Multiple Outcomes of Speculative Behavior in Theory and in the Laboratory[1]

James S. Costain
Research Division, Bank of Spain
Frank Heinemann
Technische Univ., Berlin
Peter Ockenfels
Goethe-Univ., Frankfurt am Main

First version: July 2004
This version: June 2007

Correspondence address:
División de Investigación, Servicio de Estudios
Banco de España
Calle Alcalá 48
28014 Madrid, Spain
james.costain@bde.es

[1]Thanks to Rosemarie Nagel, Dan Friedman, and seminar participants at Univ. Carlos III and the EEA-ESEM 2004 meetings for helpful discussions. The experiments described in this paper were programmed and conducted with the z-Tree software package (Urs Fischbacher, Univ. of Zurich, 1999). Financial support from the Spanish Ministry of Science and Technology (MCyT grant SEC2002-01601) is gratefully acknowledged. Errors are the responsibility of the authors. Raw data, additional graphs, and other experimental and numerical materials can be found at the following address:

http://www.wm.tu-berlin.de/~makro/Heinemann/publics/cho.html.

Abstract

This paper challenges Morris and Shin's (1998) argument that the outcome of a speculative attack is uniquely determined by macroeconomic fundamentals. We generalize Morris and Shin's model, and the experiment of Heinemann, Nagel, and Ockenfels (2004), by making decisions sequential and allowing some previous actions to be observed. We show sufficient conditions that guarantee the existence of a range of fundamentals where multiple outcomes occur. The main requirement is simply that most players must observe a sufficiently large number of previous choices.

In our experimental sessions, eight to twelve players observe signals about the aggregate state and may also observe a random subset of previous actions. Our subjects display herding behavior consistent with the unique logit equilibrium of a boundedly rational version of our game. These strategies imply a unique mapping between fundamentals and the fraction of players attacking if previous actions are unobserved. But when most previous actions are observed, they give rise to a "tripartite classification of fundamentals": there is a significant middle interval in which all players attacking, and no players attacking, both occur with more than 1% probability.

JEL classification: C62, C72, C73, C92, E00, F32

Keywords: Multiplicity, herding, global games, logit equilibrium, currency crises, experiments

1 Introduction

In an influential paper, Morris and Shin (1998; henceforth MS98) argued that because of imperfect information, speculation against a fixed exchange rate is likely to yield a unique outcome for any given state of macroeconomic fundamentals, in spite of the strategic complementarities involved in currency speculation. This paper challenges their conclusion, showing that it relies heavily on an assumption of exactly simultaneous choice. We show that if most players observe a sufficient number of previous actions before making their own decisions, then there exists a range of fundamentals over which the outcome of the speculative attack game is unpredictable.[1] In other words, sequential choice brings back the more traditional "tripartite classification of fundamentals", as illustrated in Figure 1, that occurs in models like that of Obstfeld (1996).

Considering how sequential choice alters the predictions of the MS98 model is natural, since it is well known from the herding literature (*e.g.* Banerjee 1992; Bikhchandani, Hirshleifer, and Welch 1992, 1998; Chamley 2003B; Chari and Kehoe 2003) that when players can observe previous actions, small variations in the initial choices can lead to substantially different aggregate outcomes. This variation in aggregate outcomes arises even within the context of a single equilibrium. Thus, a number of recent theoretical papers have considered coordination games where not all choices are simultaneous (Chamley 2003A; Dasgupta 2007; Heidhues and Melissas 2006; Angeletos, Hellwig, and Pavan 2006; Costain 2007), and in most of these contexts the possibility of multiple outcomes returns. This occurs in spite of fact that these models include private information like that which serves to prove uniqueness in MS98, Carlsson and van Damme (1993), Frankel, Morris, and Pauzner (2003), and other "global games" papers.

Therefore this paper proposes a laboratory experiment which generalizes the MS98 game to allow non-simultaneous choice. Up to now, experiments

[1] It is important to emphasize that the primary focus of this paper is multiplicity of macroeconomic *outcomes* conditional on aggregate fundamentals— which we believe is the relevant issue for policy makers— rather than the related but more theoretical question of multiplicity of *equilibrium*.

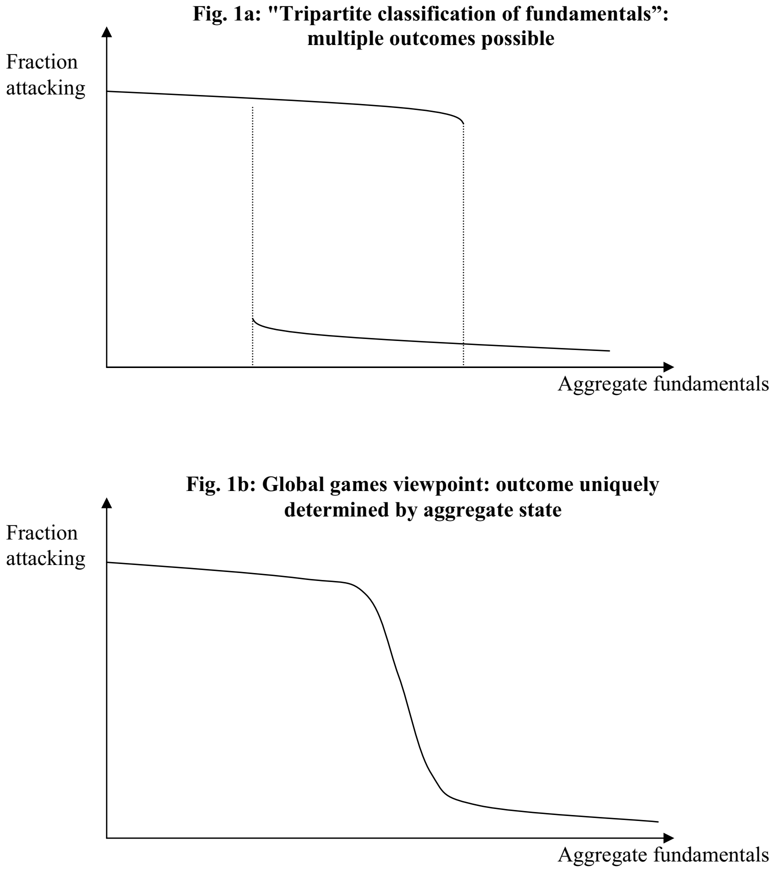

Figure 1: Two views of speculative attacks

testing the global games framework have given decidedly mixed results. Heinemann, Nagel, and Ockenfels (2004; henceforth HNO04) find uniqueness in laboratory experiments based on MS98, as do Cabrales, Nagel, and Armenter (2007) for a static game related to Carlsson and VanDamme (1993).[2] However, they find no evidence for the differences between complete and incomplete information settings that the global games framework implies. Independent from the information setting, aggregate behavior is fairly predictable and with common information, subjects coordinate on thresholds somewhere between the

[2]Other recent experiments on currency or banking crises include Cheung and Friedman (2005), Cornand (2006), Schotter and Yorulmazer (2003), Shurchkov (2007) and Duffy and Ochs (2007).

global-game solution and the payoff-dominant equilibrium.

In contrast, the mechanism underlying herding models has been much more successfully validated in the laboratory; see Anderson and Holt (1997, 1998), Kubler and Weizsacker (2004), Drehmann, Oechssler, and Roider (2005), and Alevy, Haigh, and List (2006). The main experimental failing of the original theoretical herding papers is that the degree of rationality obtained in the laboratory is lower than the theory papers assumed. Anderson and Holt (1997, p. 848) indicate that "human subjects frequently deviate from rational Bayesian inferences" and that therefore behavior can be better described by a logistic error response function. In particular, quantal response equilibria yield a good fit of data generated by laboratory experiments on herding games. Kubler and Weizsacker (2004) show that the quality of data fit can be significantly improved by allowing different error rates for different levels of reasoning. We build upon these approaches by estimating a logit equilibrium and show that it provides a significantly better fit of data in a speculative-attack model with sequential moves than the equilibrium for fully rational behavior.

Our model applies the herding structure of Costain (2007), in which a random number of previous choices is observed, to a game with is otherwise identical to that in MS98 and in the experiment of HNO04. Theoretically, we focus on a well-behaved class of equilibria we call "double monotonic herding equilibria", for which we can prove existence numerically, and which closely match our experimental observations. We prove that if when there are many rational players, most of whom observe many previous actions, then equilibria of this type exhibit a "tripartite classification of fundamentals" whenever a full-information model like Obstfeld (1996) would. That is, for some intermediate values of fundamentals, the aggregate outcome is unpredictable, in contrast to MS98. Nonetheless, equilibria of this type have well-defined and intuitive predictions and comparative statics: the probability distribution over the number of attacking agents is well defined and continuous in the fundamentals, and attacks are more likely in bad aggregate states. This feature distinguishes a rational herding model from a sunspot model, in which probabilities of reaching different equilibria are not well defined. Thereby, we avoid one of the main weaknesses of Obstfeld (1996) that was criticized by MS98.

Experimentally, we compare sessions in which previous actions are unobserved (a setup equivalent to MS98 and HNO04) to sessions in which 50%, 75%, or 100% of previous actions are observed. As our model predicts, the more previous actions are observed, the more often aggregate outcomes appear to cluster near the extremes, either with most players attacking, or most not attacking. To check how this aggregate behavior arises, we estimate our experimental subjects' strategies in two ways: first with a flexible reduced form, and then by structurally estimating a logit equilibrium of our model. The coefficients of our reduced-form estimates always have the signs implied by a double monotone herding equilibrium. The structural estimation of the logit equilibrium also fits the data very closely, and shows that subjects come close to making fully rational decisions. We then use our both our strategy estimates to reconstruct the implied probability distribution over aggregate outcomes. Whenever at least 75% of previous actions are observed, we find a "tripartite classification of fundamentals" with a significant central interval over which either extreme outcome (no players attacking, or all players attacking) can occur with at least 1% probability. However, the degree of coordination of players' decisions appears insufficient to generate the sharply bimodal distribution of aggregate outcomes that the fully rational version of our model implies.

2 The herding model

2.1 Morris and Shin's game

Our intention here is to construct a game which can be played in the laboratory, and which is as close as possible to MS98's stylized model of speculative attacks, except that choice is not exactly simultaneous. Let I be the number of players.[3] As in MS98, many traders choose whether or not to attack a currency peg on the basis of limited information about the true state Θ of

[3]Costain (2007) analyzes the limiting case $I = \infty$, which simplifies the structure of the equilibrium; but here we focus mostly on the finite I case, for compatibility with our experiments.

the macroeconomy. If the proportion of traders who choose to attack exceeds a hurdle function $a(\Theta)$, devaluation occurs, resulting in a payoff $R(\Theta) - t$ to the attackers. Otherwise, attackers lose the transactions cost t. Players who choose not to attack have payoff zero. It is assumed that $a' > 0$ and $R' < 0$, so that larger Θ represents a better state of the economy.

Ex ante, the fundamental state Θ has *c.d.f.* $G(\Theta)$, which is a uniform distribution over the support $\Omega \equiv [\underline{\theta}, \overline{\theta}]$. To make the problem interesting, we follow MS98 by assuming parameters that would guarantee a "tripartite classification of fundamentals" if there were full information.

> **Assumption 1. a.** If the state is sufficiently bad, it is worthwhile to attack: $R(\underline{\theta})/t > 1$ and $a(\underline{\theta}) < 1/I$.
> **b.** If the state is sufficiently good, it is worthwhile not to attack: $R(\overline{\theta})/t < 1$ or $a(\overline{\theta}) > 1$.
> **c.** The interval $\Omega^m \equiv [\theta_m, \theta^m] \equiv \{\theta : 1/I < a(\theta) < 1 < R(\theta)/t\}$ is nonempty. That is, for any θ in $\Omega^m \subset \Omega$, a full information game of simultaneous moves would have multiple equilibria.

Under these assumptions, the bounds on the full-information region of multiplicity are $\theta_m = a^{-1}(1/I)$ and $\theta^m = \min\{R^{-1}(t), a^{-1}(1)\}$. However, we instead assume incomplete information, as in MS98. Before deciding whether or not to attack, player i receives a signal x_i that reveals information about Θ. The signal x_i is *i.i.d.* across individuals, with a conditional *c.d.f.* $F(x_i|\Theta)$ that is uniform on $[\Theta - \epsilon, \Theta + \epsilon]$.[4]

2.2 Sequential choice

In our model and in our experiment, the players $i \in \mathcal{I}_I \equiv \{1, 2, 3, \ldots, I\}$ make their choices in numerical order. Writing actions as $\eta_i \in \{0, 1\}$, where 1 means "attack", the proportion of attacking so far, up to and including agent i, is:

$$\alpha_i \equiv \frac{1}{i} \sum_{j=1}^{i} \eta_j \in \mathcal{J}_i \equiv \left\{0, \frac{1}{i}, \frac{2}{i}, \ldots, 1\right\} \tag{1}$$

[4]Numerically and experimentally, we must restrict variables to finite, discrete supports. Thus we draw Θ from an equally-spaced grid Γ_Θ between $\underline{\theta}$ and $\overline{\theta}$, and draw x_i from an equally spaced grid Γ_x over $[\Theta - \epsilon, \Theta + \epsilon]$.

Success or failure of the attack depends on the overall proportion of attackers $\alpha_I \in \mathcal{J}_I$. The attack succeeds, and the currency is devalued, if and only if

$$\alpha_I \;\; \geq \;\; a(\Theta)$$

Obviously, sequential choice is logically equivalent to simultaneous choice if preceding choices are unknown. Therefore, the key difference between our model and that of MS98 is that we allow players to observe some previous decisions with nonzero probability. For simplicity, person i's observations are drawn *with equal probability* (but without replacement) from her set of predecessors $\mathcal{I}_{i-1} \equiv \{1, 2, \ldots, i-1\}$, so that i is just as likely to observe the choice of agent 1 as that of agent $i-1$ or that of any other preceding player. Player i can observe up to $N_i^{max} \equiv \min(N^{max}, i-1)$ predecessors. The probability of observing $n_i = n$ predecessors is assumed binomial:

$$N_i(n) \;\; \equiv \;\; prob(n_i = n) \;\; \equiv \;\; \left(\begin{array}{c} N_i^{max} \\ n \end{array} \right) q^n (1-q)^{N_i^{max}-n} \tag{2}$$

for some $q \in [0, 1]$. That is, q is the probability that any given preceding player is observed. The number of observed predecessors who attacked is denoted m_i.

Individuals know the total number of players I. However, in our model they do *not* necessarily know their position in the sequence, nor do they know the positions of the predecessors they observe. More precisely, a player places a uniform prior over her possible positions $i \in \mathcal{I}_I$, and the only clue that helps her update this assessment is that if she observes n previous choices, then her position cannot be less than $n+1$. Thus, the information set on which she bases her choice is:

$$(n_i, m_i, x_i) \tag{3}$$

Given her information set, the player must try to deduce the aggregate outcome $(\Theta, \alpha_I) \in [\underline{\theta}, \overline{\theta}] \times \mathcal{J}_I$. We denote her conditional probability assessment as:

$$\Pi(\theta, \alpha | n, m, x, \eta) \;\; \equiv \;\; \text{prob}(\Theta \leq \theta, \alpha_I \leq \alpha | n, m, x, \eta) \tag{4}$$

Note that the individual's choice η influences her perceived distribution of the aggregate outcome α_I, except perhaps in the limiting case $I = \infty$.

In general, a player's strategy can be defined by her probability $y(n,m,x)$ of playing $\eta = 1$ conditional any information set (n,m,x). If we impose the innocuous assumption that players attack when indifferent, then an optimal strategy must satisfy:

$$\begin{array}{lll} y(n,m,x) = 1 & \text{iff} & E_\Pi\left[R(\Theta)\mathbf{1}_{\alpha_I \geq a(\Theta)}|n,m,x,\eta=1\right] \ \geq \ t \\ y(n,m,x) = 0 & \text{iff} & \text{otherwise} \end{array} \tag{5}$$

Here $\mathbf{1}_{\alpha_I \geq a(\Theta)}$ is an index function taking the value 1 if there is a successful attack, and 0 otherwise. The expectation is evaluated using the probability distribution Π, conditional on observing n agents, of whom m attacked, and observing signal x, and also conditional on playing $\eta = 1$.

However, generically, players will strictly prefer 0 or 1 for almost every (n,m,x). Therefore, we typically expect to find threshold equilibria $\tau(n,m)$ of the following form:

$$\begin{array}{lll} y(n,m,x) = 1 & \text{iff} & x \leq \tau(n,m) \\ y(n,m,x) = 0 & \text{iff} & x > \tau(n,m) \end{array}$$

The threshold signal $\tau(n,m)$ must satisfy the indifference condition

$$t \ = \ E_\Pi\left[R(\Theta)\mathbf{1}_{\alpha_I \geq a(\Theta)}|n,m,\tau(n,m),\eta=1\right] \tag{6}$$

Assuming threshold strategies will be helpful for some of our analytical results, but we will not impose this assumption on our numerical simulations.

2.3 Rational herding equilibrium

Any strategy $y(n,m,x)$ or $\tau(n,m)$ induces stochastic processes η_i and α_i for $i \in \mathcal{I}_I$. The process α_i can be calculated recursively:

$$\alpha_i \ = \ \alpha_{i-1} + \frac{1}{i}(\eta_i - \alpha_{i-1}) \tag{7}$$

This representation is helpful because it shows that each history $\{\alpha_i\}_{i=1}^I$ has the structure of a *stochastic recursive algorithm*, so that we can use results from the adaptive learning literature to analyze convergence of α_i as $i \to \infty$.

To spell out the stochastic processes for η_i and α_i, recall first that the *c.d.f.* of the aggregate fundamental is $G(\Theta)$. Next, for each $i \in \mathcal{I}$, a signal x_i is drawn with distribution $F(x|\Theta)$. The number of observations n_i is drawn with distribution $N_i(n)$. These observations are drawn randomly from the set of predecessors (without replacement). Therefore, if the fraction of predecessors who have attacked is $\alpha_{i-1} = \alpha$, then the probability of observing exactly m attackers in a sample of $n_i = n$ predecessors, is

$$M_i(m|n,\alpha) \quad \equiv \quad \text{prob}(m_i = m|i,n,\alpha) \quad \equiv \quad \frac{\begin{pmatrix} \alpha(i-1) \\ m \end{pmatrix} \begin{pmatrix} (1-\alpha)(i-1) \\ n-m \end{pmatrix}}{\begin{pmatrix} i-1 \\ n \end{pmatrix}} \tag{8}$$

which goes to $M(m|n,\alpha) \equiv \begin{pmatrix} n \\ m \end{pmatrix} \alpha^m (1-\alpha)^{n-m} \equiv \frac{n!}{m!(n-m)!} \alpha^m(1-\alpha)^{n-m}$ in the limit as i goes to infinity.

Given the individual state (n_i, m_i, x_i), player i's choice is $\eta_i = 1$ with probability $y(n_i, m_i, x_i)$, and zero otherwise. This implies an explicit formula for the probability that $\eta_i = 1$, given the fraction α_{i-1} of predecessors who attacked, the index i, the state Θ, and the strategy $y()$:

$$T_i(\alpha_{i-1}, \Theta, y) \quad \equiv \quad \text{prob}(\eta_i = 1|i, \alpha_{i-1}, \Theta, y) \quad = \tag{9}$$

$$\sum_{n=0}^{N_i^{max}} N_i(n) \sum_{m=0}^{n} M_i(m|n, \alpha_{i-1}) \int_{\underline{\theta}-\epsilon}^{\bar{\theta}+\epsilon} y(n,m,x) dF(x|\Theta)$$

For large i, $N_i^{max} = N^{max}$ and $M_i(m|n,\alpha) \to M(m|n,\alpha)$, so the sequence of functions T_i approaches a limit $T(\alpha, \Theta, y)$. Note that M and therefore T are C^∞ functions of α.

The function T_i states the probability that trader i attacks the currency, given the fraction who attacked prior to him. Using T_i, we can construct all the other probabilities that are needed to solve the model. In particular, we need the following joint probability (for details, see Appendix B):

$$\text{prob}(\alpha_I, \Theta, i, n_i, m_i, x_i, |\eta_i = 1, y)$$

which means the joint probability of the event in which the the aggregate outcome is (α_I, Θ), the individual position is i, and individual information set is (n_i, m_i, x_i), assuming that all other players use strategy y and the individual chooses $\eta_i = 1$. Knowing this joint probability, the trader can construct the conditional distribution he needs to solve his maximization problem:

$$\Pi(\alpha, \theta | n, m, x, \eta = 1, y) \quad = \quad \frac{\sum_{i, \alpha_I \leq \alpha, \Theta \leq \theta} \text{prob}(\alpha_I, \Theta, i, n, m, x, |\eta = 1, y)}{\sum_{i, \alpha_I, \Theta} \text{prob}(\alpha_I, \Theta, i, n, m, x, |\eta = 1, y)} \tag{10}$$

The conditional probability distribution in (10) is all the information necessary to choose an optimal threshold strategy, so it also implies a fixed point problem that defines the equilibrium strategy. Given any strategy y, let $\Pi(\alpha, \theta | n, m, x, \eta = 1, y)$ be given by (10). Then let By be the strategy calculated from (5) for a player who makes inferences according to Π. That is, By is the best response to y, so we have:

> **Definition.** A **rational herding equilibrium** is a strategy y^* which is a fixed point of the best response mapping B:
>
> $$y^* \quad = \quad By^* \tag{11}$$

Obviously, this definition also suggests an algorithm for calculating the equilibrium:

1. Guess a strategy $y(n, m, x)$ for all possible information sets (n, m, x).
2. Using the mappings T_i, construct the conditional distribution Π over aggregate outcomes conditional on individual information and on choice $\eta = 1$, given that all others play y, as in (10).
3. For each (n, m, x), find the optimal probability of attack given Π, as in (5).
4. Return to 2 and iterate to convergence.

Steps (2) and (3) constitute the best response mapping $y' = By$. For the details of this calculation, see Appendix B.

Distribution Π is one of the most important equilibrium objects, and can in principle be observed in the laboratory, but it cannot be observed with macroeconomic data. The observable macroeconomic implications of the model are summarized by the distribution over aggregate outcomes. In particular, the macroecononomic implications of the model can be stated in terms of the probability of any given aggregate outcome $\alpha_I \in \{0, 1/i, 2/i, ...1\}$, conditional on aggregate fundamentals Θ:

$$\hat{p}_I(\alpha|\theta) \equiv prob(\alpha_I = \alpha|\Theta = \theta) \tag{12}$$

We will also use the notation $\hat{P}_I(\alpha|\theta) \equiv prob(\alpha_I \leq \alpha|\Theta = \theta)$ to refer to the associated *c.d.f.* Figures 2, 3, and 6 graph examples of the conditional outcome distribution $\hat{p}_I$ for various calibrations of the model, drawn as contour plots.

2.4 Boundedly rational decisions

The preceding model description is based on perfectly rational behavior (subject, of course, to the constraints imposed on the information set). That is, it assumes that the probability of attacking jumps from exactly zero to exactly one at the indifference threshold. However, it is unrealistic to expect such sharp calculation in the laboratory, so it is also helpful to consider boundedly rational behavior. In particular, in this subsection we restate the model under the assumption of logit choice, which is often a successful representation of discrete choice in laboratory work (Goeree and Holt 1999).

Under fully rational discrete choice of $\eta \in \{0, 1\}$, action $\eta = 1$ is chosen with probability one if its payoff is strictly higher than that of $\eta = 0$, and *vice versa.* Logit choice weakens this condition, and instead imposes the following logistic probability of playing $\eta = 1$:

$$\frac{\exp[\lambda^{-1}u(1)]}{\exp[\lambda^{-1}u(0)] + \exp[\lambda^{-1}u(1)]} \tag{13}$$

where $u(\eta)$ represents the expected payoff, in equilibrium, of action η. To translate this behavior into our herding model, we simply need to plug in the appropriate payoff function $u(\eta)$. This turns out to be especially simple since $u(0) = 0$ (the payoff of not attacking is always zero, regardless of the information set). Thus, using our previous notation, the logit probability of attacking $y(n,m,x)$ can be written as:

$$y(n,m,x) \quad = \quad \frac{1}{1+\exp\left\{t/\lambda - \lambda^{-1}E_{\Pi}\left[R(\Theta)\mathbf{1}_{\alpha_I \geq a(\Theta)}|n,m,x,\eta=1\right]\right\}} \tag{14}$$

Given this probability of attacking, the implied stochastic process for η_i and α_i can be constructed in terms of the function $T_i(\alpha,\Theta,y)$ as before. Thus we can also calculate the implied distribution of outcomes conditional on behavior, $\Pi(\alpha,\theta|n,m,x,\eta=1,y)$, following the same sequence of steps as in the fully rational case. Therefore, a *logit herding equilibrium* solves a fixed point problem analogous to that which defines a rational herding equilibrium.[5] In particular, define $B_\lambda y$ as the logit strategy defined by (14), when the conditional distribution Π is given by (10). Then:

> **Definition.** For a given λ, a **logit herding equilibrium** is a strategy y^*_λ that is a fixed point of the logit response mapping B_λ:
> $$y^*_\lambda \quad = \quad B_\lambda y^*_\lambda \tag{15}$$

Note that all the complicated steps in computing a logit equilibrium are identical to those involved in finding a rational equilibrium: the hard part is computing the conditional distribution Π implied by a given strategy y. The only difference between the two fixed point problems is that the response to a given Π is given by (14) in the logit case, and by (5) in the fully rational case. In fact, the logit equilibria nests the rational case when we set $\lambda = 0$: that is, $B_0 = B$. At the opposite extreme, logit equilibrium also nests the trivial case of fully random play. That is, $y^*_\infty = B_\infty y^*_\infty = 0.5$ for all (n,m,x): a strategy of attacking with 50% probability regardless of the information set is the unique logit equilibrium associated with $\lambda = \infty$.

[5] Logit equilibrium is a simple special case of quantal response equilibrium. See McKelvey and Palfrey (1995). Kuebler and Weizsaecker (2004) also apply quantal response equilibrium to a herding game.

3 Characterizing equilibrium

3.1 Monotonicity properties

Intuitively, we expect more attacks when the aggregate state Θ is bad. Thus a low signal x not only suggests that Θ is probably low, but also that the fraction attacking is likely to be high. Likewise, observing a high fraction attacking m/n both suggests that α_I is high and that Θ is low. In other words, in a well-behaved equilibrium, we might expect players' inferences Π to be monotonic in the following sense.

> **Definition.** The probability assessment Π exhibits **monotonic inferences** if the expectation $E_\Pi(\Theta|n, m, x)$ is increasing in x and decreasing in m, and the expectation $E_\Pi(\alpha_I|n, m, x)$ is increasing in m and decreasing in x.

The net benefit from attacking, $R(\Theta)\mathbf{1}_{\alpha_I \geq a(\Theta)} - t$, increases in α_I and decreases in Θ. Therefore, if players make monotonic inferences, their expected net payoff from attacking also varies monotonically with their observations. That is,

$$E_\Pi \left[R(\Theta)\mathbf{1}_{\alpha_I \geq a(\Theta)} | n, m, x, \eta = 1 \right] - t$$

is increasing in m and decreasing in x. Thus, when inferences are monotonic, players should choose monotonic strategies, in the following sense.

> **Definition.** The probability of attacking $y(n, m, x)$ is a **doubly monotonic strategy** if $\frac{\partial y}{\partial m} \geq 0$ and $\frac{\partial y}{\partial x} \leq 0$.

Monotonic inferences imply doubly monotonic strategies both in the case of fully rational behavior and in the more general case of logit behavior. In the fully rational case, where the probability of attacking jumps from zero to one exactly at the point of indifference, a doubly monotonic strategy can be called a **double threshold strategy**. That is, it implies thresholds both in the x and m directions: for each n and m there is a threshold $\tau(n, m)$ such that $y(n, m, x) = 1$ when $x \leq \tau(n, m)$, and is zero otherwise; while for each n and x, there is a threshold $\mu(n, x)$ such that $y(n, m, x) = 1$ when $m \geq \mu(n, x)$, and is zero otherwise.

If agents play doubly monotonic strategies, then for each player i, the conditional probability of attacking T_i is increasing in α_{i-1} (because higher α_{i-1} makes a higher m_i more likely) and is decreasing in Θ (because higher Θ makes a higher x_i more likely). Therefore, by induction over the functions T_i, we conclude that for each i, higher states Θ make a large number of attackers α_i less likely. This holds at the first step α_1 because a higher Θ makes a higher x_1 more likely, which discourages attacks. It holds at all later steps α_i because a higher Θ makes a higher x_i more likely, and because if Θ is high then α_{i-1} and therefore m_i are likely to be lower. Both these factors discourage attacks at each step i. That is, for any α, if $\theta_a > \theta_b$, then for any i,

$$\hat{P}_i(\alpha|\theta_a) \equiv prob(\alpha_i \leq \alpha|\Theta = \theta_a) \geq prob(\alpha_i \leq \alpha|\Theta = \theta_b) \equiv \hat{P}_i(\alpha|\theta_b) \tag{16}$$

Equivalently, this says that if $\theta_a > \theta_b$, then $\hat{P}_i(\alpha|\theta_b)$ first-order stochastically dominates $\hat{P}_i(\alpha|\theta_a)$.

This statement characterizes the distribution of intermediate outcomes α_i and the aggregate outcome α_I conditional on the aggregate fundamental Θ. But players can also draw analogous conclusions on the basis of their own information sets. That is, if all players other than i follow a doubly monotonic strategy, then agent i knows that $\hat{P}_{i-1}(\alpha|\theta_b)$ first-order stochastically dominates $\hat{P}_{i-1}(\alpha|\theta_a)$ whenever $\theta_a > \theta_b$. Therefore, if agent i observes a high x_i or low m_i, she should conclude that Θ is likely to be high and α_{i-1} is likely to be low. Conditional on any given value of her choice η_i, agent i should therefore also expect that the final outcome α_I is more likely to be low. But this is equivalent to saying that $E_\Pi(\Theta|n, m, x)$ is increasing in x and decreasing in m, and that $E_\Pi(\alpha_I|n, m, x)$ is increasing in m and decreasing in x: in other words, player i should make monotonic inferences.

This brings us back to the assumption that started off our chain of reasoning. Therefore, although we have not proved that equilibrium must necessarily involve monotonic inferences, we have listed a number of additional properties that must hold if and only if herding equilibrium exhibits monotonic inferences. The first proposition summarizes these findings.

Proposition 1. A herding equilibrium which has any one of the following properties has all of them.
a. The conditional distribution Π exhibits monotonic inferences.
b. Agents play a doubly monotonic strategy y.
c. For each i, mapping T_i is increasing in α_{i-1} and Θ.
d. For each i, the distribution $\hat{P}_i(\alpha|\theta')$ first-order stochastically dominates $\hat{P}_i(\alpha|\theta)$ if $\theta > \theta'$.

When we restrict consideration to threshold strategies, we will call an equilibrium that has these properties a **double threshold herding equilibrium**. In section 3.3, we will demonstrate by numerical construction that equilibria of this type exist. Furthermore, in section 4.1, we will estimate the strategy being used by our experimental subjects. The signs on the point estimates of our coefficients are always consistent with the assumption that players are using doubly monotonic strategies.

3.2 Limiting results: sufficient conditions for bimodal outcomes

To further characterize the model's behavior, it is helpful to focus on the limiting case of a large number of players, $I = \infty$. This clears away the sampling noise associated with the finite I game, offering a sharper characterization of the distribution of aggregate outcomes. Equation (7), which shows that α_i is a "stochastic recursive algorithm", tells us that in the large numbers game, the only possible outcomes $\alpha_\infty \equiv \lim_{I\to\infty} \frac{1}{I}\sum_{i=1}^{I} \eta_i$ conditional on a given state θ must be "E-stable" points in the sense of Evans and Honkapohja (2001).[6] In terms of our previous notation, the E-stable interior points are simply the points where T crosses the 45^o line from above. At these points, the probability that a given player i attacks is equal to the fraction of players who have already attacked.[7] Corner solutions can also be E-stable: $\alpha_\infty = 0$ is a solution if T is zero at $\alpha = 0$, and $\alpha_\infty = 1$ is a solution if T is one at $\alpha = 1$.

[6]The dynamics of the $I = \infty$ case of a closely related game are analyzed in greater detail in Costain (2007). In particular, Prop. 3 of that paper demonstrates the role of E-stability.

[7]More precisely, the probability that player i attacks equals the fraction who have already attacked at the crossings of T_i. Ror large i, these are approximately the crossings of the limiting function T.

Thus, for a given aggregate state $\Theta = \theta$, and given an equilibrium strategy y^*, the set of possible outcomes α_∞ of the $I = \infty$ model is a set of discrete points: those points where $T(\alpha, \theta, y^*)$ crosses the 45^o line from above, plus any appropriate corners. There may be only one such point for each θ, which means that the aggregate outcome is a well-defined function $\alpha_\infty(\theta)$ of the aggregate state. The first-order stochastic dominance property of Prop. 1d implies that $\alpha'_\infty(\theta) < 0$: less players attack when the aggregate state is better, as shown in Fig. 1b. But there may also be multiple crossings and/or corners for some θ, which gives rise to an equilibrium correspondence like that in Fig. 1a. In this case, multiple aggregate outcomes occur with positive probability at values of θ for which T has multiple crossings. The monotonicity properties of T with respect to α and θ imply that the branches of the outcome correspondence behave as Fig. 1a indicates: each branch of the correspondence is downward sloping, and the lowest point on a higher branch is higher than the highest point on a lower branch.

When do these multiple branches of the outcome correspondence arise? Given Assumption 1, which says that the full-information model has a tripartite classification of states, we can show that the $I = \infty$ herding model also has a middle range with multiple outcomes as long as N^{max} is sufficiently large and players are sufficiently rational. (We will state this fact for logit herding equilibrium, which implies that it is also true for rational herding equilibrium, since this is just logit with $\lambda = 0$.) Intuitively, if there are many players, who mostly observe many other choices, then most players will be able to guess the aggregate outcome. Therefore they should follow the crowd as if they had full information about others' decisions.

Proposition 2. Let $I = \infty$ and fix q. For sufficiently large N^{max}, in any doubly monotonic logit herding equilibrium y_λ with sufficiently small λ, there exists a nonempty interval of fundamentals (θ_*, θ^*) where the conditional distribution of aggregate outcomes $\hat{P}(\alpha_\infty|\theta, y_\lambda)$ places positive probability on at least two values of α_∞ for each $\theta \in (\theta_*, \theta^*)$.

Proof. See Appendix C.

The outcome distribution mentioned in Prop. 2 resembles that implied by a full-information model with multiple equilibria, like Obstfeld (1996), in which all players' choices depend on an exogenously-imposed, extrinsic "sunspot". In herding equilibrium, though, the first few players' actions depend mostly on their own private information; thereafter, many players can make their own decisions by simply following the crowd. In this sense, the herding equilibrium endogenously creates a "sunspot", which is just the consensus action of preceding players. The proposition shows that a region of unpredictable outcomes exists whenever the parameters permit sunspot equilibria in the full information game, and most players are sufficiently rational and observe enough previous actions. But in spite of the unpredictability of the actual outcome, players' double threshold strategies ensure that the distribution over aggregate outcomes is well behaved: bad fundamentals make it more likely that many players attack and therefore also more likely that the attack is successful. By contrast, a full-information sunspot equilibrium need not have these intuitively reasonable comparative statics properties, which is one of MS98's main criticisms of analyses based on sunspot equilibria.

3.3 Numerical results: comparative statics properties

When the number of players is finite, the set of possible outcomes is no longer limited to a small number of discrete points, because sampling error will spread out the possible realizations of α_I. Nonetheless, we will now show by simulation that the results are qualitatively quite similar with finite I. Intuitively, players' incentives to learn from others and to follow the crowd are not very different when I is large and finite from the $I = \infty$ case. Thus the functions T_i are quantitatively similar, causing the outcomes of the model with sufficiently large I to be tightly clustered around the outcomes of the infinite-player model. Thus if the $I = \infty$ model has a single possible outcome, then the finite-I model should have a unimodal distribution $\hat{p}_I$ over α for each θ. If T has two stable crossings over some range of θ for the the $I = \infty$ model, then there should be a strongly bimodal distribution of outcomes α_I conditional on roughly the same θ for large but finite I.

Results for some simulations with finite I are shown in Figure 2. In these examples we choose parameters to guarantee the existence of dominance regions at both ends of support of aggregate fundamentals, as in Assumption 1. We assume Θ lies on a grid Γ_θ from 15 to 85, and that signals x are drawn from $\Theta - 15$ to $\Theta + 15$. We assume the hurdle function $a(\Theta) = \frac{\Theta}{40} - \frac{3}{4}$, and speculation payoff function $R(\Theta) = 100 - \Theta$, and transactions cost $t = 30$. With these parameters, the full-information model would have multiple equilibria from $\theta_m = 30 + 40/I$ to $\theta^m = 70$.

The first row of the figure shows the probability distribution implied by the model for finite games with $I = 8$, 12, and 16. The graphs refer to a logit equilibrium y_λ with an intermediate level of rationality ($\lambda = 6.4$), setting $q = 1$, so that all previous actions are observed. The graphs show the contour lines of the conditional probability function $\hat{p}_I(\alpha|\theta, y_\lambda)$. What can see from the contour lines is a clear "tripartite classification of fundamentals". There is an interval of sufficiently bad aggregate states, up to almost $\theta = 40$, where players almost always attack. There is also a region of sufficiently good states, starting around $\Theta = 70$, where players almost never attack. However, over the middle range, the distribution of aggregate outcomes is sharply bimodal in the α direction.

Over the intermediate range of aggregate states, most outcomes take one of two forms: most players attack, or most players do not attack. Thus, there is not a unique relation between the state and the outcome. Nonetheless, the comparative statics of the distribution of outcomes with respect to the aggregate state are well-defined and well-behaved: the contour lines make clear that the probability of an outcome in which most agents attack is decreasing in Θ. In addition, unlike the $I = \infty$ model, it does occasionally happen that the fraction attacking is intermediate. As intuition would suggest, outcomes with roughly 50% attacking are more common when the total number of players is small. That is, while "following the crowd" is not a meaningful strategy for the

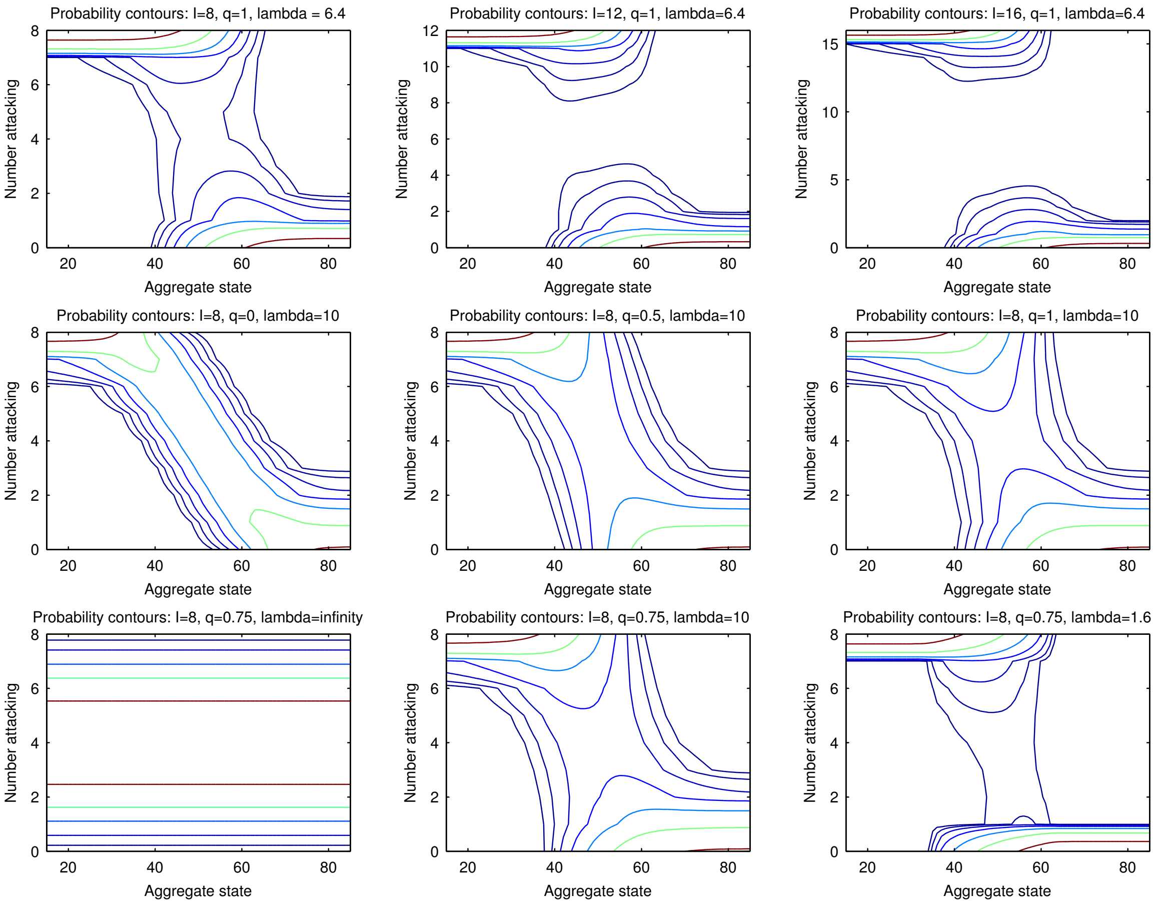

Figure 2: Contours of outcome probabilities $\hat{p}_I$: 1%, 2%, 4%, 8%, 16%, 32%, and 64% contours

first few agents, with sufficiently many players the aggregate outcome eventually snowballs towards all attacking, or none attacking, making the distribution more strongly bimodal with twelve or sixteen players than with eight.

Besides assuming an infinite number of players, the proof of a sharply bimodal outcome distribution in Prop. 2 relied on two additional assumptions: observations of a sufficiently large number of previous choices, and fully rational behavior. Relaxing these assumptions can undo our results in Prop. 2. In particular, there are obviously some circumstances in which our model generates a unique outcome as a function of fundamentals, since it nests the model of Morris and Shin (1998) when no previous choices are observed. However, our model also generates a variety of intermediate cases, as Fig. 2 illustrates.

The second row of Fig. 2 shows how the equilibrium changes with the number of previous observations. The three graphs show logit equilibria with 8 players and $\lambda = 10$, for different values of q. The graph on the left is the case $q = 0$, with no observations of previous actions, which makes our model equivalent to that of MS98 and HNO04. In this case, we see a unimodal distribution of outcomes conditional on any given fundamental θ: since players cannot condition on the choices of others, herding cannot get started. The next two graphs show that as we increase q, multiple outcomes gradually become possible. In the middle graph, with $q = 0.5$, we already see a tripartite classification of fundamentals, with almost all attacks on the left, almost no attacks on the right, and with all possible outcomes (from 0 to 8 attacks) occurring with probability above 1% in the middle region from roughly 42 to 58. In the graph on the right, with $q = 1$, players are more able to coordinate their actions, so the distribution becomes bimodal in the middle range.

Finally, the last row of Fig. 2 shows the effect of the level of rationality. (All equilibria shown assume 8 players and q=0.75.) With $\lambda = \infty$ (left graph) attacks occur randomly, with probability 0.5. Therefore the number attacking is symmetrically distributed around 4, regardless of θ. With an intermediate level of rationality, like $\lambda = 10$, coordination is already strong enough so that all outcomes occur from roughly $\theta = 38$ to $\theta = 61$, and over much of this range the distribution of α_I is bimodal. As rationality increases, to $\lambda = 1.6$ (which is almost indistinguishable from the fully rational case $\lambda = 0$), the bimodality becomes much stronger. Thus in this case, over the middle range where multiplicity occurs, almost all outcomes are predicted to have 0, 1, 7, or 8 attacks.

In summary, our model predicts a tripartite classification of fundamentals, with multiple outcomes in the middle region, when parameters are chosen so that (1) the full-information game also has a tripartite classification of fundamentals; (2) the number of players is sufficiently large; (3) players observe a sufficiently large number of previous actions; (4) players are sufficiently rational. As we reduce the number of players, or the number of observations, or the level of rationality, coordination becomes weaker. Typically this implies first that the sharp bimodality of the middle interval gradually disappears, and eventually that the multiplicity disappears too.

4 Experimental results

Our experimental sessions had eight or twelve participants. They interacted via a computer network, using z-Tree software (Fischbacher 1999). In sessions with eight participants, nine rounds were played; when we had twelve participants, only seven rounds were played. In each round, all participants played our game eight times (in parallel) so that each session yielded 72 or 56 aggregate outcomes of our game and 72*8=576 or 56*12=672 observations of individual decisions.

Appendix A translates the instructions we handed out to participants in one of our sessions at Univ. Carlos III in Madrid. In our sessions, play passed through a series of decision steps, in each of which a participant made (at most) one choice. When a participant was required to choose, the computer informed her of her information set. Specifically, the computer screen displayed her signal x, and some information on previous players' decisions: the total number of observations n, the number attacking, m, and the number not attacking, $n-m$. In some sessions, players knew their position in the sequence, because they had exactly one choice in each decision step (so they knew they were first in one sequence, second in another, etc.) In other sessions, we made the number of decision steps substantially larger than the number of players, and started some sequences later than others, so that players could not infer their positions from the timing of their choices. However, there was no noticeable difference between the sessions in which positions were known and those in which they were unknown.

Figure 3 shows the results from four of our sessions, using parameters like those in Sec. 3.3. We drew Θ from a grid over $[15, 85]$, and x from a grid over

$[\Theta - 15, \Theta + 15]$. We imposed the attack payoff function $R(\Theta) = 100 - \Theta$, hurdle function $a(\Theta) = \frac{\Theta}{40} - \frac{3}{4}$, and transactions cost $t = 30$.[8] The figure compares four sessions in which the fraction of previous plays observed was $q = 0$, 0.5, 0.75, or 1. (In each case, the session shown is the first in which we ran that particular parameter configuration.) For each session, the figure plots the contour lines of the conditional probability function $\hat{p}_8(\alpha|\theta, y^*)$, calculated numerically from the rational herding equilibrium.[9] As in Fig. 2, the contour lines shown represent 1%, 2%, 4%, 8%, 16%, 32%, and 64% probability levels. The actual observations from the experiment are plotted as stars, superimposed on the theoretical predictions. The data shown correspond to the last six of the nine rounds played, in order to remove any transitional dynamics caused by learning.[10]

The results of session 4, in which previous actions are unobserved, so that choices are logically simultaneous even if they do not occur at exactly the same time, are shown in the first panel. Our theory predicts a unimodal distribution in this case. The experimental results (the stars in the graph) show that from $\theta = 35$ to $\theta = 65$, the fraction attacking decreases gradually, smoothly, and almost linearly from one to zero. Thus the experimental distribution is always clearly unimodal conditional on θ, and it is also close to the distribution $\hat{p}_8$

[8]HNO04 found that subjects understood the game better with the definition of the aggregate state reversed so that $a' < 0$ and $R' > 0$. Therefore they described the model in terms of the transformed state $Y \equiv \Theta - 100$. Also, they wrote the hurdle function as a number of players rather than as a fraction. We follow their conventions. So when describing the game to subjects, we state that the payoff is $R(Y) = Y$, and that the hurdle function is $Ia(Y) = 8(\frac{100-Y}{40} - 3/4) = 14 - 0.2Y$ or $Ia(Y) = 12(\frac{100-Y}{40} - 3/4) = 21 - 0.3Y$, depending on the number of players. We give a neutral description of the game, calling the choices A and B, and making no reference to currency speculation.

[9]Actually, the figures show a logit equilibrium with $\lambda = 1$. Previous versions of the paper reported numerical simulations of a rational herding equilibrium, imposing threshold form; these results are available on request. However, that procedure gave virtually identical results, so rather than run two separate sets of simulation routines, we now report a logit equilibrium with low λ instead of the rational herding equilibrium.

[10]Our impression is that players are initially mostly predisposed to follow their own signals, but then learn to follow others more over the first few rounds. The learning effects appear small, though: the graphs are quite similar when we include all rounds.

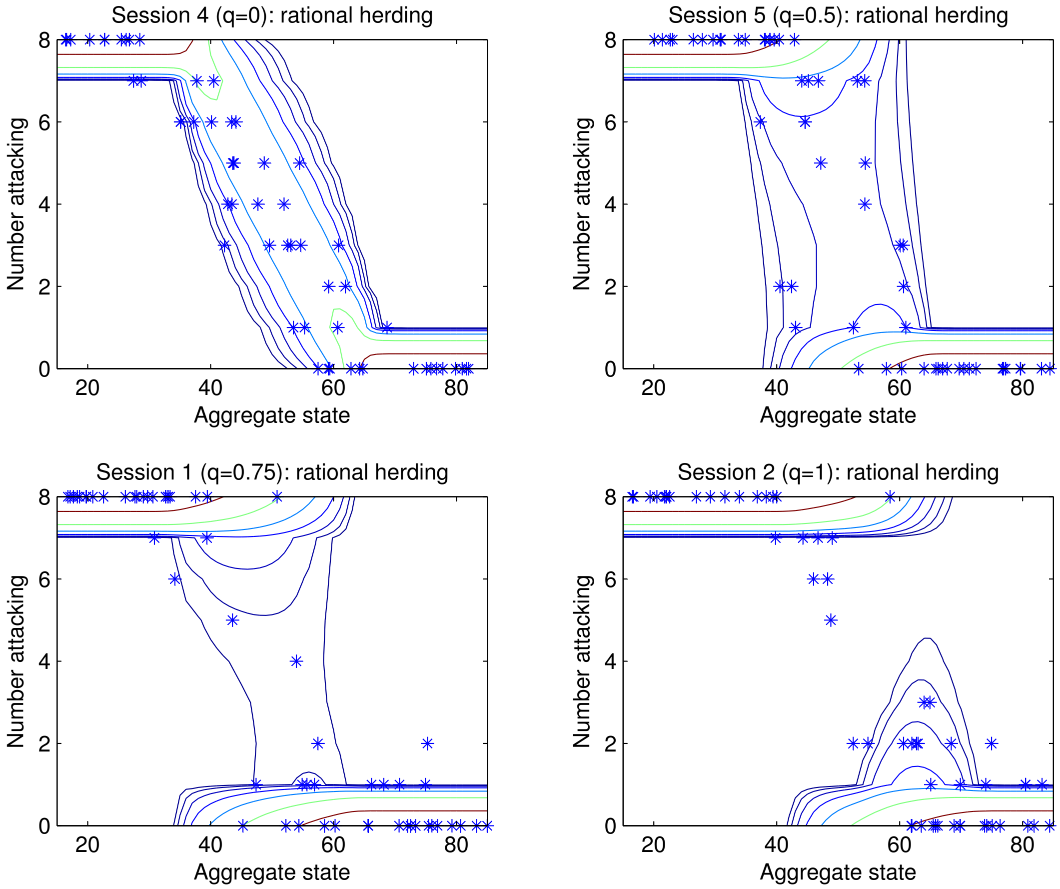

Figure 3: Contours of outcome probabilities $\hat{p}_I$ in rational herding equilibrium, with experimental observations.

implied by theory (illustrated by the contour lines). These results resemble those found by HNO04.

In the other three panels, though, the experimental results are distinctly different. In particular, we see more coordination, with outcomes of 0, 1, 7, or 8 attackers far more common than intermediate outcomes. For low θ, almost everyone attacks, and for high θ almost no one does. But there is not an intermediate range with intermediate numbers attacking. Outcomes with 3, 4, or 5 attackers are rare, especially in the treatments with $q = 0.75$ and $q = 1$, which is consistent with the theoretical predictions shown by the contour lines.

Thus, even at intermediate values of fundamentals, the outcomes tend to lie at the extremes. Moreover, in all the $q \geq 0.5$ panels, some of the extreme outcomes overlap in the "wrong" direction relative to fundamentals. In the $q = 0.5$ treatment, for example, we observe attacks by seven players twice near $\theta = 55$, but another case with only one attack when θ was only slightly over 40. Both the finding that extreme outcomes are more common than intermediate outcomes, and the finding that some of the outcomes that coordinate on attacking occur at higher fundamentals than others that coordinate on not attacking, are consistent with the strongly bimodal outcome distribution implied by the theory. So instead of interpreting these extreme outcomes vacuously as "outliers", it makes more sense to interpret them as examples of herding behavior, in which the choices of the first few agents serve to coordinate the choices of following agents.

Of course, a few experimental observations in these sessions are well outside the range predicted by the theoretical model. But given that this version of the model makes no allowance at all for human error, the fit is remarkably good. We now go on to consider all our 14 sessions, with 8 or 12 players, to see how these results hold up. Besides varying q from 0 to 0.5, 0.75, or 1, we also vary the support of the idiosyncratic signal, comparing $\theta \pm 15$ with $\theta \pm 8$. We also compare sessions in which players were informed of their position in the sequence, and sessions where they could only deduce their position from the number of previous choices observed. To summarize the results obtained across different sessions, we next estimate players' strategies in each session, and compare their implications.

4.1 Reduced form strategy estimates

A glance at Figure 3 suggests that the distribution of aggregate outcomes implied by our game differs greatly depending on whether or not previous actions are observed. But it is hard to draw firm conclusions from these graphs because of the relatively small number of aggregate outcomes observed— and the even smaller number of cases where the fundamentals lie in the most interesting range— in each experimental session. Therefore, we next estimate the strategy used by our experimental subjects, allowing us to compare their behavior with the predictions of the model. We can also use the estimated strategy to calculate the distribution of aggregate outcomes associated with each experimental session, thus obtaining a clearer, more quantitative understanding of aggregate behavior.

In our fully-rational theoretical model, players attack if and only if their signal exceeds a threshold $\tau(n,m)$. Obviously, in the laboratory, play will be more random than this. This motivates the logit equilibrium defined in Sec. 2.4. In the next subsection, we will estimate the logit equilibrium. But first, we estimate a simpler reduced-form choice model, in order to avoid imposing so much theoretical structure on players decision. Thus we model our experimental subjects' probability of attacking, conditional on their information sets, as the following linear-logistic function:

Probability of attacking:

$$prob(\eta_i = 1|x_i, n_i, m_i) \quad = \quad \frac{1}{1+\exp\{\gamma(n_i)[x_i - \delta(n_i) - \xi(n_i)(m_i/n_i - 0.5)]\}} \tag{17}$$

In this formula, the probability of attacking, conditional on the information set (x_i, n_i, m_i), is a number between 0 and 1. This probability depends on three parameters, $\gamma(n)$, $\delta(n)$, and $\xi(n)$, which may in general vary with the number of observations n. The formula is written so that the estimates of $\gamma(n)$, $\delta(n)$, and $\xi(n)$ ought to be positive. As long as γ is positive, the probability of attacking is a decreasing function of the signal x.

The parameters in the formula have straightforward interpretations. Parameter δ can be seen as an unconditional threshold signal. That is, if the observations of other actions are uninformative ($n_i = 0$ or $m_i/n_i = 0.5$), then the player attacks with probability 0.5 when she observes the signal $x_i = \delta(n_i)$.[11] When observations of other actions are informative ($m_i/n_i \neq 0.5$), then we can interpret $\delta(n) + \xi(n)(m_i/n_i - 0.5)$ as the conditional threshold signal at

[11]For the case $n = 0$, if the coefficients are allowed to vary with n, then $\delta(0)$ and $\xi(0)$ are not separately identified; in this case we preserve the interpretation of $\delta(0)$ as an unconditional threshold signal by setting $\xi(0) = 0$. On the other hand, when we do not allow the coefficients to vary with n, we preserve the interpretation of δ as an unconditional threshold signal by setting $m_i/n_i \equiv 0.5$ whenever $n_i = 0$.

which the probability of attacking is one half. Positive ξ means the threshold signal for attacking is higher when a higher proportion of attacks is observed: that is, players are more likely to attack when they observe others attacking. Finally, parameter $\gamma(n)$ indexes the precision of individual behavior. If $\gamma(n) = 0$, then behavior is always totally random; that is, the probability of attacking is always 0.5. If $\gamma(n) = \infty$, then there is no randomness in individual behavior, conditional on the individual information set, as in our fully rational theoretical model.

Using experimental data on signals x_i, observations n_i and m_i, and decisions η_i, we estimate this strategy by maximum likelihood for each experimental session separately. Using the estimated strategy y, we can then calculate the implied conditional outcome probabilities $\hat{p}_I(\alpha_I|\Theta, y)$. We can also use it to simulate additional experimental sessions. By repeating our maximum likelihood estimation on these artificial experimental data sets, we can obtain bootstrap estimates of the confidence intervals on our parameter estimates. Importantly, we can also estimate confidence intervals for any statistics of the aggregate distribution $\hat{p}_I$ that interest us.

Our parameter estimates for each session are displayed in Tables 5-4. To eliminate transitional dynamics in players' strategies due to learning, we discard the first two rounds of each session before estimation. We allow the coefficients γ, δ, and ξ to vary with n, though the number of data points is too small (especially for high n) to permit fully arbitrary variation with n. However, this appears not to be a problem. The likelihood of our sample improves greatly if we allow ξ to change between low and high n. In sessions with eight participants we estimate $\xi_{0-2} \equiv \xi(n),\ n \in \{0, 1, 2\}$ separately from $\xi_{3-7} \equiv \xi(n),\ n \in \{3, 4, 5, 6, 7\}$; with twelve players we estimate $\xi_{0-3} \equiv \xi(n),\ n \in \{0, 1, 2, 3\}$ separately from $\xi_{4-11} \equiv \xi(n),\ n \in \{4, 5, 6, 7, 8, 9, 10, 11\}$. Similarly, we check whether the likelihood is improved by estimating δ or γ (see Tables 2 and 4) separately for low and high n, or by breaking the coefficients at more than one n, but the changes rarely yield significantly different point estimates or a significant improvement in likelihood.

Estimates of the unconditional threshold signal δ range from 46.77 to 58.81; estimates over 50 imply that players attack slightly more often than they refrain from attacking. Estimates of $\xi(n)$ are robustly found to be larger when n is larger. This makes sense: it means that subjects react more to the observed fraction attacking if they have observed many actions. For example, the coefficient estimates for session F3 imply that the difference between the threshold when only attacks are observed, and when no attacks are observed, is 7.136 for $n \leq 2$, but rises significantly to 34.36 for $n \geq 3$. The pattern for γ is less clear, but in those sessions where we find significant differences, γ decreases with n.

We also report two statistics describing the distribution of aggregate outcomes. We search for values of the aggregate state such that the conditional distribution of α_I, the aggregate fraction attacking, is bimodal. As our model predicts, no bimodal region is detected for the sessions in which previous actions are unobserved, and generally the estimated bimodal region is larger when more predecessors are observed. But unlike the rational herding model, in which players act without error, only very mild bimodality is observed in the experiment. The widest estimated region of bimodality is only 3.5 units wide, and the bimodal regions are (individually) never statistically significant.

We also report the width of the interval of fundamentals in which both the lowest and the highest values of α_I occur with at least 1% probability. Again, as we would expect from our model, this region is wider when more previous actions are observed. When no previous actions are observed, no aggregate states are found in which both $\alpha_I = 0$ and $\alpha_I = 1$ occur. But when we set $q = 0.5$, so that half of the preceding actions are observed, only one out of 5 sessions fails to detect a region in which both extreme outcomes occur. For the other four sessions with q=0.5, the width of the region of multiple outcomes is estimated to be 1.5, 13.5, 6.0, or 8.5; the last three are significant. For $q = 0.75$ and $q = 1$, we always detect regions in which both extreme outcomes occur with at least 1% probability. The estimated widths of the region of multiple outcomes for these sessions are 13.5, 13.0, 9.0, 18.0, 13.0, 4.5, and 9.0, all of which are significantly different from zero. Thus sessions with $q \geq 0.5$ usually exhibit a "tripartite classification" of aggregate fundamentals, with a middle region where multiple outcomes can occur.

4.2 Estimating logit equilibrium

We next pursue a more structural empirical strategy, by directly estimating the logit equilibrium. There are a number of important advantages to this estimation procedure. First, it imposes a lot of structure on the predicted outcome distribution, since the logit equilibrium has only one free parameter: namely, the rationality index λ.

Estimating logit equilibrium is computationally intensive, because calculating a single equilibrium is already a challenging problem, and for the estimation we must do this repeatedly across values of λ. Obviously, though, the fact that we only need to search over one free parameter enormously simplifies the estimation. Moreover, it also simplifies the calculation of each equilibrium. We know immediately that the logit equilibrium associated with $\lambda = \infty$ is $y(n, m, x) = 0.5$ for all (n, m, x). We can then use this equilibrium as a starting guess for the fixed point calculation for some large but finite value of λ. Also, for sufficiently large λ, players respond weakly to any change in payoffs, which means they respond weakly to the strategies of other players, which means that the mapping $B\lambda$ is a contraction map, and that logit equilibrium is unique.

Therefore the natural way to compute the logit equilibria is to calculate them sequentially, for a decreasing sequence of values of λ, using the equilibrium associated with λ_j as a starting point for the fixed point calculation for $\lambda_{j+1} < \lambda_j$. Since the equilibrium $y_{\lambda_{j+1}}$ should be close to the equilibrium λ_j, we can numerically calculate the dominant eigenvalue of each best response mapping B_{λ_j} by looking at the local rate of convergence. This allows us to check whether each equilibrium y_{λ_j} is locally unique (though global uniqueness is not guaranteed by this routine). If local uniqueness breaks down, we propose to select an equilibrium numerically by imposing partial adjustment— that is, by finding a fixed point of $y' = \alpha B_\lambda y + (1 - \alpha) y$ instead of $y' = B_\lambda y$, for some

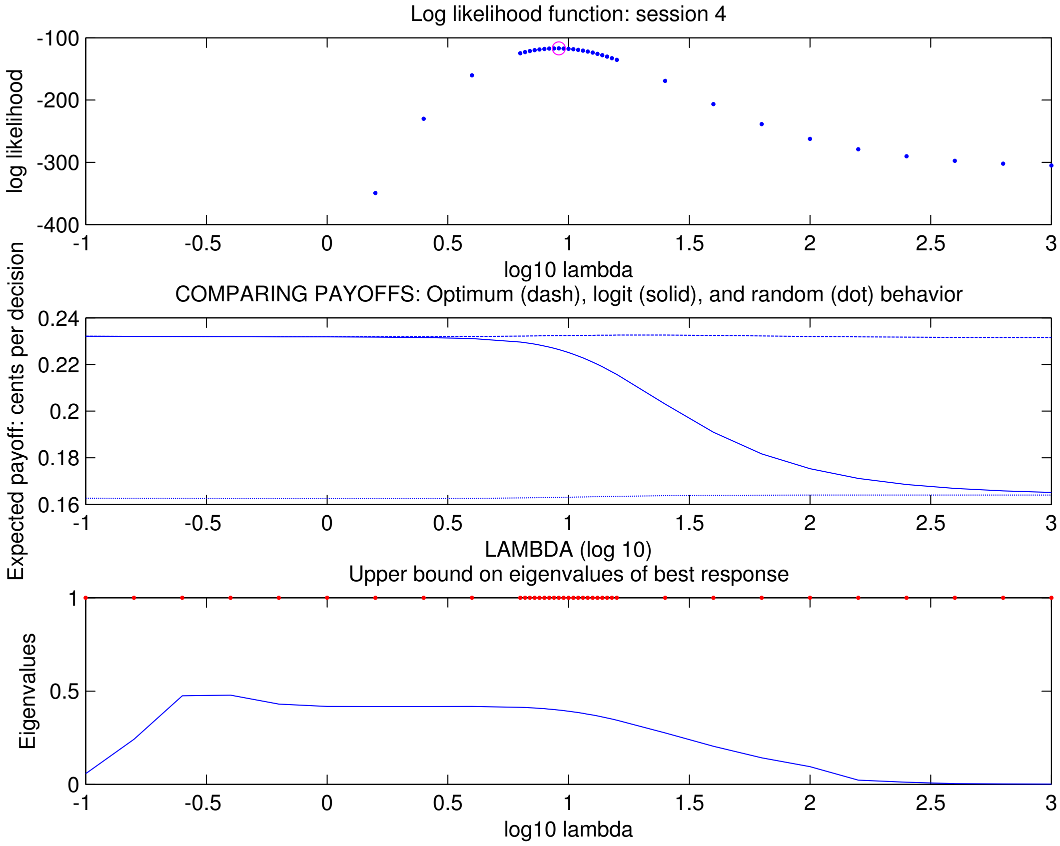

Figure 4: Log likelihood and diagnostics, Session 4 (q=0)

$0 < \alpha < 1$. Thus we can compute a whole spectrum of logit equilibria from $\lambda = \infty$ to $\lambda \approx 0$, which is arbitrarily close to a rational herding equilibrium.

Once we have computed the system of logit equilibria (which is slow, but not much slower than computing a single rational herding equilibrium from arbitrary initial conditions), we can quickly and easily evaluate the likelihood of each equilibrium, given our experimental data. Figures 4 and 5 show the log likelihood function, and associated diagnostics, for sessions 4 and 1, which are the $q = 0$ and $q = 0.75$ sessions already seen in Fig. 3. The log likelihood functions are well behaved, and we estimate $\lambda = 9.12$ for session 4, and $\lambda = 9.55$ for session 1. The standard errors on these estimates are 0.13 and 0.11, so the estimated λ differs slightly but significantly between these two sessions. Results for these and other logit equilibrium estimates are given in Table 5.

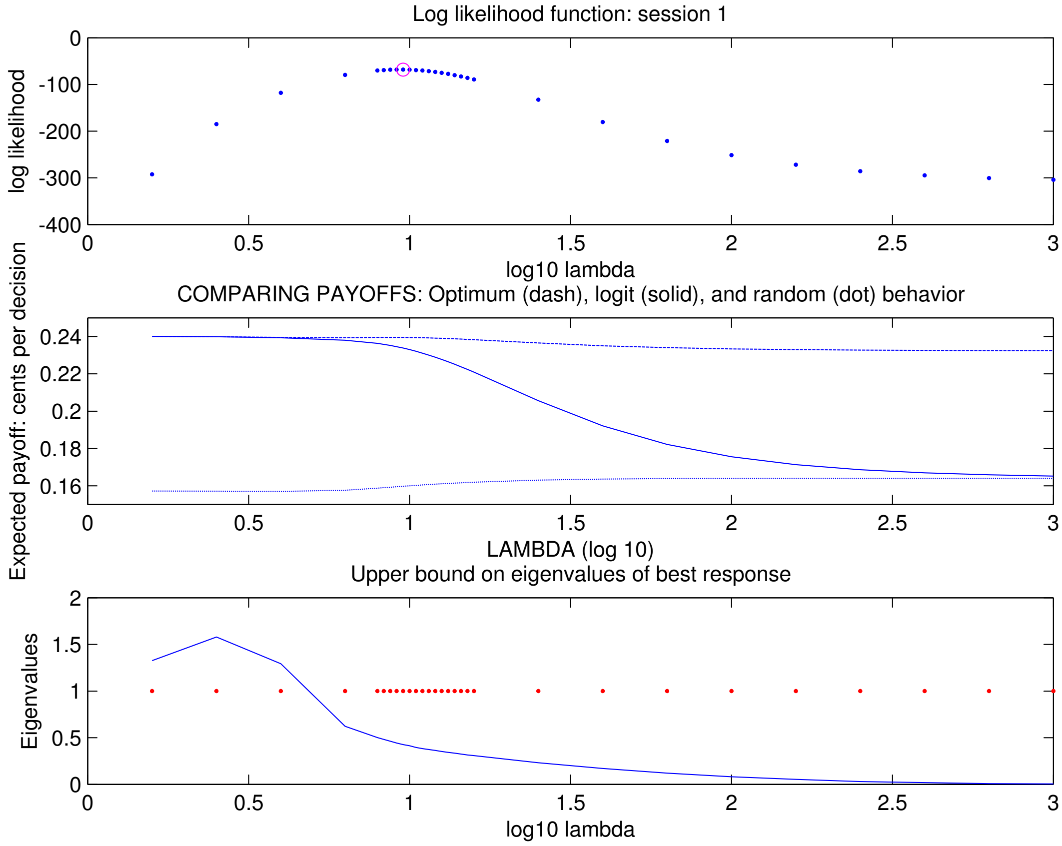

Figure 5: Log likelihood and diagnostics, Session 1 (q=0.75)

Given the estimated λ, we can graph the distribution of aggregate outcomes in the estimated equilibrium, as before. As Figure 6 shows, the equilibrium appears to predict the distribution of experimental outcomes very well. Like the fully rational case, the estimated $q = 0$ equilibrium is unimodal, with the attack fraction decreasing smoothly from $\alpha \approx 1$ to $\alpha \approx 0$. There are no values of the fundamental θ at which both extreme outcomes occur with probability 1% or more.

On the other hand, for all $q \geq 0.5$, the estimated logit equilibrium exhibits a tripartite classification of states. The width of the region of multiplicity (defined as at least 1% probability of each extreme outcome) is similar to that in the fully rational equilibria of Fig. 3. However, the fit of the estimated logit equilibria is much better than that of the corresponding rational herding

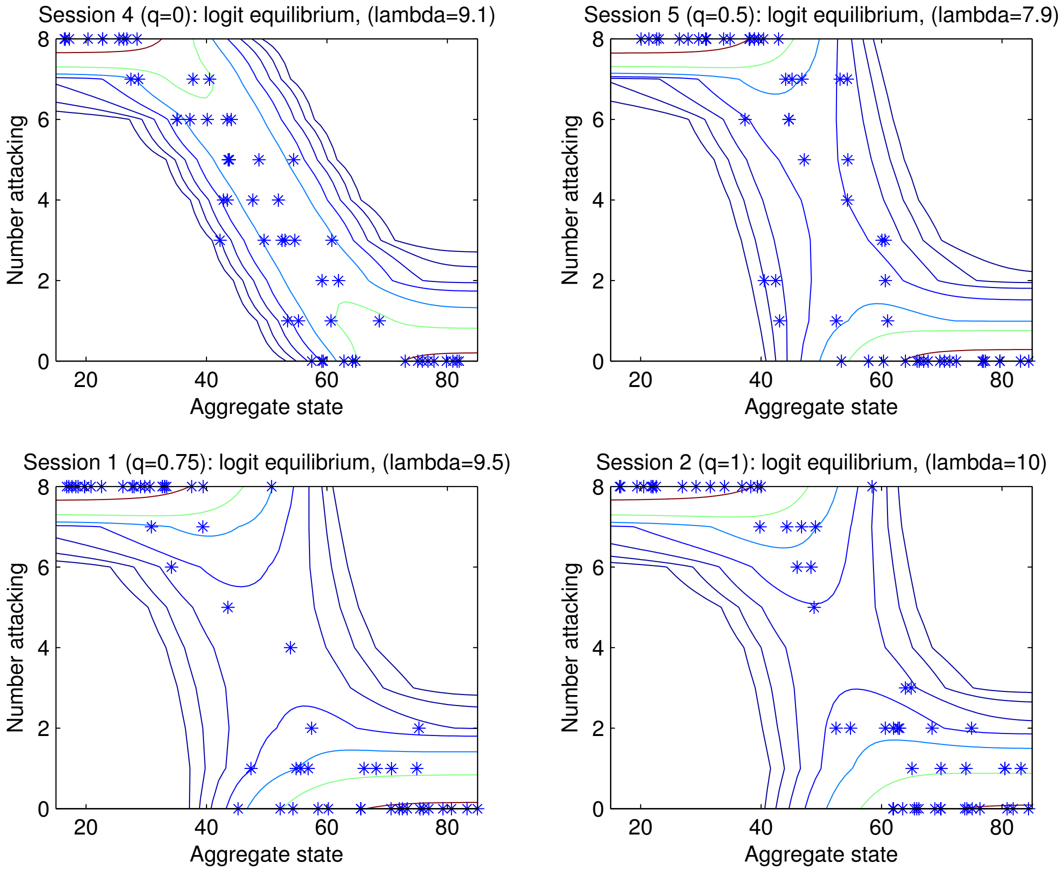

Figure 6: Contours of outcome probabilities $\hat{p}_I$ in estimated logit equilibrium, with experimental observations.

equilibria — there are no outliers with respect to the estimated equilibria. This is because logit equilibrium allows for some errors, thus making the distribution less sharply bimodal than the rational herding equilibrium. Still, it is more bimodal than the distributions implied by our reduced form strategy estimates: even with $q = 0.5$, the predicted shape of the distribution is weakly bimodal. As expected, it is even more strongly so for $q = 0.75$ and $q = 1$.

Of course, one way our estimation routine might account for all the observations, without outliers, is that we could estimate such a low degree of rationality that any outcome is possible. But this is clearly not what is happening, because our subjects' realized earnings show that they are actually highly rational. The second row of Figs. 4 and 5 shows the estimated payoffs in the logit equilibria associated with all values of λ. The middle curve in each graph is the unconditional expected payoff of a player in the estimated logit equilibrium. We compare this with the payoffs of the best response to the estimated equilibrium. This shown by the upper curve, which represents the payoff to a single perfectly rational player ($\lambda \approx 0$) who knows that her opponents use the equilibrium strategy y^*_λ. In other words, the top curve is the value of $y = B_0 y^*_\lambda$. Note that at the maximum likelihood estimators, the payoff of the estimated strategy is €0.226 per decision for the $q = 0$ session, and €0.234 for the $q = 0.75$ session. By way of comparison, the best response to these estimated equilibria would pay €0.232 per decision or €0.240 per decision, respectively. Thus even though our subjects' behavior is noticeably different from fully rational play, further improvement would only raise payoffs by around six tenths of a euro cent. Also, subjects earn much more than the value of random play, (that is, the payoff of attacking with probability 0.5 when everyone else plays the estimated strategy y^*_λ), which earns around sixteen euro cents per decision, as shown by the lower curve in the graph.

Finally, Figs. 4 and 5 also show an upper bound on the eigenvalues of the best response mapping at the logit equilibrium y_λ for each λ. Note that for session 4 ($q = 0$), the eigenvalue is less than 0.5 (in absolute value) at all the equilibria calculated, down to $\lambda = 0.1$, which is extremely close to full rationality. That is, all the $q = 0$ equilibria we calculated are locally unique. For session 1 (q=0.75), on the other hand, the eigenvalue bound does climb above 1 as rationality increases. Therefore, we cannot guarantee local uniqueness of equilibrium at all levels of rationality. Nonetheless, uniqueness only breaks down for $\lambda < 4$, substantially beyond the degree of rationality observed for our experimental subjects. The results are similar in all of the other sessions we have estimated. While the dominant eigenvalue of the best response mapping typically rises above one at sufficient rationality, this never occurs at levels of rationality observed in our experiment.

Therefore, our experimental results provide an interesting contrast in the conclusions they suggest about the theoretical issue of uniqueness of equilibrium, and the empirical issue of uniqueness of outcomes— two issues which cannot be addressed separately when a speculative attack is modelled as a simultaneous move global game. When we allow for players' imperfect rationality, equilibrium is unique at the levels of rationality estimated in the laboratory. Therefore the model provides unambiguous predictions about players' strategies and the resulting distribution of aggregate outcomes— and the predictions work very well in the laboratory. Nonetheless, what is predicted is the probability distribution over outcomes, not the outcome itself. When q is large, the predicted distribution implies that agents tend to coordinate randomly on one extreme of the outcome distribution or the other. Empirically, therefore, the model says there is no reason to expect a unique outcome as a function of fundamentals.

5 Conclusions

This paper reports an experiment on a speculative attack game identical to that in Morris and Shin (1998) and Heinemann et al. (2004), except that we allow for the possibility that some previous actions are unobserved. As our model predicts, we find that players not only condition their attack decisions on their private information, but also on their observations of previous actions. Players condition more strongly on the average previous action when the number of actions they observe is larger.

In our model, we show how this type of herding behavior affects the distribution of aggregate outcomes. In particular, we show sufficient conditions under which the aggregate outcome becomes unpredictable as a function of aggregate fundamentals. A region of fundamentals with unpredictable outcomes arises when a sufficiently large number of sufficiently rational agents play a coordination game which would have multiple equilibria under full information, and these players observe a sufficiently large number of previous actions. In this case, the private signals of the first few players can become decisive for the final outcome, as other players follow their lead.

The model nests a variety of cases in which the sharp type of outcome multiplicity from the infinite player game gradually breaks down. When the number of players is smaller, sampling error smoothes out the distribution of outcomes; and when players are less rational or have less information about previous plays, they coordinate their play less strongly with previous choices. All these effects reshape the distribution (continuously) in ways that tend to eliminate multiplicity. In particular, if we go to the "global games" case of perfectly simultaneous choice, coordination becomes impossible, and therefore variation in aggregate outcomes, if any, is due entirely to sampling error.

Our experiment allows us both to test our model and to see which of the special cases of our model best describes actual subjects' decisions. Observed play in the laboratory somewhat resembles the fully rational equilibrium for a finite number of players, but log likelihood sharply rejects full rationality. A boundedly rational version of our model, however, appears to fit our experimental observations extremely well. Our reduced form strategy estimates have all the right signs that our model implies. Our structural estimate of the logit equilibrium does an excellent job of predicting the distribution of aggregate outcomes, with a similar degree of rationality predicting outcome distributions under substantially different parameter configurations.

Even though full rationality is rejected, we calculate that individual players' losses due to irrationality are very small. The degree of rationality we estimate is sufficient to yield an outcome distribution with a "tripartite classification of fundamentals" in a game of eight or twelve players. Of course, as the model predicts, this depends on whether previous actions are observed. With no observations of previous actions, we obtain a unique outcome at any level of fundamentals, subject to some small-sample variation. With $q \geq 0.5$, we sometimes detect a significant middle range in which both extreme outcomes occur; with $q \geq 0.75$ we always detect a significant range of multiplicity. However, the level of rationality does not appear sufficient to produce bimodality of outcomes in the middle region of fundamentals, at least with the number of players we have been able to study in our laboratory work.

These results suggest a number of conclusions about the "global games" methodology, and its application to currency crises. It would definitely be wrong to conclude, on the basis of the "realism" of adding some private information to the model, that speculative attacks must have a uniquely defined outcome. Here we further increase "realism", by considering decisions that are not exactly simultaneous, and we find that a uniquely defined outcome no longer follows. But whether or not self-fulfilling speculative behavior often randomizes over aggregate outcomes in practice remains an empirical question, because it depends (among other things) on how large the number of significant players in financial markets is, and on the degree of information that these players have about others' moves when they make their decisions.

Methodologically, using a small amount of heterogeneity (such as private information) in order to smooth players' average responses enough that the theoretical model makes a well-defined prediction remains a very useful idea. But this paper suggests that other forms of heterogeneity may be more useful modeling devices for this purpose. In this paper, the noisiness of choices implied by the small amount of bounded rationality we observed in the laboratory was sufficient to guarantee uniqueness of our equilibrium (and to match our experimental observations). Using bounded rationality as a modeling trick to generate uniqueness allowed us to extend our speculative game to nonsimultaneous choice, which has proved challenging in models where the preferred modeling trick was instead private information. And extending to nonsimultaneous choice was crucial for us in characterizing what sorts of situations are likely to permit self-fulfilling speculative behavior to occur.

References

Alevy, Jonathan E.; Michael S. Haigh; and John List (2006), "Information Cascades: Evidence from an Experiment with Financial Market Professionals." NBER Working Paper #12767.

Anderson, Lisa R., and Charles A. Holt (1997), "Informational Cascades in the Laboratory." *American Economic Review* 87 (5), pp. 847-62.

Anderson, Lisa R., and Charles A. Holt (1998), "Informational Cascade Experiments." In *Handbook of Results in Experimental Economics*, Charles Plott, ed., forthcoming 2004, Elsevier Science Ltd.

Angeletos, George-Marios; Christian Hellwig; and Alessandro Pavan (2006), "Signaling in a Global Game: Coordination and Policy Traps." *Journal of Political Economy* 114 (3), pp. 452-84.

Angeletos, George-Marios, and Ivan Werning (2006), "Crises and Prices: Information Aggregation, Multiplicity, and Volatility." *American Economic Review* 96 (5), pp. 1720-36.

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

Bikhchandani, Sushil; David Hirshliefer, and Ivo Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy* 100 (5), pp. 992-1026.

Bikhchandani, Sushil; David Hirshliefer, and Ivo Welch (1996), "Informational Cascades and Rational Herding: An Annotated Bibliography." Available at: *http://welch.som.yale.edu/cascades/*

Bikhchandani, Sushil; David Hirshliefer, and Ivo Welch (1998), "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades." *Journal of Economic Perspectives* 12 (3) pp. 151-70.

Cabrales, Antonio; Rosemarie Nagel, and Roc Armenter (2007), "Equilibrium Selection through Incomplete Information in Coordination Games: an Experimental Study." Forthcoming, *Experimental Economics.* Also Univ. Pompeu Fabra Economics Working paper #601, 2002.

Carlsson, Hans and Eric van Damme (1993), "Global Games and Equilibrium Selection." *Econometrica* 61 pp. 989-1018.

Chamley, Christophe (2003A), "Dynamic Speculative Attacks." *American Economic Review* 93 (3), pp. 603-21.

Chamley, Christophe (2003B), *Rational Herds: Economic Models of Social Learning*, Cambridge University Press.

Chari, V.V., and Patrick Kehoe (2003), "Financial Crises as Herds: Overturning the Critiques." *Journal of Economic Theory* 119 (1), pp. 128-50.

Cheung, Yin-Wong, and Daniel Friedman (2005), "Speculative Attacks: A Laboratory Study in Continuous Time." UCSC Working Paper #606.

Cornand, Camille (2006), "Speculative Attacks and Information Structure: An Experimental Study." *Review of International Economics* 14 (5), pp. 797-817.

Costain, James (2007), "A Herding Perspective on Global Games and Multiplicity." *B.E. Journal of Theoretical Economics*, Vol. 7, Issue 1 (Contributions), Art. 22.

Dasgupta, Amil (2007), "Coordination and Delay in Global Games." Forthcoming, *Journal of Economic Theory.*

Drehmann, Mathias; Jorg Oechssler; and Andreas Roider (2005), "Herding with and without Payoff Externalities – An Internet Experiment." Manuscript, Univ. of Bonn.

Duffy, John, and Eric Fisher (2005), "Sunspots in the Laboratory." *American Economic Review* 95 (3), pp. 510-29.

Duffy, John, and Jack Ochs (2007), "Equilibrium Selection in Entry Games: An Experimental Study." Manuscript, Univ. of Pittsburgh.

Evans, George W., and Seppo Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.

Fischbacher, Urs (1999), "z-Tree: Zurich Toolbox for Readymade Economic Experiments— Experimenter's Manual." Working paper #21, Institute for Empirical Research in Economics, University of Zurich. Additional information available at *http://www.iew.unizh.ch/ztree.*

Frankel, David M., Stephen Morris, and Ady Pauzner (2003), "Equilibrium Selection in Global Games with Strategic Complementarities." *Journal of Economic Theory* 108, pp. 1-44.

Goeree, Jacob K., and Charles A. Holt (1999), "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory." *Proc. Natl. Acad. Sci. USA* 96, pp. 10564-67.

Heidhues, Paul, and Nicolas Melissas (2006), "Equilibria in a Dynamic Global Game: the Role of Cohort Effects." *Economic Theory* 28, pp. 531-57.

Heinemann, Frank; Rosemarie Nagel; and Peter Ockenfels (2004), "The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information." *Econometrica* 72 (5), pp. 1583-99.

Hellwig, Christian (2002), "Public Information, Private Information, and the Multiplicity of Equilibria in Coordination Games." *Journal of Economic Theory* 107, pp. 191-222.

Kubler, Dorothea, and Georg Weizsacker (2004), "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory." *Review of Economic Studies* 71, pp. 425-42.

McKelvey, Richard, and Thomas Palfrey (1995), "Quantal Response Equilibria for Normal Form Games." *Games and Economic Behavior* 10, pp. 6-38.

Morris, Stephen, and Hyun Song Shin (1998), "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88 (3), pp. 587-97.

Morris, Stephen, and Hyun Song Shin (2003), "Heterogeneity and Uniqueness in Interaction Games." Cowles Foundation Discussion Paper #1402, Yale Univ.

Obstfeld, Maurice (1996), "Models of Currency Crises with Self-Fulfilling Features." *European Economic Review* 40 (3-5), pp. 1037-47.

Schurchkov, Olga (2007), "Coordination and Learning in Dynamic Global Games: Experimental Evidence." Manuscript, MIT.

Vives, Xavier (1993), "How Fast Do Rational Agents Learn?" *Review of Economic Studies* 60, pp. 329-347.

Appendix A: Instructions for participants

The next three pages are an English translation of the instructions handed out to experimental subjects in our first session at the Univ. Carlos III in Madrid. The original instructions handed out in Spanish or German are available upon request.

General information
Thank you for participating in an economic experiment, in which you will have a chance to earn money. Please do not talk to the other participants from now on. If you have a question, please raise your hand, and one of the instructors will come to you.

You are one of **8 participants** who will interact with each other in the experiment. The rules are the same for all participants. The experiment consists of **9 rounds**. Each round consists of **8 independent situations** in which you must make a decision.

Decision situations
The important fact about each situation is a number, called Y, between 15 and 85, which will be chosen randomly by the experimenter's computer. This number will be the same for all participants. All numbers between 15 and 85 are equally probable. When you make your decision, you will **not know Y.**

However, even though you do not know Y, you will receive a **hint** about its value. This hint will be another random number, between Y-15 and Y+15. All numbers between Y-10 and Y+10 are equally probable. Every participant will receive a hint, but their hints will be independent, so they will not necessarily be equal.

After receiving your hint, you will choose one of **two actions**, **A or B**.

If you choose **action A**, you will receive a payoff of **30 points**. This payoff is the same in every situation, in every round, for every participant.

If you choose **action B**, you will receive a payoff of **Y points**, as long as a sufficiently large number of other participants choose B too. To be precise, action B will be successful--- that is, it will pay you Y points--- if the total number of participants choosing B is at least 14-Y/5. There is a table on the next page to clarify this formula. According to this formula, action B is more likely to be successful if Y is large, and/or if many participants choose B. But if action B is not successful, then anyone who chose B will receive a payoff of **0 points.**

In summary, action A always pays 30 points. Action B can pay Y points (if many players choose B), or it can pay 0 points (if few players choose B). Keep in mind that you will *not* know exactly how big Y is.

(Please note: the software used in this experiment indicates decimal numbers with a decimal point, instead of a comma.)

If the unknown number *Y* (which is at least 15, and at most 85) lies in the interval mentioned below,	then at least this number of participants must choose B in order for B to be successful, implying a payoff of *Y* points.
Y is less than 30:	B cannot be successful
Y is between 30.00 and 34.99:	8 participants must choose B
Y is between 35.00 and 39.99:	7 or more must choose B
Y is between 40.00 and 44.99:	6 or more must choose B
Y is between 45.00 and 49.99:	5 or more must choose B
Y is between 50.00 and 54.99:	4 or more must choose B
Y is between 55.00 and 59.99:	3 or more must choose B
Y is between 60.00 and 64.99:	2 or more must choose B
Y is 65 or greater:	If you choose B, then B will be successful

In each situation, all participants will make a choice. In each situation, they will choose one after another, but in each situation, the order of their choices is different. The number of participants who have made their choices before you, in a given situation, is equal to the number of decisions you have already made. For example, if you are making your fourth decision in a given round, you will know that three participants already made decisions in the situation you are currently facing.

When you need to make a decision in a given situation, some information about that situation will appear, in red, on your computer screen. You will see the numerical hint which tells you, roughly but not exactly, the value of Y. You may also see some information about the decisions made by other participants in the same situation.

The computer may, randomly, tell you some of the choices made in the same situation before your decision. To be precise, each previous decision has probability 0.75 of being revealed to you. Therefore, you might not receive any information about previous decisions, even though some other participants have already chosen.

Therefore, in general you will not observe all the choices made before your decide in a given situation. What you will know is how many previous choices the computer has revealed to you. You will know how many of these decisions were A, and how many were B. If you observe some previous choices, you will not know which participants made the choices you observe.

The random set of previous decisions that are revealed is different for each participant and for each situation. You will never know who has learned about your own decisions, or whether they have been revealed at all.

How the computers work
When the red information about a given situation appears on your screen, you should then choose one of the two options A or B. You will make your choice by clicking on

one of the two buttons labelled “A” and “B”. Next, you must click “OK” (at the bottom of the screen) to confirm your choice. You can change your choice before you click “OK”, but thereafter your choice is fixed.

When it is your turn to choose, a clock will appear at the top of your computer screen, showing a total of 30 seconds available to make your decision. Please try to decide during this time, so that the experiment moves ahead rapidly. However, there is no penalty for exceeding the time limit; even after the 30 seconds you will still have the opportunity to make your choice.

Information after each round
A round ends after all participants have made their choices in all 8 situations of that round. When the round ends, the following information will be displayed on your computer screen:

(1) the true value of the number Y in each situation;
(2) the total number of participants who chose A, and the number who chose B, in each situation;
(3) the number of points you earned in each situation.

You will also be informed about the total number of points you have earned in the rounds that have finished so far. Obviously, this total can grow over time, and cannot decrease.

This information will be visible on your screen for up to 120 seconds, indicated by a clock at the top of your computer screen. During this time, you may take notes, if you wish. After 120 seconds, the information will disappear from your screen, but if you wish you can click on the “OK” button to erase it earlier. Once the information disappears, there is no way to retrieve it.

Once all participants have clicked “OK”, or after 120 seconds, the next round will begin. The rules are the same in all rounds.

The questionnaire
At the end of the experiment, we will ask that you fill out a questionnaire. The personal information collected in the questionnaire will be treated with the strictest confidentiality, and will be used only for research purposes.

Final payment
After the experiment, you will be paid, in euros, the value of the points you earned during the experiment. Each point is worth one half of a cent, that is, 200 points are equivalent to €1.

Practice quiz
In order to make sure that all participants understand the rules of the experiment, we now ask that you complete a short practice exercise. The practice questions will appear shortly on your computer screen. When all participants have correctly finished all the practice questions, the first round of the experiment will begin.

Please raise your hand to ask a question if you have any doubts about the rules of the experiment or if you have trouble with any of the practice questions.

Appendix B: Simulation details

For numerical and experimental purposes, we assumed that the distributions $G(\Theta)$ and $F(x|\Theta)$ place positive probability only on a discrete grid; let $g(\Theta)$ be the probabilities associated with the grid points of the aggregate state. For a given state Θ, if the first player uses strategy y, then she will attack (implying $\eta_1 = \alpha_1 = 1$) with probability $\int_{\underline{\theta}-\epsilon}^{\bar{\theta}+\epsilon} y(0,0,x)dF(x|\Theta)$. Starting here, we can calculate the probability of any fraction of attackers α_i up to and including individual i, conditional on some aggregate state Θ:

$$\hat{P}_i(\alpha|\Theta, y) \equiv \text{prob}(\alpha_i = \alpha|i, \Theta, y) \quad \text{for any } \alpha \in \mathcal{J}_i$$

These probabilities can be calculated recursively, using the functions T_i:

$$\begin{aligned}\text{prob}(\alpha_i = \alpha|i, \Theta, y) &= \text{prob}\left(\alpha_{i-1} = \frac{i\alpha}{i-1}|i-1, \Theta, y\right)\left[1 - T_i\left(\frac{i\alpha}{i-1}, \Theta, y\right)\right] \\ &\quad + \text{prob}\left(\alpha_{i-1} = \frac{i\alpha - 1}{i-1}|i-1, \Theta, y\right) T_i\left(\frac{i\alpha - 1}{i-1}, \Theta, y\right)\end{aligned}$$

Next we can easily calculate the joint probability

$$\begin{aligned}&\text{prob}(\alpha_{i-1}, \Theta, i, n, m|y) = \\ &\qquad \text{prob}(\alpha_{i-1}|i-1, \Theta, y)\text{prob}(i)\text{prob}(n|i)\text{prob}(m|n, \alpha_{i-1}, i)\text{prob}(\Theta)\end{aligned}$$

This is the joint probability that the player is the ith in the sequence, that the state is Θ, that the fraction of predecessors attacking is α_{i-1}, and that she observes n predecessors, of whom m attacked, given that all other agents are playing strategy y. All the probabilities in this product are known from our description of the model: $\text{prob}(i) = 1/I$, $\text{prob}(n|i) = N_i(n)$, $\text{prob}(m|n, \alpha_{i-1}, i) = M_i(m|n, \alpha_{i-1})$, and $\text{prob}(\Theta) = g(\Theta)$.

Now if trader i plays $\eta_i = 1$, then $\alpha_i = ((i-1)\alpha_{i-1} + 1)/i \in \{\frac{1}{i}, \frac{2}{i}, \ldots, 1\}$. Thus for $\alpha_i = ((i-1)\alpha_{i-1} + 1)/i$, we have

$$\text{prob}(\alpha_i, \Theta, i, n, m|y, \eta_i = 1) = \text{prob}(\alpha_{i-1}, \Theta, i, n, m|y)$$

From here, we go on updating, assuming that other agents play strategy y, to calculate probability distributions over α_j for $j > i$. In the end, we need the probabilities over α_I, the aggregate fraction attacking; in particular, we must know

$$\text{prob}(\alpha_I, \Theta, i, n, m|y, \eta_i = 1)$$

which is the joint probability of the event in which the aggregate outcome is (α_I, Θ), the player is the ith individual, and the player observes n predecessors of whom m attack, given that others play strategy y and the individual plays

$\eta_i = 1$. This can be calculated by updating with T_i, as we did before. For any $j > i$ and $\alpha_j \in \{\frac{1}{j}, \frac{2}{j}, \ldots, 1\}$,

$$\begin{aligned}\text{prob}(\alpha_j, \Theta, i, n, m|y, \eta_i = 1) \;=\; & \text{prob}\left(\alpha_{j-1} = \frac{i\alpha_j}{i-1}, \Theta, i, n, m|y, \eta_i = 1\right)\left[1 - T_j\left(\frac{i\alpha_j}{i-1}, \Theta, y\right)\right] \\ & + \text{prob}\left(\alpha_{j-1} = \frac{i\alpha_j - 1}{i-1}, \Theta, i, n, m|y, \eta_i = 1\right) T_j\left(\frac{i\alpha_j - 1}{i-1}, \Theta, y\right)\end{aligned}$$

Next, for any signal x, we can multiply by $\text{prob}(x|\Theta)$ to calculate

$$\text{prob}(\alpha_I, \Theta, i, n, m, x|y, \eta_i = 1)$$

which is the player's distribution over the aggregate outcome conditional on his information set and his action (we only need this information for the case $\eta_i = 1$, since if $\eta_i = 0$ then the payoff is zero regardless of the aggregate outcome). This, at last, is the probability that enters into formula (10) from which we calculate the conditional probability $\Pi(\alpha_I, \Theta|n, m, x, \eta_i = 1, y)$ that the player must know in order to choose his optimal strategy.

Plugging this probability Π in to either the rational first-order condition (5) or the logit choice equation (14) for some given λ, we obtain the best response (or logit response) to the strategy y. That is, we have calculated the best or logit response $y' = By$ or $y' = B_\lambda y$. We iterate on the mapping B or B_λ until the average absolute change in y across all information sets $y(n, m, x)$ is less than 10^{-7}.

Appendix C: Proof of Proposition 2

Proof.

We start by considering the rational herding equilibrium case $\lambda = 0$, in which case doubly monotonic strategies are double threshold strategies.

Pick $\gamma \in (0, 0.5)$ and a small $\delta > 0$. For any equilibrium, define

$$\begin{aligned}\theta_* &\equiv \inf\{\theta : \alpha_\infty < a(\theta) \text{ with probability } > \gamma\} \\ \theta^* &\equiv \sup\{\theta : \alpha_\infty \geq a(\theta) \text{ with probability } > \gamma\}\end{aligned}$$

Thus θ_* is the lowest aggregate state such that there is not a successful attack with at least probability γ, while θ^* is the highest aggregate state such that there is a successful attack with at least probability γ.

Note that if there are sometimes multiple outcomes, then $\theta_* \leq \theta^*$. On the other hand, if there is a unique outcome for all θ (that is, if $\alpha_\infty(\theta)$ is a well-defined function of θ, rather than a correspondence) then $\theta_* = \theta^* \equiv \theta_0$, and $\alpha_\infty(\theta_0) = a(\theta_0) \equiv \alpha_0$. Let us assume that there is a unique outcome for all θ, and try to show a contradiction.

Equation (8) shows that if i and n are large, then $M_i(m|n, \alpha) \equiv prob(m_i = m|i, n_i, \alpha)$ goes very quickly to zero for any m_i not approximately equal to αn. Moreover, when N^{max} is large, most players will observe large samples of previous actions (since q is fixed). On the other hand, $f(x|\theta)$ is fixed, independent of N^{max}. Therefore, we can pick N^{max} large enough so that receiving a signal x far from its mean, θ, is vastly more probable than receiving a sample m_i/n_i far from its mean α_{i-1}.

Now, when $I = \infty$, Assumption 1 guarantees that there are just two conceivable configurations: (a) $\theta_m \leq \theta_0 \leq \theta^m$ and (b) $\theta_m < \theta^m \leq \theta_0$. We begin by considering case (a).

Consider a player who observes a large sample of previous actions n_i, with $\frac{m_i}{n_i} \geq \alpha_0 + \delta$ and $x_i = \theta_0 + \epsilon - \delta$. Since $\frac{m_i}{n_i}$ is overwhelmingly likely to be close to its mean for sufficiently large n_i, such a player (knowing that in equilibrium, the outcome $\alpha_\infty(\theta)$ is a well-defined, single-valued function of the aggregate state θ) should conclude that the true state is a θ slightly less than θ_0. Therefore he should conclude that a successful attack is occurring, and since $\theta_0 < \theta_m$, he should also conclude that it is profitable to join a successful attack.

On the other hand, if x_i were just slightly higher ($x_i = \theta_0 + \epsilon + \delta$), he should conclude that $\theta > \theta_0$, implying that a successful attack is impossible, and therefore that attacking is undesirable. Therefore, when n_i is large, and $m/n \geq \alpha_0 + \delta$, the optimal threshold is $\tau(m, n) \approx \theta_0 + \epsilon$.

Similarly, a player who observes many previous actions, with $\frac{m_i}{n_i} \geq \alpha_0 - \delta$ and $x_i \approx \theta_0 - \epsilon$, will conclude that he is observing an unsuccessful attack if x_i is slightly greater than $\theta - \epsilon$, but that he is observing a successful attack if x_i is slightly less than ϵ; therefore the optimal threshold when n_i is large and $m/n \geq \alpha_0 - \delta$ is $\tau(m, n) \approx \theta_0 - \epsilon$.

Therefore, consider what happens when the true underlying state is θ_0. For sufficiently large N^{max}, we have $T(\alpha_0 - \delta, \theta_0, \tau) \approx 0$: if the fraction of initial players attacking prior to i is less than or equal to $\alpha_0 - \delta$, then player i's threshold is extremely likely to be $\tau(m, n) \approx \theta_0 - \epsilon$, so his signal is extremely likely to exceed his threshold, so his probability of attacking will be roughly 0.

Likewise, $T(\alpha_0 + \delta, \theta_0, \tau) \approx 1$: if the fraction of previous players attacking is at least $\alpha_0 + \delta$, then player player i's threshold is extremely likely to be $\tau(m, n) \approx \theta_0 + \epsilon$, so his signal is extremely likely to fall below his threshold, so his probability of attacking will be roughly 1.

Therefore, $T(\alpha, \theta_0, \tau)$ crosses the 45^o line in an upward direction at $\alpha \approx \alpha_0$. By continuity, there exists an interval around θ_0 where T has multiple crossings. But this means multiple outcomes are possible in an interval around θ_0, which contradicts our initial assumption.

Moreover, at θ_0, all the initial signals can be arbitrarily close to $\theta_0 + \epsilon$, or arbitrarily close to $\theta_0 - \epsilon$, with positive probability. Since we have assumed that outcomes are unique, any player who observes no previous choices should choose a threshold near θ_0, because she knows that the eventual outcome will be a successful attack if $\theta < \theta_0$, while attacks will be unsuccessful if $\theta > \theta_0$. Players with a small but nonzero sample have a similar incentive, but they also have an incentive to follow previously observed actions.

Thus if the first signals are sufficiently close to $\theta_0 + \epsilon$, no initial players should attack. Analogously, if the first signals are sufficiently close to $\theta_0 - \epsilon$, all initial players should attack. Therefore there is positive probability of an arbitrarily large number of initial attacks; or an arbitrarily large number of initial nonattacks. Thereafter, the shape of the T function for large n comes into play, so that the outcome may converge to either of the crossings of $T(\alpha, \theta_0, \tau)$ with positive probability. This contradicts the hypothesis that outcomes are unique in case (a).

We must still consider case (b), in which $\theta_m < \theta^m \leq \theta_0$. If this configuration holds, it means successful attacks occur in equilibrium for all θ between $\underline{\theta}$ and θ_0, but that when $\theta \in (\theta^m, \theta_0)$, the attackers are making a mistake: for these θ, attacking is undesirable because $R(\theta) < t$. Now define $\alpha^m \equiv \alpha_\infty(\theta^m)$. It is straightforward to show, with arguments like those used above, that in this case, a player observing many n_i and $\frac{m_i}{n_i} \approx \alpha^m + \delta$ will choose a threshold $\tau(m, n) \approx \theta^m + \epsilon$, while a player observing many n_i and $\frac{m_i}{n_i} \approx \alpha^m - \delta$ will choose a threshold $\tau(m, n) \approx \theta^m - \epsilon$. Following our previous arguments, we can conclude that $T(\alpha, \theta_0, \tau)$ crosses the 45^o line from below near α^m, and that its upper and lower stable crossings are both outcomes that occur with positive probability. Again, this contradicts our initial assumption of unique outcomes.

A doubly monotonic logit herding strategy with λ strictly positive but still sufficiently close to zero is arbitrarily close to a double threshold strategy. The arguments used in this proof still hold for this case.

QED.

Table 1: Estimated strategies: comparative statics on q.

Session	I	q	ϵ	*Position known?*	γ	δ	ξ_{1-2}	ξ_{3-7}	*Width bimodal region*	*Width region all possible*
F4	8	0	15	*yes*	**0.2281**	**46.77**	*n.a.*	*n.a.*	0	0
					(0.0214)	(0.685)			(0)	(0)
F7	8	0.5	15	*no*	**0.3875**	**58.78**	3.571	**13.14**	0	1.5
					(0.0765)	(0.633)	(2.59)	(5.53)	(0)	(1.60)
F5	8	0.5	15	*yes*	**0.2450**	**51.52**	**20.11**	34.81	0	**13.5**
					(0.0442)	(1.18)	(3.62)	(195)	(2.37)	(2.92)
M4	8	0.5	15	*no*	**0.2519**	**49.83**	5.289	30.86	0	**6.0**
					(0.0391)	(1.04)	(3.34)	(287)	(0.224)	(2.69)
M2	8	0.5	15	*yes*	**0.2730**	**48.59**	**9.110**	**35.07**	0	**8.5**
					(0.0753)	(1.0568)	(2.85)	(13.0)	(0)	(2.60)
F1	8	0.75	15	*no*	**0.1752**	**48.92**	**14.00**	**29.21**	0	**13.5**
					(0.0333)	(1.15)	(4.69)	(9.25)	(4.19)	(4.46)
F3	8	0.75	15	*yes*	**0.3452**	**56.83**	**7.136**	**34.36**	0	**13.0**
					(0.0663)	(0.851)	(2.99)	(7.30)	(2.74)	(2.46)
M3	8	0.75	15	*no*	**0.2855**	**54.18**	**11.63**	**19.61**	0	**9.0**
					(0.101)	(1.13)	(4.64)	(3.85)	(0.919)	(2.38)
M1	8	0.75	15	*yes*	**0.1535**	**51.27**	**21.63**	**33.76**	3.5	**18.0**
					(0.0259)	(1.24)	(6.06)	(10.8)	(5.17)	(4.97)
F2	8	1	15	*yes*	**0.1378**	**55.90**	2.330	**33.29**	0.5	**13.0**
					(0.0195)	(1.41)	(5.89)	(8.11)	(3.57)	(4.34)

Bootstrap standard errors in parentheses; coefficients significant at 5% shown in bold.

Data of first two rounds of each session deleted to eliminate nonstationary behavior due to learning.

Coefficients γ and δ assumed independent of n; ξ is estimated separately for $n \in \{1,2\}$ and $n \in \{3,4,5,6,7\}$.

Table 2: Estimated strategies: comparative statics on q.

Session	I	q	ϵ	*Position known?*	γ_{0-2}	γ_{3-7}	δ	ξ_{1-2}	ξ_{3-7}	*Width bimodal region*	*Width region all possible*
F4	8	0	15	*yes*	**0.2281**	*n.a.*	**46.77**	*n.a.*	*n.a.*	0	0
					(0.0214)		(0.685)			(0)	(0)
F7	8	0.5	15	*no*	**0.3756**	0.4354	**58.81**	3.707	**12.37**	0	1.0
					(0.0825)	(21.6)	(0.607)	(2.37)	(5.56)	(0)	(1.44)
F5	8	0.5	15	*yes*	**0.3019**	**0.1366**	**51.55**	**18.51**	51.92	1.5	**15.5**
					(0.0524)	(0.0562)	(1.21)	(3.35)	(30.0)	(1.88)	(2.33)
M4	8	0.5	15	*no*	**0.2451**	0.2957	**49.71**	5.429	28.27	0	6.0
					(0.0375)	($5.1*10^{12}$)	(1.09)	(3.11)	(14.8)	(0)	(3.17)
M2	8	0.5	15	*yes*	**0.2796**	0.2476	**48.57**	**8.97**	35.90	0	**8.5**
					(0.0863)	(14.9)	(1.04)	(2.81)	(113)	(0.559)	(2.34)
F1	8	0.75	15	*no*	**0.1912**	**0.1422**	**48.97**	**12.71**	36.56	0	**14.5**
					(0.0406)	(0.0481)	(0.716)	(3.30)	(45.6)	(1.80)	(3.34)
F3	8	0.75	15	*yes*	**0.2963**	1.155	**56.81**	**7.899**	**32.99**	1.5	**12.5**
					(0.0768)	($2.2*10^{13}$)	(0.7058)	(3.639)	(6.192)	(5.371)	(2.821)
M3	8	0.75	15	*no*	0.3493	**0.2194**	**54.00**	**10.61**	**23.54**	0	**10.5**
					(0.191)	(0.0888)	(1.24)	(4.07)	(7.12)	(0.638)	(2.95)
M1	8	0.75	15	*yes*	**0.1444**	**0.1802**	**51.02**	**22.87**	28.71	2.5	**16.5**
					(0.0315)	(0.0495)	(1.19)	(7.12)	(19.7)	(4.89)	(5.11)
F2	8	1	15	*yes*	**0.1602**	**0.1046**	**55.80**	0.8077	**46.49**	3.0	**15.0**
					(0.0338)	(0.0233)	(1.25)	(5.24)	(16.0)	(4.09)	(2.95)

Bootstrap standard errors in parentheses; coefficients significant at 5% shown in bold.

Data of first two rounds of each session deleted to eliminate nonstationary behavior due to learning.

Coefficient δ assumed independent of n; γ is estimated separately for $n \in \{0, 1, 2\}$ and $n \in \{3, 4, 5, 6, 7\}$, while ξ is estimated separately for $n \in \{1, 2\}$ and $n \in \{3, 4, 5, 6, 7\}$.

Table 3: Estimated strategies: comparative statics on ϵ and I.

Session	I	q	ϵ	*Position known?*	γ	δ	ξ_{1-2}	ξ_{3-7}	*Width bimodal region*	*Width region all possible*
F9	8	0	8	*yes*	**0.4712**	**51.57**	*n.a.*	*n.a.*	0	0
					(0.105)	(0.859)			(0)	(0)
F10	8	1	8	*yes*	**0.5090**	**55.99**	6.420	**9.114**	0	**4.5**
					(0.1498)	(0.5461)	(3.7775)	(2.2647)	(1.3139)	(1.8057)

Session	I	q	ϵ	*Position known?*	γ	δ	ξ_{1-3}	ξ_{4-11}	*Width bimodal region*	*Width region all possible*
F6	12	0.5	15	*yes*	**0.1682**	**50.33**	**14.35**	**12.25**	0	0
					(0.0233)	(1.1846)	(3.8697)	(5.0404)	(0)	(0)
F8	12	0.75	15	*yes*	**0.1963**	**50.72**	**16.77**	**35.02**	1.0	**9.0**
					(0.0413)	(1.5183)	(2.7207)	(5.9482)	(1.0894)	(1.8909)

Bootstrap standard errors in parentheses; coefficients significant at 5% shown in bold.

Data of first two rounds of each session deleted to eliminate nonstationary behavior due to learning.

Coefficients γ and δ assumed independent of n; ξ estimated separately for low and high n.

Table 4: Estimated strategies: comparative statics on ϵ and I.

Session	I	q	ϵ	*Position known?*	γ_{0-2}	γ_{3-7}	δ	ξ_{1-2}	ξ_{3-7}	*Width bimodal region*	*Width region all possible*
F9	8	0	8	*yes*	**0.4712**	*n.a.*	**51.57**	*n.a.*	*n.a.*	0	0
					(0.105)		(0.859)			(0)	(0)
F10	8	1	8	*yes*	0.9200	0.3488	**56.47**	4.799	12.67	0	**5.5**
					($1.5*10^{13}$)	(0.373)	(0.953)	($4.0*10^{10}$)	(6.67)	(4.93)	(2.42)
Session	I	q	ϵ	*Position known?*	γ_{0-3}	γ_{4-11}	δ	ξ_{1-3}	ξ_{4-11}	*Width bimodal region*	*Width region all possible*
F6	12	0.5	15	*yes*	**0.2847**	**0.0799**	**51.04**	**8.861**	40.68	0	0
					(0.0663)	(0.0209)	(1.22)	(2.14)	(20.9)	(0)	(0.671)
F8	12	0.75	15	*yes*	**0.4188**	**0.1153**	**50.40**	**12.28**	**50.84**	1.0	**8.0**
					(0.162)	(0.0393)	(1.29)	(2.62)	(22.7)	(2.36)	(2.30)

Bootstrap standard errors in parentheses; coefficients significant at 5% shown in bold.

Data of first two rounds of each session deleted to eliminate nonstationary behavior due to learning.

Coefficient δ assumed independent of n; γ and ξ are estimated separately for low and high n.

Table 5: Selected logit equilibrium estimates.

Session	I	q	ϵ	*Position known?*	λ	*Eigenvalue bound*	*Loss per decision*
F4	8	0	15	*yes*	**9.12**	0.40	€0.0060
					(0.13)		
F7	8	0.5	15	*no*	**10.00**	0.34	€0.0067
					(0.11)		
F5	8	0.5	15	*yes*	**7.42**	0.38	€0.0026
					(0.13)		
M4	8	0.5	15	*no*	**10.47**	0.34	€0.0075
					(0.10)		
M2	8	0.5	15	*yes*	**9.12**	0.35	€0.0052
					(0.11)		
F1	8	0.75	15	*no*	**9.55**	0.43	€0.0057
					(0.11)		
F2	8	1	15	*yes*	**10.00**	0.46	€0.0069
					(0.11)		

Standard errors in parentheses; coefficients significant at 5% shown in bold.

Data of first two rounds of each session deleted to eliminate nonstationary behavior due to learning.