Naive Herding
(PRELIMNARY VERSION FOR UCL CONFERENCE)

Erik Eyster
Department of Economics
LSE

Matthew Rabin
Department of Economics
University of California—Berkeley

May 21, 2008

Abstract

Within a social-learning environment, we investigate the implications of the assumption (formalized in various ways) that players may naively believe that a previous player's actions solely reflect that player's private information. This can lead players to become very confident about the true state of the world, even in environments where rational players would never become certain, as well as to herd on incorrect actions, even in environments where fully rational players never incorrectly herd.

Keywords: Innit?, Didn't I?
JEL Classification: B49

Acknowledgments: We thank Zack Grossman for valuable research assistance, Ignacio Esponda, Guy Mayraz, Marco Ottaviani, Peter Sorensen, and Georg Weizsäcker for helpful conversations, and seminar participants at Berkeley, Birkbeck/UCL, Bocconi, Brescia, Davis, ESEM 2007, ESSET Gerzensee, Harvard/MIT, Hebrew, Lausanne, Naples, Tel Aviv, Toulouse, and Wash U for helpful comments, and the NSF and [[[XXX]]] for financial support. Eyster thanks Berkeley for hospitality while part of this research was undertaken.

Contact: e.eyster@lse.ac.uk; rabin@econ.berkeley.edu

1 Introduction

Beginning with Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), a theoretical literature has explored the role of inference in herding behaviour. In the simplest model, a group of people arranged in sequence chooses one of two options. They have common preferences over the two choices but do not know which one is preferable. Instead, each person receives an independent and equally strong private binary signal about which option is better, and each can observe all of her predecessors' choices. In this setting, rational agents herd: if the first two choose A, then so too will all who follow, regardless of their private information. This is because two A choices mean that (given any natural tie-breaking rule for the second mover) each of the first two players must hold a signal indicating A preferable; the next mover should realize this and, even if she gets a B signal, choose action A. And all subsequent players will follow suit. Although everyone chooses A, they understand that all but the first two movers ignore their signals. Hence, they never become highly confident that A is better. Generalizing this result, the literature finds that rational herders will—when the action space and the signal space both have finite support—always converge in behaviour, albeit it with positive probability to the wrong action. Rational players find it individually optimal to herd, despite the social inefficiency inherent in possibly herding on the wrong action.

One recognized limitation of this literature is that the logic of rational herding relies crucially on features of the environment that may be neither realistic nor intuitively central to herding. Finer action spaces *or* finer signal spaces reduce the scope for herding, and in either of the limiting cases where actions fully reveal beliefs or where people can receive arbitrarily strong signals equilibrium actions converge to the best response to the true state with probability one (Smith and Sorensen, 2001).

In addition, the rationality assumption maintained by these models seems neither entirely realistic nor necessary for—nor even conducive to—herding behaviour. In this paper, we consider the possibility of herding by players who are more limited in their rationality. We find that a simple form of naivety in inference can lead players to stronger and wronger herding than rational inference. Such naivety also leads to wrong herding in a much wider range of situations than rational inference. We introduce two particular types of departure from purely rational play that both reflect general types of errors that people may make across a wide variety of strategic situations and capture specific errors that they make in social learning contexts, as suggested by evidence from laboratory experiments. First, players may be "cursed"—and underappreciate the degree to which previous movers' actions reflect these movers' information.[1] Players who commit this error choose actions that depend only upon their own private signal and not upon their observation of previous play. More plausibly, players might be "inferentially naive"—and realize that previous movers' actions reflect these movers' own signals without appreciating that these previous movers themselves make informational inferences from their own predecessors. Players who commit this error believe that their predecessors' actions depend upon their private signals alone, essentially mistakenly believing all of their predecessors cursed. Indeed, this form of naive inference may seem especially natural to anyone who has struggled to learn or teach informational cascades: the intuition that herders' actions may bear very little resemblance to their private information is a difficult one to grasp; those whose actions we are trying to model may themselves find it equally challenging.

It is this latter error that we emphasize in this paper and that we show to have dramatic implications. To consider the example above, inferential naivety leads each mover to interpret every previous A choice as another bit of evidence favouring A. Unlike rational players, each successive naive player herding on A becomes more convinced that A is in fact the true state of the world, so that players eventually come to believe with certainty that A is the true state of the world. This is especially striking because, just as in the rational case, the herd may be "mistaken"—a few unlikely signals early on can move the players to believe in the wrong theory, which they then may subsequently believe more and more strongly.

While naive players hold beliefs that differ markedly from their rational counterparts, in the simple binary-action-binary-signal-binary-state model outlined above they behave in the same way.

[1] This would be a simple application of "cursed equilibrium" as modeled in Eyster and Rabin (2005) and also captured by versions of Jehiel and Koessler's (forthcoming) "analogy-based expecations" in Bayesian games. As argued in Eyster and Rabin (2009), however, social-learning environments like the one considered in this paper provide enough scope for observing opponents' play that they should not be particularly conducive to extreme "cursedness".

To differentiate inferentially naive from rational play in actions as well as beliefs, we turn to a model of social learning with much richer action and signal spaces. In our primary model, each player receives a signal drawn from the continuum that corresponds to the probability that the true state is A—and then chooses an action from the continuum that reflects her beliefs. In this informationally rich environment, rational players' beliefs and actions converge with probability one to the correct state; indeed, either a rich signal or action space suffices for the result, and this model has both. Strikingly, under a very weak assumption on the distribution of signals—somewhat stronger than that "magic signals" revealing the true state have zero probability—with positive probability naive players converge to the wrong beliefs and actions. Indeed, we show by simulation that this probability is approximately eleven percent when the distribution of signal strengths is uniform.

What is the intuition for this? Not realizing that the second mover conditions his action on both his private signal and his inference about the first mover's private signal (from her move), the third mover combines her own private signal with the moves of the first and second mover. But since the second move already incorporates the first mover's private signal, this procedure double counts the first mover's signal. Naive inference by the fourth mover, in turn, leads her to count the first mover's signal four-fold: thinking she is using the actions of the first three movers to count each of the first three signals once, she is in fact counting the first signal four times (once from the first action, once from the second mover's action, and twice from the third mover's) and the second signal twice (once each from the first and second mover). In general, by naively interpreting each predecessor's action as reflecting solely that mover's signal, naive players come to massively overweight early signals—mover k counts the first signal 2^{k-1} times, the 2nd signal 2^{k-2} times, etc. If the first signal happens to be misleading, limit herds may so over-use it as to outweigh an infinite sequence of further signals, including an infinity of arbitrarily strong signals of the truth.

In Section 2 we formalize our assumptions about cursedness and inferential naivety and relate them to existing more general models. In Section 3, we establish our formal results for our primary model. We establish that with two possible states of the world, inferentially naive players will converge to certain beliefs about the state, and with positive probability this will be on the wrong state. While we prove the result in a model with a continuum of actions and a continuum of signals, it holds equally well in models with a finite number of actions and signals. We also establish a curious result: if beliefs for some time weakly favour one state over the other, then the other state is more likely to be the true state. The intuition behind the result relates to its limited applicability: it is, in fact, very unlikely that beliefs will remain weak after a long sequence signals. But the most likely way that a long sequence of players can be unsure of the true state is if they start out believing the wrong thing and then receive a long sequence of contradictory evidence—all of which still adds up to less than the herders implicitly infer from the earlier signals.[2]

In Section 4 we consider the implications of inferential naivety in several further environments. In particular, we consider the case of limited observability where, instead of observing all the previous moves, each participant observes the immediately preceding k movers, for some finite k. Interestingly, if $k = 1$, meaning that each herder observes solely the move before him, then inferentially naive players behave just as fully rational players do—despite having the wrong model of their predecessors' decision making. Here, each player mistakenly believes that her action depends on only two signals—her predecessor's and her own—when in fact her predecessor's action depends on his own predecessor's action and signal, and so forth. When $k = 2$, however, once more there is positive probability that players converge to the wrong limiting belief and action; as k becomes larger, the probability of a wrong herd increases and converges in the limit to the probability in the full-observability model. In this sense, we show that the wrongness result is robust to limited observability of prior moves.

In Section 5 we return to the possibility that players may not be solely inferentially naive but also "cursed" or rational. Blending cursedness with inferential naivety may be of special interest because of their potential for making opposite predictions: cursed players are too little influenced by their predecessors' actions, while inferentially naive players are seemingly too much so. While there is truth to this contrast, the results of Section 5 help emphasize that inferential naivety does not amount to "over" inference and indeed that some of its implications are not undone by

[2] The intuition is similar to that developed in Rabin and Schrag's (1999) model of confirmatory bias, who assumed an individual processor of information tended to misread later signals as reinforcing earlier signals. Indeed, much of the intuition of our model reflects the sort of "social-inference confirmatory bias" that inferential naivety generates.

cursedness. Inferential naivety says not just that prior actions are overweighted too much, but that a) herders interpret them too literally as reflecting new information, and b) that (unbeknownst to the herders themselves) far too much weight accrues to early relative to more recent signals. Combining cursedness with inferential naivety may, in fact, explain precisely the mix of behaviours observed in some laboratory experiments, where herders appear to underweight the sum total of others' signals relative to their own whilst reading more into herds that do form than warranted. To capture this intuition, we show in a variant of our primary where players choose actions in $\{0, 1\}$ rather than the continuum that if everybody is inferentially naive and not too cursed that in fact limiting public beliefs may be closer to the wrong state than the right one. The wedge between this and the stronger results in Section 3 derives from the fact that cursed players' beliefs and actions do not converge because they overweight their own signals so much. Section 5 also considers the case where there are a mixture of types—naive, cursed, and even rational—and shows an interesting form of robustness. In fact, with the assumption that the rational types are aware of the precise distribution and identity of the other types, then in the same environments where rational people converge to the truth when rationality is common knowledge, they will also do so when they know others are not rational. Essentially, rational herders who observe a lot always figure out the truth so long as previous behaviour richly reflects the distribution of signals.[3] Cursed players simply ignore others' signals, and inferentially naive players may again anchor with full confidence on the wrong state entirely if the first signals are wrong. This can be true even though they will be behaving opposite as the rational people they observe. In environments where there is full support of actions that fully reveal beliefs, in fact, this anchoring can occur even when the naive herders are a small fraction of the herders.

[3] Note, in fact, that rational herders can figure out the truth even in the binary signal/binary signal case when there are sufficiently many sufficiently cursed folk around, so that in this sense the presence of not-fully-rational types in fact enhances the ability of fully rational people to extract the truth from a herd.

1.1 Experimental Evidence

While the primary purpose of this article is to clarify the logical implications for herding of a very simple form of inferential error rather than directly model existing data, in Section X we briefly discuss how our analysis relates to existing evidence. Recently, such models have been tested extensively in the laboratory. While of course some predictions of the rational model have found support, several systematic discrepancies have been uncovered. Kübler and Weizsäcker (2004) show in a variant of the model (described in more detail below) that subjects' beliefs become too extreme to be well explained by rational play. Kübler and Weizsäcker (2005) report the related finding that longer cascades are more stable. Intuitively because after a long string of A choices, people come to believe A more and more likely, reducing the likelihood that anyone will break the cascade by choosing B.

2 Informational Naivety

In this section, we formally define the solution concepts that we shall use throughout the paper. We begin with Eyster and Rabin's (2005) concept of cursed equilibrium, which we weaken in a way that preserves the form of informational misinference in a cursed equilibrium yet relaxes the equilibrium assumption that players correctly predict the distribution of one another's actions. We use this notion of "cursed best response" to underpin our definition of "best-response trailing naive inference" (BRTNI) play, the primary focus of this paper.

2.1 Cursed Equilibrium

Consider a finite Bayesian game $G = (A_1, \ldots, A_N; T_0, T_1, \ldots, T_N; p; u_1, \ldots, u_N)$ played by players $k \in \{1, ..., N\}$. A_k is the finite set of Player k's actions; in a sequential game, an action specifies what Player k does at each of her information sets. T_k is the finite set of Player k's "types", each type representing different information that Player k can have. For conceptual and notational ease, we include a finite set of types for "nature", T_0. $T \equiv T_0 \times T_1 \times ... \times T_N$ is the set of type profiles, and p is the prior probability distribution over T, which we assume to be commonly known.

Player k's payoff function $u_k : A \times T \to \mathbb{R}$ depends on all players' actions as well as their types. A (mixed) strategy σ_k for Player k specifies a probability distribution over actions for each type: $\sigma_k : T_k \to \triangle A_k$. Let $\sigma_k(a_k|t_k)$ be the probability that strategy σ_k assigns to type t_k playing action a_k.

The common prior probability distribution p puts positive weight on each $t_k \in T_k$ and fully determines the conditional probability distributions $p_k(t_{-k}|t_k)$, Player k's beliefs about the types $T_{-k} \equiv \underset{j \neq k}{\times} T_j$ of other players (including nature) given her own type $t_k \in T_k$. Let $A_{-k} \equiv \underset{j \neq 0,k}{\times} A_j$ be the set of action profiles for players $j \neq k$ (excluding nature, who takes no action), and $\sigma_{-k} : T_{-k} \to \underset{j \neq 0,k}{\times} \triangle A_j$ be a strategy of Player k's opponents (where they mix independently), where $\sigma_{-k}(a_{-k}|t_{-k})$ is the probability that type $t_{-k} \in T_{-k}$ plays action profile a_{-k} under strategy $\sigma_{-k}(t_{-k})$.

The standard solution concept in such games is Bayesian Nash equilibrium:

Definition 1 *The strategy profile σ is a* Bayesian Nash equilibrium *if for each Player k, each type $t_k \in T_k$, and each a_k^*, $\sigma_k(a_k^*|t_k) > 0$ only if*

$$a_k^* \in \arg\max_{a_k \in A_k} \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \cdot \sum_{a_{-k} \in A_{-k}} \sigma_{-k}(a_{-k}|t_{-k}) u_k(a_k, a_{-k}; t_k, t_{-k}).$$

In a Bayesian Nash equilibrium, each player correctly predicts both the probability distribution over the other players' actions and the correlation between those actions and the other players' types.

Before defining cursed equilibrium, we define for each type of each player the "average mixed strategy" of other players, averaged over the other players' types. Formally, for each strategy for Players $-k$, σ_{-k}, and type of Player k, t_k, define $\overline{\sigma}_{-k}(\cdot|t_k)$ by

$$\overline{\sigma}_{-k}(a_{-k}|t_k) \equiv \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \cdot \sigma_{-k}(a_{-k}|t_{-k}).$$

When Player k is of type t_k, $\overline{\sigma}_{-k}(a_{-k}|t_k)$ is the probability that strategy profile σ_{-k} assigns to Players $-k$ playing the action profile a_{-k}. A type t_k of Player k who (perhaps mistakenly) believes that each type profile of the other players plays the same mixed action profile—while correctly predicting the distribution over their actions—believes that each type profile t_{-k} of the other players mixes over action profiles according to $\overline{\sigma}_{-k}(\cdot|t_k)$ rather than according to their true strategy profile $\sigma_{-k}(\cdot|t_{-k})$. Different types of Player k may hold different beliefs about the probability distribution over Players $-k$'s actions, as reflected in the fact that $\overline{\sigma}_{-k}(a_{-k}|t_k)$ depends on t_k. Let $\overline{\sigma}_{-k}(t_k) : T_{-k} \rightarrow \underset{j\neq 0,k}{\times} \triangle A_j$ denote t_k's beliefs about the average strategy of players $j \neq k$, namely $\overline{\sigma}_{-k}(t_k)$ is the strategy players $j \neq k$ would play if each type profile t_{-k} played a_{-k} with probability $\overline{\sigma}_{-k}(a_{-k}|t_k)$.

Definition 2 *The mixed-strategy profile* σ *is a* fully cursed equilibrium *if for each Player* k*, type* $t_k \in T_k$*, and action* a_k^**,* $\sigma_k(a_k^*|t_k) > 0$ *only if*

$$a_k^* \in \arg\max_{a_k \in A_k} \sum_{t_{-k}\in T_{-k}} p_k(t_{-k}|t_k) \cdot \sum_{a_{-k}\in A_{-k}} \overline{\sigma}_{-k}(a_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k}).$$

The strategy profile σ is a fully cursed equilibrium if each type t_k of each Player k plays a best response to the other players' playing $\overline{\sigma}_{-k}(\cdot|t_k)$ rather than their true, type-contingent fully cursed equilibrium strategy $\sigma_{-k}(\cdot|t_{-k})$.

Like in Bayesian Nash equilibrium, players in a fully cursed equilibrium correctly predict the distribution over others players' actions; unlike in a Bayesian Nash equilibrium, they ignore the correlation between those other players' actions and types. We interpret such "cursedness" not as Player k being certain that the other players' behave according to $\overline{\sigma}_{-k}$ rather that σ_{-k} but instead as Player k not properly thinking through the logic of other players' strategies. By neglecting how other players' types or private information influence their actions, players fail to perform proper inference from observing or conditioning upon other players' actions. For instance, suppose that the strategy profile σ specifies a pure, separating strategy for every player: Player k plays the pure strategy $a_k(t_k) : T_k \rightarrow A_k$, and $a_k(t_k) = a_k(t'_k)$ implies $t_k = t'_k$. A fully rational Player j understands that if $a_k = a_k(t_k)$, then Player k must have type t_k. By contrast, a fully cursed Player j infers nothing about Player k's type from observing or conditioning upon $a_k = a_k(t_k)$ and maintains her priors, $p(t_k|t_j)$.

Definition 3 *The mixed-strategy profile* σ *is a* χ- cursed equilibrium *if for each Player* k*, type*

$t_k \in T_k$, and action a_k^*, $\sigma_k(a_k^*|t_k) > 0$ only if

$$a_k^* \in \arg\max_{a_k \in A_k} \sum_{t_{-k} \in T_{-k}} p_k(t_{-k}|t_k) \cdot \sum_{a_{-k} \in A_{-k}} \left[(1-\chi)\sigma_{-k}(a_{-k}, t_{-k}) + \chi\overline{\sigma}_{-k}(a_{-k}|t_k) u_k(a_k, a_{-k}; t_k, t_{-k})\right].$$

The strategy profile σ is a χ-cursed equilibrium if each type t_k of each Player k plays a best response to the other players' playing $(1-\chi)\sigma_{-k}(\cdot|t_{-k}) + \chi\overline{\sigma}_{-k}(\cdot|t_k)$ rather than their true, type-contingent fully cursed equilibrium strategy $\sigma_{-k}(\cdot|t_{-k})$. When $\chi = 0$, this reduces to Bayesian Nash equilibrium; when $\chi = 1$, it reduces to fully cursed equilibrium. For intermediate values of χ, players partially, but but not fully, appreciate how others' actions depend upon their types.

2.1.1 Best Response to Cursed Beliefs

Cursed equilibrium maintains the feature of Bayesian Nash equilibrium that players correctly predict one another's actions while dropping the feature that the players understand the mapping from one another's types into actions. This seems reasonable in many contexts where players may learn the distribution of one another's actions without ever observing other players' types. For example, in government oil or timber auctions, a single set of bidders who compete for many similar tracts of land or sea may come to learn the distribution of bids without ever learning their competitors' estimates of the assets' values.[4]

In other contexts, players may neither know the distribution nor understand the information content of their opponents' actions. One example may be players who encounter a strategic setting for the first time. In still other contexts like the model of social learning studied in this paper, players' beliefs about the distribution of their opponents' actions do not affect their best responses; hence, none of our results depends on any assumption about players' beliefs about one another's actions. This prompts us to take the most parsimonious approach by assuming nothing about these beliefs.

[4] In some settings, in addition to the distribution of other players' actions, players may learn the distribution of their own payoffs as function of their type and action as well as other players' actions. Esponda (forthcoming) develops a solution concept that joins cursed inference to this richer information feedback.

To develop our concept formally, we first define

$$u_k^{\chi}(a;t_k,t_{-k}) = (1-\chi)u_k(a;t_k,t_{-k}) + \chi\left(\sum_{t'_{-k}\in T_{-k}} p(t'_{-k}|t_k)u_k(a;t_k,t'_{-k})\right),$$

the χ-weighted average of type t_k's payoffs from the action profile a when facing type t_{-k} and when averaging over all types $t'_{-k} \in T_{-k}$. Eyster and Rabin (2005) show that when type t_k of Player k maximises the "χ-virtual payoffs" $u_k^{\chi}(a;t_k,t_{-k})$ given beliefs that the other players use strategy $\sigma_{-k}(\cdot|t_{-k})$ she also maximises her original payoffs $u_k(a;t_k,t_{-k})$ given beliefs that the other players use strategy $(1-\chi)\sigma_{-k}(\cdot|t_{-k}) + \chi\overline{\sigma}_{-k}(\cdot|t_k)$.[5] Hence a type t_k of Player k who plays an action a_k in the set

$$CBR_k^{\chi}\left(t_k;\widehat{A}_{-k}\right) = \left\{\begin{array}{c} a_k \in A_k : \exists\sigma_{-k} : T_{-k} \to \triangle\left(\times_{j\neq k}\hat{A}_j\right), a_k \in \arg\max_{a'_k\in A_k} \\ \sum_{t_{-k}\in T_{-k}}\sum_{a_{-k}\in\hat{A}_{-k}} p(t_{-k}|t_k)\sigma_{-k}(a_{-k}|t_{-k})u_k^{\chi}(a'_k,a_{-k};t_k,t_{-k}) \end{array}\right\},$$

plays a "χ-cursed best response" to beliefs that the other players play some strategies with support $\hat{A}_{-k}$. The set $CBR_k^1(t_k;A_{-k})$ contains those actions in A_k that are "fully-cursed best responses" by t_k to beliefs that the other players play some strategies with support A_{-k}. The definition of $CBR_k^1(t_k;A_{-k})$ emphasizes how a fully cursed player can be of two minds about her opponents' behaviour, both discerning when forming beliefs about opponents' strategies that they condition their actions upon type and neglecting that correlation when attempting a best response to those beliefs. An interpretation is that Player k is fully rational except for erring in computing best responses: whatever her beliefs about the joint distribution of her opponents' types and actions, she computes best responses using the correct marginals and their product distribution rather than their true joint distribution. Note that our definition of $CBR_k^1(t_k;A_{-k})$ allows different types of Player k to hold different beliefs about other players' strategies.

[5] Intuitively, the former averages Player k's original payoffs, $u_k(a;t_k,t_{-k})$, over her opponents' types while using correct beliefs about their strategies; the latter averages Player k's beliefs about her opponents' types while using the original payoffs.

2.2 BRTNI Play

Players who play cursed best responses fail to appreciate the information content inherent in others' play. In this section, we introduce a solution concept where—in contrast to cursed inference—players do understand that others' actions depend on their information but misunderstand how that dependence works. In particular, every player understands that the other players' types influence their actions yet disregards that other players understand the same thing. Consequently, each player believes that she uniquely perceives information content in play. Because fully cursed players do not recognise the information content in play, we can model this type of naive inference by having each player misapprehend all other players as being fully cursed.

First, define

$$CBR_k^{\chi}\left(t_k; CBR_{-k}^{\chi'}\right) = \left\{ \begin{array}{c} a_k \in A_k : \exists \sigma_{-k} : T_{-k} \rightarrow \triangle\left(\times_{j \neq k} A_j\right), \sigma_{-k}(a_{-k}|t_{-k}) > 0 \\ \Rightarrow \forall j \neq k, \forall t_j \in T_j,\ a_j \in CBR_j^{\chi'}(t_j; A_{-j}), a_k \in \arg\max_{a'_k \in A_k} \\ \sum_{t_{-k} \in T_{-k}} \sum_{a_{-k} \in \hat{A}_{-k}} p(t_{-k}|t_k)\sigma_{-k}(a_{-k}|t_{-k}) u_k^{\chi}(a'_k, a_{-k}; t_k, t_{-k}) \end{array} \right\},$$

the set of type t_k of Player k's χ-cursed best responses to beliefs that all other players play χ'-cursed best responses to some beliefs about others' strategies.

Definition 4 *The action profile* $(a_1, \ldots, a_N)$ *is* (best response trailing naive inference) BRTNI play *if for each type* t_k *of each Player* k, $a_k \in CBR_k^0\left(t_k; CBR_{-k}^1\right)$.

In BRTNI play, each player plays a true best response to beliefs that all other players play fully cursed best responses. As with cursed equilibrium, we interpret BRTNI not so much as Player k being convinced that Player $i \neq k$ is cursed as much as Player k not thinking through how Player $i \neq k$ makes informational inferences from Player $j \neq i$'s actions. Since j may equal k, in general this means that Player k does not think through how Player $i \neq k$ makes inference from her own or a third player's actions.

Eyster and Rabin (2005) show that for each $\chi \in [0, 1]$, χ-cursed equilibrium exists in any finite Bayesian game. This implies that the larger set of cursed best responses is non-empty, and hence BRTNI play exists in all finite Bayesian games. BRTNI play is far too weak a solution concept for most context; for instance, in complete-information games it merely restricts players to playing best responses to best responses, far less than rationalizability (Bernheim, 1984, and Pearce, 1984). However, in a multi-stage game with observable actions where players move one at a time and do not care about their successors' actions, refining players' beliefs about their opponents actions would not affect their best responses. Thus, in the social-learning environment studied in this paper, BRTNI is as strong a solution concept as we need.

3 Cascades

We now consider the role of inferential naivety in information cascades. Beginning with Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), a theoretical literature has explored the role that rational inference plays in herding behaviour. When we observe that people making choices in sequence come to imitate their predecessors, it does not imply that they have a taste for conformity; instead, people who believe that their predecessors base their choices on private evidence rationally learn about their own preferences by observing previous choices.

In the simplest model, a sequence of people choose one of two options in turn. They have common preferences over the two choices but do not know which they prefer. Rather, each person receives private information about which option is better; all private information is of the same quality. In addition, each person can observe all of her predecessors' choices. In this setting, rational agents herd: if the first two choose A, then so too will the third, regardless of her private information. (Because the first has private information favouring A, and the second on average does too—given any tie-breaking rule that rules out the event that an agent with private information favouring B chooses A for sure —the third disregards her private information and chooses A.) And all subsequent players will follow suit. Although everyone in a rational herd may choose A, because everyone understands that a herd forms after the first two people choose A they never become certain that A is the better alternative.

If B is the true state of the world, yet the first two signals favour A, then a rational herd forms on the wrong action. More generally, any finite-action, finite-signal model (with no signal strong enough to perfectly reveal the true state) allows for such inefficient herding. In many settings, though, richer action spaces allow people to express the intensity of beliefs. For instance, a trader may choose not only whether to buy or sell some asset but also how many shares to trade. The finer the action space, the lower the scope for rational herding, and in the limiting case where no two distinct beliefs give rise to the same best response, Bayesian Nash equilibrium beliefs (and actions) converge to the (best response to) the true state with probability one.

Nevertheless, we suspect that people in such informationally-rich settings may wind up making wrong choices even in the long-run. In this section, we show how BRTNI captures a very intuitive error in reasoning that people may make in herding contexts that can lead to incorrect herding in even this informationally rich setting. We then

Since BRTNI players best respond to other players' best responding to "cursed beliefs", we begin by exploring cursed best response. In a fully cursed best response, players do not appreciate the informational content of their predecessors' actions. Consequently, their actions depend only upon their own signals and not upon previous play. In BRTNI play, each player believes that her predecessors' actions depend upon their signals alone, and so each A choice is interpreted as another piece of evidence favouring A. Unlike rational players, each successive BRTNI player herding on A becomes more convinced that A is in fact the true state of the world.

Recently, such models have been tested extensively in the laboratory. While some predictions of the rational model have found support, several systematic discrepancies have been uncovered. Anderson and Holt (1997) and Hung and Plot (2001) find that when rational inference requires subjects to ignore their own signals they do so more often that not but far less than all the time. Nöth and Weber (1999), Kariv (2003) and Goeree, Palfrey, Rogers and McKelvey (2004) report that subjects tend to follow their own private information more often than they should. Kübler and Weizsäcker (2004) show in a variant of the model (described in more detail below) that subjects' beliefs become too extreme to be well explained by rational play. Kübler and Weizsäcker (2005) report the related finding that longer cascades are more stable, intuitively because after a long string of A choices, people come to believe A more and more likely, reducing the likelihood that anyone will break the cascade by choosing B.

One widely recognized limitation of this literature is that the action space is too coarse to represent many contexts where people's actions reveal the intensity of their beliefs. Finer action spaces reduce the scope for rational herding, and in the limiting case where no two distinct beliefs give rise to the same best response, equilibrium actions converge to the best response to the true state with probability one. Nevertheless, we suspect that people may herd in such settings. In this section, we show how BRTNI play can lead to incorrect herding in even this informationally rich setting.

3.1 A Rich Model with BRTNI Herding

There are two possible states of the world, $\omega \in \{0, 1\}$, each equally likely *ex ante*. Each player k receives the signal $s_k \in [0, 1]$; signals are independent and identically distributed conditional on the state. When $\omega = 0$, signals have the density function f_0; when $\omega = 1$, they have density f_1. Each player observes her signal and the actions of all previous players before choosing an action in $[0, 1]$. For simplicity, we assume that the model is symmetric: for each $s \in [0, 1]$, $f_0(s) = f_1(1 - s)$, and we normalize signals such that $s = \Pr[\omega = 1|s]$. In addition, we assume that the likelihood ratio $L(s) \equiv \frac{f_0(s)}{f_1(s)}$ is continuously differentiable with image $\mathbb{R}_+$ and derivative $L'(s) < 0$. Let $a_k(a_1, \ldots, a_{k-1}; s_k)$ be the action taken by the kth player as a function of previous players' actions and her own private information, and let $a \equiv (a_1, a_2, \ldots) \in [0, 1]^{\mathbb{N}}$ be the profile of all players' actions. Every Player k has payoff function $g_k(a; \omega) = -(a_k - \omega)^2$, which is maximized by setting $a_k = E[\omega|F_k]$, where F_k is all the information available to Player k.

This social-learning environment provides players with two sources of rich information. First, an unbounded likelihood ratio of players' private signals means that sometimes players receive arbitrarily strong signals about the identity of the true state. Second, by choosing actions in the continuum, players reveal to their successors their posteriors. Either of these features suffices to guarantee that rational players form beliefs and choose actions that converge to the true state with probability one (Smith and Sorenson, 2001); that is, rational players would converge with probability one to correct beliefs and actions with binary signals and a continuum of actions or a continuum of signals and binary actions.

In BRTNI play, each agent thinks that all previous agents are fully cursed, whereas in reality no agent is cursed. Clearly $a_1(s_1) = \Pr[\omega = 1|s_1] = s_1$: the first agent follows her signal. For the remainder of the analysis, it is easier to work with the log odds ratios $\ln\left(\frac{a}{1-a}\right)$, the log of the ratio the agent's beliefs that $\omega = 1$ versus $\omega = 0$. For the first agent, $\ln\left(\frac{a_1}{1-a_1}\right) = \ln\left(\frac{s_1}{1-s_1}\right)$. Because the first player would follow her signal whatever her cursedness, the second player correctly infers the first player's signal from her action and chooses

$$\begin{aligned}\ln\left(\frac{a_2}{1-a_2}\right) &= \ln\left(\frac{a_1}{1-a_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right) \\ &= \ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right),\end{aligned}$$

just as she would in a Bayesian Nash Equilibrium. The third player believes the second to be fully-cursed. Because a fully-cursed Player 2 would choose $\ln\left(\frac{a_2}{1-a_2}\right) = \ln\left(\frac{s_2}{1-s_2}\right)$, Player 3 mistakenly ignores the effect that a_1 has on a_2. Hence he chooses

$$\begin{aligned}\ln\left(\frac{a_3}{1-a_3}\right) &= \ln\left(\frac{a_1}{1-a_1}\right) + \ln\left(\frac{a_2}{1-a_2}\right) + \ln\left(\frac{s_3}{1-s_3}\right) \\ &= \ln\left(\frac{s_1}{1-s_1}\right) + \left(\ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right)\right) + \ln\left(\frac{s_3}{1-s_3}\right) \\ &= 2\ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right) + \ln\left(\frac{s_3}{1-s_3}\right).\end{aligned}$$

The third player's action differs from the optimal choice by overweighting the first signal. Intuitively, because Player 3 ignores how Player 2's action depends upon Player 1's action and hence upon Player 1's signal, Player 3 weights Player 1's signal more than she should. Similarly,

$$\begin{aligned}\ln\left(\frac{a_4}{1-a_4}\right) &= \ln\left(\frac{a_1}{1-a_1}\right) + \ln\left(\frac{a_2}{1-a_2}\right) + \ln\left(\frac{a_3}{1-a_3}\right) + \ln\left(\frac{s_4}{1-s_4}\right) \\ &= \ln\left(\frac{s_1}{1-s_1}\right) + \left(\ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right)\right) \\ &\quad + \left(2\ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right) + \ln\left(\frac{s_3}{1-s_3}\right)\right) + \ln\left(\frac{s_4}{1-s_4}\right) \\ &= 4\ln\left(\frac{s_1}{1-s_1}\right) + 2\ln\left(\frac{s_2}{1-s_2}\right) + \ln\left(\frac{s_3}{1-s_3}\right) + \ln\left(\frac{s_4}{1-s_4}\right).\end{aligned}$$

In general,

$$\ln\left(\frac{a_t}{1-a_t}\right) = \sum_{\tau<t} 2^{\tau-(t-1)} \ln\left(\frac{s_\tau}{1-s_\tau}\right) + \ln\left(\frac{s_t}{1-s_t}\right).$$

(By contrast, rational players give all signals equal weight.) Relative to rational play, BRTNI players overweight early signals, giving the first signal half the weight of all signals, the second half of what remains, etc. Players exhibit a social confirmatory bias, as early signals exert undue influence on subsequent beliefs and actions.

Because BRTNI play weights early signals so heavily, it seems possible that even an arbitrarily large number of players may fail to learn the true state in the event that the first few players have inaccurate signals. On the other hand, the fact that the likelihood ratio goes to infinity at $s \in \{0, 1\}$ allows players to receive arbitrarily strong signals of the state. If arbitrarily strong signals occur frequently enough, then players should learn the true state. If not, then they may "herd" on wrong beliefs and actions.

The following proposition shows that BRTNI players may herd on wrong beliefs and actions not only in the example but in the more general model outlined above, where the key assumption is the lack of any positive-probability "magic signal" that reveals the true state with certainty.

Proposition 1: Suppose that $var\left(\ln\left(\frac{S}{1-S}\right)\right) = \sigma^2 < +\infty$. Then in BRTNI play for each $k \in \left(\frac{1}{2}, 1\right)$ there exists $\delta > 0$ such that $\Pr[a_t > k \ \forall t | \omega = 0] > \delta$.

Proposition 1 establishes that when $\omega = 0$ there is positive probability that BRTNI players all choose actions that exceed any given threshold lying above one-half. The result is striking because the information structure allows players to receive arbitrarily strong signals that the state is $\omega = 0$ as well as to transmit their posteriors to succeeding players. Yet if the first couple of agents receive signals high enough to take actions above k, then with positive probability no agent ever takes an action below k. This occurs because of the speed with which BRTNI players come to believe that $\omega = 0$ is the true state.

The assumption in the proposition that the log likelihood ratio of signals can take on any real value implies that BRTNI players never observe actions that they deem impossible. (See Footnote 8 on this topic.) The assumption of finite variance rules out positive-probability "magic signals" that reveal the true state with probability one; clearly this would preclude incorrect herding in many models—fully or boundedly rational—becuase eventually some player would receive a private signal that reveals the true state and act accordingly; succeeding players should then follow suit, precluding incorrect herding.

Proposition 1 does not tell us that BRTNI beliefs converge. But since public beliefs form a bounded submartingale when above than one-half and a bounded supermartingale when below one-half, they do converge with probability one to a limiting distribution that can only assign positive probability to zero or one. Consequently, BRTNI players converge to the wrong actions and beliefs with positive probability.

While we prove the result in a model with a continuum of actions and a continuum of signals, it holds equally well in models with a finite number of actions and signals. (Of course, with finite actions, rational play can also be consistent with herding on the wrong state.) [6]

To illustrate Proposition 1, consider the case where the densities are $f_0(s) = 2(1-s)$ and $f_1(s) = 2s$ as illustrated in Figure 1.

[6] Finite models complicate our solution concept since BRTNI players observe actions inconsistent with their model of the world. To see this, consider a modification of our model that leaves the action space intact but reduces the signal space to a finite set. Whatever the quality of the strongest signal that $\omega = 1$, there exists an action $\widehat{a}$ higher than Bayesian posteriors following observation of that signal alone. While BRTNI players believe no action $a \geq \widehat{a}$ will ever be played, some must as actions converge to one. This raises the question of what beliefs BRTNI should adopt after observing their predecessors make choices ruled out by their model of the world. We could address this limitation of our concept by extending it to assume that a player who observes her predecessor behave in a way inconsistent with full cursedness believes that this predecessor is χ-cursed for as small a χ as possible consistent with the history of play. Because this predecessor is in fact uncursed, such a χ must exist. Intuitively, this extension forces a player who observes her predecessor choose an action too high to be consistent with naive inference to believe that her predecessor received the highest possible signal. It is then easy to see how this allows for positive probability of false herding.

Because exteding BRTNI in this way would lengthen the paper more than it would shed any light on irrational herding, we choose to leave this sort of foundational issue for future work.

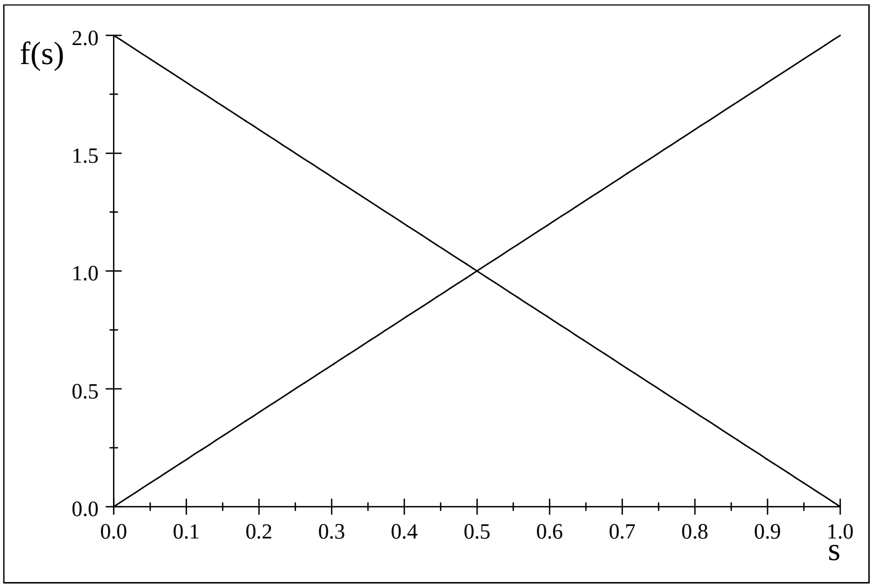

When $\omega = 0$, signals come from a triangular distribution with mode at $s = 0$, and when $\omega = 1$ they come from a triangular distribution with mode $s = 1$. Note that the signals $s \in \{0, 1\}$ reveal the state but occur with probability zero. The table below reports simulations of BRTNI as well as Bayesian Nash play for these distributions when $\omega = 1$.

Player	$\Pr[a_t < 0.05]$ BRTNI	$\Pr[a_t < 0.05]$ BNE	$\Pr[a_t > 0.95]$ BRTNI	$\Pr[a_t > 0.95]$ BNE
1	0.0025	0.0026	0.0977	0.0976
2	0.0058	0.0060	0.3030	0.3035
3	0.0216	0.0070	0.5965	0.4871
4	0.0483	0.0069	0.7640	0.6247
5	0.0739	0.0060	0.8332	0.7232
6	0.0914	0.0051	0.8623	0.7954
7	0.1016	0.0041	0.8754	0.8477
8	0.1068	0.0033	0.8815	0.8857
9	0.1098	0.0026	0.8845	0.9148
10	0.1115	0.0020	0.8856	0.9356

The table reports the probabilities of the various players choosing actions that are either very high or very low under the two different solution concepts. Since BRTNI and BNE coincide for the first two players, these should be the same; the small differences appearing in the table derive from the fact that different signals were simulated for the two solution concepts. For each, the likelihood that the second player chooses a very low action is about 0.006. A rational Player 3 more likely than not chooses a higher action than Player 2 since when $\omega = 1$, most signals move posteriors in that direction. Indeed, for rational players, the likelihood that Players 2 and 3 choose low actions is similar. BRTNI Player 3's, however, are more than three times as likely as their predecessors to choose a low action. Intuitively, this is because she interprets Player 1 and 2's low actions as two very strong pieces of evidence in favour of $\omega = 0$, in which case she needs a very high signal to choose an action above 0.05. Moving down Column 2 to examine later players' actions suggests that BRTNI player converge to $a = 0$ when $\omega = 1$ with probability around 11 percent. Column 4 suggests that this cannot occur with rational players, who are by Player 10 only 2 percent as likely as BRTNI players to be choosing low actions.

Another interesting feature of BRTNI play is the speed of convergence. Summing entries in the last row from Columns 2 and 4 gives that 99.71% of BRTNI Player Tens play actions below 0.05 or above 0.95. By contrast, only 93.58% or rational Player Tens do. Although we have no formal result along these lines, the observation suggests that BRTNI play converges faster than rational play.

Since BRTNI play appears to converge fast—and we know with positive probability to the wrong thing—a natural question is what the world looks like when play is slow to converge.

Proposition 2: For each interval $[c, d] \subset \mathbb{R}_{++}$ there exists $T \in \mathbb{N}$ such that if for each $t \in \{1, \ldots, T\}$, $\ln\left(\frac{a_t}{1-a_t}\right) \in [c, d]$ under BRTNI play, then

$$\Pr[\omega = 0 | a_1, \ldots, a_T] > \Pr[\omega = 1 | a_1, \ldots, a_T].$$

The Proposition establishes that if all players for many periods believe $\omega = 1$ more likely than $\omega = 0$—they take actions a above one-half such that $\ln\left(\frac{a}{1-a}\right) > 0$—then in fact it is more likely that $\omega = 0$ than that $\omega = 1$. In BRTNI play, each player believes that all of her predecessors' actions coincide with their signals. A player at the end of a long run of high actions believes that her predecessors must all have high signals. So, the only reason why she would not conclude that

$\omega = 1$ with virtual certainty must be that she receives a very low signal herself. Hence, the way that players can take actions above one-half for many periods without any one of them taking an action sufficiently close to one is that if after a while all of those players receive low signals; this event indicates that zero is more likely than one to be the true state.

A limitation of Proposition 2 is that it only applies to settings where each of the first T players chooses an action suggesting that the state is more—but not too—likely to be one than zero. A stronger result would be that for any $d > 0$ there exists some integer T such that if we know nothing other than that Player T chooses an action whose log odds is in $(0, d)$, then it is more likely that zero is the true state. While we lack any formal result along those lines, simulation results for the parametric example introduced above suggest it to be true. The following table presents the probabilities of BRTNI Player Tens through Player Twenties playing actions in various intervals *when* $\omega = 1$:

P	$a_t^{BRTNI} \in (0.05, 0.50)$	$a_t^{BRTNI} \in (0.50, 0.95)$
10	0.00143435	0.00144879
11	0.00071807	0.00071808
12	0.00036055	0.00035963
13	0.00017987	0.00017923
14	0.00008881	0.00008972
15	0.00004537	0.00004535
16	0.00002306	0.00002254
17	0.00001168	0.00001131

Players 14 through 17 choose actions that are bounded away from the limit and below one-half more frequently than they do actions bounded away from the limit and above one-half. Because the model is symmetric, this implies that when, say, Player 17 chooses an action in $(0.05, 0.50)$ it is more likely that $\omega = 1$ than $\omega = 0$; that is, Player 17's weak beliefs that $\omega = 0$ is more likely than $\omega = 1$ are on average wrong. Of course, the fact that BRTNI players' beliefs converge so quickly with such high probability limits the empirical relevance of this result as well as that contained in Proposition 3.

4 Limits on Observation

BRTNI players overweight the signals of each mover by counting them again and again through each predecessor's action as well as directly. For instance, Player 3 double-counts Player 1's action by counting it once through Player 2's action and then once again directly. Naturally, a Player 3 who cannot observe Player 1's action cannot double count Player 1's information in this way. This suggests that BRTNI players who only observe their immediate predecessor's action behave in the same way as rational players. Indeed, if each BRTNI player can only observe her immediate predecessor, then BRTNI and rational play coincide. However, milder limits on players' observation of past play do not overturn our main result of positive probability herding on the wrong action and beliefs.

Proposition 3: When BRTNI players can only observe their two immediate predecessors' actions, the result of Proposition 1 holds.

In the benchmark model, the signal of player t comes to have twice the weight of the signal of player $t+1$ in actions in period $t+2$ and beyond. When BRTNI players can observe the previous two players' actions, as t approaches infinity this ratio converges to φ, the golden ratio. Intuitively, limiting every player to observing only her immediate two predecessors only takes effect with Player 4. Since Player 3 has already overweighted Player 1's signal, Player 4 will as well. And when Player 4 overweights Player 2's action, he also overweights Player 1's signal. The proof for Proposition 1 goes through almost directly when the weight 2 is replaced by φ. Likewise, having players observe their three predecessors' actions changes the limiting ratio to ρ, the ratio of terms in the generalized Fibonacci sequence $s_t = s_{t-1} + s_{t-2} + s_{t-3}$, or a root of the cubic $x^3 - x^2 - x - 1$. In fact, the conclusion of Proposition 1 goes through when players can observe their immediate k predecessors for any $k > 1$.

The result that BRTNI players may come to hold wrong limiting beliefs and take the wrong limiting actions holds equally well when players cannot observe the order of their predecessors' play.

Proposition 4: When BRTNI players can observe all their predecessors' actions, but not the order in which those actions were taken, the result of Proposition 1 holds.

Rational players' actions depend upon the order of their predecessor's moves. For instance, Player 3 would like to combine Player 2's action with her own signal and ignore Player 1's action. But a Player 3 who cannot observe the order of her predecessors' moves cannot do that in our model where every action corresponds to an optimal action given some priors and signal realization. BRTNI players, however, do not attend to the order of their predecessors play because they believe that each of their predecessors simply follows his signal and that all signals have equal validity. Since observing the order of predecessors' play does not affect BRTNI players, clearly any result that we establish with observability holds equally well without. For instance, Proposition 3 would continue to hold if BRTNI players could not observe the order of their two immediate predecessors' moves.

We believe that Propositions 3 and 4 demonstrate an important robustness of our results. It is hard to imagine a setting where people know that they are in a social-learning environment but can observe no more than the action of their immediate predecessor. But in many settings it seems unrealistic to know which predecessor moved when, and people may simply receive summary statistics of their predecessors actions.

5 Cursing BRTNI

BRTNI play converges quickly—sometimes the wrong beliefs and action—because public beliefs become very extreme very quickly. When players are in fact partially cursed, then they will underinfer the predecessors' information from their actions—rightly or wrongly—preventing public beliefs from becoming so extreme. Hence, cursedness in this information-cascade setting works to counteract BRTNI inference. This section explores the extent to which cursedness overturns our main results in a simplified version of the previous model where actions are binary: $a \in \{0, 1\}$. As before, we allow for continuous signals.

Definition 5 *The strategy profile* $(\sigma_1, \ldots, \sigma_N)$ *is* (χ-cursed best response trailing naive inference) χ-cursed BRTNI play *if for each type* t_k *of each Player* k*,* $\sigma_k(a_k|t_k) > 0$ *only if* $a_k \in CBR_k^{\chi}\left(t_k; CBR_{-k}^{1}\right)$.

In χ-cursed BRTNI play, players play a χ-cursed best response to beliefs that others play a fully cursed best response to some beliefs. The concept of χ-cursed best response captures the idea that players underappreciate the information content in others' play, while BRTNI means that to the extent that players appreciate that there is information content in others' play they misconstrue it by failing to perceive that other players make the informational inference that they do.

Proposition 5: In the binary-action model, there exists $\underline{\chi} \in (0,1)$ such that for each $\chi \in \left(0, \underline{\chi}\right)$, when $\omega = 0$ public beliefs in $\chi-$cursed BRTNI play converge to $1 - \frac{\chi}{2}$ with positive probability.

Unlike BRTNI and rational play, cursed play need not converge as players overweight their own private information. For instance, when $\chi = 1$, players follow their signals, and, hence, play does not converge. Public beliefs, however, do converge, in this case to one-half as no player infers anything from her predecessors' moves. Proposition 5 establishes that limited cursedness does not prevent public beliefs from converging to something close to the wrong state. In the binary-action model, public beliefs in period t depend only on the number of players before period t playing $a = 1$ minus the number playing $a = 0$. Suppose that many players choose $a = 1$ such that public beliefs are close to $1 - \frac{\chi}{2}$. If not too cursed, then a χ-cursed BRTNI player chooses $a = 1$ with probability greater than one-half with probability greater than one-half. Hence the number of players choosing $a = 1$ less those choosing $a = 0$ is a random walk with positive drift, which with positive probability never returns to its current position.

While Proposition 5 is a limiting result in χ, the degree of cursedness that is compatible with wrong limiting beliefs can be substantial. In our example above where $f_0(s) = 2(1-s)$ and $f_1(s) = 2s$—signals are uniformly distributed—χ need only be smaller than $2 - \sqrt{2} \simeq 0.59$.

6 Heterogeneity and Other Extensions

Sections 3 and 4 treated the case where all players are BRTNI, and Section 5 showed that a limited form of our main result holds when all players are cursed and BRTNI. In this section, we explore play when some players differ in their strategic sophistication by assuming that some fraction $\frac{1}{n}$ of players are BRTNI with the remainder rational. For simplicity, suppose that the rational players know which of their predecessors are rational and which BRTNI. As usual, BRTNI players believe that all predecessors are rational. Let $\mathcal{B}$ denote the set of players who are BRTNI.

Proposition 6: Suppose that $var\left(\ln\left(\frac{S}{1-S}\right)\right) = \sigma^2 < +\infty$ and that for some $n \in \mathbb{N}$, Players $n, 2n, \ldots$ are BRTNI and the remainder rational. Then there exists some $\delta > 0$ such that $\Pr\left[\lim_{t \in \mathcal{B}, t \to \infty} a_t = 1 \middle| \omega = 0\right] > \delta$.

Proposition 6 asserts that our main qualitative result—with positive probability, BRTNI players wind up with wrong limiting beliefs and actions—remains valid no matter their share of the population. It does not, however, imply that rational players among BRTNI players come to hold misleading beliefs and choose wrong actions. Indeed they will not. Yet we abstract from the issue of how the presence of BRTNI players slows rational learning with the extreme and implausible assumption that rational players know which of their predecessors are rational and which BRTNI.

7 Discussion and Conclusion

BRTNI players believe that each of their predecessors' actions reflects that predecessor's private information alone; if true, then those action follow the same distribution as players' signals do in one of the two possible states of the world. But BRTNI players' actions very much *do not* conform to either of these distributions; instead, they converge either to zero or one. Hence, a BRTNI player with some meta-awareness that she might hold the wrong model of the world should come to recognize the unlikeliness of observed play and consequently update to become "less BRTNI." Intuitively, BRTNI players far enough along in the sequence might ask themselves, say, why their one hundred immediate predecessors all chose actions very close to one. Could this really happen when each followed her signal? Since the likelihood of this event is small, BRTNI players may conclude that they likely hold the wrong model of the world. This suggests a potential non-robustness of our results.

The more cursed players are, however, they more they behave as their successors expect them to behave. In the binary-action model treated above, cursed BRTNI players do not herd: because the likelihood ratio has full support, starting from any point in the game, both actions are played in the next period with positive probability. Hence, BRTNI players' observations of their predecessors' play may not differ significantly enough from their theory as to how those predecessors should play to cause them to abandon that theory. This suggests some robustness to our result in the case where BRTNI players are cursed, namely that cursed BRTNI players' public beliefs converge with positive probability to the "wrong" limit.

8 Appendix

Proof of Proposition 1: Choose $k \in \left(\frac{1}{2}, 1\right)$ and let $K = \ln\left(\frac{k}{1-k}\right)$. Let P_t be the log likelihood of public beliefs in period t, and note that with BRTNI play $P_{t+1} = 2P_t + \ln\left(\frac{s_t}{1-s_t}\right)$. When $\omega = 0$, with positive probability $P_2 \geq 3K$. If $\ln\left(\frac{s_t}{1-s_t}\right) > -tK$ for each t, then $P_3 = 2P_2 + \ln\left(\frac{s_2}{1-s_2}\right) > 2{\cdot}3K - 2K = 4K$, and then $P_4 = 2P_3 + \ln\left(\frac{s_3}{1-s_3}\right) > 2{\cdot}4K - 3K = 5K$, etc. In general, $P_t > (t+1)K$, and since $\ln\left(\frac{a_t}{1-a_t}\right) = P_t + \ln\left(\frac{s_t}{1-s_t}\right) > (t+1)K - tK = K$ as desired. From Chebyshev's Inequality, $\Pr\left[\ln\left(\frac{s_t}{1-s_t}\right) > -tK\right] > \frac{t^2K^2-\sigma^2}{t^2K^2}$, and hence

$$\begin{aligned}
\Pr\left[\forall t, \left(\frac{s_t}{1-s_t}\right) > e^{-tK}\right] &> \prod_t \frac{t^2K^2-\sigma^2}{t^2K^2} = \exp\left\{\sum_t \ln\left(\frac{t^2K^2-\sigma^2}{t^2K^2}\right)\right\} \\
&= \exp\left\{\sum_t -\frac{\sigma^2}{z_t}\right\},
\end{aligned}$$

for $z_t \in \left(t^2K^2 - \sigma^2, t^2K^2\right)$, by the Mean-Value Theorem. Hence

$$\Pr\left[\forall t, \left(\frac{s_t}{1-s_t}\right) > e^{-tK}\right] > \exp\left\{\sum_t -\frac{\sigma^2}{t^2K^2}\right\} = \exp\left\{-\frac{\sigma^2\pi}{6K^2}\right\} > 0.$$

Q.E.D.

Proof of Proposition 2: Let $[c, d] \subset \mathbb{R}_{++}$ be given. Define $T_1 = \left\lfloor \frac{d}{c} + 1 \right\rfloor$, so that $(T_1 - 1)c \leq d <$

$T_1 c$. Choose $\delta \in (0, T_1 c - d)$. For Player $T_1 + 1$,

$$\begin{aligned}
\ln\left(\frac{a_{T_1+1}}{1-a_{T_1+1}}\right) &= \sum_{\tau < T_1+1} \ln\left(\frac{a_\tau}{1-a_\tau}\right) + \ln\left(\frac{s_{T_1+1}}{1-s_{T_1+1}}\right) \\
&> T_1 c + \ln\left(\frac{s_{T_1+1}}{1-s_{T_1+1}}\right).
\end{aligned}$$

If $\ln\left(\frac{a_{T_1+1}}{1-a_{T_1+1}}\right) \leq d$, then $\ln\left(\frac{s_{T_1+1}}{1-s_{T_1+1}}\right) < -\delta$. The same is true for Player T_1+2 and so forth. Now pick T_2 such that $T_1 d - \delta T_2 < 0$ and set $T = T_1 + T_2$. We claim that if $\ln\left(\frac{a_t}{1-a_t}\right)$ for each $t \in \{1, \ldots, T\}$, then $\Pr[\omega = 0|\,(a_1, \ldots, a_T)] > \Pr[\omega = 1|\,(a_1, \ldots, a_T)]$. To see that, note that the first T_1 players have signals with log likelihoods no larger than d (otherwise one would choose an action with log odds above d), and the next T_2 have signals with log likelihoods no larger than δ. Since $T_1 d - \delta T_2 < 0$, Bayesian beliefs after T periods ascribe higher probability to the state being zero than one. **Q.E.D.**

Proof of Proposition 3: Similar to Proof of Proposition 1 and hence omitted.

Proof of Proposition 4: Since BRTNI players do not heed the order of their predecessors' moves, it follows from Proposition 1.

Intuition for Proposition 5: Assume that $f_0(s) = 2(1-s)$ and $f_1(s) = 2s$. Suppose first K players get signals above $\frac{1}{2}$, with K large enough that public beliefs are $1 - \frac{\chi}{2} - \varepsilon$ for $\varepsilon > 0$ small. Player $K+1$ chooses $a = 1$ with

$$\Pr\left[s : \frac{s\left(1 - \frac{\chi}{2} - \varepsilon\right)}{s\left(1 - \frac{\chi}{2} - \varepsilon\right) + (1-s)\left(\frac{\chi}{2} + \varepsilon\right)} \geq \frac{1}{2} \middle| \omega = 0 \right].$$

Letting $\varepsilon \to 0$, this $> \frac{1}{2}$ when $\chi < 2 - 2^{1/2} = 0.586$. So, after K instances of $a = 1$, $a_{K+1} = 1$ with probability $p > \frac{1}{2}$.

Public beliefs at time t depend only on Markov process

$$n(t) = \#\{\tau < t : a_\tau = 1\} - \#\{\tau < t : a_\tau = 0\},$$

and since

$$\Pr[a_t = 1 | n(t) > K] > \Pr[a_t = 1 | n(t) = K] = p > \frac{1}{2},$$

$\Pr\left[\exists \widehat{t} > K : n\left(\widehat{t}\right) = K\right] <$ under random walk with $\Pr[a_t = 1] = p \ \forall t$. Since this < 1, there is positive probability that $n(t) > K \ \forall t > K$. In this case, public beliefs cannot converge to $\frac{\chi}{2}$ and must converge with probability one to $1 - \frac{\chi}{2}$. To conclude, when $\omega = 0$ public beliefs converge to $1 - \frac{\chi}{2}$ with positive probability.

Proof of Proposition 6: To be completed.

References

[1] Banerjee, Abhijit. "A Simple Model of Herd Behaviour." *Quarterly Journal of Economics* 107(3), August 1992: 797-817.

[2] Bernheim, B. Douglas. "Rationalizable Strategic Behavior". *Econometrica*, Vol. 52(4), July 1984: 1007-1028

[3] Bikhchandani, S., D. Hirshleifer, and I. Welch. "A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades." *Journal of Political Economy* 100, 1992: 992-1026.

[4] Esponda, Ignacio. "Behavioral Equilibrium in Economies with Adverse Selection." Forthcoming, *American Economic Review.*

[5] Eyster, Erik and Matthew Rabin. "Cursed Equilibrium." *Econometrica* 73(5), September 2005: 1623-1672.

[6] Goeree, Jakob, Thomas R. Palfrey, Brian W. Rogers, Richard D. McKelvey. "Self-Correcting Information Cascades" *Review of Economic Studies* 74 (3), 2007: 733–762

[7] Hung, Angela A. and Charles R. Plott. "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding." *American Economic Review*, 2001: 1508-1520.

[8] Jehiel, Philippe and Frederic Koessler. "Revisiting Games of Incomplete Information with Analogy-Based Expectations," *Games and Economic Behavior*, forthcoming.

[9] Kübler, Dorothea and Georg Weizsäcker. "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory." *Review of Economic Studies* 71 (2), 2004: 425–441

[10] Kübler, Dorothea and Georg Weizsäcker. "Are Longer Cascades More Stable?" *Journal of the European Economic Association*, 3(2-3), 2005.

[11] Pearce, David G. "Rationalizable Strategic Behavior and the Problem of Perfection". *Econometrica*, Vol. 52(4), July 1984: 1029-1050

[12] Smith, Lones and Peter Sorensen. "Pathological Outcomes of Observational Learning." *Econometrica* 68, 2001: 371-398.