Social Learning with Payoff Complementarities*

Amil Dasgupta†
Yale University

November 1999; Revised May 2000
Preliminary. Comments welcome.

REVISED VERSIONS WILL BE AVAILABLE AT
http://www.econ.yale.edu/~amil/research.html

Abstract

We incorporate strategic complementarities into a multi-agent sequential choice model with observable actions and private information. In this framework agents are concerned with learning from predecessors, signalling to successors, and coordinating their actions with those of others. Coordination problems have hitherto been studied using static coordination games which do not allow for learning behavior. Social learning has been examined using games of sequential action under uncertainty, but in the absence of strategic complementarities (herding models). Our model captures the strategic behavior of static coordination games, the social learning aspect of herding models, and the signalling behavior missing from both of these classes of models in one unified framework. In sequential action problems with incomplete information, agents exhibit herd behavior if later decision makers assign too little importance to their private information, choosing instead to imitate their predecessors. In our setting we demonstrate that agents may exhibit either strong herd behavior (complete imitation) or weak herd behavior (overoptimism) and characterize the informational requirements for these distinct outcomes. We also characterize the informational requirements to ensure the possibility of coordination upon a risky but socially optimal action in a game with finite but unboundedly large numbers of players.

Key words: Learning, Coordination, Herding, Cascades, Strategic Complementarities.

---

*Stephen Morris provided invaluable guidance for this project. I also gratefully acknowledge useful conversations with David Pearce, Ben Polak, Giuseppe Moscarini, Dirk Bergemann, Hyun Song Shin, Jonathan Levin, Felix Kubler, and participants at the Yale seminar on game theory. Part of this project was supported by the NSF and the Cowles Foundation. A previous version of this paper was circulated under the title: "Learning, Signalling, and Coordinating: A Rational Theory of "Irrational Exuberance.""

†Department of Economics, Yale University, P. O. Box 208268, New Haven, CT 06520-8268, USA. E-mail: `amil.dasgupta@yale.edu`, `http://www.econ.yale.edu/~amil`

1 Introduction

Observers of financial market booms and busts, both casual and experienced, will often note that behavior in the market is characterized by an excess of optimism or pessimism. There appears to be a tendency for market participants to "jump on the bandwagon." They get so carried away by the decisions of others around them that they simply imitate their predecessors, paying no attention to any information about fundamentals that they may receive, or making no effort to gather such information. It is as though they were all moving in a herd. When the market tanks, traders tend to exit the market quicker then they would have if they took into account the fundamentals of the economy. When the market booms, traders get excessively optimistic compared to levels that are justified by the underlying fundamentals. This, perhaps, is what Alan Greenspan was referring to in his now famous "irrational exuberance" speech of December 5, 1996, at the heart of the stock market boom of the 1990s. Given the pertinence of such market herd behavior in both good times and bad, there is clearly a need to analyze the probem carefully. To begin, let us try to separate the central stylized characteristics of situations such as stock market booms and panics, currency crises, or bank runs.

The salient features of such situations are as follows. A number of market participants are called upon to make similar decisions (buy/sell, long/short, withdraw/remain etc.) at about the same time. Since they are all in the same market, they can observe each other's actions. Each participant has non-trivial private information (ideas, intuition, acquired knowledge) about the fundamentals of the situation. These are, after all, educated financial traders. In order to make their decisions, partipants may use either their private information, or the public information generated by observing their predecessors' actions, or both. However, participants also have to worry about their successors, because each person's payoff depends upon the actions of everybody else. Even if a few predecessors have chosen to go short on the market, a trader may worry that his successors will not, thus preventing a market downturn and leaving him stranded. In short, there are *strategic complementarities*[1]. The observation that agents seem to herd, then, amounts to noting that later agents pay "too much" attention to the choices of their predecessors and "too little" attention to their own private information.

In this paper, we propose a game theoretic model to study such situations. There are n risk neutral agents in our model who act in an exogenous sequence and choose either to invest or not. There are two states of the world, a state that is conducive to investment, and a state that is not. Agents receive signals that are informative about the state of the world in a stochastic sense: very roughly, higher signals increase the likelihood of the state being good. Conditional on the state, the signal generating process is independent and identical across agents. At the point when they have to choose their actions, agents are able to observe the choices of their predecesors and their own private signals (but not the signals of their predecessors). Finally, there are strong strategic complementarities. Investment leads to a positive net payoff only if the state is good and all other agents also choose to invest. Otherwise, it generates negative net

[1]A term coined by Bulow, Geanakoplos, and Klemperer (1985), otherwise refered to as positive payoff externalities,network externalities, supermodularities etc. in various specific contexts.

return. Not investing costs and pays nothing, independent of the state of the world.

In this set up, we show that it is inevitable that agents shall become progressively more optimistic as more and more predecessors choose to invest (Proposition 5). This is natural and to be expected. However, it turns out that such optimism can take excessive forms, depending on the properties of the information system of the game. If the information system has the property that likelihood ratios for individual agents are bounded (i.e. agents can exhibit only limited amounts of personal skepticism based upon their available information), then agents may literally start to imitate others and ignore their own payoff relevant information (*strong herd behavior*). Indeed, under these circumstances, such "irrational exuberance" is the *only* outcome of rational behavior (Proposition 6). We are able to tightly characterize the informational requirements that would lead to such strong herd behavior for linear information systems (Propositions 7 and 8). However, we also show that if the information system is rich enough to allow agents to exhibit unbounded personal skepticism, i.e., possesses the unbounded likelihood ratio property, such extreme forms of exuberance are ruled out. The exclusion of strong herd behavior does not mean that overoptimism vanishes. In fact it is quite possible that agents do not ignore their private information but are overoptimistic in comparison to the case where information is aggregated efficiently in the market. We call such phenomena *weak herd behavior* and lay down informational conditions necessary and sufficient for weak herding to occur. We show that for the important class of Gaussian information systems weak herding occurs with positive probability.

It is apparent from the structure of the model that if players do not exhibit strong herd behavior, it shall be harder and harder to persuade the first player to invest as the number of agents gets progressively larger. In order to address these concerns, we characterize the informational requirements that shall create the possibility of coordinated investment even in games with unbounded likelihood ratios when the number of players is arbitrarily large (Propositions 9 and 10).

The study of situations where people's decisions are influenced by those of others around them is not new. Stylized versions of situations similar to ours have been extensively studied in the literature. The pioneering papers are by Banerjee and by Bikhchandani, Hershleifer, and Welch, both in 1992. Variations, generalizations, and applications have also been studied. Lee (1993) provides conditions on the action choices of agents in a generalized herding model that guarantee herding. Gul and Lundholm (1995), Chamley and Gale (1994), and Chari and Kehoe (2000) examine similar models but allow the order of action to be endogenous. Froot, Scharfstein, and Stein (1992), Chari and Kehoe (1997), Avery and Zemsky (1998), and Lee (1998), among others, apply herding models to study various financial situations. For a recent selective survey of this literature see Bikhchandani, Hershleifer, and Welch (1998). However, all these models are characterized by a common feature: *individual payoffs are unaffected by the actions of others*. The only externality present in these models is an informational one. Agents are concerned about each other's choices only to the extent that prior actions generate information about the state of the world. There are no strategic complementarities. Therefore, agents in these models exhibit only backward-looking behavior. As a result, it becomes much harder to apply these models to real financial situations.

In many settings, in addition to the informational externality, it is essential to incorporate direct payoff externalities. The situations discussed above are but a few of a plethora of possible examples. When payoff complementarities exist, agents must be concerned not only with the actions of their predecessors but also with those of their successors. Thus, in situations such as these, agents would exhibit both backward-looking (learning) and forward-looking (strategic) behavior. This strategic component complicates the arguments in the models of Banerjee, and Bikhchandani, Hershliefer, and Welch. Games with payoff complementarities that capture strategic behavior by agents have been studied in the literature with the goal of explaining situations similar to the ones above. For example, Obstfeld (1986), Cole and Kehoe (1996), and Morris and Shin (1998) model currency crises in various degrees as static coordination games under uncertainty with payoff complementarities. However, the static nature of these games excludes the learning behavior seen in sequential action models. Finally, the interaction of sequential action with strategic complementarities creates signalling behavior in our model, an effect that is missing from both herding models and static coordination games. Agents are concerned about the signals that their action choices send to their successors. We are, therefore, able to capture the learning behavior of herding models, the strategic behavior of static coordination games, and the signalling behavior absent from both of these previous classes of models in one unified framework.

Our analysis also helps to understand better the way in which information plays a role in creating strong herding in markets. We provide two versions of the model featuring qualitatively different information systems, one in which the private information of agents is rich enough to allow them to exercise unlimited personal skepticism (unbounded likelihood ratios) and one in which this is not possible. We demonstrate that that latter is necessary (but not sufficient) for strong herd behavior and characterize the precise conditions under which strong herding takes place under additional assumptions. This provides a foundation upon which to build a theory of optimal information structure in such games, paving the way for mechanism design in situations where market participants must be prevented from herding or persuaded to herd upon risky but socially productive alternatives.

In an important recent contribution, Smith and Sorensen (1999) provide similar characterizations of the informational prerequisites for herd behavior. Their model generalizes the traditional herding literature by allowing for heterogeneous preferences and makes explicit the conditions under which Bayesian learning may be incomplete as opposed to confounded. Our setting retains the identical preferences of traditional herding models, but adds in payoff complementarities with the goal of capturing the other relevant strategic aspects of a market boom or bust, thereby unifying the literature on herding with static models of coordination.

Choi (1997) builds network externalities into a model of sequential action under uncertainty. His model is one of strategic technology choice by firms. Firms choose between two competing technologies with unknown values. It is beneficial for firms to choose the technology that shall be adopted by most other firms because of network externalities. While this model is ostensibly similar to ours, it is significantly different in spirit. First, once a technology is used by a firm, its true value becomes common knowledge amongst participants in the game. Thus, after the first player has chosen a technology, the rest of the game is effectively one of complete information. Second, since firms receive no private signals about the alternative technologies, there is no private information in Choi's model. Herding happens purely due to the network effect and risk aversion. Herding, in the traditional sense, is simply the phenomenon by which followers may pregressively (suddenly or gradually) disregard their private information in favour of already available public signals. A proper analysis of herding requires a fully-specified model that explicitly distinguishes between private and public information. Our model provides such a framework.

Two other recent papers that contain elements of strategic complementarities and herding are Jeitschko and Taylor (1999) and Corsetti, Dasgupta, Morris, and Shin (2000). In the former, agents play pairwise coordination games due to random matching, but learning is not "social" since agents observe only their own private histories. In the latter, a sequential coordination game is set up to explore the influence of a large trader in a model of speculative currency attacks with private information. When the large trader is arbitrarily better informed in comparison to the rest of the market, smaller traders exhibit strong herd behavior in the sense of our model.

The rest of the paper is organized as follows. In section 2 we lay out the model. Section 3 demonstrates two important properties of the equilibria of this game. Section 4 defines strong and weak herding in our setting and provides informational requirements for their occurence. In section 5 we characterize the informational requirements to ensure the possibility of coordination in numerous-player versions of our game. Section 6 discusses our results and section 7 provides a simple illustrative application. Section 8 discusses caveats and potential extensions of the model.

2 The Model

2.1 The Structure of the Game

There are n agents who choose whether to invest (I) or not (N). We write $a_i \in A_i = \{I, N\}$ for $i = 1, 2, ..., n$, and $A = \times_{i=1}^{n} A_i$. There are two states of the world: a state G which is good for investment, and a state B which is bad for investment. Nature selects which state of the world occurs. Investing is risky. For an agent to get positive net return (of 1) from investing, it is necessary that the state is conducive to investment, i.e., G, and that all other agents also choose to invest. If even one of these conditions are violated, then investment generates negative net return of $-c$. Not investing is safe. It generates a constant return of 0 independent of the actions of other agents and the state of the world. Agents' payoffs can thus be represented by the mappings $(u_i : \{G, B\} \times A \to \mathbb{R})_{i=1}^{n}$ defined for each i by:

$$u_i(G, a_i, a_{-i}) \quad = \quad \begin{cases} 1 & \text{when } a_i = I \text{ and } a_j = I \text{ for all } j \neq i, \\ -c & \text{when } a_i = I \text{ and } a_j = N \text{ for some } j \neq i, \\ 0 & \text{when } a_i = N \end{cases}$$

$$u_i(B, a_i, a_{-i}) \quad = \quad \begin{cases} -c & \text{when } a_i = I \\ 0 & \text{when } a_i = N \end{cases}$$

Agents act sequentially, in the order $1, 2, ..., n$. Each agent observes the actions of those who have preceded her. In addition, each agent receives a private signal (her type), which summarizes her private information about the state of the world. In particular, agent i receives signal $s_i \in S = [\underline{s}, \bar{s}] \subset \mathbb{R}$ or $S_i = \mathbb{R}$ for all i.[2] Conditional on the state, the signals are *independent* and *identically distributed*. For each i, s_i is distributed according to some *continuous, state-dependent* density, $f(.)$. We require that in state G, private signals have full support.[3] $f(.)$ satisfies the following (strict) *monotone likelihood ratio property* (MLRP): $\frac{f(s|B)}{f(s|G)}$ is *strictly* decreasing in s. We shall sometimes refer to these stochastic processes as making up the *information system* for the game, and write $f = \{f(.|G), f(.|B)\}$ to denote it.

Agents share a common prior over the state of the world: $Pr(G) = 1 - Pr(B) = \pi \in [0, 1]$. They are expected utility maximizers.

For future reference, we shall denote the game we have just described by $\Gamma(n)$, where the argument refers to the number of players in the game. Unless otherwise stated, we shall assume that $n \in Z_{++}$, i.e. the number of players is finite. In what follows, we consider Weak Perfect Bayesian Equilibria of $\Gamma(n)$ which are defined below. We preface our analysis by some brief remarks about the information system.

2.2 A Note on Likelihood Ratios

As we have noted above, the signals in $\Gamma(n)$ can be generated either from some closed subinterval of $\mathbb{R}$, or from $\mathbb{R}$ itself. This distinction is made to explicitly distinguish between two versions of the game: the case with *bounded likelihood ratios* and the case with *unbounded likelihood ratios*.[4] We denote the likelihood ratio by $r(s) = \frac{f(s|B)}{f(s|G)}$. The full support assumption on $f(s|G)$ ensures that $r(s)$ is well defined on S. When $S = [\underline{s}, \bar{s}] \subset \mathbb{R}$, the MLRP property and the boundedness of probability density functions implies that there exist bounds $B \geq 0$ and $T < \infty$ such that $r(s) \in [B, T]$ for $s \in S$. $B = 0$ when $f(s|B)$ is not full support. When $S = \mathbb{R}$, the MLRP property implies that $r(s)$ is unbounded above and asymptotes to 0 below. Conversely, when $r(s)$ is unbounded above or below, the boundedness of probability density functions that $S = \mathbb{R}$.

Intuitively, the case with unbounded likelihood ratios can be thought to be the version of $\Gamma(n)$ when players exhibit unbounded personal skepticism, i.e. may observe some private information that reverses any level of optimism they may have enjoyed ex ante. The case with bounded likelihood ratios is the reverse: players are only boundedly skeptical. A certain level of ex ante

[2]For notational convenience, we shall use $\underline{s}$ and $\bar{s}$ below to denote lower and upper bounds for S even when $S = \mathbb{R}$, assuming implicitly that $\bar{s} = \infty$ and $\underline{s} = -\infty$ when this is the case.

[3]This is done to eliminate the trivial case where an agent may discover for sure that the state is B. In such a case, there is no strategic content left in the game.

[4]While the distinction is formal, i.e., represents alternative modelling strategies, it is useful in classifying the results of the game.

optimism cannot be reversed by any private information, however discouraging. The properties of $r(.)$ shall turn out to be crucial to our analysis of $\Gamma(n)$, and we shall return to this point again below.

2.3 Possible Strategy Profiles

How does an agent, say i, decide whether to invest or not? When agent i is called upon to act, she knows only what her predecessors have done and the value of her own signal. Hence, her strategies take the form of mappings from her predecessors' actions and her own private signal to her action set. Formally, for each i, $\sigma_i : (\times_{j<i} A_j) \times S \rightarrow \{I, N\}$. Given this notation, we first provide a useful definition:

Definition 1 *Player i follows a* trigger strategy *in $\Gamma(n)$ if she chooses her actions according to the map*

$$\sigma_i(s_i, (a_j)_{j<i}) = \begin{cases} I & \text{when } s_i \geq t_i \text{ and } a_j = I \ \forall j < i \\ N & \text{otherwise} \end{cases}$$

for some $t_i \in \Re$ where $\Re = \mathbb{R} \cup \{-\infty\} \cup \{\infty\}$, the augmented real line.

We call t_i player i's trigger. An equilibrium in which each player follows a trigger strategy is called a *trigger equilibrium*.

It is important to note that while players' *signals* are drawn from some subset of the real numbers $\mathbb{R}$, *triggers* are drawn from the augmented real line, $\Re$, because players may follow strategies of "always invest" (corresponding to a trigger of $-\infty$ if $S = \mathbb{R}$) or of "never invest" (corresponding to a trigger of ∞ if $S = \mathbb{R}$).

Given the payoff complementarities, it is clear that in any equilibrium if $a_j = N$ for some $j < i$, $\sigma_i((a_j)_{j<i}, s_i) = N$ for all s_i, since investing is a strictly dominated action. Thus, agent i's decision problem is interesting only in the instance that $a_j = I$ for all $j < i$. In this instance, since we know by definition $(a_j)_{j<i} = (I, ..., I)$, an agent's equilibrium strategy is formally just some function of her private signal. Thus, for notational convenience, we can now drop the explicit dependence of the strategies σ_i on the observed history of actions $(a_j)_{j<i}$. When the argument is suppressed, it is tacitly assumed that $a_j = I$ for all $j < i$.

If she observes investment by all her predecessors, then agent i has some beliefs (posterior) about the state of the world, say $\pi_i \in [0, 1]$. Her expected utility from investing depends upon this posterior belief, her private signal, and the strategies of her successors. Formally, $EU_i(\pi_i, s_i, (\sigma_j)_{j>i}) = (1)P_i + (-c)(1 - P_i)$, where

$$\begin{aligned} P_i &= Pr(G, (\sigma_j(s_j) = I)_{j>i} | s_i) \\ &= \frac{Pr(G)Pr((\sigma_j(s_j) = I)_{j>i}, s_i | G)}{Pr(G)Pr(s_i|G) + Pr(B)Pr(s_i|B)} \end{aligned}$$

$$= \frac{\pi_i Pr((\sigma_j(s_j) = I)_{j>i}|G) f(s_i|G)}{\pi_i f(s_i|G) + (1-\pi_i) f(s_i|B)}$$
$$= \frac{\pi_i Pr((\sigma_j(s_j) = I)_{j>i}|G)}{\pi_i + (1-\pi_i)\frac{f(s_i|B)}{f(s_i|G)}}$$

where the third equality follows from the conditional independence of the signals.[5] Given this notation, we define a Weak Perfect Bayesian Equilibrium for $\Gamma(n)$.

Definition 2 *A* Weak Perfect Bayesian Equilibrium *of* $\Gamma(n)$ *is a tuple of strategies* $(\sigma_1, ..., \sigma_n)$ *and a tuple of posterior beliefs* $(\pi_1, ..., \pi_n)$ *where for each* i, $\pi_i : (a_j)_{j<i} \to [0,1]$ *which satisfy the following conditions:*

1. *Given* π_i, σ_i *is a best response to* σ_{-i} *after every possible history.*
2. *If the observed history of play can happen with positive probability in the equilibrium path prescribed by* σ, *then* π_i *is derived from the original priors by Bayesian updating. If not, then* π_i *is any member of* $[0,1]$.

In the specific setting of our model, these conditions translate into the following:

1. For each i, if $a_j = N$ for any $j < i$, then $\sigma_i = N$. If $a_j = I$ for all $j < i$, then $\sigma_i = I$ if and only if
$$\frac{\pi_i Pr((\sigma_j(s_j) = I)_{j>i}|G)}{\pi_i + (1-\pi_i)\frac{f(s_i|B)}{f(s_i|G)}}(1+c) - c \geq 0$$
2. If $Pr(\sigma_j(s_j) = I \ \forall j < i) > 0$, then π_i is obtained by updating $\pi_1 = \pi$ using Bayes' rule. If not, then π_i is *any* member of the interval $[0,1]$.

It turns out that in any Weak Pefect Bayesian Equilibrium, each player in $\Gamma(n)$ will follow a simple trigger strategy as we demonstrate below.

Proposition 1 *Any Weak Perfect Bayesian Equilibrium of* $\Gamma(n)$ *is a trigger equilibrium.*

Proof: Let $(\sigma_1, ..., \sigma_n)$ be any WPBE of $\Gamma(n)$. We shall show that each player follows a trigger strategy. We already know that each player i, conditional on having observed a history of investment, follows $\sigma_i = I$ if and only if

$$EU_i = \frac{\pi_i Pr((\sigma_j(s_j) = I)_{j>i}|G)}{\pi_i + (1-\pi_i)\frac{f(s_i|B)}{f(s_i|G)}}(1+c) - c \geq 0$$

Since $EU_i(s_i)$ is clearly increasing and continuous in s_i, player i will adopt then invest only if $s_i \geq t_i$ where t_i is defined by $EU_i(t_i) = 0$. Upon not observing a history of investment, player i will not invest for sure. Thus, player i follows:

$$\sigma_i(s_i, (a_j)_{j<i}) = \begin{cases} I & \text{when } s_i \geq t_i \text{ and } a_j = I \ \forall j < i \\ N & \text{otherwise} \end{cases}$$

[5]Note that we could divide by $f(s|G)$ above because of our assumption of full support in state G.

which is exactly a trigger strategy as defined above. But this immediately implies that any WPBE of $\Gamma(n)$ is a trigger equilibrium. $\diamond$

Proposition 1 allows us to restrict our attention to trigger equilibria. Thus, the Weak Perfect Bayesian Equilibria of $\Gamma(n)$ are n-tuples, $(t_1, ..., t_n) \in \Re^n$ where player i follows a trigger strategy with trigger t_i. Henceforth, we shall refer to equilibria of $\Gamma(n)$ simply as trigger equilibria. Before proceeding further, it may be helpful to consider a simple example. We present one below.

2.4 A Simple Example with Trigger Equilibria

Four players play $\Gamma(4)$, where their signals are generated by a Gaussian (Normal) information system. $S = \mathbb{R}$. The information system is as follows: In state B, signals are generated by a Gaussian process with mean 0 and standard deviation 5. In state G signals are generated by a Gaussian process with the same standard deviation but a mean of 5. It is easy to see that this information system satisfies the strict MLRP property required above. Miscoordinated investment has costs: $c = 1$. Priors are mildly optimistic: $\pi = 0.6$. Given this set up, we use *Gauss* to numerically determine the (in this case unique) trigger equilibrium. Rounded to two decimal places, the triggers are:

```
t1 = 1.71
t2 = -2.57
t3 = -4.37
t4 = -5.46
```

Each player invests if and only if her private signal is above their equilibrium trigger conditional on a history of investment. It is interesting to compare this equilibrium to the hypothetical first best that could be achieved via an informed social planner: i.e., when all players invest if and only if the state is good. In this equilibrium, the probability of coordinated investment in state G is 65% and in state B is 18%. Thus this equilibrium is clearly not pareto efficient, a feature common to many equilibria of strategic games. We shall return to this example later to illustrate other aspects of $\Gamma(n)$.

This example shows that for a very specific realization of $\Gamma(n)$ there is a trigger equilibrium. We now address the question of existence for the general game.

2.5 The Existence of Trigger Equilibria

2.5.1 Case 1: Bounded Signal Support

To demonstrate the existence of pure strategy equilibria, we define the best response function:

Definition 3 *Consider a set of triggers $t \in [\underline{s}, \bar{s}]^n$. Denote the* best response mapping *by $\beta : [\underline{s}, \bar{s}]^n \to [\underline{s}, \bar{s}]^n$ and the ith component of $\beta(t)$ by $\beta_i(t)$. Then, $\beta_i(t) = r^{-1}(E_i)$, where*

$$E_i = \frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j>i} Pr(s \geq t_j|G) - 1\right]\prod_{j<i}\frac{Pr(s \geq t_j|G)}{Pr(s \geq t_j|B)}$$

if $E_i \in [B, T]$*. If* $E_i < B$*,* $\beta_i(t) = \bar{s}$ *and if* $E_i > T$*,* $\beta_i(t) = \underline{s}$*.*

The strict MLRP property implies that $\beta(.)$ is single valued, and the continuity of $f(.|G)$ and $f(.|B)$ imply that $\beta(.)$ is continuous. Thus, $\beta(.)$ is a continuous function that maps $[\underline{s}, \bar{s}]^n$, a compact and convex set, into itself. Therefore, by Brouwer's Fixed Point Theorem, there is a $t^* \in [\underline{s}, \bar{s}]^n$ such that $\beta(t^*) = t^*$. Thus a trigger equilibrium of $\Gamma(n)$ exists.

However, since the argument above admits the possibility that t^* lies on the boundary of S, this leaves open the possibility that the only equilibrium of $\Gamma(n)$ is the trivial equilibrium in which $t_j = \bar{s}$ for all j, and thus nobody invests in equilibrium.

In fact, the extreme form of strategic complementarities embodied in $\Gamma(n)$ ensures that there is always a trivial trigger strategy equilibrium in which nobody ever invests. Let us construct such an equilibrium. Consider the problem of a player, say i, with posterior belief π_i upon observing investment by her predecessors, who is sure that all her successors (if any) will invest. Let $t(\pi_i)$ be the trigger selected by such a player. Clearly, $t(\pi_i)$ is decreasing in π_i. Let π^* be defined by $Pr(s \geq t(\pi^*)|G) = \frac{c}{1+c}$. Note that for some $\epsilon \in (0, \pi^*]$, $Pr(s \geq t(\pi^* - \epsilon)|G) < \frac{c}{1+c}$.

Now consider the problem of player $i-1$ with posterior beliefs π_{i-1}, who observes signal s_{i-1}. She knows that upon observing investment by her, player i will have posterior beliefs $\pi^* - \epsilon$. She will certainly not invest if $Pr(G, s \geq t(\pi^* - \epsilon)|s_{i-1}) < \frac{c}{1+c}$, i.e., if she does not assign sufficient probability to the event that the state is good and (at least) her immediate successor invests (if her immediate successor doesn't invest, it matters not to player $i-1$ what later player do). But notice that

$$Pr(G, s \geq t(\pi^* - \epsilon)|s) = \frac{\pi_{i-1} Pr(s \geq t(\pi^* - \epsilon)|G)}{\pi_{i-1} + (1 - \pi_{i-1})\frac{f(s|B)}{f(s|G)}} \leq Pr(s \geq t(\pi^* - \epsilon)|G) < \frac{c}{1+c}$$

where the first inequality corresponds to the case where $s_{i-1} = \bar{s}$. This means that if player $i-1$ knew that upon observing her invest player i would have beliefs $\pi^* - \epsilon$, then she would assign probability strictly less than $\frac{c}{1+c}$ to the event that the state is good and that player i will invest, *regardless of her own prior belief*. So, player $i-1$ will not invest.

Now it is easy to see that the strategy set $(\bar{s}, \bar{s}, ..., \bar{s})$ is a Perfect Bayesian Equilibrium if upon seeing investment by a predecessor each player has probabilistic beliefs given by $\pi^* - \epsilon$ for any $\epsilon \in (0, \pi^*]$. Given these beliefs off the equilibrium path, the first player will never find it profitable to deviate from her equilibrium strategy to "never invest." This is because even if she believed that players $3, ..., n$ would invest for sure conditional upon investment by their predecessors, she would still assign too low a probability to the event that the state is good and that player 2 (with beliefs given by $\pi^* - \epsilon$ upon seeing player 1 invest) will invest. Thus, the first player will not invest, and so all her successors will set their triggers optimally to infinity (i.e., never invest) [6].

[6] These out of equilibrium beliefs are sufficient but not necessary to support the $(\bar{s}, ..., \bar{s})$ strategy profile as an equilibrium.

However, this is not a very interesting equilibrium, and it is natural to wonder if there is a non-trivial trigger equilibrium of $\Gamma(n)$. In such an equilibrium players would choose interior triggers, and therefore allow for the possibility of coordinated investment. Formally, we refer to these equilibria as investment equilibria.

Definition 4 *Trigger equilibrium* $(t_1, ..., t_n) \in \Re^n$ *of* $\Gamma(n)$ *is a* investment equilibrium *if* $t_j < \bar{s}$ *for all* $j = 1, ..., n$.

Investment equilibria allow for the possibility of coordinated investment.

In order to ensure the existence of investment equilibria, we must lay down some sufficient conditions on the information system of the game. This requires a preamble.

Consider the following situation. Players 1 through $n-1$ choose to invest blindly, i.e., $t_1 = ... = t_{n-1} = \underline{s}$. Consider player n's best response to such strategies (upon observing investment by all her predecessors), t_n. Given our definition of the best response function above, In other words,

$$t_n = r^{-1}\left[\frac{\pi}{1-\pi}\frac{1}{c}\right]$$

This uniquely defines t_n in terms of the parameters (c), the prior (π), and the information system (f). We write $t_n = U_n(c, \pi, f)$, or U_n for short. Further, we require, that

$$Pr(s \geq U_n(c, f)|G) > \frac{c}{1+c}$$

Call this *Condition* Ψ_n.

Now consider the situation where players 1 through $n-2$ choose to invest blindly, i.e., $t_1 = ... = t_{n-2} = \underline{s}$, while player n plays according to trigger U_n. Now, player $n-1$ will choose her trigger according to:

$$t_{n-1} = r^{-1}\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c}Pr(s_n \geq U_n|G) - 1\right)\right]$$

Note that Condition Ψ_n ensures that t_{n-1} is well defined in terms of the parameters, and we write $t_{n-1} = U_{n-1}(c, \pi, f)$ or U_{n-1} for short. Now we require that

$$Pr(s \geq U_{n-1}|G)Pr(s \geq U_n|G) > \frac{c}{1+c}$$

Call this *Condition* Ψ_{n-1}.

We continue iteratively in this way, defining $U_{n-2}(c, \pi, f), ..., U_1(c, \pi, f)$, and conditions $\Psi_{n-2}, ..., \Psi_1$. Now we are ready to define the useful properties of the information system promised above.

Definition 5 *Let* $U_1(c, \pi, f), ..., U_n(c, \pi, f)$ *and the conditions* $\Psi_1, ..., \Psi_n$ *be defined as above. We say* Property Ψ *holds if conditions* Ψ_2 *through* Ψ_n *hold simultaneously, i.e., if*

$$\prod_{j=2}^{n}\int_{U_j(c,\pi,f)}^{\infty} f(x|G)dx > \frac{c}{1+c}$$

Given c and π, the property defined above imposes a restriction on the *information system,* i.e., on the stochastic process generating the private information processes of the agents. Stated in words, Property Ψ simply says that in the *good* state, the information system must be *reliable* enough, i.e., generate signals above predetermined levels (the U_j's) with sufficient probability. This property turns out to be useful for the case with unbounded likelihood ratios.[7] However, for the present case with bounded likelihood ratios, we need a slightly stronger condition. Finally, therefore, a last definition.

Definition 6 *We say that f satisfies* Property $\Psi+$ *in $\Gamma(n)$ if it satisfies Property Ψ and if*

$$\frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j>1} Pr(s \geq U_j|G) - 1\right] > B$$

Then the following result holds:

Proposition 2 *When Property $\Psi+$ holds, there exist $L, U \in S$ with $L < U$ and $U < \bar{s}$ such that for $t \in [L, U]$, $\beta(t) \in [L, U]$.*

Proof: Let $L = (L_1, L_2, ..., L_n)$, where $L_i = \beta((U_1, ..., U_{i-1}), (L_{i+1}, ..., Ln))$ where U_i is defined as above. Let $U = (U_1, U_2, ..., U_n)$. Clearly, $L < U$. Let $t \in [L, U]$. Since Property $\Psi+$ holds, $U_1 < \bar{s}$, thus t is interior in $[\underline{s}, \bar{s}]$ and $\beta(t) \in (B, T)$. Thus,

$$\beta_i(t) = r^{-1}\left(\frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j>i} Pr(s \geq t_j|G) - 1\right]\prod_{j<i}\frac{Pr(s \geq t_j|G)}{Pr(s \geq t_j|B)}\right)$$

Notice that

$$L_i = r^{-1}\left(\frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j>i} Pr(s \geq L_j|G) - 1\right]\prod_{j<i}\frac{Pr(s \geq U_j|G)}{Pr(s \geq U_j|B)}\right)$$

and since $L \leq t \leq U$, $\beta_i(t) \geq L_i$. Similarly, notice that

$$U_i = r^{-1}\left(\frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j>i} Pr(s \geq U_j|G) - 1\right]\right)$$

and thus $\beta_i(t) \leq U_i$. $\diamond$

Corollary 1 *When Property $\Psi+$ holds, there is an investment equilibrium in $\Gamma(n)$ with bounded signal support.*

[7] As we shall see below, Property Ψ turns out to be sufficient to guarantee existence of investment equilibria in $\Gamma(n)$ when $S = \mathbb{R}$.

Proof: When Property $\Psi+$ holds, there exists a compact and convex set $[L, U] \subset [\underline{s}, \bar{s}]$ with $L < \bar{s}$ such that $\beta(t) \in [L, U]$ for all $t \in [L, U]$. Observation of the best response mapping establishes immediately that $\beta(.)$ is continuous on $[L, U]$. Thus, by Brouwer's Fixed Point Theorem, $\beta(.)$ has a fixed point in $[L, U]$. $\diamond$

We now turn to the case for existence of trigger equilibria in the case with unbounded signal support (thus unbounded likelihood ratios).

2.5.2 Case 2: Signals drawn from $\mathbb{R}$

Since we do not have to worry about endpoint problems when $S = \mathbb{R}$, the best response mapping is more simply defined than above. Given a set of triggers $t \in \mathbb{R}^n$, the best response is defined to be $\beta(t) \in \mathbb{R}^n$, where

$$\beta_i(t) = r^{-1}\left[\frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j>i} Pr(s \geq t_j|G) - 1\right]\prod_{j<i}\frac{Pr(s \geq t_j|G)}{Pr(s \geq t_j|B)}\right]$$

Given this definition, we are ready to examine the existence of trigger equilibria for the game with $S = \mathbb{R}$.

As in the case with bounded signal support, a trivial trigger equilibrium with infinite triggers exists. Such an equilibrium can be constructed with an argument identical to the one above. However, the no-investment equilibrium is not very interesting, and we naturally turn to the existence of investment equilibria. The question of existence of investment equilibria is more involved when signals are drawn from $\mathbb{R}$, since the underlying signal generating process has unbouded support, making it impossible to directly appeal to a fixed point theorem. However, a subtler argument establishes that if Property Ψ holds, then investment equilibria exist even in this case. The following result is crucial.

Proposition 3 *If Property Ψ holds, then there exist $\underline{t} \in \mathbb{R}^n$ and $\bar{t} \in \mathbb{R}^n$ such that for all $t \in \Re^n$, $\underline{t} \leq \beta(t) \leq \bar{t}$.*

The proof of this result is involved. It is relegated to the appendix.

Corollary 2 *When Property Ψ holds, there is an investment equilibrium in $\Gamma(n)$ with $S = \mathbb{R}$.*

Proof: Proposition 3 tells us that there exists $[\underline{t}, \bar{t}\,] \subset \mathbb{R}^n$ such that for all $t \in \Re^n$, $\beta(t) \in [\underline{t}, \bar{t}\,]$. Thus, in particular, for all $t \in [\underline{t}, \bar{t}\,]$, $\beta(t) \in [\underline{t}, \bar{t}\,]$. Inspection of the best response mapping establishes that for all $t \in [\underline{t}, \bar{t}\,]$, $\beta(.)$ is single-valued and continuous. Clearly $[\underline{t}, \bar{t}\,]$ is compact and convex. Now, by a simple application of Brouwer's fixed point theorem, we note that there exists $t^* \in [\underline{t}, \bar{t}\,]$ such that $t^* = \beta(t^*)$. Thus, a bounded equilibrium of $\Gamma(n)$ exists. $\diamond$

Next we examine some pertinent properties of these investment equilibria.

3 Properties of Investment Equilibria

In this section we demonstrate two key structural properties of investment equilibria. The properties apply to investment equilibria of $\Gamma(n)$ in general, *regardless* of whether the underlying signals have bounded or unbounded support. Thus, in order to ensure the existence of these equilibria we tacitly assume that Properties Ψ or $\Psi+$ hold, depending on which version of $\Gamma(n)$ we are considering.

The first of these properties encapsulates a simple relation between the relative magnitudes of triggers in any investment equilibrium of $\Gamma(n)$.

Proposition 4 *Suppose* $(t_1, t_2, ..., t_n)$ *is any investment equilibrium of* $\Gamma(n)$*. Then,* t_i *is* decreasing *as a function of* t_j *for* $j < i$*, and* increasing *as a function of* t_j *for* $j > i$*. In other words, an agent's equilbrium triggers is increasing in the triggers of her successors, and decreasing in the triggers of her predecessors.*

Proof: The proof follows directly upon examination of the best response correspondence. Let $(t_1, t_2, ..., t_n)$ be any investment equilbrium of $\Gamma(n)$. Then, by the definition of the best response mapping:

$$t_i = r^{-1}\left[\frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j>i} Pr(s \geq t_j|G) - 1\right]\prod_{j<i}\frac{Pr(s \geq t_j|G)}{Pr(s \geq t_j|B)}\right]$$

Now it is apparent that t_i is increasing in t_j for $j > i$ and decreasing in t_j for $j < i$ since $r(s)$ is decreasing and $\frac{Pr(s \geq x|G)}{Pr(s \geq x|B)}$ is increasing in x. $\diamond$

The intuition behind this result is simple. In equilibrium, conditional upon observing investment by a predecessor, the higher the predecessor's trigger, the higher the signal the predecessor must have observed. The MLRP property of the information system of $\Gamma(n)$ implies that higher private signals make players more optimistic about the state of the world. Large signals are "good news" for would-be investors. They make it likelier that the state is G, which in turn makes it likelier that other players will receive relatively high signals. Thus, observing investment by a predecessor with a high trigger conveys more "good news" for a player, and makes her more optimistic about the state of the world, and about the probability that her successors will also invest. Agents in this model have two sources of information, both of which affect their level of optimism: the public information encapsulated in the observed decisions of their predecessors, and the private information contained in their signals. Thus, when an agent observes more encouraging public information (investment by predecessors with high triggers), she requires less persuasive private information in order to choose to invest. Thus, she picks a lower trigger. Similarly, if an agent believes that her successors have extremely high triggers, then she may be concerned that they shall not invest with higher probability, and "leave her stranded" if she chooses to invest. Thus, she will be inclined in equilibrium to require more persuasive private evidence for the fact that the state is G before deciding to invest. In other words, she will choose a higher trigger.

The above proposition has an immediate consequence for the two player version of our game, as we note below:

Corollary 3 *There is a unique investment equilibrium in* $\Gamma(2)$.

In an investment equilibrium, the process by which agents become more optimistic (or pessimistic) about the state of the world is by Bayesian learning. Agents update their priors about the state of the world by Bayes' Rule upon observing their predecessor's actions. As we have just argued, observation of investment by a predecessor with a high trigger makes an agent more optimistic about the state of the world than the observation of investment by a predecessor with a low trigger. Intuitively, it seems also likely that observing investment by two predecessors makes an agent (at least weakly) more optimistic about the state of the world than observing investment by one predecessor.[8] Thus, upon observing investment by more and more predecessors, later players will require less and less persuasive private information in order to invest. In other words, the greater the mass of public evidence in favour of a good state, the lower the level of private evidence required to make investors take potentially productive but risky actions. The following result captures this intuition.

Proposition 5 *In any Investment Equilibrium of* $\Gamma(n)$, $(t_1, ..., t_n)$, $t_j \geq t_{j+1}$ *for* $j = 1, .., n-1$.

Proof: In equilibrium $(t_1, ..., t_n)$, consider the magnitudes of t_i and t_{i+1}. We know from the definition of the best response mapping:

$$\frac{f(t_i|B)}{f(t_i|G)} = \frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j=i+1}^{n} Pr(s_j \geq t_j|G) - 1\right]\prod_{j=1}^{i-1}\frac{Pr(s_j \geq t_j|G)}{Pr(s_j \geq t_j|B)}$$

$$\frac{f(t_{i+1}|B)}{f(t_{i+1}|G)} = \frac{\pi}{1-\pi}\left[\frac{1+c}{c}\prod_{j=i+2}^{n} Pr(s_j \geq t_j|G) - 1\right]\prod_{i=1}^{i}\frac{Pr(s_j \geq t_j|G)}{Pr(s_j \geq t_j|B)}$$

Note that $Pr(s \geq t_{i+1}|G) \leq 1$ and $\frac{Pr(s \geq t_i|G)}{Pr(s \geq t_i|B)} \geq 1$ due to the MLRP property of f. Note that the equalities follow in both cases if and only if $t_{i+1} = \underline{s} or - \infty$ and $t_i = \underline{s} or - \infty$ respectively. Thus, $\prod_{j=i+1}^{n} Pr(s_j \geq t_j|G) \leq \prod_{j=i+2}^{n} Pr(s_j \geq t_j|G)$ and $\prod_{j=1}^{i-1}\frac{Pr(s_j \geq t_j|G)}{Pr(s_j \geq t_j|B)} \leq \prod_{j=1}^{i}\frac{Pr(s_j \geq t_j|G)}{Pr(s_j \geq t_j|B)}$. This means that $\frac{f(t_i|B)}{f(t_i|G)} \leq \frac{f(t_{i+1}|B)}{f(t_{i+1}|G)}$, and thus $t_i \geq t_{i+1}$. $\diamond$

This proposition implies that conditional upon observing investment by predecessors, later players shall tend to invest more easily, i.e., for larger ranges of private information. Very roughly speaking, this means that later agents are less concerned about the content of their private information than earlier agents. This is because the observation of investment by predecessors make later players progressively more optimistic. Let us consider an instance of such optimism by returning to the example of section 2.4 with a Gaussian information system. Here, there

[8] If the second predecessor has a trigger of $\underline{s}$ or $-\infty$, then her decision to invest does not affect the optimism of succeeding players. Hence the relation if weak.

are only 4 players. The first player pays some attention to her private information. She picks a trigger of 1.71 and the probability that she shall not invest based upon discouraging private information is about 26%. However, the fourth player pays very little attention to her private information. She picks a trigger of -5.46 and the probability that she shall not invest based upon discouraging private information is only about 2%. So, even upon observing investment by only three predecessors, it is possible for the fourth player to invest for extremely large ranges of her own private information.

The extreme optimism of player 4 in this example raises the natural question: is it possible that Bayesian learning has made players "too optimistic" relative to some (as yet unspecified) superior social alternative? Could there be versions of $\Gamma(n)$ where successors completely ignore their private information upon observing predecessors invest, and thereby clearly act suboptimally in a social sense? These questions are addressed in the following section.

4 Herding

In settings with sequential decision making in the presence of uncertainty, private information, and observed public actions, agents are said to "herd" if they blindly imitate their predecessors' choices without heed to their own private information. In other words, herd behavior occurs when, in the words of Douglas Gale (1996) "imitation dominates information." Interpreting this concept literally in terms of our trigger equilibria, agents herd in equilibrium if one or more successors set their triggers to $\underline{s}$, the lower bound of the signal generating process. This means that conditional upon observing investment by their predecessors, some agents choose to invest for *all* possible values of their private signals. The strongest version of such her behavior is if *all* successors choose to blindly imitate the first player. We shall call this type of behavior *strong herding*. For our purposes, we shall also define a much weaker form of herd behavior. When strong herding occurs, later agents become so optimistic that they pay no attention whatsoever to their private information. However, it is not difficult to imagine situations in which agents do not become optimistic enough that to imitate blindly, but still become overoptimistic compared to a situation where private information was aggregated efficiently in the market. We shall call such phenomena *weak herding*. The idea is formally defined later in the paper.

It turns out that strong herding can occur in $\Gamma(n)$ only if signals are drawn from a bounded support, i.e. the likelihood ratios are bounded. However, in the case where signals are drawn from $\mathbb{R}$ and likelihood ratios are unbounded, there can still be weak herding. In what follows, we lay down the informational prerequisites for strong and weak herding in $\Gamma(n)$.

4.1 Bounded Likelihood Ratios: Strong Herd Behavior

We begin with a definition.

Definition 7 *An investment equilibrium t of $\Gamma(n)$ exhibits* strong herding *if $\underline{s} < t_1 < \bar{s}$ and $t_j = \underline{s}$ for $j \geq 2$.*

In words, this simply means that all followers choose to completely ignore their private information. This definition, taken together with Proposition 4 leads to a very useful property of strong herding equilibria.

Proposition 6 *If investment equilibrium t of $\Gamma(n)$ exhibits strong herding, it is unique.*

Proof: Let $t = (t_1, \underline{s}, ..., \underline{s})$ where $\underline{s} < t_1 < \bar{s}$ be a strong herding equilibrium. Suppose it is not unique. Let z be an investment equilibrium, with $z \neq t$. Since $z \neq t$, clearly, there is a j, $j \geq 2$ such that $z_j > t_j$. Proposition 5 implies that if this is so, $z_2 > t_2$. For simplicity, let $z_j = t_j$ for $j \geq 3$. Let $z_2 > t_2$. Then, by Proposition 4 we know $z_1 > t_1$ (equilibrium triggers are increasing in those of successors). But this in turn implies, also by Proposition 4, that $z_2 < t_2$ (equilibrium triggers are decreasing in those of predecessors). This is a contradiction. $\diamond$

Proposition 6 tells us that if $\Gamma(n)$ has a strong herding equilibrium, it is unique in the class of investment equilibria. Corollary 3 tells us that the two player game has a unique equilibrium. Together the two imply that in order to analyze whether $\Gamma(n)$ has a unique strong herding equilibrium, it is sufficient to look at $\Gamma(2)$. In what follows, therefore, we lay down the conditions under which $\Gamma(2)$ has a strong herding equilibrium.

Let (t_1, t_2) be any investment equilibrium of $\Gamma(2)$. Recall that $r(s) \in [B, T]$. Then,

$$r(t_1) = \frac{\pi}{1-\pi}\left[\frac{1+c}{c}Pr(s \geq t_2|G) - 1\right]$$

$$r(t_2) = \frac{\pi}{1-\pi}\frac{1}{c}\frac{Pr(s \geq t_1|G)}{Pr(s \geq t_1|B)}$$

For t to be a strong herding equilibrium, we need $t_2 = \underline{s}$. So $t_1 = r^{-1}\left(\frac{\pi}{1-\pi}\frac{1}{c}\right)$. But if $t_2 = \underline{s}$, then $r(t_2) \geq T$. This implies

$$\frac{Pr(s \geq r^{-1}\left(\frac{\pi}{1-\pi}\frac{1}{c}\right)|G)}{Pr(s \geq r^{-1}\left(\frac{\pi}{1-\pi}\frac{1}{c}\right)|B)} \geq \frac{1-\pi}{\pi}\ c\ T$$

This condition defines precisely the informational requirements for strong herding in $\Gamma(2)$, and therefore for $\Gamma(n)$. It is apparent that there is no unique way of characterizing the information systems that satisfy such a condition. However, it is possible to provide tight characterizations over broad classes of information systems. Below, we provide such a characterization for all information systems where the signal generating processes are linear in the signals.

In considering the case for linear signal generating processes we limit attention without loss of generality to a support of $[0, 1]$. In addition, since it is the ratio of densities and not the individual densities that are important in our model, we normalize the density in state G to be uniform ($U[0, 1]$) also without loss of generality. Given the strict MLRP, this means that $f(s|B)$ is decreasing and linear in s. We are now ready to provide two results characterizing when strong herding will occur in $\Gamma(n)$.

Proposition 7 *Suppose $S = [0,1]$. Let $f(s|G) = 1$. Consider the class of densities in state B that are linear in the signal and that do not have full support. Then, for a given c and for any prior π we can construct an information system such that $\Gamma(n)$ has a unique equilibrium with the strong herding property.*

Proof: Let $f(s|B) = a - bs$, where $a > 0$ and $b > 0$. Since $f(s|B)$ is a density, it must integrate to 1 over its support. The support is given by $[0, \frac{a}{b}]$ where since $f(s|B)$ is not full support we require that $a \leq b$. Also, $\int_0^{a/b}(a - bs)ds = 1$ implies $b = \frac{a^2}{2}$. Putting this two relations together, we get $a \geq 2$. Thus, $f(s|B) = a - \frac{a^2}{2}s$, and since $f(s|G) = 1$, $r(s) = f(s|B)$. Thus, clearly, $0 \leq r(s) \leq a$. Let $k = \frac{\pi}{1-\pi}\frac{1}{c}$. By the conditions defining a strong herding equilibrium (t_1, t_2), we require that $r(t_1) = k$ which implies $t_1 = \frac{2}{a^2}(a - k)$. To ensure that this is an investment equilibrium, we require that $t_1 < 1$, i.e., $a - \frac{a^2}{2} < k$. To ensure that the equilibrium is not trivial, we require $t_1 > 0$, i.e., $a > k$. Finally, in order to make $r(t_2) \geq a$, we require $k\,\frac{Pr(s \geq t_1|G)}{Pr(s \geq t_1|B)} \geq a$, which, upon algebraic simplification implies that we require $\frac{a-2}{1-\frac{2}{a}} \geq k$. Thus we want $a \geq 2$, $a > k$, $a - \frac{a^2}{2} < k$, and $\frac{a-2}{1-\frac{2}{a}} \geq k$. Clearly, for any given k, by picking a large enough, we can satisfy these conditions, which are necessary and sufficient for the existence of a unique investment equilibrium with strong herding. $\diamond$

Proposition 8 *Suppose $S = [0,1]$. Let $f(s|G) = 1$. Consider the class of densities in state B that are linear in the signal and full support. Then there is no equilibrium of $\Gamma(n)$ with the strong herding property.*

Proof: Let $f(s|B) = a - bs$ where $a > 0$ and $b > 0$. Since $f(s|B$ is a density it must integrate to 1 over it's support. The support in this case is given by $[0,1]$ and in order to ensure this full support, we require that $a > b$. So, $\int_0^1(a - bs)ds = 1$ which implies $b = 2(a-1)$. Together, these imply $a < 2$ and, since $b > 0$, $a > 1$. Thus, $1 < a < 2$. So, $f(s|B) = a - 2(a-1)s$ and $r(s) = f(s|B)$. Clearly, $2 - a \leq r(s) \leq a$. Let $k = \frac{\pi}{1-\pi}\frac{1}{c}$. (t_1, t_2) is a strong herding equilibrium if and only if $t_1 = r^{-1}(k) = \frac{a-k}{2(a-1)}$. Also, since $0 < t_1 < 1$ by definition of strong herding, $a > k$ and $a > 2 - k$. Finally, $t_2 = 0$ if and only if $k\,\frac{Pr(s \geq t_1|G)}{Pr(s \geq t_1|B)} \geq a$. Upon algebraic simplification, this yields, $a \leq k$. But we have already required $a > k$. This is a contradiction. $\diamond$

Thus, in the class of linear information systems with bounded likelihood ratios, $\Gamma(n)$ can have a strong herding equilibrium if and only if the the signal generating process in state B is not full support. This provides a criterion for mechanism design in contexts where players need to be coordinated upon some socially productive risky action (or prevented from coordinating upon a socially unproductive one). If it was possible for the mechanism designer to provide private information via linear stochastic processes to the players, she would know exactly how to make all but one player ignore their own private information, or, by the same token, how to force players to pay more attention to their private information.

We now address the question of herd behavior when likelihood ratios are unbounded.

4.2 Unbounded Likelihood Ratios: Weak Herd Behavior

It is apparent by inspection of the best response mapping that when $r(s)$ is unbounded, i.e., when $S = \mathbb{R}$, it is impossible to have strong herding in $\Gamma(n)$. However, the lack of extreme informational inefficiencies in such instances does not mean that there aren't any. Inefficiencies in the aggregation of information can lead to "excessive optimism" in $\Gamma(n)$. How can we measure such excessive optimism?

Information about the state of the world in $\Gamma(n)$ is generated by the sequence of payoff relevant private signals that the players receive. The problem is that the signals are *private*, i.e., only the original recipient of the signal can observe its true value. Others must be satisfied with simply observing the *actions* chosen by the original recipient and guessing from this action what the recipient's signal may have been. If information was efficiently aggregated, each agent would be able to observe the equivalent of all signals that had been received (by herself or others) at the point of time they she is called upon to act. This could be achieved, for example, by a social planner, who could observe each agent's signal and announce it to the rest of the group.[9] Let us denote this variation of $\Gamma(n)$ with observed signals by $\hat{\Gamma}(n)$. We can now define weak herding in $\Gamma(n)$.

Definition 8 *An investment equilibrium t in $\Gamma(n)$ is said to exhibit* weak herding *if there exists $i > 1$ such that with positive probability $t_i < \hat{t}_i$, where $\hat{t}$ is the unique equilibrium in $\hat{\Gamma}(n)$.*

In other words, an investment equilibrium of $\Gamma(n)$ exhibits weak herding if at least one follower becomes excessively optimistic with positive probability. In what follows we present in brief the game with observed signals. For brevity, we simply consider the two-player case.

4.2.1 $\Gamma(2)$ with Observed Signals

Let π_2 denote player 2's updated prior after she has observed her predecessor's signal. Clearly,

$$\frac{\pi_2}{1-\pi_2} = \frac{\pi}{1-\pi}\frac{f(s_1|G)}{f(s_1|B)},$$

where s_1 is the realization of player 1's signal. Upon observing player 1 invest, player 2's expected utility from investing is given by

$$EU_2(I) = \frac{\pi_2}{\pi_2 + (1-\pi_2)\frac{f(s_2|B)}{f(s_2|G)}}(1+c) - c$$

if she observes private signal s_2. Her expected utility from not investing is 0. Since $EU_2(I)$ is clearly increasing and continuous in s_2, player 2 shall choose a trigger strategy, where her trigger

[9]Note that it is not easily possible to simply get each agent to simply announce their signals, since there are significant credibility problems inherent in such announcements. Once an agent has chosen to invest, she has a clear incentive to get her successors to invest, *regardless* of the actual state of the world, and thus has motive to overstate her signal.

$\hat{t}_2$ is defined by $\hat{t}_2 = r^{-1}(\frac{\pi_2}{1-\pi_2}\frac{1}{c})$, i.e.,

$$\hat{t}_2 = r^{-1}(\frac{\pi}{1-\pi}\frac{1}{c}\frac{f(s_1|G)}{f(s_1|B)})$$

.

We are now ready to provide a characterization of weak herding in $\Gamma(2)$. Recall that by definition of the best response mapping in $\Gamma(2)$, $t_2 = r^{-1}(\frac{\pi}{1-\pi}\frac{1}{c}\frac{Pr(s\geq t_1|G)}{Pr(s\geq t_1|B)})$. The investment equilbrium t possesses the weak herding property if with positive probability, $t_2 < \hat{t}_2$, i.e. with positive probability

$$\frac{Pr(s \geq t_1|G)}{Pr(s \geq t_1|B)} > \frac{f(s_1|G)}{f(s_1|B)}$$

It is apparent that this property shall hold for large classes of full support distributions on $\mathbb{R}$, particularly those with thin tails. A natural example of this is the Gaussian Distribution. In what follows, we present a few examples of how weak herding occurs with positive probability in $\Gamma(2)$ when the information system is Normal.

The results are presented in Table 1. We assume the same parameter values as in the example in Section 2.4, but vary the standard deviation of the signal generating processes. In each case, in the table below, we provide, the unique equilibrium triggers of $\Gamma(2)$, the ranges of signals for which weak herding occurs, as well as the corresponding ex ante probability of weak herd behavior. We call the upper and lower bounds of the weak herding range wh_U, and wh_L respectively.

StDev	Equilibrium		WH Range		Pr(Herding)
σ	t_1	t_2	wh_L	wh_U	(*ex ante*)
1	2.419	1.449	2.419	2.908	1%
2	2.199	0.645	2.199	3.477	12%
3	1.910	-0.344	1.910	4.187	22%
4	1.559	-1.478	1.559	4.937	27%
5	1.123	-2.735	1.123	5.712	32%
10	-3.011	-10.471	-3.011	9.647	46%

Table 1: Signal Ranges for Weak Herding

Thus, when $c = 1$, $\pi = 0.6$, $f(s|G) = N(5,5)$, and $f(s|B) = N(0,5)$, agents will exhibit weak herd behavior whenever signals are anywhere between 1.123 and 5.712, which implies an ex ante probability of weak herd behavior of about 32%.[10]

[10]Naturally, in each of the cases above, it is also possible that weak herd behavior shall not occur, since the signals received shall be high enough to justify, or even dwarf, the optimism inherent in the trigger equilibrium. The point of this exercise with unbounded likelihood ratios is to demonstrate that overoptimism is *likely*, not inevitable, as the outcome of rational behavior in $\Gamma(2)$.

The comparison of $\Gamma(2)$ with the game with observed signals raises an obvious related question. If inappropriate aggregation of information in $\Gamma(2)$ represents a source of potential inefficiency, so do the strategic complementarities built into the payoffs. Conditional upon investment by their predecessors, when agents choose whether to invest or not, they take into account only their personal gains and losses from investing, not the gains and losses to society as a whole. When agent 2 chooses not to invest (conditional upon agent 1 having already invested) she imposes an immediate cost of $-c$ on Agent 1. Agent 2 does not take this cost into account, and thereby may possibly be too conservative in her investment strategy relative to the social optimum. In order to compare $\Gamma(2)$ with the alternative game that incorporates *both* observed signals and awareness of social costs, we have to consider the single-agent decision problem. In this, one agent chooses successively to invest or not, and gains the sum of the payoffs to individual players in $\Gamma(2)$. This means that she earns 2 at the end if she chooses to invest twice and the state is G, $-2c$ if the state is B, $-c$ if she chooses to invest exactly once, and 0 if she does does not invest at all. The agent remembers her history of signals when choosing to invest. The comparison of the equilibrium of this single agent decision problem with the trigger equilibria of $\Gamma(2)$ turns out to be similar to the comparison with the case of observed signals above. The broad conclusions are that the trigger equilibrium of $\Gamma(2)$ is inefficient for sure, and can exhibit both overoptimism and overpessimism with positive probability. The single agent version of $\Gamma(2)$ is worked out in the appendix.

The lack of the strong herding property in equilibria of $\Gamma(n)$ when $S = \mathbb{R}$ raises another interesting question. When an equilibrium of $\Gamma(n)$ exhibits the strong herding property, it is possible to coordinate an infinite number of players upon risky investment. However, when the equilibria of $\Gamma(n)$ lack the strong herding property it is natural to wonder whether it is possible to coordinate larger and larger numbers of players upon investment in equilibrium. Intuitively, since all players in the game choose finite triggers, as the number of players get larger and larger, it may be harder and harder to convince the first player to take a risk and invest, since there are more and more later players who could "leave him stranded" by choosing to not invest after he does so. What would happen in $\Gamma(n)$ with $S = \mathbb{R}$ as n grew larger and larger? Is it still possible to make Player 1 invest in equilibrium?

Recall that in our discussion to date, we have assumed that the information system in $\Gamma(n)$ satisfied Property Ψ, which guaranteed the existence of a investment equilibrium when $S = \mathbb{R}$. However, a cursory glance at the definition of Property Ψ might lead one to believe that as one increased the number of players, the information system may no longer satisfy Ψ. However, note that Ψ is only sufficient and not necessary for the existence of equilibria in this model. Whether we can satisfy Ψ in $\Gamma(n)$ as n grows larger, and whether an equilibrium might exist even if Ψ is violated, depends crucially on the information structure chosen for the game. Thus, in considering the effects of increasing n, we are effectively proposing an exercise in comparative information systems. In order to give some structure to such a comparative exercise, it is necessary to parametrize the information system. For this purpose, we henceforth consider Gaussian (Normal) information systems for arbitrary (general) parameter values. This is simply an analytical simplification. Several of our results will not be contingent on the precise functional form of the Gaussian distribution, and we shall point out generalizations in due course. With this in mind, we progress to characterizing the informational requirements for creating the possibility of coordinated investment in $\Gamma(n)$ when $S = \mathbb{R}$.

5 On the Possibility of Coordination

The existence of investment equilibria in $\Gamma(n)$ ensures the possibility of coordinated investment by the participants in the game. It is intuitive that it should be harder to coordinate progressively larger numbers of players upon a risky but socially productive action in our setting. In this section, we explore the informational conditions that will allow us to ensure that it is at least *possible* to coordinate a large number of players upon a productive risky action.

Our new setting specializes the original setting in one sense. The information system $f = \{f(.|G), f(.|B)\}$ is specified to be Gaussian. In the good state, signals are generated by some arbitrary Gaussian process with mean $\mu > 0$ and standard deviation $\sigma > 0$. In the bad state, signals are generated by a Gaussian process with mean 0 and standard deviation σ. Choosing the standard deviation to be identical in both states ensures the strict MLRP property of f. The choice of 0 as the mean of the signal generating process in the bad state is without loss of generality. It is easy to see that what matters is the *difference* between the two means. Thus, μ the mean in the good state could also be viewed as the difference of means between the two states: $\mu = \mu_G - \mu_B$. μ_B is set to 0 for notational simplicity. In sum, therefore, $f = \{N(\mu, \sigma), N(0, \sigma)\}$. We denote this information system by $f(\mu, \sigma)$.

In this new setting consider what happens in $\Gamma(n)$ as n gets bigger. It is easy to see that as the number of players gets large, it becomes harder for f to satisfy Property Ψ. Whether Ψ is violated or not turns out to depend on the initial level of optimism of the players. We shall show below that for initial priors above a certain cutoff point *determined solely by the parameters*, we can always find an information system to satisfy Ψ, and thus ensure the existence of an investment equilibrium. On the other hand, for priors below this cutoff point, there is some finite n for which Ψ is violated in $\Gamma(n)$. However, Ψ is only sufficient and not necessary for the existence of an investment equilibrium. Thus, the violation of Ψ does not exclude the possibility of investment equilibria. In fact, we shall show that for *any* nondegenerate prior on the states, we can find an information system that ensures the existence of an investment equilibrium for $\Gamma(n)$ where n is as large as we please in the set of integers. The following propositions explicate these points.

Proposition 9 *If $\pi > \frac{c}{1+c}$, then for any $n \in Z_{++}$ there exists a $f(\mu, \sigma)$ with σ large and finite such that Property Ψ is satisfied in $\Gamma(n)$. If $\pi < \frac{c}{1+c}$ there exists $n \in Z_{++}$ such that Ψ is violated for $\Gamma(n)$.*

Proof: Adapting the best response mapping to the specific context of the Gaussian information system $f(\mu, \sigma)$, we know that t_i is defined by:

$$\exp(\frac{\mu^2 - 2\mu t_i}{2\sigma^2}) = \frac{\pi}{1-\pi}\left(\frac{1+c}{c}\prod_{j=i+1}^{n} Pr(s_j \geq t_j|G) - 1\right)\prod_{j=1}^{i-1}\frac{Pr(s_j \geq t_j|G)}{Pr(s_j \geq t_j|B)}$$

This implies,

$$t_i = \frac{\mu}{2} - \frac{\sigma^2}{\mu}\ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c}\prod_{j=i+1}^{n} Pr(s_j \geq t_j|G) - 1\right)\prod_{j=1}^{i-1}\frac{Pr(s_j \geq t_j|G)}{Pr(s_j \geq t_j|B)}\right] \tag{1}$$

Now, recalling the definitions of $U_j(c, \pi, f)$ for $j = 2, ..., n$ from section 2.5, we can write

$$U_n = \frac{\mu}{2} - \frac{\sigma^2}{\mu}\ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c} - 1\right)\right] \tag{2}$$

For $i = 2, ..., n-1$

$$U_i = \frac{\mu}{2} - \frac{\sigma^2}{\mu}\ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c}\prod_{j=i+1}^{n} Pr(s_j \geq U_j|G) - 1\right)\right] \tag{3}$$

Now consider the case where $\pi > \frac{c}{1+c}$. This means that $\ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c} - 1\right)\right] > 0$. Equation (2) now implies that U_n is *decreasing* as a function of σ. In particular, $U_n = \frac{\mu}{2} - J_n\sigma^2$ for some $J_n \in \mathbb{R}_{++}$. This means that by picking σ high enough we can ensure that

$$\ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c}Pr(s \geq U_n|G) - 1\right)\right] = \ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c}Pr(z \geq -J_n\sigma - \frac{\mu}{2}) - 1\right)\right] > 1$$

where z is the standard Gaussian variable.

But from equation (3), this in turn ensures that U_{n-1} is also decreasing in σ. In particular, $U_{n-1} = \frac{\mu}{2} - J_{n-1}(\sigma)\sigma^2$ for some with $0 < J_{n-1}(\sigma) < J_n$ for $\sigma < \infty$. So we can pick σ high enough to ensure that

$$\ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c}Pr(z \geq -J_n\sigma - \frac{\mu}{2})Pr(z \geq -J_{n-1}\sigma - \frac{\mu}{2}) - 1\right)\right] > 1$$

where z is the standard normal variable. This ensures that U_{n-2} is decreasing in σ, $U_{n-2} = \frac{\mu}{2} - J_{n-2}(\sigma)\sigma^2$ with $J_{n-2}(\sigma) < J_{n-1}(\sigma) < J_n$ for any $\sigma < \infty$.. Notice that as σ gets arbitrarily large, the J_i's get arbitrarily close to each other. Therefore, continuing in this way, for any finite n, we can clearly choose σ high enough (but finite) to satisfy Property Ψ in $\Gamma(n)$.

Next consider the case where $\pi < \frac{c}{1+c}$. By analogy to the above, we know that this means that U_n is *increasing* as a function of σ. Thus, U_n is bounded below by $\frac{\mu}{2}$ (corresponding to $\sigma = 0$). But notice that U_{n-1}, U_{n-2},...etc. are all bounded below by $\frac{\mu}{2}$ because by equation (3) we observe that they are all of the form

$$U_i = \frac{\mu}{2} - \frac{\sigma^2}{\mu}\ln\left[\frac{\pi}{1-\pi}\left(\frac{1+c}{c}J - 1\right)\right]$$

where $J < 1$. But clearly $\frac{1+c}{c}J - 1 < \frac{1+c}{c} - 1 < 1$. Thus, the product $\prod_{j=2}^{n} Pr(s_j \geq U_j|G)$ is bounded below by $[Pr(s \geq \frac{\mu}{2}|G)]^{n-1}$. This clearly vanishes as $n \to \infty$. Thus, there exists an $n \in Z_{++}$ for which Ψ is violated in $\Gamma(n)$. ⋄

This result is ostensibly counterintuitive but powerful. It says that if sufficiently optimistic players are offered sufficiently *garbled* information about the state of the world, then there exists the possibility for coordinated investment regardless of how large the number of players is, as long as the number is finite and known. What makes the result strong is that the level of initial optimism is *independent of the number of players.*

What if the beliefs of potential players fell below the range that guarantees the existence of an investment equilibrium from Proposition 9? It turns out that there is an alternative way to ensure the possibility of coordinated investment that is independent of the initial priors: to provide sufficiently accurate (instead of sufficiently garbled) information to players.

Proposition 10 *Fix $n \in Z_{++}$ and $\mu > 0$. Let $M > 0$ be such that $[1 - \Phi(-M - \mu)]^{n-1} > \frac{c}{1+c}$. Then we can find $(\epsilon, \sigma) > 0$ such that $t = (t_1, ..., t_n)$ where $t_1 \in [\frac{\mu}{2} - \epsilon, \frac{\mu}{2} + \epsilon]$, and $t_j \leq -M$ for $j = 2, ..., n$ is a trigger equilibrium of $\Gamma(n)$*

Proof: Recall from equation (1) that

$$t_1 = \frac{\mu}{2} - \frac{\sigma^2}{\mu} \ln \left[\frac{\pi}{1-\pi} \left(\frac{1+c}{c} \prod_{j=2}^{n} Pr(s_j \geq t_j|G) - 1 \right) \right]$$

Our choice of M above means that t_1 is well defined for $\sigma \leq 1$.

$$t_2 = \frac{\mu}{2} - \frac{\sigma^2}{\mu} \ln \left[\frac{\pi}{1-\pi} \left(\frac{1+c}{c} \prod_{j=2}^{n} Pr(s_j \geq t_j|G) - 1 \right) \frac{Pr(s \geq t_1|G)}{Pr(s \geq t_1|B)} \right]$$

Notice that as $\sigma \to 0$ and resultantly $t_1 \to \frac{\mu}{2}$, the thin tailed property of the Gaussian distribution ensures that $\frac{Pr(s \geq t_1|G)}{Pr(s \geq t_1|B)} \to \infty$. In particular, this term explodes much faster than $\sigma^2 \to 0$ so that $t_2 \to -\infty$. Thus, for a given $M > 0$, we can clearly find $(\epsilon, \sigma) > 0$ such that $t_1 \in [\frac{\mu}{2} - \epsilon, \frac{\mu}{2} + \epsilon]$ and $t_2 \leq -M$. But notice that in this setting $t_2 \geq t_3 \geq ... \geq t_n$. Thus clearly $t_j \leq -M$ for $j = 2, ..., n$ and we have found a $f(\mu, \sigma)$ to rationalize t as a bounded equilibrium of $\Gamma(n)$. ⋄

Proposition 9 and Proposition 10 lay down sufficient informational conditions to create the possibility of coordination in $\Gamma(n)$ for any $n \in Z_{++}$, no matter how large. Jointly they imply that for initially optimistic individuals, either very good or very bad quality information creates the possibility of coordinated investment. Intermediate quality information does not guarantee the possibility of coordination. Sufficiently good quality information always creates the possibility of coordinated investment regardless of whether agents are initially optimistic or not.

6 Discussion

The results presented in the preceding sections provide a general framework within which to view phenomena associated with herd behavior such as market panics, bank runs, and currency crashes. We unify two prior strands of the literature that address such phenomena: sequential choice models with Bayesian learning without payoff complementarities (herding models), and static coordination games with payoff complementarities without Bayesian learning. It is apparent that agents involved in bank runs and financial panics in general have to simultaneously solve coordination problems (captured by the static games literature), learn from their predecessors (captured by the herding literature), and send effective signals to other market participants (not captured by either prior class of models). This model, albeit in extremely stylized form, captures all these three aspects to the behavior of agents within one framework.

An important message that emerges from our analysis is that information matters. In particular, the stochastic properties of the information that agents receive is relevant in determining outcomes. In the preceding sections, we have demonstrated that herd behavior emerges in equilibrium in varying degrees depending on the properties of the information system. In particular, if the information system is such that agents can exhibit only limited amounts of personal skepticism (bounded likelihood ratios), we have demonstrated that very extreme forms of herd behavior can emerge as the unique outcome of rational behavior. If information systems are such that agents exhibit unbounded personal skepticism (unbounded likelihood ratios), however, the outcomes of the model are less extreme, but by no means represent the efficient aggregation of information. Agents may still exhibit the excessive optimism (or pessimism) that is often observed in the market.

While extremely stylized, our model can be used to analyze several potential applications. For example, it can be applied to the problem of technology choice under uncertainty in the presence of network externalities, much as in Choi (1997). Consider, for example, a group of firms choosing between a "safe" old technology and a "risky" new one. The older technology provides a low payoff that is independent of the state of the world and the choices of other firms. The new technology provides a potentially high payoff, but requires that the various other firms also adopt it, and that the industrial environment as a whole is conducive to change. If other firms do not adopt the new technology, the firm that adopts suffers adjustment costs.

The model can also be used to analyze market booms and busts or currency crises, the traditional subjects of herding models. Since we began our discussion with reference to such phenomena, it is fitting to conclude with a simple example of how this model can be used to examine these. We present such an example below.

7 A Simple Application

In what follows, we present a simple example of a model of coordinated attacks on a currency peg under uncertainty. It contains features of the original example of currency crises as coordination games presented by Obstfeld in 1986, and of more recent work building on similar models by Morris and Shin (1998) and Obstfeld (1998).

We assume that there are two traders who have two choices each: Attack (A) or Don't Attack (D). There are two states of the world: either the economy is strong S, and conditions are not conducive for an attack, or the economy is weak (W) and conditions are ripe for an attack. By not attacking, each trader receives a payoff of 0 regardless of the state of the economy. If the economy is strong, then even a coordinated attack is useless for the traders (even though it may cause significant stress on the economy not modelled here) and thus in the state S, A produces a payoff of $-c$ to each player. When the economy is weak, however, a coordinated attack can pay off: the currency crashes and each trader realizes a payoff of 1. On the other hand, even if the economy is weak, one trader by herself cannot bring down the currency. If one trader attacks in state W and the other does not, she realizes a payoff of $-c$.

The states S and W occur with equal probability. Agents receive private signals about the state of the world. Signals can be low (L), medium (M), or high (H), indicating the strength of the economy. When the economy is strong the government has the resources to provide relatively accurate information to agents and they receive signal H with probability $\frac{2}{3}$, and signal M with probability $\frac{1}{3}$, but never receive signal L. When the economy is weak, the government has no control over the signals it sends to agents, and agents receive signals L, M, and H with equal probability. The information system is summarized in Table 2.

	Signals		
	H	M	L
p_W	1/3	1/3	1/3
p_S	2/3	1/3	0

Table 2: Information System in Currency Attack Game

Trader 1 acts first and trader 2 follows after observing her act. In this setting, it is easy to see based upon our earlier discussion, that there is a *unique* trigger equilibrium with the strong herding property. $t_1 = M$, and $t_2 = H$. This means that trader 1 chooses to attack if and only if she receives signal either M or L. Conditional on trader 1's attack, trader 2 attacks for sure, exhibiting strong herding.

The important conclusion of this model is that even when the state of the economy is strong, there is a 33% chance of a coordinated currency attack. Even though the assumptions of the model preclude a currency crash in this setting, it is quite likely that such a coordinated attack will weaken the economy by causing stresses (high interest rates, lowered reserves levels due to defense of the peg) that shall push the economy to state W and increasing the likelihood of a future successful currency attack. However, that is not our central point. The startling conclusion is that unanticipated (and fundamentally unjustified) coordinated currency attacks can occur with high probability in the unique outcome of rational behavior. To the casual observer, such behavior can appear seemingly irrational and be termed herd behavior.

8 Extensions

While our model takes an essential step towards appropriately modeling the strategic aspect of market booms and busts, and extends the herding literature by incorporating forward-looking behavior on the part of agents, it admits several caveats. One of these is that we require complete coordination on the part of agents to achieve positive payoff from investment. A richer model would allow for payoffs from investment to be a continuous and increasing function of the number of investors. The arguments in such a model would be complicated by combinatorial considerations, because the order of action matters in models such as ours. However, we conjecture that the broad results will be similar to ours. In particular, in a set-up similar to this, we antipate that agents will still follow trigger strategies, and that later agents will choose smaller triggers exhibiting strong or weak herding along the lines of this model. In particular, since requiring complete agreement encourages greater conservatism on the part of players in their choices of action, we conjecture that herd behavior would occur more easily in games where positive payoffs may be earned even without complete agreement.

Other potential extensions of the model would allow players to choose their time of entry into the market, i.e., endogenize the order of actions, or allow for imperfect observation of prior choices. Richer models such as these would more closely approximate the reality of financial market booms and panics. This model provides a benchmark against which to compare such future models.

9 Appendix

9.1 Proof of Proposition 3

Consider the largest and smallest possible triggers that can be chosen by the players. Suppose initially we allow players to choose triggers in an unrestricted way, i.e., anywhere in $\Re$. Call the bounds corresponding to this $\underline{t}^0$ and $\bar{t}^0$. We denote the jth component of $\underline{t}^0$ by $\underline{t_j}^0$ and of $\bar{t}^0$ by $\bar{t_j}^0$. So, $\underline{t}^0 = (-\infty, ..., -\infty)$ and $\bar{t}^0 = (\infty, ..., \infty)$.

Now consider best responses to triggers chosen in $[\underline{t}^0, \bar{t}^0]$. What range would these best responses lie in? Call this range $[\underline{t}^1, \bar{t}^1]$. What is $\underline{t_1}^1$? $\underline{t_1}^1$ is a best response to $(\underline{t_2}^0, ..., \underline{t_n}^0) = (-\infty, ..., -\infty)$. Thus, in the spirit of the computations above, $\underline{t_1}^1$ is defined by the solution to the equation, $Pr(G|x) = \frac{c}{1+c}$, which implies $\frac{f(\underline{t_1}^1|B)}{f(\underline{t_1}^1|G)} = \frac{1}{c}$. Thus, $\underline{t_1}^1 \in \mathbb{R}$[11] How about $\underline{t_2}^1$? It is a best response to $\bar{t_1}^0$ and $(\underline{t_3}^0, ..., \underline{t_n}^0)$. Since $\bar{t_1}^0 = \infty$ Player 2 assumes upon observing investment that Player

[11] In particular, $\underline{t_1}^1 = U_n(c, \pi, f)$ defined above.

1 must have received an infinitely high signal, or, equivalently, that the state must be G.[12] So $\underline{t_2}^1 = -\infty$. By the same token $\underline{t_j}^1 = -\infty$ for $j = 3, ..., n$. How about $\bar{t_1}^1$? This is a best response to $(\bar{t_2}^0, ..., \bar{t_n}^0) = (\infty, ..., \infty)$. So, Player 1 knows that Players 2 through n will *never* invest. So, her best response must be to set her own trigger to ∞. Thus, $\bar{t_1}^1 = \infty$. Similarly, $\bar{t_j}^1 = \infty$ for $j = 2, ..., n-1$. However, $\bar{t_n}^1$ is a best response to $(\underline{t_1}^0, ..., \underline{t_{n-1}}^0) = (-\infty, ..., -\infty)$, and so, by reasoning identical to the case of $\underline{t_1}^1$ that $\bar{t_n}^1$ is defined as the solution to $Pr(G|x) = \frac{c}{1+c}$, and thus, $\bar{t_n}^1 = \underline{t_1}^1 \in \mathbb{R}$. Thus,

$$(\underline{t}^1, \bar{t}^1) = ((\underline{t_1}^1, -\infty, ..., -\infty), (\infty, ..., \infty, \bar{t_n}^1))$$

Now consider best responses to triggers in $[\underline{t}^1, \bar{t}^1]$. What range would these best responses lie in? Call this range $[\underline{t}^2, \bar{t}^2]$. Notice first that $\underline{t_1}^2 = \underline{t_1}^1 \in \mathbb{R}$ because they are best reponses to the same triggers. Also, by an argument identical to that constructed for computing $\underline{t_j}^1 = -\infty$ for $j = 2, ..., n$, we observe that $\underline{t_j}^2 = -\infty$ for $j = 2, ..., n$. Similarly, for $j = 1, ..., n-2$, each of $\bar{t_j}^2$ is a best response to at least one successor trigger of ∞. So, for $j = 1, ..., n-2$, $\bar{t_j}^2 = \infty$. However, $\bar{t_{n-1}}^2$ is a best response to predecessors $(\underline{t_1}^1, -\infty, ..., -\infty)$ and successor $\bar{t_n}^1 \in \mathbb{R}$. Thus, $\bar{t_{n-1}}^2$ is chosen to solve the equation $Pr(G, s_n \geq \bar{t_n}^1 | s_1 \geq \underline{t_1}^1, x) = \frac{c}{1+c}$, and thus clearly, $\bar{t_{n-1}}^2 \in \mathbb{R}$. Finally, note that since triggers are decreasing as best responses to their predecessors and since $\underline{t_1}^0 < \underline{t_1}^1$ while $\underline{t_j}^0 = \underline{t_j}^1$ for $j = 2, ..., n$, $\bar{t_n}^2 \leq \bar{t_n}^1$. Thus,

$$(\underline{t}^2, \bar{t}^2) = ((\underline{t_1}^2, -\infty, ..., -\infty), (\infty, ..., \infty, \bar{t_{n-1}}^2, \bar{t_n}^2))$$

By iterating this argument it is easy to see that

$$(\underline{t}^3, \bar{t}^3) = ((\underline{t_1}^3, -\infty, ..., -\infty), (\infty, ..., \infty, \bar{t_{n-2}}^3, \bar{t_{n-1}}^3, \bar{t_n}^3))$$

and so on, until finally

$$(\underline{t}^{n+1}, \bar{t}^{n+1}) = ((\underline{t_1}^{n+1}, ..., \underline{t_n}^{n+1}), (\bar{t_1}^{n+1}, ..., \bar{t_n}^{n+1}))$$

i.e., $(\underline{t}^{n+1}, \bar{t}^{n+1}) \in \mathbb{R}^n \times \mathbb{R}^n$.

It is clear that the iterative process defined above satisfies the property that $\underline{t}^j \geq \underline{t}^{j-1}$ for all j and $\bar{t}^j \leq \bar{t}^{j-1}$ for all j. Thus, since $(\underline{t}^{n+1}, \bar{t}^{n+1}) \in \mathbb{R}^n \times \mathbb{R}^n$ and $(\underline{t}^j, \bar{t}^j) \in [\underline{t}^{n+1}, \bar{t}^{n+1}]^2$ for all $j \geq n+1$, the trigger bound sequence is monotonic and bounded. So it must converge. Thus there exist $\underline{t} \in \mathbb{R}^n$ and $\bar{t} \in \mathbb{R}^n$ such that for all $t \in \Re^n$, $\underline{t} \leq \beta(t) \leq \bar{t}$. ⋄

9.2 The Single Player Version of $\Gamma(2)$

Consider a modification to $\Gamma(2)$ in which a single player makes all choices, in order, constrained by the same signal structure. She remembers her past signals. Her payoffs are given by the total

[12] This is an ad hoc refinement that we introduce for this iterative process. It is important to note that the refinement is irrelevant for investment equilibria.

payoffs at the end of the game contingent upon *both* her choices in periods 1 and 2, by summing the final payoffs from these choices. Thus, payoffs are given by the following:

$$u(G, a_1, a_2) = \begin{cases} 2 & \text{when } a_1 = I \text{ and } a_2 = I, \\ -c & \text{when } a_1 = I \text{ and } a_2 = N, \\ -c & \text{when } a_1 = N \text{ and } a_2 = I, \\ 0 & \text{when } a_1 = N \end{cases}$$

$$u(B, a_1, a_2) = \begin{cases} -2c & \text{when } a_1 = I \text{ and } a_2 = I, \\ -c & \text{when } a_1 = I \text{ and } a_2 = N, \\ -c & \text{when } a_1 = N \text{ and } a_2 = I, \\ 0 & \text{when } a_1 = N \end{cases}$$

We denote the decision maker's updated beliefs about the state after observing the signal in period 1 by π_2. Clearly

$$\frac{\pi_2}{1-\pi_2} = \frac{\pi}{1-\pi}\frac{f(s_1|G)}{f(s_1|B)}.$$

Now, the expected utility to the decision maker of investing in period 2, conditional upon having invested in the previous period is given by

$$EU_2(I) = \frac{\pi_2}{\pi_2 + (1-\pi_2)\frac{f(s_1|B)}{f(s_1|G)}}(2+2c) - 2c.$$

The expected utility to the decision maker of not investing in period 2 conditional upon having invested in period 1 is simply $-c$. Thus, by employing arguments made repeatedly above, the decision maker shall choose a trigger $t_2^{'}$ given by

$$t_2^{'} = r^{-1}\left(\frac{\pi}{1-\pi}\frac{f(s_1|G)}{f(s_1|B)}\frac{2+c}{c}\right).$$

Clearly then, the unique trigger equilibrium of $\Gamma(2)$ is socially optimal only if

$$\frac{Pr(s_1 \geq t_1|G)}{Pr(s_1 \geq t_1|B)} = (2+c)\frac{f(s_1|G)}{f(s_1|B)}.$$

This, of course, happens with zero probability given the assumptions on the signal generating processes. The trigger equilibrium shall be characterized by both overoptimism and underoptimism compared to the socially optimal case with positive probability (for finite c).

References

[1] Avery, C. and P. Zemsky, (1998) "Multidimensional Uncertainty and Herd Behavior in Financial Markets," *American-Economic-Review*, 88(4), pp. 724-48.

[2] Banerjee, A., (1992) "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107(3), pp. 797-818.

[3] Bikhchandani, S., D. Hirshleifer, and I. Welch, (1998) "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *Journal of Economic Perspectives*, 12(3), pp. 151-70.

[4] Bikhchandani, S., D. Hirshleifer, and I. Welch, (1992) "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100(5), pp. 992-1026.

[5] Bulow, J., J. Geanakoplos, and P. Klemperer, (1985) "Multimarket Oligopoly: Strategic Substitutes and Complements," *Journal of Political Economy*, 93(3), pp. 488-511.

[6] Chamley, C. and D. Gale, (1994) "Information Revelation and Strategic Delay in Irreversible Decisions," *Econometrica*, 62(5), pp. 1065-85.

[7] Chari, V., and P. Kehoe, (1997) "Hot Money," March 1997, Mimeo.

[8] Chari, V. and P. Kehoe (2000) "Financial Crises as Herds," Mimeo.

[9] Choi, J., (1997) "Herd Behavior, the Penguin Effect, and the Suppression of Informational Diffusion: An Analysis of Informational Externalities and Payoff Interdependency," *The Rand Journal of Economics*, 28(3), pp. 407-25.

[10] Cole, H. and T. Kehoe, (1996) "A Self-Fulfilling Model of MExico's 1994-1995 Debt Crisis," *Journal of International Economics*, 41(3-4), pp. 309-30.

[11] Corsetti, G., A. Dasgupta, S. Morris, and H. Shin (2000) "Does One Soros Make a Difference? A Model of Currency Crises with Large and Small Traders", Mimeo, Yale University.

[12] Froot, K., D. Scharfstein, and J. Stein, (1992) "Herd on the Street: Informational Inefficiencies in a Market with Short-Term Speculation," *Journal of Finance*, 47(4), pp. 1461-84.

[13] Gale, D., (1996) "What Have We Learned From Social Learning", *European Economic Review*, 40(3-5), pp. 617-28.

[14] Gul, F. and R. Lundholm, (1995) "Endogenous Timing and the Clustering of Agents' Decisions," *Journal of Political Economy*, 103(5), pp. 1039-66.

[15] Jeitschko, T. and C. Taylor (1999) "Local Discouragement and Global Collapse: A Theory of Coordination Avalanches", forthcoming *American Economic Review.*

[16] Lee, I., (1993) "On the Convergence of Informational Cascades," *Journal of Economic Theory*, 61(2), pp. 396-411.

[17] Lee, I., (1998) "Market Crashes and Information Avalanches," *Review of Economic Studies*, 65(4), pp. 741-59.

[18] Morris, S. and H. Shin, (1998) "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," *American Economic Review*, 88(3), pp. 587-97.

[19] Obstfeld, M., (1986) "Rational and Self-Fulfilling Balance of Payments Crises," *American Economic Review*, 76(1), pp. 72-81.

[20] Obstfeld, M., (1998) "Open-Economy Macroeconomics: Developments in Theory and Policy," *Scandinavian Journal of Economics*, 100(1), pp. 247-75.

[21] Smith, L. and P. Sorensen, (1999) "Pathological Outcomes of Observational Learning," forthcoming in *Econometrica*.