met,*

$$i^* (y + 1) > i^* (y) \qquad \text{for all } y \in Y \backslash \{N\},$$

and there exist a unique switching state $y^* (z)$ such that the phase is high if and only if $y < y^* (z)$.

Proof. The second part of the statement follows since a sufficient condition for the existence of $y^* (z)$ is that $i^* (y)$ is monotone increasing in y. For example, if $i^* (y_1) > z^*$, resulting in a low phase, then the phase is low for all $y > y_1$. Suppose, for convenience, that y was a continuous variable, and implicitly differentiate 6.3 with respect to y to get

$$\left[\frac{\partial \mathbf{H} (i^* (y))}{\partial i^* (y)} + \frac{1}{\theta} \right] \frac{\partial i^* (y)}{\partial y} = \frac{1}{\theta}.$$

The condition 6.5 guaranteeing uniqueness precisely requires that $\partial \mathbf{H} (i) / \partial i$ is greater than $-1/\theta$ for all $i \in [y - \theta, y]$, implying that $\partial i^* (y) / dy > 0$ for any y. Using the same argument for successively higher continuous values of y, it follows that $i^* (y + 1) > i^* (y)$

for $y \in Y\backslash\{N\}$, and hence there indeed exists a unique switching state $y^*(z)$, provided there is a unique switching equilibrium. □

The unique switching state is characterised by

$$y^*(z) = \min\{y \in Y|\, i^*(y) > z^*\}. \tag{6.6}$$

Note that $y^*(z)$ depends on κ via two channels. A higher κ raises z^* directly by lowering the cost of the decisive agent and increases $y^*(z)$ for an exogenous reason. In addition, $i^*(y)$ is decreasing in κ as agents are willing to invest more when a high regime is easier to support. This is the endogenous channel through which a higher κ raises $y^*(z)$. See also the uniform prior example in section 8.

7 Sequential Learning

Consider the full information set

$$\mathcal{I}_t^i = \left\{c_i, \{y_s\}_{s\leq t}, \{b_s\}_{s<t}\right\}$$

as given by 2.3 where agents have memories and may use the history to update their information about the payoffs of others. Private information includes an agent's own payoff and the history of play which is a record of the states that have been visited so far and the corresponding regimes. In this set-up, agents therefore learn from private information as well as sequential observations, and the analysis for this case must take into account the dynamic evolution of beliefs.

The qualitative implications of the sequential learning case can be intuitively described as follows. Suppose that (1) state y_t is visited for the first time in a high regime, (2) no higher state has been visited, and (3) there is no regime switch at the end of period t. By observing that the regime has not switched at state y_t, agents learn that the average cost cannot be too high, and accordingly increase their posteriors that the regime will switch at some state higher than y_t. In other words, agents' beliefs about the fundamentals (average cost) become more optimistic in the absence of a regime switch over time (as y_t rises), since agents rule out the average cost being higher than a successively lower threshold. Agents then invest with probability one at the visited states, as their belief support is truncated by learning sequentially. The belief density over the remaining support can still take an arbitrary shape, however, depending on the prior as well as the sequential truncations. Therefore, any unvisited state can potentially be a regime switching state depending on the skewness of the belief density. The objective of this section is to determine the switching strategy and the corresponding switching state as a function of the model's parameters.

To fix ideas, suppose throughout this section that the economy starts in a high phase, which is determined by the unique Bayesian Nash equilibrium in monotone strategies as described in the previous section. The analysis for starting in a low regime is identical as the model is symmetric.

As before, not all agents need to be considered at any given time, as some will have behaviour dictated by dominant strategies. At time t and corresponding state y_t, the (assumed non-null) set of agents without a dominant strategy is given by

$$I_t \equiv [y_t - \theta, y_t] \cap [z, z+n] .$$

The fact that there is no phase switch in period $t-1$ implies that the decisive agent must at least not have a dominant strategy to be inactive, i.e. $z^* \geq y_{t-1} - \theta$. Since for any fixed y_{t-1}, y_t is at most $y_{t-1}+1$, we have $z^* \geq y_t - \theta - 1$, or $z+n \geq y_t + (n - \kappa - \theta - 1)$. By assumption 1, that $n - \kappa > \theta + 1$, we can conclude that $z+n \geq y_t$. Assumption 1 also implies $n > \theta$, i.e. $[y_t - \theta, y_t]$ may be a strict subset of $[z, z+n]$. Thus, I_t is given by

$$I_t = [\max\{y_t - \theta, z\}, y_t] .$$

It follows immediately that

$$\sup I_t = y_t \leq N \tag{7.1}$$

$$\inf I_t \geq z. \tag{7.2}$$

7.1 Simple beliefs

Let the first period when agents have no memories be $t=0$. The equilibrium phase at the end of this period is determined by the (assumed unique) Bayesian Nash equilibrium in monotone strategies as described in the previous section. Consider the next period $t=1$, when it is observed that $b_0 = b_h$. Because every agent follows a monotone strategy, and period 0 is known to generate a high phase, everyone can deduce that the decisive agent z^* found it optimal to invest at time 0, i.e. $z^* \geq i^*(y_0)$, or $z \geq i^*(y_0) - \kappa$. The support of z in the information set of agent i in period 1 is then truncated to

$$z \in [i-n, i] \cap [i^*(y_0) - \kappa, \overline{z}] .$$

Lemma 4. *For $i \in I_1$ where I_1 is non-null,*

$$[i-n, i] \cap [i^*(y_0) - \kappa, \overline{z}] = [i^*(y_0) - \kappa, i] .$$

Proof. Three inequalities need to be proved. Using assumption 2 that $\overline{z} \geq N$ and the fact that $\sup I_t \leq N$, one obtains the first inequality $i \leq \overline{z}$ for any $i \in I_1$. Next, since a high phase in period 0 implies $z \geq i^*(y_0) - \kappa$, the inequality $\inf I_t \geq z$ must imply $i \geq i^*(y_0) - \kappa$ for any $i \in I_1$. Lastly, note that $\min[i^*(y_0) - \kappa] = y_0 - \theta - \kappa$ whereas $\max_{y_1} \max_{i \in I_1}(i - n) = \max_{y_1}(y_1 - n) = y_0 - n + 1$. Since $n - \kappa > \theta + 1$ by assumption 1, we have the last inequality $i^*(y_0) - \kappa \geq i - n$ for any $i \in I_1$, and the result follows. □

Hence at the beginning of period 1, each agent $i \in I_1$ holds an updated belief on the support of the form

$$z \in [i^*(y_0) - \kappa, i]. \tag{7.3}$$

In other words, agents whose strategic decisions are nontrivial all learn additional information from the first observation of phase history. In particular, all agents $i \in I_1$ hold the same belief about the lower bound for z, having observed an identical binding public signal. They continue to hold diverse opinions about the upper bound however.

Shortly we will establish conditions under which the equilibrium play at the end of period 1 may be determined using iterated dominance as in the previous section. The beliefs that agents hold at the beginning of period 2, provided that there has not been a regime switch, can then be determined by combining the existing and new (if any) information generated by observing the play outcome of period 1. It is always feasible to derive the dynamics of beliefs and the corresponding equilibrium play by repeating this procedure. However note that learning from the equilibrium outcome in the first period is one-sided in the sense that it provides a publicly observed lower bound to the fundamental z, leading to updated beliefs of the form expressed in equation 7.3. Given that this one-sided property continues to hold until there is a regime switch, there is a reason to suspect that this particular form of beliefs may be held by agents more generally in the subsequent stages of the learning process. A natural way to verify the conjecture would be to check if this class of beliefs is a fixed point of the learning update mapping on the space of all classes of beliefs. A more general definition of this class of beliefs will therefore prove useful. Suppose that all agents observe a public signal at time t that $z \geq \widehat{z}_t$, i.e. $\widehat{z}_t$ is the commonly observed lower bound for z.

Definition 1. *Posterior beliefs about the value of z at time t are 'simple' if $\widehat{z}_t \in [y_t - n, \inf I_t]$, so that each agent $i \in I_t$ believes that $z \in [\widehat{z}_t, i]$.*

If $\widehat{z}_t < y_t - n$, agents $i \in [\widehat{z}_t + n, y_t]$ use $i - n$ from their own signals ($z \in [i - n, i]$) as the lower bound for z instead of $\widehat{z}_t$, whereas $\widehat{z}_t > \inf I_t$ would imply contradictions between private signals of agents $i \in [\inf I_t, \widehat{z}_t]$ and the public $\widehat{z}_t$. In other words, agents' beliefs are simple if there is a public signal on the lower bound for z that is informative to all agents who do not have a dominant strategy at the current state. For example, lemma 4 shows that beliefs are always simple in period 1, with $\widehat{z}_1 = i^*(y_0) - \kappa$.

We ask the following general question; given that agents initially hold simple beliefs, which is true in the first period by lemma 4, what is the resulting equilibrium outcome?

And given this equilibrium outcome, and supposing that the regime does not switch, what form of beliefs will agents hold after observing such outcome? By answering these questions, we establish whether the class of simple beliefs is a fixed point of the learning mapping.

Under simple beliefs, the conditional posterior of z is given by

$$\Pr\left(z < z' \middle| \mathcal{I}_t^i\right) = \frac{H(z') - H(\widehat{z}_t)}{H(i) - H(\widehat{z}_t)}, \tag{7.4}$$

for $z' \in [\widehat{z}_t, i]$.

7.2 Iterated dominance under simple beliefs

Under simple beliefs, any agent $i \in I_t$ holds a non-degenerate belief density that z lies on the interval $[\widehat{z}_t, i]$, thus each agent assigns a positive probability to herself being z, and thus at least one dominance region exists. Hence, lemma 1 again applies. For sufficiently low $\widehat{z}_t$, the other dominance region may also exist, but this is not essential for iterated dominance argument.

Define for any $i \in I_t$

$$\mathbf{H}(i, \widehat{z}_t) \equiv \frac{H(i) - H(i-\kappa)}{H(i) - H(\widehat{z}_t)}. \tag{7.5}$$

$\mathbf{H}(i, \widehat{z}_t)$ is strictly positive under the assumption that $\widehat{z}_t < z$ (so that $\widehat{z}_t < z \leq \inf I_t$), which will be shown to always hold true later (claim 1). However $\mathbf{H}(i, \widehat{z}_t)$ is monotonic increasing in $\widehat{z}_t$ and exceeds 1 when $\widehat{z}_t > i - \kappa$. Despite $\mathbf{H}(i, \widehat{z}_t)$ not always being a probability measure, it is easy to see from equation 7.4 that $\Pr\left(z^* > i \middle| \mathcal{I}_t^i\right) = \mathbf{H}(i, \widehat{z}_t)$ whenever $\mathbf{H}(i, \widehat{z}_t) \leq 1$. Thus we have

$$\Pr\left(z^* > i \middle| \mathcal{I}_t^i\right) = \min\left\{\mathbf{H}(i, \widehat{z}_t), 1\right\}.$$

Intuitively, when $\widehat{z}_t$ is high, agents may be positive that the decisive agent has a lower cost than she does. The following corollary is immediately implied by proposition 1, showing that the equilibrium outcome under simple beliefs can still be constructed using the iterated dominance argument.

Corollary 2. *Suppose that agents hold simple beliefs at the beginning of time t, and that $\widehat{z}_t < z$ (beliefs are 'rational'). Then iterated dominance leads to a unique equilibrium if and only if there is a unique solution i^* to*

$$\min\left\{\mathbf{H}(i^*, \widehat{z}_t), 1\right\} = \frac{y_t - i^*}{\theta}. \tag{7.6}$$

The unique equilibrium is a switching equilibrium defined by i^*, where every $i < i^*$ chooses

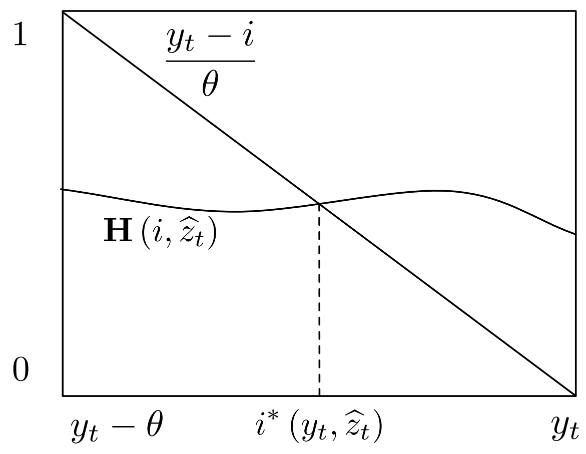

Figure 3: Equilibrium switching strategy

$x = 0$ *and every* $i > i^*$ *chooses* $x = 1$.

The proof is again omitted. It is not unimaginable that $\mathbf{H}(i, \widehat{z}_t)$ may be equal to 1 for an interval of (lower values of) i for some y_t, resulting in a switching equilibrium with $i^* = y_t - \theta$, i.e. all agents without a dominant strategy decide to invest. The possibility of a corner solution highlights the general implications of sequential learning; if agents hold sufficiently optimistic beliefs about the costs of others (i.e. high $\widehat{z}_t$) and y_t is low, then they may be absolutely certain that the regime will be high. On the other hand, given that the economy starts in a high regime, the complementary corner solution where all agents $i \in [y_t - \theta, y_t]$ choose to be inactive can never occur. This is because $\mathbf{H}(i, \widehat{z}_t)$ remains strictly positive under the assumption that $\widehat{z}_t < z \leq \inf I_t$. Existence of an equilibrium is still ensured, as $\mathbf{H}(i, \widehat{z}_t)$ is continuous. The conditions for uniqueness will be detailed shortly.

Denote the unique solution to equation 7.6, if one exists, by $i^*(y_t, \widehat{z}_t)$. Note that $\mathbf{H}(i, \widehat{z}_t)$ is increasing in $\widehat{z}_t$, implying that $i^*(y_t, \widehat{z}_t)$ is decreasing in $\widehat{z}_t$. Clearly $i^*(y_t, \widehat{z}_t)$ is increasing in y_t. See Figure 3. Note also that $\mathbf{H}(i, \widehat{z}_1) > \mathbf{H}(i)$ for all $i \in I_1$ by virtue of lemma 4. Having characterised the equilibrium play generated by simple beliefs, we now turn our attention to the next-period beliefs generated by this equilibrium play.

7.3 Beliefs dynamics

Assume for the moment that at time t the belief posteriors take a simple form. Given that there has not been a phase switch at the end of period t, agents can infer from observing a high phase that $i^*(y_t, \widehat{z}_t) < z^*$, i.e. that $z > i^*(y_t, \widehat{z}_t) - \kappa$. Since the existing beliefs are that $z \in [\widehat{z}_t, i]$, this new inference is a binding update only if $i^*(y_t, \widehat{z}_t) - \kappa > \widehat{z}_t$. Thus, the dynamics of $\widehat{z}_t$ conditional on belief being simple at time t is given by

$$\widehat{z}_{t+1} = \max\{\widehat{z}_t, i^*(y_t, \widehat{z}_t) - \kappa\}. \tag{7.7}$$

At time $t+1$, the state may have evolved, and there may have have been a new set of agents without dominant strategies. What beliefs will this new set of agents hold about z, given the new public signal $\widehat{z}_{t+1}$?

Lemma 5. *Assume that in period t equilibrium is unique, the phase is high, and beliefs are simple. Then in period $t+1$ the posterior beliefs remain simple.*

Proof. The objective is to prove that given simple beliefs at date t, the set of beliefs about the possible values of z at date $t+1$ is sufficiently bounded, so that it also takes a simple form. Specifically we desire to show that $\widehat{z}_{t+1} \in [y_{t+1} - n, \inf I_{t+1}]$. Consider each bound in turn.

Upper bound: Since $\inf I_t \geq z$ for any t, $\widehat{z}_t < z$ implies $\widehat{z}_t < \inf I_t$. Hence it is sufficient to prove that $\widehat{z}_t < z$ implies $\widehat{z}_{t+1} < z$. Suppose that $\widehat{z}_t < z$. Conditional on the belief being simple and the phase being high at time t, it follows that $i^* (y_t, \widehat{z}_t) < z^*$ and that $\widehat{z}_{t+1}$ is given by equation 7.7. From equation 7.7, either $\widehat{z}_{t+1} = \widehat{z}_t < z$ (by assumption), or $\widehat{z}_{t+1} = i^* (y_t, \widehat{z}_t) - \kappa < z$ (since a high phase at time t implies $i^* (y_t, \widehat{z}_t) < z^*$). Thus $\widehat{z}_t < z$ implies $\widehat{z}_{t+1} < z$.

Lower bound: The targeted result is that $\widehat{z}_t > y_t - n$ implies $\widehat{z}_{t+1} > y_{t+1} - n$. Since $\widehat{z}_t$ is nondecreasing as a function of t, the result obtains trivially for $y_{t+1} = y_t - 1$ or y_t. Thus only the case where $y_{t+1} = y_t + 1$ needs to be considered. Next, if $\widehat{z}_t > y_t - n + 1$, then $\widehat{z}_{t+1} \geq \widehat{z}_t > y_t - n + 1 = y_{t+1} - n$, and the result holds. Hence the case that needs to be checked is where $y_t - n < \widehat{z}_t < y_t - n + 1$. Under this constraint, the result holds if and only if $\widehat{z}_{t+1} = i^* (y_t, \widehat{z}_t) - \kappa > y_t - n + 1 > \widehat{z}_t$. But $\min [i^* (y_t, \widehat{z}_t) - \kappa] = y_t - \theta - \kappa > y_t - n + 1$ by assumption 1.

Combining the recursions for the two bounds, the lemma is proved. □

Lemmas 4 and 5 together imply, by induction, that attention can be restricted to the class of simple beliefs throughout the learning process up to the period when there is a phase switch. Thus, a unique switching equilibrium in each period except the first is determined by the solution to equation 7.6, with the corresponding belief posteriors being completely characterised by the dynamics of $\widehat{z}_t$ according to equation 7.7. This property will greatly simplify our analysis.

The proof to the lemma above also provides a useful result that $\widehat{z}_t < z$ implies $\widehat{z}_{t+1} < z$. Since $\widehat{z}_1 = i^* (y_0) - \kappa < z$, we have another small result that verifies the rationality of our learning process.

Claim 1. *Under any simple beliefs, $\widehat{z}_t < z$ for all t.*

Since $\widehat{z}_t$ is increasing, $\widehat{z}_t$ edges closer to the true z over time. In section 7.5 below, we shall show how the difference $z - \widehat{z}_t$ is related to the property of the equilibrium play. In particular, it will be shown that there exists a limiting equilibrium where $|z - \widehat{z}_t| \to 0$, which implies that agents delay switching their actions to a maximum extent. However, we first deal with the equilibrium uniqueness issue in the next section.

7.4 Uniqueness

In the class of static games with incomplete information, a unique equilibrium can always be ensured by requiring the density $h(z)$ to be sufficiently flat. It is natural as a starting point to follow a similar contraction argument in this case. Recall that the method used in such argument involves constructing an appropriate benchmark discontinuous density.

Differentiating $\mathbf{H}(i, \widehat{z}_t)$ yields

$$-\frac{\partial \mathbf{H}(i, \widehat{z}_t)}{\partial i} = \frac{h(i-\kappa) - h(i) + \mathbf{H}(i, \widehat{z}_t)\, h(i)}{H(i) - H(\widehat{z}_t)}.$$

For fixed $h(i-\kappa)$ and $h(i)$, the numerator is at its maximum when $\mathbf{H}(i, \widehat{z}_t) = 1$. Thus

$$\sup\left[h(i-\kappa) - h(i) + \mathbf{H}(i, \widehat{z}_t)\, h(i)\right] \leq \sup h(i-\kappa).$$

Note that for any fixed $\inf h(z)$, $\sup h(z) = \inf h(z) + \epsilon_h$ by the uniform boundedness assumption. In other words, when $\mathbf{H}(i, \widehat{z}_t) = 1$, the numerator is maximised only if the uniform boundedness condition is binding. Consider a benchmark case obeying such restriction, where $h(z) = \overline{h} - \epsilon_h$ for $z \in [\widehat{z}_t, i] \setminus \{i-\kappa\}$, and $h(z) = \overline{h}$ otherwise. Defining

$$\widehat{i}_t \equiv i - \widehat{z}_t$$

for notational simplicity, we have $\overline{h} = \left(1 + \widehat{i}_t \epsilon_h\right) / |z|$ so that the probability density is well-defined. Supposing that only the uniform boundedness condition is binding and not the atomlessness assumption, i.e. $\inf h(z) > \alpha / |z|$, this benchmark serves as an upper bound for $-\partial \mathbf{H}(i, \widehat{z}_t) / \partial i$, since the numerator $h(i-\kappa) = \overline{h}$ is maximised subject to $\inf h(z) = \overline{h} - \epsilon_h$, and the denominator is minimised subject to the uniform boundedness binding. Thus when $\sup h(z) = \inf h(z) + \epsilon_h$ and $\inf h(z) > \alpha / |z|$, we have

$$\begin{aligned} -\frac{\partial \mathbf{H}(i, \widehat{z}_t)}{\partial i} &< \frac{\overline{h}}{\widehat{i}_t \left(\overline{h} - \epsilon_h\right)} \\ &= \frac{1 + \widehat{i}_t \epsilon_h}{\widehat{i}_t \left(1 - \epsilon_h \left(|z| - \widehat{i}_t\right)\right)}. \end{aligned}$$

On the other hand, the atomlessness condition may also be binding, in which case we consider the benchmark where $h(z) = \alpha / |z|$ for $z \in [\widehat{z}_t, i] \setminus \{i-\kappa\}$, $h(i-\kappa) = \alpha / |z| + \epsilon_h$, and $h(z) = \left(|z| - \alpha\widehat{i}_t\right) / |z| \left(|z| - \widehat{i}_t\right)$ otherwise. Under this scenario we have

$$-\frac{\partial \mathbf{H}(i, \widehat{z}_t)}{\partial i} < \frac{\alpha + \epsilon_h |z|}{\alpha \widehat{i}_t}.$$

To summarise we have

$$-\frac{\partial \mathbf{H}(i, \widehat{z}_t)}{\partial i} < \begin{cases} \frac{1+\widehat{i}_t\epsilon_h}{\widehat{i}_t\left(1-\epsilon_h\left(|z|-\widehat{i}_t\right)\right)} & \text{if } \epsilon_h < \frac{1-\alpha}{|z|-\widehat{i}_t} \\ \frac{\alpha+\epsilon_h|z|}{\alpha\widehat{i}_t} & \text{otherwise.} \end{cases}$$

As before, the contraction argument requires the upper bound to be less than $1/\theta$, to ensure a unique equilibrium.

The functional form of this upper bound differs fundamentally from the static counterpart. Most importantly, the upper bound may now evolve over time. It can easily be shown that the upper bound is strictly decreasing in $\widehat{i}_t$. Since the belief dynamics $\widehat{z}_t$ is non-decreasing, $\widehat{i}_t$ is decreasing with time, and hence the upper bound is an increasing function of time. For instance, in the extreme case of $\widehat{i}_t \to 0$ (i.e. when i and $\widehat{z}_t$ tend to z, which cannot be ruled out), the upper bound tends to infinity. The sufficient condition for uniqueness may therefore only be met up to a period, after which indeterminacy surfaces. Intuitively sequential learning allows agents to accumulate more information over time about the game that they are playing. As the posteriors become less diffuse, iterated dominance may eventually fail.

Given that the sufficient condition becomes more demanding over time, it is natural to ask if a unique equilibrium can be guaranteed for any set of parameters at all. Although the bound is monotonic increasing in ϵ_h so that the bound may be made small by lowering ϵ_h, its limit as $\epsilon_h \to 0$ is given by $1/\widehat{i}_t$, which may tend to infinity as mentioned above. With sequential learning, $-\partial \mathbf{H}(i, \widehat{z}_t)/\partial i$ is not bounded everywhere, and a uniqueness argument based on contraction of $\mathbf{H}(i, \widehat{z}_t)$ does not suffice. It will be noticed immediately however, that since the L.H.S of the equilibrium condition 7.6 is $\max\{\mathbf{H}(i, \widehat{z}_t), 1\}$ and not $\mathbf{H}(i, \widehat{z}_t)$, $-\partial \mathbf{H}(i, \widehat{z}_t)/\partial i$ only needs to be bounded for $\mathbf{H}(i, \widehat{z}_t) < 1$ for a unique equilibrium to obtain, as $\max\{\mathbf{H}(i, \widehat{z}_t), 1\}$ does not vary with i for $\mathbf{H}(i, \widehat{z}_t) \geq 1$. By definition, $\mathbf{H}(i, \widehat{z}_t) < 1$ if and only if $\widehat{i}_t > \kappa$. Thus, when $\mathbf{H}(i, \widehat{z}_t) < 1$, the limiting upper bound for $-\partial \mathbf{H}(i, \widehat{z}_t)/\partial i$ as $\epsilon_h \to 0$ cannot be greater than $1/\kappa$. Invoking assumption 1, it is clear that the limiting upper bound is strictly less than $1/\theta$, and hence a unique equilibrium is always ensured in the limit. We summarise this conclusion by the following proposition.

Proposition 3. *Under assumption 1, as $\epsilon_h \to 0$ there exists a unique equilibrium in every stage of the sequential learning until there is a regime switch.*

It should be stressed that the upper bound obtained for $-\partial \mathbf{H}(i, \widehat{z}_t)/\partial i$, which itself may not be the tightest possible, only serves as a sufficient condition for uniqueness. Proposition 3 does not rule out a unique equilibrium when $-\partial \mathbf{H}(i, \widehat{z}_t)/\partial i > 1/\theta$; for example if $\mathbf{H}(i, \widehat{z}_t) \geq 1$ for all $i \in [y_t - \theta, y_t]$, then a unique corner solution exists regardless of $\partial \mathbf{H}(i, \widehat{z}_t)/\partial i$. The strong requirement that $\epsilon_h = 0$ is not always necessary for the uniqueness. The result above merely states that regardless of other features of the prior distribution, indeterminacy can be ruled out solely by a sufficiently small ϵ_h, zero being sufficient for all circumstances, a conclusion that is analogous to proposition 2 of the static case.

7.5 Hysteresis

Under static games of incomplete information, it is found that the switching state is given by $y^*(z)$ in equation 6.6, where it is not unreasonable given moderate κ that $y^*(z)$ should lie comfortably in the interior of the interval $[z^*, z^* + \theta]$. The qualitative difference between this benchmark case and the sequential learning is not hard to see. By implication of lemma 4, we have $\mathbf{H}(i, \widehat{z}_t) > \mathbf{H}(i)$ which implies that $i^*(y_t, \widehat{z}_t) < i^*(y_t)$, so that it is possible for the regime to be high when agents observe $\widehat{z}_t$, but low when agents can only use the prior. In this sense, being able to observe a high regime in the past is more favourable for the regime to be high in the current period. Moreover, since $i^*(y_t, \widehat{z}_t)$ is decreasing in $\widehat{z}_t$ and $\widehat{z}_t$ is increasing over time, for any fixed y we also have $i^*(y, \widehat{z}_t) \leq i^*(y, \widehat{z}_{t+1})$, i.e. more learning tends to delay a regime switch. This effect is what we mean by *hysteresis*; having been in a particular regime, agents rationally hold beliefs about the fundamental favouring that regime, and accordingly coordinate on an action that tend to sustain the status quo.

Having observed that the equilibrium regime exhibits a hysteresis property, it is natural to quantify its effect. How long can a high regime survive under sequential learning? We know that a high regime cannot continue after state $y_t = z^* + \theta$ is reached, as by then the decisive agent would stop investing as a dominant strategy. As a starting point, can one rule out the extreme outcome of 'maximum hysteresis' where the high regime still remains at the corner state $y_t = z^* + \theta$ (ignoring the integer constraint for simplicity)? Assuming that a unique switching equilibrium is guaranteed, and that there has been no regime switch earlier, the regime will remain high at $y_t = z^* + \theta$ if and only if agent z^* invests, i.e. $\mathbf{H}(z^*, \widehat{z}_t) \geq 1$, so that the unique solution is a corner solution. By definition of $\mathbf{H}(z^*, \widehat{z}_t)$, this inequality is satisfied when $\widehat{z}_t \geq z$. As claim 1 forbids $\widehat{z}_t > z$, the condition is given by $\widehat{z}_t = z$. In fact when $\widehat{z}_t = z$, the unique equilibrium is a corner solution for all $y_t \leq z^* + \theta$. This follows because, in our model, once a corner solution is obtained as an equilibrium, it will remain an equilibrium until there is a regime switch. In sum, when learning is exhausted so that agents hold a lower bound belief $\widehat{z}_t$ that is identical to the true z, there could be a maximum delay in the regime switching, provided there has not been a regime switch previously.

Claim 2. *$\widehat{z}_t = z$ guarantees a high regime at state $y_t = z^* + \theta$, provided there has not been a regime switch before period t.*

The requirement that there has not been a regime switch is an important one. For instance, $\widehat{z}_t$ may be equal to z potentially causing a maximum delay at time t if it was reached, but if $\widehat{z}_{t-1}$ is far smaller than $\widehat{z}_t$, there might well be a regime switch at time $t-1$. In other words, the actual degree of hysteresis may depend on the frequency of belief updates, with infrequent updates potentially reducing the degree of hysteresis. Thus even if in general a higher $\widehat{z}_t$ results in a longer delay other things being equal, it is difficult to fully address the issue of hysteresis without considering the dynamics of the system as a whole. To assess the quantitative effect of hysteresis, one must track the evolution of beliefs. We now introduce a simple device for this purpose.

Assume that a unique equilibrium always obtains, and consider the dynamics of $i^*(y_t, \widehat{z}_t)$. We know that $i^*(y_t, \widehat{z}_t)$ is increasing in y_t and decreasing in $\widehat{z}_t$, but both y_t and $\widehat{z}_t$ tend to grow over time given a high regime leaving the net effect on $i^*(y_t, \widehat{z}_t)$ ambiguous. Intuitively, as y_t rises, agents are less willing to tolerate a high investment, but they also hold a more positive view about the fundamental from past learning. Which effect dominates? In answering this question, it is helpful to explicitly work out how y_t and $\widehat{z}_t$ coevolve. In a manner analogous to lemma 3, it is possible to reformulate the learning dynamics $\widehat{z}_t$ in terms of the dynamics of the underlying state variable y_t, i.e. $\widehat{z}_t$ and y_t are perfectly coupled. The dynamical rule 7.7 says that additional information is learned from observing the phase b_t if and only if $\widehat{z}_{t+1} > \widehat{z}_t$.

Claim 3. *A necessary condition for $\widehat{z}_{t+1} > \widehat{z}_t$ is that $y_t > y_{t-1}$.*

Proof. By lemma 4, there is always a new information from observing b_0, with $\widehat{z}_1 = i^*(y_0) - \kappa$. Consider first whether there will be a learning of extra information from observing b_1. By lemma 4, $\mathbf{H}(i, \widehat{z}_1) > \mathbf{H}(i)$, and thus $i^*(y, \widehat{z}_t) < i^*(y)$. Observing b_1 yields a new information if $i^*(y_1, \widehat{z}_1) > i^*(y_0)$, which may hold only if $y_1 > y_0$ as $i^*(y_t, \widehat{z}_t)$ is increasing in y_t.

There is a new information learned from observing b_t only if $i^*(y_t, \widehat{z}_t) > i^*(y_{t-1}, \widehat{z}_{t-1})$. According to the beliefs dynamics 7.7, $\widehat{z}_t \geq \widehat{z}_{t-1}$ (with strict inequality if there was a learning of additional information from observing b_{t-1}), so that $i^*(y_t, \widehat{z}_t) \leq i^*(y_t, \widehat{z}_{t-1})$ as $i^*(y_t, \widehat{z}_t)$ is decreasing in $\widehat{z}_t$. Since $i^*(y_t, \widehat{z}_t)$ is increasing in y_t, $i^*(y_t, \widehat{z}_t) > i^*(y_{t-1}, \widehat{z}_{t-1})$ may hold only if $y_t = y_{t-1} + 1$. $\square$

Claim 3 suggests that we can construct two increasing sequences $\widehat{z}^0 < \widehat{z}^1 < ... < \widehat{z}^j$, and $y^0 < y^1 < ... < y^j$, where

$$\begin{aligned}
\widehat{z}^0 &= \widehat{z}_1 = i^*(y_0) - \kappa \\
\widehat{z}^j &= i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) - \kappa, \text{ for } j = 1, 2, ...
\end{aligned} \tag{7.8}$$

and

$$\begin{aligned}
y^0 &= \min\left\{ y \in Y \middle| \, i^*\left(y, \widehat{z}^0\right) > i^*(y_0) \right\} \\
y^j &= \min\left\{ y \in Y \middle| \, i^*\left(y, \widehat{z}^j\right) > i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) \right\}, \text{ for } j = 1, 2, ...
\end{aligned} \tag{7.9}$$

Thus the sequence $\{y^j\}$ lists all the states at which there is a learning update, whilst $\{\widehat{z}^j\}$ records all the distinct beliefs formed over the learning process. These sequences or stages exist independently of the dynamic paths that y_t and $\widehat{z}_t$ may take in any realisation. They express beliefs as a function of states visited in the past rather than time itself. Define the switching stage j^* as

$$j^* = \min_j \left\{ j = 0, 1, ... \middle| \, i^*\left(y^j, \widehat{z}^j\right) > z^* \right\}. \tag{7.10}$$

It is clear that a phase switch takes place when state y^{j^*} is reached from below for the first time. In other words, y^{j^*} is the switching state under sequential learning. Clearly by the construction of the sequences $\{\widehat{z}^j\}$ and $\{y^j\}$, $i^*(y^j, \widehat{z}^j)$ is an increasing function of the index j.

The extent of hysteresis may now be reassessed quantitatively. In general, the size of hysteresis grows with $\widehat{z}^j$, and $\widehat{z}^j$ rises because successively higher y^j yields more information via higher $i^*(y^j, \widehat{z}^j)$. Note that from equations 7.8 we have $\widehat{z}^j$ and $i^*(y^{j-1}, \widehat{z}^{j-1})$ being linearly related with slope one for each j, where

$$\begin{aligned}
\widehat{z}^{j+1} - \widehat{z}^j &= i^*\left(y^j, \widehat{z}^j\right) - i^*\left(y^{j-1}, \widehat{z}^{j-1}\right), \qquad j = 1, 2, ... \\
\widehat{z}^1 - \widehat{z}^0 &= i^*\left(y^0, \widehat{z}^0\right) - i^*(y_0).
\end{aligned}$$

That is, $\widehat{z}^j$ and $i^*(y^{j-1}, \widehat{z}^j - 1)$ increase in each stage exactly by the same amount. In fact, it can be seen from equations 7.8 that as $i^*(y^{j-1}, \widehat{z}^{j-1})$ approaches the threshold z^*, $\widehat{z}^j$ tends to its upper bound z. Thus, as a regime switch becomes more imminent, agents' beliefs also tend to the most optimistic level possible under rational learning, which can support a high regime up to the dominant-strategy range. The degree of hysteresis can therefore potentially reach its maximum size provided that $i^*(y^j, \widehat{z}^j)$ approaches z^* sufficiently slowly.

The actual degree of hysteresis depends on how large $i^*(y^j, \widehat{z}^j)$ can jump in each stage. If $i^*(y^j, \widehat{z}^j)$ is much larger than $i^*(y^{j-1}, \widehat{z}^{j-1})$, it is possible that there is a phase

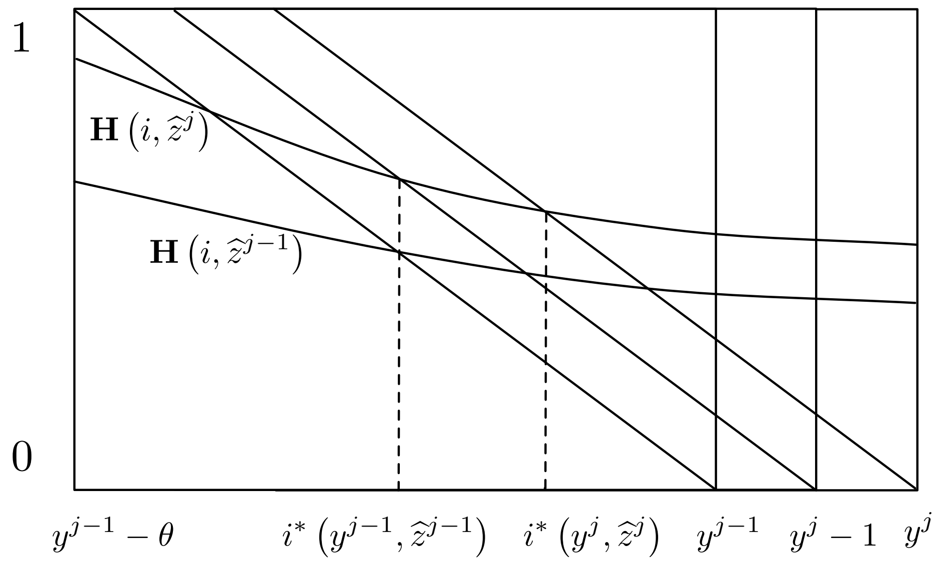

Figure 4: Maximum hysteresis

switch at stage j although $i^*(y^{j-1}, \widehat{z}^{j-1})$ is much lower than the threshold z^* (i.e. $\widehat{z}^j$ much lower than z). In such case, hysteresis is short-lived, as $i^*(y^j, \widehat{z}^j)$ jumps beyond the threshold before $\widehat{z}^{j+1}$ gets a chance to grow closer to z to sustain the high regime.

To determine how large can $i^*(y^j, \widehat{z}^j)$ jump, let us consider

$$\sup i^*\left(y^j, \widehat{z}^j\right)$$

taken over all possible realisations of sequences $\{y^j\}$ and $\{\widehat{z}^j\}$, conditional on fixed y^{j-1} and $\widehat{z}^{j-1}$ where $i^*(y^{j-1}, \widehat{z}^{j-1}) < z^*$, i.e. conditional on fixed $\widehat{z}^j < z$. Since $\widehat{z}^j$ is fixed, $\sup i^*(y^j, \widehat{z}^j)$ is reached by maximising y^j (across realisations), subject to equations 7.9, which define y^j as the minimum state such that there is a learning update. The maximum y^j under such constraint solves

$$i^*\left(y^j - 1, \widehat{z}^j\right) = i^*\left(y^{j-1}, \widehat{z}^{j-1}\right).$$

Namely, y^j reaches its maximum when the adjacent lower state $y^j - 1$ just almost results in a learning update. Figure 4 depicts the maximum y^j for given y^{j-1} and $\widehat{z}^{j-1}$.

Figure 4 shows that the maximum difference between $i^*(y^j, \widehat{z}^j)$ and $i^*(y^{j-1}, \widehat{z}^{j-1})$ depends on the slope of $\mathbf{H}(i, \widehat{z}^j)$ relative to $1/\theta$, and the size of θ relative to 1. If $\mathbf{H}(i, \widehat{z}^j)$ is steep relative to $1/\theta$, then it is more likely that $i^*(y^j, \widehat{z}^j) - i^*(y^{j-1}, \widehat{z}^{j-1})$ will be larger, as $\mathbf{H}(i, \widehat{z}^j)$ could dip and intersect $(y^j - i)/\theta$ at a higher i. However, note that the parameter scale has been implicitly defined relative to 1, and there is no *a priori* restriction on how large or small the scale must be. If the parameter scale specifies θ, κ and n to be very large relative to 1, on the diagram the functions $(y^j - i)/\theta$ will be very close to each other, whilst $\mathbf{H}\left(i,\widehat{z}^j\right)$ appears unchanged (as it is defined in terms of κ and n). With such scale, $i^*\left(y^j,\widehat{z}^j\right)-i^*\left(y^{j-1},\widehat{z}^{j-1}\right)$ will be small. In the limit of infinite scale (but finite ratios between θ, κ and n), $i^*\left(y^{j-1},\widehat{z}^{j-1}\right)\uparrow i^*\left(y^j,\widehat{z}^j\right)$, implying that $i^*\left(y^{j^*-1},\widehat{z}^{j^*-1}\right)$ is arbitrarily close to z^* leading to maximum hysteresis. Viewed in this light, the possibility of a premature switch is merely an artifact of the discreteness of the space Y. Thus the mechanism underlying our earlier observation that infrequent learning updates may result in less hysteresis is operating in our model, but can be directly controlled by the choice of parameter scale. In general, a near maximum amount of hysteresis is expected under sequential learning, with the only real hindrance being a technical one. The following proposition summarises the results regarding hysteresis.

Proposition 4. *As θ, κ and $n\to\infty$,*

$$y^{j^*}\to z^*+\theta$$

i.e. there is a maximum hysteresis in the limit where there is an infinite number of learning updates..

8 Uniform Prior Example

Some explicit results for uniform prior distribution will now be derived. The example corresponds to $\epsilon_h=0$, where a unique equilibrium is always guaranteed by proposition 3.

Under the uniform prior

$$H\left(i\right)=\frac{i-z}{|z|}$$

implying

$$\mathbf{H}\left(i\right)=\frac{\kappa}{n} \tag{8.1}$$

$$\mathbf{H}\left(i,\widehat{z}_t\right)=\frac{\kappa}{i-\widehat{z}_t}. \tag{8.2}$$

Consider first the static case without sequential learning. Using 8.1, the solution to equation 6.3 is given by

$$i^*\left(y\right)=y-\frac{\kappa}{n}\theta \tag{8.3}$$

with the corresponding switching state

$$\begin{aligned} y^*\left(z\right) &= \min\left\{y\in Y|\, i^*\left(y\right)>z^*\right\}\\ &= \min\left\{y\in Y|\, y>z^*+\frac{\kappa}{n}\theta\right\}\\ &= \min\left\{y\in Y|\, y>z+\kappa\left(1+\frac{\theta}{n}\right)\right\}. \end{aligned} \tag{8.4}$$

When κ is high, so that a high regime only requires a few agents coordinating on investing, a high regime is easier to sustain and is more likely for exogenous reasons. The exogenous effect of κ is captured by z^*, i.e. a higher κ means the decisive agent has a lower cost for any fixed z. In addition, a higher κ also induces more agents to invest, as the equilibrium cut-off agent given by 8.3 is decreasing in κ, so that the high regime is also more likely for endogenous reasons. The factor κ/n in the equilibrium solution 8.3 may be interpreted as a measure of tolerance; at any state y, there are θ agents without a dominant strategy, but a proportion κ/n of which chooses to invest by iterated dominance. The parameter κ therefore has both exogenous and endogenous implications. The total effect of κ is summarised in the equilibrium switching state 8.4 by the multiplier term $1+\theta/n$.

Next, we turn our attention to the sequential learning. An equilibrium, guaranteed to be unique by proposition 3, is given by the solution to

$$\max\left\{\frac{\kappa}{i-\widehat{z}_t},1\right\}=\frac{y_t-i}{\theta}.$$

A corner solution is obtained if $\mathbf{H}\left(y_t-\theta,\widehat{z}_t\right)\geq 1$, i.e. if $\widehat{z}_t\geq y_t-\theta-\kappa$. Otherwise the unique equilibrium is an interior solution, given by the greater solution[4] to the quadratic equation

$$\frac{\kappa}{i-\widehat{z}_t}=\frac{y_t-i}{\theta}$$

namely

$$i^*\left(y_t,\widehat{z}_t\right)=\frac{y_t+\widehat{z}_t+\sqrt{\left(y_t-\widehat{z}_t\right)^2-4\kappa\theta}}{2}. \tag{8.5}$$

Given this solution, we can attempt to construct the sequences of belief updates and states at which there is a learning update, as defined by 7.8 and 7.9. From sequence 7.8, we have

$$\begin{aligned}\widehat{z}^0 &= i^*\left(y_0\right)-\kappa\\ &= y_0-\kappa\left(1+\frac{\theta}{n}\right)\\ \widehat{z}^j &= i^*\left(y^{j-1},\widehat{z}^{j-1}\right)-\kappa, \qquad j=1,2,...\end{aligned} \tag{8.6}$$

To construct the sequence $\{y^j\}$ according to sequence 7.9, let us solve for the minimum y satisfying the following inequality

$$i^*\left(y,\widehat{z}^j\right)>i^*\left(y^{j-1},\widehat{z}^{j-1}\right).$$

[4]It can be checked that the smaller solution is always less than $y_t-\theta$, and can never be a fixed point over $[y_t-\theta,y_t]$.

Substitute the solution from equation 8.5 for the L.H.S, use $\widehat{z}^j = i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) - \kappa$, and simplify the inequality to get

$$\begin{aligned} \sqrt{\left(\kappa + y - i^*\left(y^{j-1}, \widehat{z}^{j-1}\right)\right)^2 - 4\kappa\theta} &> \kappa - y + i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) \\ y &> i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) + \theta. \end{aligned}$$

Thus we have

$$\begin{aligned} y^0 &= \min\left\{y \in Y | \, y > i^*\left(y_0\right) + \theta\right\} \\ &= \min\left\{y \in Y | \, y > y_0 + \theta\left(1 - \frac{\kappa}{n}\right)\right\} \\ y^j &= \min\left\{y \in Y | \, y > i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) + \theta\right\}, \qquad j = 1, 2, ... \end{aligned} \tag{8.7}$$

Proposition 4 obviously still applies in the present case. Following sequence 8.7, which shows how the discreteness of Y imposes restrictions on the learning, let us define $\epsilon^j < 1$ by

$$\epsilon^j \equiv y^j - i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) - \theta.$$

Combining this with sequence 8.6 and substituting in 8.5, we get

$$\begin{aligned} i^*\left(y^j, \widehat{z}^j\right) &= \frac{y^j + \widehat{z}^j + \sqrt{\left(y^j - \widehat{z}^j\right)^2 - 4\kappa\theta}}{2} \\ &= i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) + \frac{1}{2}\left[\theta - \kappa + \epsilon^j + \sqrt{\left(\theta + \kappa + \epsilon^j\right)^2 - 4\kappa\theta}\right]. \end{aligned}$$

Clearly as ϵ^j becomes negligible, i.e. as θ and κ become large, then we have $i^*\left(y^{j-1}, \widehat{z}^{j-1}\right) \uparrow i^*\left(y^j, \widehat{z}^j\right)$, and $i^*\left(y^{j^*-1}, \widehat{z}^{j^*-1}\right) \to z^*$ so that the equilibrium approaches a corner solution of maximum hysteresis.

9 Conclusion

This paper investigates the potential link between information availability and fluctuations. The model studied is essentially a dynamic coordination game with Bayesian learning, whereby the equilibrium outcome is sequentially observed over time. The global game technique is shown to be applicable, and is used to derive a sufficient condition for a unique equilibrium. Sequential learning is found to have two key implications. Firstly, at the technical level, it is shown that the global game technique when cast in a dynamic context may become less powerful if learning is allowed, as the posteriors may become less diffuse over time. Secondly, sequential learning serves as a strong propagation mechanism in the dynamic coordination game considered. In our model, learning and equilibrium play reinforce each other, causing the equilibrium to exhibit a strong hysteresis property.

The fact that the economy is in a boom sends a signal to agents that the economy is healthy, encouraging them to keep investing thereby sustaining the boom and so on.

It is instructive to contrast our approach with Caplin and Leahy (1994), who also model market crashes in an environment where agents learn about the unobserved fundamental in deciding whether to invest. They study the process by which agents learn about the fundamental by observing the actions of others who receive a stream of independent private signals about the fundamental. When one agent stops investing, she signals her private information to others. The market crashes when a sufficient number of agents stop investing, sending a strong signal to others who then follow suit. Before such a crash, the market is in a 'business as usual' phase, where most agents continue investing. There is a built-in persistence in Caplin and Leahy (1994), which is generated by assuming that investment is irreversible, so that switching action is costly (cancellation and reinvestment involves installation costs). In this regard, their model is similar to the conventional approach of generating persistence or hysteresis via the switching cost (employed, for example, in Dixit (1989)). As agents delay changing their actions due to private switching costs, they also delay transmitting their private information to the public, leading to aggregate persistence.

In contrast, agents in our model can change their actions costlessly at any period, and therefore there is no frictions in the decision making *per se*. There is a friction in the information processing however, as agents in our model only learn about an aggregate summary of private signals (i.e. the regime), and not the individual details thereof. This crucially leads to the one-sided nature of our learning, and it leads to persistence of actions via a mechanism that is different from that in Caplin and Leahy (1994). The novel feature of our model is that persistence is made possible by the presence of strategic complementarity, which involves direct strategic interactions between players, rather than individual switching costs. This feature allows for our usage of global game technique in solving for a unique equilibrium. On the other hand, the interactions between agents in Caplin and Leahy (1994) are purely informational. Indeed, there are in general many equilibria in the model of Caplin and Leahy (1994) (the authors characterise the equilibrium set in terms of lower and upper bounds on the exit time), and it is unclear how to select among these potentially many equilibria.

Furthermore, in Caplin and Leahy (1994), the market needs not crash, provided that the fundamental is sufficiently good. If agents in their model learn that the profitability is high, then they remain invested until the final period, thus the issue of delay or hysteresis does not always arise. In our model, the market necessarily crashes, and the only question is when. Our model therefore places a much greater emphasis on the strategic delay, and its interaction with the information generating process, using a dynamic coordination game as the building block. As our interest lies on the timing of the crash, it is desirable to characterise the equilibrium set as sharply as possible, and the applicability of iterated dominance as a tool is an attractive feature of our model.

References

Abreu, D. and M. K. Brunnermeier (2003): "Bubbles and crashes," *Econometrica*, 71, 173–204.

Caplin, A. and J. Leahy (1994): "Business as usual, market crashes and wisdom after the fact," *American Economic Review*, 84, 548–565.

Carlsson, H. and E. van Damme (1993): "Global games and equilibrium selection," *Econometrica*, 61, 989–1018.

Chamley, C. (1999): "Coordinating regime switches," *The Quarterly Journal of Economics*, 114, 869–905.

Dixit, A. K. (1989): "Entry and exit decisions under uncertainty," *Journal of Political Economy*, 97, 620–638.

Rungcharoenkitkul, P. (2006): "Complementarities and fluctuations," Ph.D. thesis, University of Oxford.

Shapley, L. S. (1953): "Stochastic games," *Proceedings of the National Academy of Sciences of the USA*, 39, 1095–1100.