Naive Herding

Erik Eyster
Department of Economics
LSE

Matthew Rabin
Department of Economics
University of California—Berkeley

27 August 2008

Abstract

In a social-learning environment, we investigate implications of the assumption that players naively believe that a previous player's actions reflect solely that player's private information. This error leads players to inadvertently over-weight early players' private signals by neglecting that interim players' actions also embed these signals. Consequently, naive players can become extremely and wrongly confident about the state even where rational players never become confident and may herd on incorrect actions even where rational players never do. Wrong herding can happen when naive players observe no more than two previous actions and also when other players are rational or under-infer.

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

Acknowledgments: This model originally appeared as part of a more general paper entitled "Inferential Naivety", first presented in February 2006 and circulated in November 2006. On this or the earlier paper, we thank Zack Grossman for valuable research assistance and Ignacio Esponda, Guy Mayraz, Marco Ottaviani, Peter Sørensen, and Georg Weizsäcker for helpful conversations. We also thank seminar participants in 2006 at Berkeley, Birkbeck-UCL, Davis, ESSET Gerzensee, Harvard/MIT, and Lausanne; in 2007 at Bocconi, Brescia, ESEM 2007, Toulouse, and Wash U.; and in 2008 at Hebrew, Naples, Tel Aviv and UCL. Rabin thanks the NSF for financial support, and the Trento, Italy Police Station for hospitality when some of the work on this paper was completed.
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 sequence of people each choose in turn one of two options, A or B, each person observing all of her predecessors' choices. They have common preferences over the two choices but do not know which is better. Rather, they receive independent and equally strong private binary signals that provide information about the right choice. In this setting, rational agents herd: if the first two choose A, then so too do all who follow, regardless of their private information. This is because two A choices mean (given a natural tie-breaking rule for the second mover) that each party must have received a signal indicating that A is better; the next mover should realize this and, even if she gets a B signal, choose action A. And all subsequent players follow suit. Although everyone chooses A, each understands that all but the first two movers ignore their signals and, hence, never becomes highly confident that A is better. Generalizing this result, the social-learning literature finds that when the action space and signal space are both finite, and each signal is imperfect, rational people almost always eventually herd, but with positive probability on an action that is suboptimal.

Despite its many insights, this literature has some important limitations. The logic of when rational people herd relies crucially on features of the environment that seem neither realistic nor intuitively central to herding. As is widely recognized, *either* richer action spaces *or* richer signal spaces reduce the scope for herding; in the limiting cases where a person's action fully reveals her beliefs or people can receive fully revealing signals of the state, equilibrium actions almost surely converge to the best one. And although intuition and evidence suggest that the rationality assumption maintained in these models may be importantly unrealistic, relatively little is known about the implications of more realistic assumptions. In this paper, we investigate herding by players who are more limited in their rationality. In particular, we investigate the implications of the assumption that players naively believe that a previous player's actions reflect solely that player's private information. This error leads players to inadvertently over-weight early players' signals by failing to appreciate that interim players' actions also embed those signals. As a consequence, naive players can become extremely and wrongly confident about the state of the world even in environments where rational players never become confident, as well as to herd on incorrect actions even in environments where fully rational players always herd on correct actions.

In Section 2, we formalize two types of departure from purely rational play that correspond to two distinct general errors in strategic inference that people may make and also accord with intuition and some experimental evidence on herding. First, players may be "cursed"—and fail to appreciate the link between previous movers' actions and those movers' private signals.[1] A cursed player bases her action upon her own signal alone and not upon previous play. More plausibly in social-learning environments, players might be "inferentially naive"—and realize that previous movers' actions reflect these movers' own signals but fail to appreciate that these previous movers themselves also perform inference from still earlier movers' actions. Players making this error believe that each predecessor's action depends upon that predecessor's signal alone; inferentially naive players believe that their predecessors are cursed. For each of these two types of error, we define both weak and strong solution concepts, and in the remainder of the paper primarily apply only the weaker versions.

In Section 3, we study this latter error of inferential naivety and show it to have dramatic implications. We begin with the classic two-state, finite-action, finite-signal model already introduced. In the example above, inferential naivety leads herders to interpret each prior

[1] In this paper we define and study a weaker form of "cursedness" that is in the spirit of Eyster and Rabin's (2005) "cursed equilibrium" and similar to the "coarse analogy partition" variant of Jehiel's (2005) "analogy-based expecations equilibrium" as applied to Bayesian games by Jehiel and Koessler (2008).

A choice as more evidence favouring A. Unlike rational players, each successive naive player following a herd on A becomes more convinced that A is better, and eventually herders come to believe with certainty that A is better. More generally, we show that inferential naivety almost surely leads players to converge to fully confident beliefs that their actions are optimal. 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. We construct an example, in fact, where inferential naivety leads players to converge with probability arbitrarily close to one to fully confident beliefs—that are wrong.

To illustrate more starkly how inferential naivety may lead to erroneous confident beliefs, we continue Section 3 by constructing a related model where actions and signals both can take on any value in the continuum. In such a setting, each player takes an action that fully reveals her beliefs; we assume that each signal can be anything from virtually fully revealing of the state to virtually uninformative. In this informationally rich environment, rational herders almost surely converge to the correct beliefs about the state. But even here, with positive probability naive herders converge to certain belief in the wrong state. Indeed, we show by simulation of a simple model where each player receives a signal whose strength is uniformly distributed that approximately 11% of the time the herd converges to beliefs and actions that represent *certainty* in the wrong state.

What is the intuition for this? Not realizing that the second mover's action reflects beliefs that combine the first and second movers' signals, the third mover's attempt to incorporate both predecessors' signals into her own action leads her in fact to count the first mover's signal twice. Naive inference by the fourth mover, in turn, causes him to count the first mover's signal four-fold: once from the first action, once from the second, and twice from the third. He counts the second signal twice: once each from the second and third mover. Naive herders in this model are massively over-influenced by the 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 mislead, limit herds may so over-use it as to outweigh an infinite sequence of further signals—even if this sequence includes an infinity of very strong signals of the truth.

We conclude Section 3 with 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 one. The intuition behind the result relates to its limited applicability: it is, in fact, very unlikely that beliefs remain weak after a long sequence signals, but the most likely way is if herders 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.

In Section 4 we briefly consider the implications of inferential naivety in a couple variants of this social-learning setting. First, we consider the case of limited observability where, instead of observing all previous moves, each player observes only her immediate $k < \infty$ predecessors. Interestingly, if $k = 1$, meaning that each player observes solely the move before him, then—despite having the wrong model of others' decision making—inferential naivety does not lead players astray. Each player mistakenly believes her predecessor's action reflects solely that predecessor's signal when in fact it depends on his own predecessor's action and, hence, signal. But because in this case each person correctly infers her predecessor's beliefs (despite having the wrong theory of those beliefs' provenance) and lacks the observations to double count earlier movers' signals, all signals are counted rationally. When $k = 2$, however, players once more can and do double count, and again with positive probability they 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 Section 4 we illustrate how inferential naivety also can lead players to overestimate how much information their actions embody. We do this in a binary-state model where the state follows a Markov process, changing with some fixed probability each period. Each player chooses whether to pay some cost to receive a fully-revealing signal of the state before choosing her action, both after observing her predecessors' actions but not purchasing decisions. Rational players in this model follow cyclical strategies: Player 1 buys a signal, and succeeding players herd until enough periods have elapsed that Player 1's action provides a sufficiently weak signal of the current state to warrant another player to purchase a signal, and so forth. Each naive player, by contrast, erroneously believes that all of her predecessors have bought signals and therefore refrains from buying one herself. When the state is very persistent, naive players far enough along in the sequence of players believe that they know they know the state with probability close to one, whereas the true probably hardly exceeds one-half.

In Section 5, we explore how inferential naivety may combine with other types of play. We first explore what happens when all players are inferentially naive and to some degree "cursed". 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, inferential naivety says not just that herders infer too much from previous actions, but more specifically that far too much weight accrues to early relative to late signals. Indeed, cursedness does not undo our core result than inferential naivety may lead to certain but wrong beliefs. To capture this intuition, we show in a binary-action variant of our primary model that if everybody is inferentially naive and not too cursed, then their limiting public beliefs may be closer to the wrong state than the right one. Next we introduce a mixture of types—naive, cursed, and even rational—into our primary model and establish an interesting form of robustness. When fully informed about others' types, rational players' ability to eventually learn the true state does not depend upon their share of the population. Essentially, rational players who observe a lot always uncover the truth, so long as previous behaviour richly reflects the distribution of signals.[2] Cursed players simply ignore others' signals, and inferentially naive players may again anchor with

[2] Indeed, rational herders learn the truth even in the binary signal/binary signal case, meaning that the presence of not-fully-rational types in fact enhances rational players' ability to infer information from play and thereby improves efficiency.

full confidence on the wrong state if the first few signals mislead.

In part because our main predictions lie in settings with little experimental evidence, and in part because this paper aims to introduce a model that isolates one form of error from other realistic ones, we neither attempt to tightly fit existing evidence nor to compare our predictions systematically to other theories of departure from Bayesian-Nash equilibrium. But in Section 6 we provide a brief and speculative discussion of how our model may help to interpret existing experimental evidence and how further evidence could be gathered to test it. We also briefly compare our model to other theories of non-Bayesian play.

2 Solution Concepts

In this section, we formally define the solution concepts that we apply in Sections 3 to 5 and intuitively discuss their predictions relative to those of other solution concepts. 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.

It is worth noting that all the variants of limited rationality that we explore in this paper involve players who can be thought of as maximizing their expected payoffs in games given some beliefs about others' strategies. Formally, all the solution concepts we develop specify correct beliefs about the game's structure and specific beliefs about other players' strategies. Our models of non-rational play simply comprise particular theories as to how players form the incorrect beliefs against which they optimize.

2.1 Preliminaries

Let $G = (A_1, \ldots, A_N; T_0, T_1, \ldots, T_N; p; u_1, \ldots, u_N)$ be a finite Bayesian game 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 $u_k : A \times T \rightarrow \mathbb{R}$ 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: $\sigma_k : T_k \rightarrow \triangle A_k$. 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 $T_{-k} \equiv \times_{j \neq k} T_j$ of other players (including nature) given her own type $t_k \in T_k$. Let $A_{-k} \equiv \times_{j \neq 0,k} A_j$ be the set of action profiles for players $j \neq k$ (excluding nature, who takes no action), and $\sigma_{-k} : T_{-k} \rightarrow \times_{j \neq 0,k} \triangle A_j$ 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.

2.2 Cursedness

Eyster and Rabin (2005) define the concept of *cursed equilibrium* to capture the idea that even players who come to correctly anticipate others' actions may not fully attend to the informational content that those actions convey. A bidder in a common-values auction, for instance, may suffer from the winner's curse: she may not optimally account for the fact that when others bid lower than her—a necessary condition for her to win—it indicates that they possess negative information about the value of the good.

To define cursed equilibrium, we first 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 anticipates others' actions but when choosing an optimal response fails to account for the informational content of play behaves as if she 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, type-contingent strategy profile $\sigma_{-k}(\cdot|t_{-k})$. Different types of Player k may hold different beliefs about Players $-k$'s strategies, 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 $\sigma_{-k}(\cdot|t_{-k})$.

As 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 than 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 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 $a_k(t_k) : T_k \to A_k$, 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, $p(t_k|t_j)$.

Eyster and Rabin (2005) also define a solution concept where players partially, but not fully, appreciate the informational content in others' behaviour. For $\chi \in [0, 1]$,

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 not fully, appreciate how others' actions depend upon their types. Whatever χ, players hold correct beliefs about the distribution of others' actions: their only mistake in a χ-cursed equilibrium takes the form of mis-inferring others' types from their actions.

Cursed equilibrium's combination of assuming that players have correct expectations about others' actions while under-appreciating the relationship between others' information and actions seems reasonable in many contexts where players may learn the distribution of one another's actions without ever observing other players' types.[3] 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, like when encountering a novel strategic situation, because players may not know the distribution of others' actions, cursed equilibrium may not be an appropriate solution concept, even if people do fail to take into account the informational content of others' actions.

In many settings, a far weaker notion of "cursedness" suffices for sharp predictions. When players care about the information content in others' actions—but not about the distribution of others' actions *per se*—then the strong equilibrium assumptions built into cursed equilibrium has no bite. Such is the case with the basic social-learning models considered in this paper, where others' actions exert no direct externalities: players care about others' actions only insofar as those actions signal private information. This prompts us to define the concept of "cursed best response"—the cursed analog of looking at the "dominant" actions by rational players in a game. To formally develop our concept, we first define

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

[3] This learning story is far less convincing as a justification for partially cursed equilibrium.

[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.

the χ-weighted average of type t_k's payoffs from the action profile a when facing type t_{-k} and of 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 $(1-\chi)\sigma_{-k}(\cdot|t_{-k}) + \chi\overline{\sigma}_{-k}(\cdot|t_k)$.

Definition 4 *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 type t_k of Player k to beliefs that the other players play some strategies with support A_{-k}. (Note that different types of Player k may hold different beliefs about other players' strategies.) The definition of $CBR_k^1(t_k; A_{-k})$ suggests 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; this neglect shows up in a Player k who contemplates her opponents' playing σ_{-k} best responding to $(1-\chi)\sigma_{-k}(\cdot|t_{-k}) + \chi\overline{\sigma}_{-k}(\cdot|t_k)$ rather than to σ_{-k}. An interpretation is that Player k is fully rational except for erring in computing best responses: whatever her beliefs about the correlation between her opponents' types and actions, she computes best responses without attending fully to this correlation.[5]

[5] The alternative definition

$$CBR_k^1(t_k; A_{-k}) = \left\{ \begin{array}{c} a_k \in A_k : \exists \sigma_{-k} : T_k \to \triangle(\times_{j \neq k} A_j) \text{ s.t. } \forall a_{-k} \in \widehat{A}_{-k}, \forall t_{-k}, t'_{-k} \in T_{-k}, \sigma(a_{-k}|t_{-k}) = \sigma(a_{-k}|t'_{-k}), \\ a_k \in \arg\max_{a'_k \in A_k} \sum_{t_{-k} \in T_{-k}} \sum_{a_{-k} \in A_{-k}} p(t_{-k}|t_k)\sigma_{-k}(a_{-k}|t_{-k}) u_k(a'_k, a_{-k}; t_k, t_{-k}) \end{array} \right\}$$

explicitly restricts Player k to best respond to cursed beliefs—only contemplate strategies by the other players

2.3 Naive Inference

Players who are cursed fail to appreciate the information content inherent in others' play. We now 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. (As we discuss in the next section, we believe that this type of error may be more likely than cursedness in the context of the social-learning situations we study in this paper.)

Because fully cursed players neglect the information content in play, we can model naive inference by having each player misapprehend all other players as being fully cursed. As with cursedness, we develop different models embedding inferential naivety. Although a weaker form of inferential naivety building from cursed best response suffices for our results in Sections 3-5, we begin by introducing a stronger variant that we use to compare inferential naivety to other solution concepts in Section 6. Let $\Sigma^{FCE}(G)$ be the set of fully cursed equilibria of the Bayesian game G.

Definition 5 *The mixed-strategy profile* $\hat{\sigma}$ *is* inferentially-naive information transmission (INIT) *if for each* k*, there exists* $\sigma^{FCE} \in \Sigma^{FCE}$ *such that for each* $t_k \in T_k$ *and each* a_k^* *for which* $\hat{\sigma}_k(a_k^*|t_k) > 0$,

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

In INIT play, each player best responds to the beliefs that all other players play their parts in some fully-cursed equilibrium. In games with a unique fully-cursed equilibrium, INIT players have common beliefs about which fully-cursed equilibrium the other players are

with the property that their actions are independent of their types. It lends itself less straightforwardly to the interpretation that players may appreciate the type-contingent nature of others' strategies when forming expectations about others' behaviour but fail to account for that type contingency when best responding.

playing. More generally, however, INIT play imposes no restriction on the joint distribution of players' beliefs beyond that each believes all the others to be playing in the manner prescribed by some fully-cursed equilibrium.

Despite being defined with respect to cursed equilibrium rather than cursed best responses, INIT is not in any sense an equilibrium concept: unlike cursed equilibrium, INIT players may have incorrect beliefs about even the probability distributions of over other players' actions. Indeed, whenever no fully cursed equilibrium is a Bayesian Nash equilibrium, this must be the case.[6]

To define a weaker form of inferential naivety, we use cursed best responses instead of cursed equilibrium. First, define

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

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. From this, we can define our main solution concept used in the next section.

Definition 6 *The strategy profile* $(\sigma_1, \ldots, \sigma_N)$ *is* (best response trailing naive inference) BRTNI play *if for each type* t_k *of each Player* k*,* $\sigma_k(a_k|t_k) > 0$ *only if* $a_k \in CBR_k^0\left(t_k; CBR_{-k}^1\right)$.

In BRTNI play, each player plays a true best response to beliefs that all other players play fully cursed best responses.[7] As with cursed equilibrium, we interpret BRTNI not so much

[6] Cursed equilibrium and Bayesian Nash equilibrium differ only in games that include a "common-value" component to payoffs, where one player would benefit from learning another's private information. Eyster and Rabin (2005) provide a formal statement of this condition as well as another sufficient condition for cursed and Bayesian Nash equilibrium to coincide.

[7] "BRTNI" should be pronounced "Britainy," meaning "that which resembles what you'd see in (Great) Britain," except it should be pronounced in a Britainy way, of dropping the middle syllable. Or you could pronounce it like Britney Spears, but using all the syllables (of the first name).

as Player k being convinced that Player $i \neq k$ is cursed but rather as Player k not thinking through how Player i 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 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).[8] 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.

Cursed players under-appreciate the information content in others' actions while inferential naive players misconstrue it. These two errors may coexist.

Definition 7 *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 fully cursed best responses. The concept of χ-cursed best response captures the idea that players under-appreciate 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

[8] INIT differs from BRTNI play by embedding much stronger notions of rationality. For instance, in a complete-information game with a unique Nash equilibrium, only this Nash equilibrium is INIT, while BRTNI requires only that players play best responses to best responses.

misconstrue it by failing to perceive that other players make the informational inference that they do.

3 Social Learning

We now apply the concept of inferential naivety to social learning. While we return to a discussion of the implications of both rationality and cursedness in the settings we examine, our main emphasis is on inferential naivety (in particular, BRTNI) for two reasons. First, we believe that inferential naivety may be relatively more prevalent than cursedness in sequential social-inference settings. One reason is that seeing somebody take an action that manifestly reflects her beliefs makes it salient that the action depends upon private information. Observing someone choose one restaurant over another provokes the theory that she prefers the chosen restaurant. This notion that seeing-is-inferring seems to accord with the psychology of contingent thinking; insofar as cursedness as modelled by Eyster and Rabin (2005) and analogy-based equilibrium as modelled by Jehiel (2005) reflect failures of contingent thinking, observing actions may mitigate cursedness while very much not overcoming the form of naivety embodied in BRTNI play. Indeed, while this intuition steps outside formal models, most of Eyster and Rabin's (2005) evidence for cursedness comes in the context of simultaneous-move games, and the intuition is strongest there: players may more severely neglect the informational content of others' behaviour when preparing for all contingencies—e.g., when contemplating the full range of bids by others in a sealed-bid auction—than when reacting to these others' behaviour—e.g., when responding to others' dropping out of a sequential auction.

We also emphasize inferential naivety for its striking and novel results, which seem to fit some of the patterns we observe in herding behaviour. And as Section 5 suggests, the patterns of herding it produces in social-learning environment hold in the presence of other types of players as well. Hence, the behaviour of the inferentially naive may be important even when they do not constitute the dominant type.

3.1 Convergence in the Finite Model

We begin by exploring the classical model of social learning introduced by Bikhchandani, Hirshleifer, and Welch (1992). They show that given finite action and signal spaces, rational players almost surely eventually form an information cascade—ignore their private information—and, hence, their public beliefs (a player's beliefs from her predecessors' actions but excluding her private signal) cannot converge to certainty. BRTNI players also converge to limiting actions and public beliefs. In addition, with a rich enough information structure, their public and private beliefs both almost surely do converge and to certain beliefs; in such settings, their actions almost surely converge to extreme ones.

There are two possible states of the world, $\omega \in \{0,1\}$, each equally likely *ex ante*. Let $A = \left\{a^1, \ldots, a^K\right\}$ be the finite set of real-valued actions available to each of a countably infinite sequence of players, where for each $l \in \{1, \ldots, K\}$, $a^l \in [0,1]$ and , $a^l < a^{l+1}$. Let S be the finite set of available signals, where $\pi_\omega(s)$ be the probability that any player receives the signal $s \in S$ when the state is $\omega \in \{0,1\}$; signals are independent conditional on the state. Assume that for each $s \in S$, $\frac{\pi_1(s)}{\pi_0(s)} \notin \{0, +\infty\}$, i.e. no signal perfectly reveals the state. Let $\underline{s} = \arg\min_{s \in S} \frac{\pi_1(s)}{\pi_0(s)}$, the strongest signal in favour of $\omega = 0$, and $\overline{s} = \arg\max_{s \in S} \frac{\pi_1(s)}{\pi_0(s)}$, the strongest signal in favour of $\omega = 1$. For simplicity, assume that no two distinct signals have the same likelihood ratio: for each s,$s' \neq s$, $\frac{\pi_1(s)}{\pi_0(s)} \neq \frac{\pi_1(s')}{\pi_0(s')}$. With this assumption, we can without loss of generality normalize $s = \frac{\pi_1(s)}{\pi_0(s)+\pi_1(s)} \in [0,1]$; a player with even priors who receives signal s forms posteriors s. It is well known that in this context rational players almost surely eventually cascade. Because they stop learning about the state after a finite number of imperfect signals, rational players cannot become certain that they have learned the true state. Let $b_k(a_1, \ldots, a_{k-1}; s_k)$ denote Player k's belief in the state $\omega = 1$ given her predecessors' actions as well as her private signal. Each Player k chooses an action $a_k \in A$ to maximize the payoff function $u_k(a_1, \ldots, a_k, a_{k+1}, \ldots; \omega) = -(\omega - a_k)^2$; she wishes to choose her action as close as possible to the state of the world and does not care directly about any other player's action.

Let $P(a) = \{p \in [0,1] : a \in \arg\max_A(-p(1-a)^2 - (1-p)a^2\}$, the set of beliefs over which the action a is optimal. In order to avoid situations where BRTNI players observe actions by their predecessors that contradict their model of the world, we assume that every action in A is a fully cursed best response to some signal in S.

Proposition 1: Suppose that for each $a \in A$, $E[s|s \in S \cap P(a)] \neq \frac{1}{2}$ and that there exists some $s \in S$ such that $\frac{\pi_1(s)}{\pi_0(s)+\pi_1(s)} \in P(a)$. Then for BRTNI players, with probability one, $\lim_{k\to\infty} a_k \in \{a^1, a^K\}$ and $\lim_{k\to\infty} b_k \in \{0,1\}$.

Proposition 1 uses two assumptions: first, that every action is a cursed best response to some signal, and second that every action reveals beliefs that one of the two states is more likely than the other: there is no action that is as likely to be chosen by a player with beliefs above one-half as one with beliefs below one-half.[9] When these hold, BRTNI players converge to certain beliefs and to choosing one of the two most extreme actions. To gain an intuition, suppose that actions converge to a unique limit $a^* \in A$. Because BRTNI players believe that all their predecessors follow their private signals, those far enough in the sequence believe that a proportion close to one of their predecessors receive signals $s \in P(a^*)$, each indicating the same one of the two states strictly more likely than the other; hence, they converge to certain beliefs in that state and best respond to those beliefs by choosing an extreme action. The proof in the appendix demonstrates that actions converge to a unique limit point with probability one.

[9] The second assumption is generic in the following sense: identify each action a_l with $\sup\{p : p \in P(a_l)\}$, the highest beliefs that an agent could hold and find a_l optimal. The action space A then corresponds to $K-1$ points in $[0,1]$. For any probability measure on $[0,1]$ that is absolutely continuous with respect to Lebesgue measure, the probability of choosing an action space A that satisfies this condition is one.

Given the assumption that every action is a cursed best response to some signal, rational players also almost surely converge to playing an extreme action. Unlike BRTNI players, however, they do not come to hold certain beliefs. Yet rational and BRTNI players differ more substantially than in their limiting beliefs. To draw out this difference, we modify the assumptions in Proposition 1 to allow for the presence of an action that is not a cursed best response to any signal. Because BRTNI players expect that none of their predecessors will ever play this action, we suppose that the game ends the first time it is played. For instance, suppose that $S = \{0.4, 0.6\}$ and $A = \{0.20, 0.40, 0.61\}$. In this model, rational players never play the extreme action $a = 0.20$. Rational players who receive two successive high signals (0.6) before the total number of low signals (0.4) exceeds that of high signals herd on 0.61. But those who receive one more low than high signal before receiving two successive high signals herd on 0.40. For example, if the first signal is low, then $a_1 = 0.4$. If the second signal is also low, then the second player forms posterior $\frac{(0.4)^2}{(0.4)^2+(0.6)^2} = \frac{4}{13}$, which, because it is closer to 0.4 than to 0.2, leads him to choose $a_2 = 0.4$. If the second signal is high, then the second player forms posterior 0.5, which, because it is closer to 0.4 than to 0.61, leads him to choose $a_2 = 0.4$. Hence, the second and all succeeding players herd on $a = 0.4$, implying that rational players never play $a = 0.2$. By contrast, BRTNI players interpret every $a = 0.4$ action as emanating from a new low signal and, hence, cannot herd on this action and, eventually, herd on $a = 0.2$. In short, the result that BRTNI players almost surely herd on extreme actions is more widely applicable than for rational players.[10]

In other settings, BRTNI players do not converge to extreme actions or certain beliefs. For instance, when $A = \{0.1, 0.6\}$ and $S = \{0.4, 0.6\}$, then all BRTNI players choose $a = 0.6$, and public beliefs never depart from the priors. BRTNI players who watch the same action played again and again do not necessarily conclude that it is the optimal action. Unlike concepts like "persuasion bias" (DeMarzo, Vayanos and Zwiebel, 2003), where people who

[10] Were the game not to end with the first play of $a = 0.2$, then the next player would observe a history inconsistent with her model of play. We postpone discussion of this issue to the next subsection.

hear the same information again and again eventually conclude that it must be true, BRTNI players in some settings can recognize that others' actions convey no information.

A central finding of the finite herding literature is that players who collectively possess sufficient information to identify the true state may end up herding on the wrong action. This holds for BRTNI players too. For instance, if the first two players in the example get high signals, then all BRTNI players choose 0.6; this happens with positive probability when $\omega = 0$. Yet because they misread the information content in their predecessors' actions, BRTNI players typically are more prone than rational players to herding on the wrong action. The next subsection illustrates this in a stark way.

3.2 Extreme Example of Mistakes

In certain settings, BRTNI players may wind up with the wrong limiting beliefs with probability close to one. Once again, there are two states of the world, $\omega \in \{0, 1\}$; private signals belong to the set $S = \{s_l, s_m, s_h\}$. Conditional on the state, the signals are distributed as follows:

$\Pr[s \vert \omega]$	s_l	s_m	s_h
$\omega = 1$	$\varepsilon_a \varepsilon_b$	$1 - \varepsilon_b$	$(1 - \varepsilon_a)\varepsilon_b$
$\omega = 0$	$2\varepsilon_a \varepsilon_b$	$1 - \varepsilon_b$	$(1 - 2\varepsilon_a)\varepsilon_b$,

where $\varepsilon_a > 0$ and $\varepsilon_b > 0$ both are small. The signal s_l provides weak evidence in favour of $\omega = 0$, s_h weak evidence in favour of $\omega = 1$, and s_m no evidence at all. The signal s_m occurs far more frequently than s_h, which in turn occurs far more more frequently than s_l.

Players choose actions from the set $A = \{a_l, a_m, a_h\}$, where $P(a_l) = \left[0, \frac{4}{9}\right]$, $P(a_m) = \left[\frac{4}{9}, \frac{1-\varepsilon_a}{2-3\varepsilon_a}\right]$, and $P(a_h) = \left[\frac{1-\varepsilon_a}{2-3\varepsilon_a}, 1\right]$. These sets of beliefs for which the three actions are, respectively, optimal have been chosen in such a way that BRTNI players choose a_m until some one of them observes a signal other than s_m, after which all BRTNI players choose that action. Characterizing BRTNI play following the first non-s_m signal forms the heart of our result.

What is the likely course of events? Given that $\varepsilon_a > 0$ and $\varepsilon_b > 0$ are both small, irrespective of the state the expected signal sequence comprises a long series of s_m signals eventually followed by an s_h signal, followed in turn by more s_m signals, all before the first s_l signal occurs. Rational players best respond with a sequence of a_m actions followed by a sequence of a_h actions; rational herders infer little from the play of a_h. If the first s_l appears before too long, its owner plays a_l. But if the streak of a_h continues for long enough, then rational herders infer that no s_l signal has been observed, and eventually actions may settle on a_h, irrespective of later signals, a classic herd. Rational herders ultimately settle on one of the two extreme actions without converging to certain beliefs. For small ε_a, rational players eventually herd on a_h with probability $\frac{1}{4}$, in which case public beliefs converge to $\Pr[\omega = 1|a_1, a_2, \ldots] = \frac{2}{3}$; with the remaining $\frac{3}{4}$ probability, they herd on a_l, in which case public beliefs are $\Pr[\omega = 1|a_1, a_2, \ldots] = \frac{4}{9}$.[11]

What will BRTNI do? BRTNI players believe that each in a sequence of a_h actions indicates an s_h signal rather than a far-more-likely s_l signal. After sufficiently many a_h actions, no BRTNI player deviates from a_h, even one who holds an s_l signal. Unlike rational herders, who must observe a very long sequence of a_h actions before ignoring their own signals, BRTNI players far more readily accept that $\omega = 1$ from a sequence of a_h's. In fact, if BRTNI players observe $k(\varepsilon_a)$ repetitions of a_h before observing any play of a_l, where $k(\varepsilon_a)$ solves

$$\frac{1}{1 + 2\left(\frac{1-2\varepsilon_a}{1-\varepsilon_a}\right)^{k(\varepsilon_a)}} = \frac{1 - \varepsilon_a}{2 - 3\varepsilon_a},$$

then all succeeding players choose a_h, irrespective of signal. For any $\varepsilon_a > 0$ and $\pi < 1$, there exists some $\varepsilon_b > 0$ small enough such that even when the true state is $\omega = 0$ the

[11] For ε_a small, public beliefs when herding on $a \in \{a_l, a_h\}$ have the property that receiving an extreme signal against the herd leaves the herder indifferent between continuing to herd on a and breaking the herd onto a_m. This condition generates the formulas for public beliefs, $\mu_l = \frac{4-8\varepsilon_a}{9-7\varepsilon_a}$ in an a_l herd and $\mu_h = \frac{2-2\varepsilon_a}{3-4\varepsilon_a}$ in an a_h herd. The martingale property of rational beliefs then delivers the likelihoods of the two kinds of herd.

probability of receiving a s_h followed by $k(\varepsilon_a) - 1$ s_m or s_h signals before receiving a single s_l signal is at least π. By choosing ε_a and then ε_b small enough, we can make the probability of BRTNI converging to play on a_h arbitrarily high. Formally, for any $\pi < 1$, if $\varepsilon_a < \frac{1-\sqrt{\pi}}{2}$ and $\varepsilon_b < 1 - \pi^{\frac{1}{2} \frac{\ln \frac{1-\varepsilon_a}{1-2\varepsilon_a}}{\ln 2 + \ln \frac{1-\varepsilon_a}{1-2\varepsilon_a}}}$, then the probability that BRTNI actions take the form of a sequence of a_m's followed by an infinite sequence of a_h's exceeds π.

Although the formulas above were developed for equal priors over the two states, the result extends to general priors: for any non-degenerate priors, BRTNI players come to believe with near certainty that $\omega = 1$ even when $\omega = 0$.

3.3 Naive Convergence in the Continuous Model

As discussed in the introduction, it is well understood that the basic logic driving the rational-herding literature centres around the "coarseness" of the model's action and signal spaces. While in some settings players' private information may not be easily extractable from their actions, in others the scope for observation and inference seem far too rich for fully rational players to herd inefficiently. To explore some of the more striking differences between naive social inference and rational social inference in richer settings, we develop a continuous-signal, continuous-action model of the sort discussed in the introduction.[12]

[12] A problem that our continuous model allows us to avoid—dismissed in our finite analysis by the assumption in Proposition 1 that every action is the cursed best response to some signal—is how to characterize BRTNI players' beliefs following an action that contradicts their model of play. For instance, consider a modification of our model that leaves the action space intact but reduces the signal space to a finite set. Let $\overline{s} < 1$ be the strongest signal that $\omega = 1$. While BRTNI players believe no action $a > \overline{s}$ will ever be played, whenever actions converge to one this will prove false. 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 extend our solution concept 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 the largest χ 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

Once more 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.[13] 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 assume that the likelihood ratio $L(s) \equiv \frac{f_1(s)}{f_0(s)}$ is continuously differentiable with image $\mathbb{R}_+$ and derivative $L'(s) > 0$. The assumption that the likelihood ratio is unbounded and takes every positive value implies that players may receive signals of every possible level of informativeness. These assumptions allow us to normalize signals such that $s = \Pr[\omega = 1|s]$. 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 some players receive arbitrarily strong signals about the identity of the true state. Second, by choosing actions in the continuum, players reveal their posteriors to their successors. Either of these features suffices to guarantee that rational players form beliefs and choose actions that converge almost surely to the true state; in particular, rational players converge with probability one to correct beliefs and actions with informative binary signals and a continuum of actions, or a continuum of signals and binary actions.

---

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 extending BRTNI in this way would lengthen the paper more than it would shed any light on irrational herding, we have not done so.

With this extension to our solution concept, we believe that both rational and BRTNI play in very-rich-but-finite social-learning models converge to the continuous case we explore.

[13]We use S_k to denote the Player k's signal as a random variable and s_k its realization.

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 follows 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 over-weighting the first signal. Intuitively, because Player 3 ignores how Player 2's action depends upon Player 1's action and, hence, signal, Player 3 unwittingly uses Player 1's signal twice—once when learning from Player 1, and again when learning from Player 2. Player 4, in turn, will, by dint of interpreting actions 1, 2, and 3 as signals 1, 2, and 3, in fact end up weighting the first signal four-fold:

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

More generally, player t's actions are described by

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

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.

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.

Proposition 2 shows that BRTNI players may herd on wrong beliefs, where the key assumption is the lack of any positive-probability "magic signal" that reveals the true state with certainty.

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

Proposition 2 establishes that even when $\omega=0$ there is positive probability that every single BRTNI in an infinite sequence chooses an action that exceeds any given threshold.

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 a sequence of actions that they deem impossible. The assumption of finite variance rules out positive-probability "magic signals" that reveal the true state with probability one. Proposition 2 does not tell us that BRTNI beliefs converge. But since public beliefs form a bounded submartingale when above one-half and a bounded supermartingale when below one-half, they do converge almost surely to a limiting distribution that assigns positive probability to only zero or one. Consequently, BRTNI players converge to the wrong actions and beliefs with positive probability.

To illustrate Proposition 2, consider the case where the densities are $f_0(s) = 2(1 - s)$ and $f_1(s) = 2s$. 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 $s = 1$. The two extreme signals $s \in \{0, 1\}$ fully reveal the state but occur with probability zero. Table 1 below reports simulations of BRTNI as well as Bayesian Nash play for these distributions when $\omega = 1$.

Player	$\Pr[a_t^{BRTNI} < 0.05]$	$\Pr[a_t^{BNE} < 0.05]$	$\Pr[a_t^{BRTNI} > 0.95]$	$\Pr[a_t^{BNE} > 0.95]$
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

Table 1: Simulated probabilities of BRTNI and BNE actions given $\omega = 1$.

Table 1 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 are an artifact of the simulation techniques. 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 players converge to $a = 0$ when $\omega = 1$ with probability around 11 percent. Column 4 reflects that this cannot occur with rational players, who, by Player 10, are only 2 percent as likely as BRTNI players to choose 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 there is a 99.7% chance of BRTNI Player 10 playing an action below 0.05 or above 0.95; a rational Player 10 will do so with only 93.6% chance. While we have no formal result along these lines, the observation suggests that BRTNI play converges faster than rational play.

3.4 Long-Held Weak Beliefs Are Probably Wrong

Although BRTNI converges fast, the next proposition establishes an interesting result in the rare event that beliefs converge slowly. In particular, if players' beliefs remain stable but do not converge, then they are probably wrong.

Proposition 3: 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

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

Because $\ln\left(\frac{a}{1-a}\right) > 0$ if and only if $a > \frac{1}{2}$, Proposition 3 establishes that if BRTNI for many periods believes $\omega = 1$ more likely than $\omega = 0$, 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.[14]

[14] The intuition is similar to the comparable weak-beliefs-are-probably-wrong result developed in Rabin

A limitation of Proposition 3 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 zero is more likely than one to be 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 Players 10 through 20 playing actions in various intervals *when* $\omega = 1$:

Player	$a_t^{BRTNI} \in (0.05, 0.50)$		$a_t^{BRTNI} \in (0.50, 0.95)$	$\Pr\left[\omega = 1 \,\middle|\, a_t^{BRTNI} \in (0.50, 0.95)\right]$
10	0.00143435	$<$	0.00144879	0.5025
11	0.00071807	$\approx$	0.00071808	0.5000
12	0.00036055	$\approx$	0.00035963	0.4994
13	0.00017987	$<$	0.00017923	0.4991
14	0.00008881	$<$	0.00008972	0.5026
15	0.00004537	$\approx$	0.00004535	0.4999
16	0.00002306	$>$	0.00002254	0.4943
17	0.00001168	$>$	0.00001131	0.4920

Table 2: Simulated probabilities of intermediate BRTNI actions

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

and Schrag's (1999) model of "confirmatory bias", which assumes that an individual tends to misread later signals as reinforcing earlier signals.

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 Naive Inference in Broader Settings

In this section we study the implications of naive inference outside the canonical herding setting studied in the previous section, illustrating how the propensity of BRTNI players to form extreme beliefs and wrong beliefs play out more generally. The two modifications of the environment we consider both involve limitations to how much information players have available to them.

4.1 Learning with Limited Observation

In the models of the previous section, 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 do not observe all previous play may behave more rationally. Indeed, if each BRTNI player can only observe her immediate predecessor, then BRTNI and rational play coincide:

Proposition 4: Suppose that BRTNI players can only observe their immediate predecessor's action. Then BRTNI play coincides with Bayesian-Nash-equilibrium play.

The intuition is that BRTNI players make no mistake because each correctly extracts information from the one thing she observes, her predecessor's action. Note, interestingly, however, that BRTNI players have the wrong theory of where their predecessors' beliefs come from; each believes that the action she observes contains only her predecessor's signal, when in fact it contains all the prior signals.

While Proposition 4 establishes the ironic result that severely limiting inferentially naive players' observations can lead these players to make correct inferences, milder limits on observability of past play do not qualitatively overturn our main result. When two or more actions are observed, there is still a positive probability of herding on the wrong action and beliefs:

Proposition 5: Suppose that BRTNI players can only observe their two immediate predecessors' actions and that $var\left(\left[\ln\left(\frac{S}{1-S}\right)|\omega=0\right]\right)<\infty$. Then in BRTNI play, for each $k<1$ there exists $\delta>0$ such that $\Pr[a_t>k \text{ for all } t|\omega=0]>\delta$.

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 over-weighted Player 1's signal, Player 4 will as well. And when Player 4 over-weights Player 2's action, he also over-weights Player 1's signal. The proof for Proposition 2 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 when players cannot observe the order of their predecessors' play.

Proposition 6: Suppose that BRTNI players can observe all their predecessors' actions (but not necessarily the order) and that $var\left(\left[\ln\left(\frac{S}{1-S}\right)|\omega=0\right]\right)<\infty$. Then in BRTNI play, for each $k<1$ there exists $\delta>0$ such that $\Pr[a_t>k$ for all $t|\omega=0]>\delta$.

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, any result that we established the order of moves was observable and common knowledge holds equally well without. For instance, Proposition 5 would continue to hold if BRTNI players could not observe the order of their two immediate predecessors' moves.

We believe that Propositions 5 and 6 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.

4.2 Naive Inference When Information is Costly

A central finding of the herding literature is that rational people may fail to efficiently aggregate private information: players who collectively possess sufficient information to identify the true state fail to do so when individually they have incentive to herd. In settings where rational players recognize that they have not learned the true state, we have seen above that BRTNI players may become certain that they have.

Through a variant of our basic model in which each player chooses whether to pay to receive a private signal of the state or to remain uninformed, we illustrate here a related point: BRTNI players may come to hold confident yet mistaken beliefs not only in settings where they collectively possess enough information to identify the state but also in settings where they lack such information. We assume that players observe neither whether their predecessors purchased private information nor the nature of that information when purchased; as usual, they do observe the history of their predecessors' actions. In this context, BRTNI players commit a new form of error: each believes that her predecessors—because they are "cursed"—acquire their own signals rather than rely on the behaviour of their own predecessors.

To illustrate most simply how BRTNI might be led astray, we modify our previous payoff and information structure. In each period t, there are two possible states of the world, $\omega_t \in \{0, 1\}$, each equally likely *ex ante*. From period to period, the state of the world can drift: each period, with probability $\theta \in (0, 1)$, the state "changes", in which case the new state is drawn with independent 0.5 probability. If $\theta = 1$, then the state is i.i.d.; if $\theta = 0$, then the state is fixed as before. Formally, for each t and $\omega_t \in \{0, 1\}$, $\Pr[\omega_{t+1} = \omega_t] = (1 - \theta)1 + \theta\frac{1}{2} = 1 - \frac{\theta}{2}$; the state changes in a given period with probability $\frac{\theta}{2}$.

Each Player k (with one player per period of time) can pay cost $c > 0$ to receive a *perfect* signal $s_k \in \{0, 1\}$ of the state ω_k. She then chooses an action $a_k \in \{0, 1\}$. She observes all previous players' actions (although our primary result requires only that she observe her immediate predecessor's action).

Player k has the payoff function $g_k(a_k; \omega_k) = -(a_k - \omega_k)^2 - c$ if she buys a signal and $g_k(a_k; \omega_k) = -(a_k - \omega_k)^2$ if not. Since buying a signal allows her to take the right action, a player's *de facto* payoff from buying a signal is $-c$. Her payoff from not buying a signal depends upon her likelihood of choosing the right action without buying a signal. If she chooses the right action with probability π without buying a signal, then her payoff from not buying a signal is $-(1 - \pi)$. Each player optimally purchases a signal if and only if $\pi < 1 - c$.

Rational players use time-dependent strategies. If $c > \frac{1}{2}$, then no player ever purchases a signal because, starting from a flat prior, the cost of information exceeds the benefit of becoming fully informed. Suppose instead that $c < \frac{1}{2}$. In this case, the first player purchases a signal. Succeeding players imitate this initial action until the probability that the state has departed from its initial position crosses some threshold; at this point, another player finds it optimal to purchase a signal. In particular, after the maximum integer n for which $(1-\theta)^n + [1-(1-\theta)^n]\frac{1}{2} > 1-c$, another player will purchase despite seeing all n previous players play the same action. (In the event that some player observably deviates from the equilibrium path, BNE does not impose any restrictions on subsequent behaviour.) In general, rational players efficiently trade the cost of obtaining additional information off against the benefit of learning the state.[15] When $c < \frac{1}{2}\theta$, each player purchases a signal, while when $c > \frac{1}{2}\theta$ players only periodically buy new information.

In this highly stylized setting with known and equal costs of information, players herd cyclically, knowing within each herd that only the first action conveys any information about the state. This illustrates the classical logic of rational herding: players know that the "herd" represents no new information at all and is merely imitation.

BRTNI play exhibits an altogether different dynamic. Like rational players, when $c < \frac{1}{2}\theta$ each BRTNI player purchases her own signal. But when $c > \frac{1}{2}\theta$, no BRTNI player except the first ever purchases a signal; instead, all simply follow the (infinite) herd begun by Player 1. This happens because each BRTNI player incorrectly believes that all previous players are cursed and have therefore purchased signals. When θ is small, each player believes that she knows the state with near certainty, whereas in fact players far enough along in the sequence have correct beliefs with probability only approximately equal to one-half. Like our example in Section 3.1 where BRTNI players do not herd, this result clarifies that our solution concept consists of more than the hypothesis that players take others' actions at

[15] Without characterizing the exact likelihood of correct actions or expenditures on information as a function of θ and c, it can be seen that when either θ or c is small, play will be close to efficient.

face value: here, BRTNI players incorrectly believe that their predecessors are in fact certain about the optimal action.

5 Diverse Inference

In this section, we return to the basic setting of Section 3, and explore what happens when players exhibit cursed or rational behaviour in addition to inferential naivety.

We found above that BRTNI play converges quickly—sometimes to the wrong beliefs and action—because public beliefs become very extreme very quickly. When all players are in fact partially cursed, then they under-infer predecessors' information from their actions—rightly or wrongly—preventing public beliefs from becoming so extreme. Hence, cursedness in this kind of social-learning environment counteracts BRTNI "over-inference". Rational players who recognize the types of errors that their fellow players commit, by contrast, may be able to extract a great deal of information from their predecessors' actions and, indeed, converge to correct beliefs and actions. In this section we explore the "robustness" of our predictions for inferentially naive players to the presence of other types of inference. We do so in two ways: first, we return to our solution concept "χ-cursed BRTNI play" defined in Section 2, where players are both cursed and inferentially naive, and characterize limiting beliefs and actions as they depend on the extent of their cursedness; second, we explore the behaviour of a mix of inferentially naive, cursed, and rational players.

5.1 Cursedly Naive Herding

This section explores the extent to which cursedness overturns our main results. As before, we allow continuous signals, but here we simplify the previous model by restricting actions to be binary, namely $A = \{0, 1\}$. This combination of assumptions maintains the feature that rational players almost surely form correct, certain limiting beliefs and play the appropriate action.

Recall that 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 under-appreciate 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 inferences that they do.

Unlike BRTNI and rational play, cursed play need not converge, for players overweight their own private information. For instance, when $\chi = 1$, each player follows her own signal, and, hence, play cannot converge with positive probability. Public beliefs, which we define here in the simplest way—what each player would believe about the state without her private signal—do converge, in this case to one-half as no player infers anything from her predecessors' moves. More generally, Proposition 7 characterizes what happens to public beliefs in the limit. Let m_0 be the median signal when $\omega = 0$, i.e. $\Pr[s \geq m_0 | \omega = 0] = \frac{1}{2}$.

Proposition 7: In the binary-action model, when $\chi < 2m_0$ public beliefs in χ-cursed BRTNI play converge to $1 - \frac{\chi}{2}$ with positive probability.

Proposition 7 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.

Because $m_0 < \frac{1}{2}$, Proposition 7 does rule out cases where players are very cursed. Nevertheless, 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 (unconditional on the state) uniformly distributed—χ need only be smaller than $2m_0 = 2-\sqrt{2} \simeq 0.59$ for there to be positive probability that public beliefs are wrong.

5.2 Heterogeneous Types

Sections 3 and 4 treated the case where all players are BRTNI, and the previous subsection showed that a limited form of our main result holds when all players are both inferentially naive and (boundedly) cursed. We now turn to an exploration of play when some players differ in their strategic sophistication by assuming that some fraction $\frac{1}{n}$ of players are BRTNI—with the remainder cursed or rational. For simplicity, we make the strong assumption that the rational players know which of their predecessors are rational, which are cursed, and which are BRTNI. No assumption about what cursed or BRTNI players believe has content or implication, for cursed players ignore others' actions and BRTNI players believe that all predecessors' actions directly reflect their signals.

Let $\mathcal{B}$ denote the set of players who are BRTNI.

Proposition 8: Suppose that $var\left[\ln\left(\frac{S}{1-S}\right)|\omega = 0\right] < \infty$ and that for some $n \in \mathbb{N}$, Players $n, 2n, \ldots$ are BRTNI and the remainder cursed or rational. Then there exists some $\delta > 0$ for which $\Pr\left[\lim_{t\in\mathcal{B},t\to\infty} a_t = 1|\,\omega = 0\right] > \delta$.

Proposition 8 asserts that our main qualitative result—that 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 do not. Although more generally the presence of unidentifiable BRTNI or cursed players would surely slow down rational learning, with the extreme and implausible assumption that rational players know which of their predecessors are rational, cursed, and BRTNI, they in fact will at all stages play actions that exactly reflect all prior signals and, hence, converge to the truth. Cursed players follow their own signals throughout. This combination of non-converging cursed players and differently-converging rational and BRTNI players seems an interesting implication of the underlying logic of the model, although probably unrealistic. If rather than know the type of each of her predecessors, rational players only knew the proportions of the different types, it would hinder their inference their predecessors' signals from actions. While we suspect that the same limit result would obtain, the speed of convergence would likely fall dramatically, limiting the relevance of limit results.

6 Discussion and Conclusion

The purpose of this paper is to explore the implications of a new theory of error in inference to social-learning environments. We have not concentrated on drawing out its implications in detail for existing experiments, in part because to the best of our knowledge the continuous model in which we develop our most striking predictions has never been tested in the laboratory, and in part because our simple model omits other kinds of error in strategic reasoning for which there is ample intuition and evidence. Yet we conclude the paper with a brief discussion that relates our concepts to existing evidence, speculating on how they might combine and compare with other theories of boundedly rational play in herding and other settings.

Recent laboratory tests of social-learning models have found support for some predictions of the rational model while also uncovering several systematic discrepancies. Anderson and Holt (1997), Goeree, Palfrey, Rogers and McKelvey (2007), Hung and Plott (2001), and Nöth and Weber (2003) all find that when rational inference requires players to ignore their own signals subjects do so more often than not—but far less than all the time. Researchers have attributed this pattern to "overconfidence", whereby subjects overuse their own private signals. While the psychological interpretation or foundation for such overconfidence is not always clear (nor the primary emphasis of research), we find it useful to clarify that in many circumstances what has been framed or formalized as "overconfidence" is more likely to be something akin to cursedness.

Despite their different predictions about beliefs, cursedness and overconfidence make similar predictions about actions in the finite-action models tested in the laboratory: cursedness says that subjects under-use the others' signals, whereas overconfidence says that subjects over-use their own signals; both lead to relative over-weighting of one's own signal. But "overconfidence" in the sense of players mistakenly perceiving their private information as more accurate than it is (or more accessible than another player's information) seems in these contexts to have little *a priori* psychological plausibility, and we are unfamiliar with any direct evidence for it. In a typical social-learning experiment, subjects observe signals in the form of single draws from an urn or the like. While the large psychology literature on inference identifies settings in which people over-interpret their private information, thinking it is more valuable than it is, we know of no direct evidence that people have any general propensity to regard their own random draws as superior other people's identically generated random draws. Nor are we familiar with evidence that people *per se* over-infer from a single draw from an urn in the type of symmetric-priors situations studied in the lab.[16] In future social-learning experiments with rich enough action and signal spaces to identify the first mover's beliefs, overconfidence in one's own signal should show up just strongly and more cleanly in first movers' actions than in later players' actions (where cursedness and other strategic errors act as confounds); cursedness predicts no systematic error by first movers.

Much of the experimental evidence on social learning fits neither overconfidence nor cursedness but rather seems more in line with inferential naivety. Kübler and Weizsäcker

[16] Indeed, while overconfidence seems quite plausible in many real-world social-learning environments, the fact that subjects behave in ways consistent with both overconfidence and cursedness in laboratory settings where independent evidence suggests that subjects are not overconfident should lead researchers to conclude that overconfidence is less likely to drive real-world departures from the rational model.

(2004) show in a variant of the model 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 over the course of a long string of A choices people come to believe A more and more likely, reducing the likelihood that anyone will break the herd by choosing B. Çelen and Kariv (2005), in another social-learning environment, find evidence that some players suboptimally ignore their own private information perhaps because they read too much into their predecessors' actions, exactly the prediction of BRTNI play in their setting.

The model designed to explain departures from rational play that most closely resembles ours in herding settings is Goeree, Palfrey, Rogers and McKelvey (2007), which combines Quantal-Response Equilibrium (QRE), whereby players play a noisy best response to their predecessors' actual play–with more costly actions played less frequently, with an "*ad hoc* belief-updating rule" that functions like the overconfidence described above. In the traditional finite-signal, finite-action model, Goeree, Palfrey, Rogers and McKelvey (2007) show that in a QRE players' beliefs converge to certainty. To gain an intuition, consider again the binary model first described in the introduction, and imagine that the first two players choose A; a rational third player herds on A, and rational public beliefs never budge from knowing that the first two signals are A. Because breaking the herd to choose B costs more in expected payoff terms for someone with an A signal than someone with a B signal, QRE predicts that a Player 3 with a B signal does so more frequently. Thus, a Player 3 who plays A provides successors with information that she more likely holds an A signal than a B signal. Since this same logic applies to all players, beliefs converge to certainty on the right action. The result that QRE leads players to correct limiting beliefs stands in marked contrast to the main results of our paper, where with positive probability BRTNI play leads to certainty on the wrong state.

Another natural comparison for our model is "persuasion bias" as modelled in DeMarzo, Vayanos, and Zwiebel (2003), who study how a form of naive or automatic inference from mere repetition of messages. Translated into the social-learning setting here, the logic of their model naturally predicts the same growing confidence in a herd as does inferential naivety. One might conjecture that a simple heuristic of being more and more persuaded that an action is good by seeing more and more people do it explains many anomalies. While we suspect that this simple intuition indeed plays out independent of inferential naivety, our model offers sharply different predictions across different settings about people's propensity to infer too much. For instance, a BRTNI player who shares a public signal supporting one action and observes all her predecessors take that action would *not* come to believe more and more strongly in the correctness of that action, for she understands that others' actions depend upon the public information and correctly infers that the others lack any additional information. Only when actions depend upon private information does she makes any inference. Consequently, naive inference can be viewed as almost a refinement or elaboration of generalized persuasion bias or the propensity to be convinced by repetition.

This last comparison points to a more general feature of our approach in this paper that we hope enhances its usefulness: we define a general model of errors applicable to a broad range of different games—with "zero degrees of freedom". Although the model surely omits much that could be usefully incorporated on an *ad hoc* basis, in principle the various solution concepts in Section 2 can be applied to any strategic setting without ancillary assumptions. Indeed, our stronger-than-needed solution concepts in Section 2 provide a guide as to how to embed inferential naivety in broader settings where the very weak assumptions about rationality built into BRTNI play do not suffice for sharp predictions.

The "portability" and generality of our models in Section 2 prove especially useful for comparisons of our models to others in the literature. A leading non-rational model of behaviour in other environments that is directly relevant to our model is the "Level-k Model" introduced in complete-information games by Nagel (1995) and Stahl and Wilson (1994) and extended to Bayesian games by Crawford and Iriberri (2007). In it, all players are in fact Level k, who best respond to beliefs that all other players are of Level $k-1$; Level-0 types randomize uniformly over all available actions, regardless of their private information. In Bayesian games, this implies that there is no relationship between Level-0 actions and types, so Level-1 types, who best respond to beliefs that all other players are Level 0, infer nothing about type from action. Thus, Level 1's play cursed best responses to the particular theory that their opponents' actions are uniformly distributed; cursed best response is a weaker solution concept than Level-1. Level-2 types best respond to beliefs that all other players are Level 1's, meaning that they best respond to particular cursed best responses; BRTNI play is a weaker solution concept that Level-2. Yet in all the settings we explore in this paper BRTNI makes unique predictions, so they coincide with both Level-2 and INIT predictions.[17]

While our results provide a new set of implications for Level-k models, it is worth comparing predictions of the two theories in wider contexts. Cursed equilibrium differs from Level-1 in that it predicts that players only make errors in inference and not in predicting the distribution of their opponents' play. (Cursed best response makes no prediction about players' beliefs about the distribution of their opponents' actions.) Inferential naivety also predicts that "Level-2 inferential errors" occur more frequently than "Level-2 non-inferential errors".Yet the two models also can differ in their predictions about inference.

Consider, for instance, a slight variant on the classical herding model and its cover story: people observe others sequentially entering one of two restaurants in London's West End, and infer quality from choice. But they don't observe the behaviour of patrons inside the restaurants. Once inside a restaurant, each patron can behave rationally in a civil manner or irrationally in some uncivil manner, with common knowledge that no one wishes to be

[17] Camerer, Ho and Chong's (2004) "Cognitive-Hierarchy Model of Games" extends the Level-k Model to allow Level-k players to best respond to beliefs that their opponents' levels are drawn from some distribution on $\{1, \ldots, k-1\}$, with Level k and $k-1$ sharing beliefs about the relative frequencies of levels $k-2$ and below. For instance, a Level-2 might best respond to beliefs that half of her opponents are Level-1 and the remainder Level-0. While making somewhat different predictions than BRTNI or Level-2, this model also delivers our main result that players in the continuous model come to hold wrong yet fully confident limiting beliefs with positive probability.

uncivil. The catch is that the uncivil behaviour—be it boisterous drink and loud talk, or stripping off one's clothes—exerts an externality on the following patron but not (to keep the story simple) subsequent patrons. So now if person $t + 1$ sees person t enter restaurant A she must form not only beliefs about the quality of the two restaurants but also about t's behaviour inside. To take the simplest case, suppose that uncivil behaviour generates a very large positive externality—perhaps people enjoy watching others land themselves in trouble through irresponsible partying or undressing—that outweighs the quality of the food. Then Level-1 restaurant goers will herd fully on whatever the first patron does, each predicting a 50% chance of enticingly irrational behaviour by her immediate predecessor, and ignoring her own signal. Not so a cursed-equilibrium restaurant-goer: since nobody actually acts irrationally, each only follows suit if her private information indicates good food. Because cursed best responders do not necessarily attribute rationality to their predecessors, they might or might not herd. Certainly, cursed best response does not make the unique Level-1 prediction of herding.

These different predictions generate different Level-2, INIT, and BRTNI predictions. Level-2 players do not assume any informational content in observed herds, and therefore follow their own signals. Adding appropriate assumptions about heterogeneity in tastes or off-equilibrium-path responses to make further predictions, we are likely to observe very little herding. INIT play, by contrast, predicts herding for exactly the same reason that underlies it in the models of this paper. BRTNI play is indeterminate but does not exclude herding.

Of course, such contrived examples do not form the core of the social-learning literature, and many predictions of inferential naivety hold equally well for Level-k models. But in more complicated and realistic models—such as, for instance, cases where crowding or other externalities matter—the models' predictions do differ, and so it seems important to explore the frequency and implications of informational versus non-informational mistakes.

We conclude by noting a problematic and interesting feature that our model shares with many models of bounded rationality, namely that sufficiently reflective players might suspect from their predecessors' behaviour that something is amiss with their world view. BRTNI players believe that each predecessor takes an action based on her private information alone; if so, then actions follow the same distribution as players' signals. But BRTNI players' actions very much *do not* conform to the distribution of signals in either state of the world; 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 might 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 each truly have 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—to the extend that such reflection occurs.

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.

7 Appendix

Proof of Proposition 1: Let L^* be the set of limit points of BRTNI play; for each $l \in L^*$, let $s(l) = E[s|s \in S \cap P(l)]$, which by assumption cannot equal one-half, and let $S^* = \{s(l) : l \in L^*\}$. First note that if $S^* \subset \left(0, \frac{1}{2}\right)$ or $S^* \subset \left(\frac{1}{2}, 1\right)$, then $L^* = \{a_1\}$ or $L^* = \{a_K\}$. To see this, note that if $S^* \subset \left(0, \frac{1}{2}\right)$, then each play of an action in L^* provides future BRTNI players with further evidence that $\omega = 0$, and for any $C > 1$ there exists some $N \in \mathbb{N}$ such that any player $n > N$ believes $\omega = 0$ at least C times as likely as $\omega = 1$. For C sufficiently high, she best responds by choosing a_1, and hence $L^* = \{a_1\}$.

Now suppose that S^* contains points on either side of one-half; let $a = \max\left\{l \in L^* : s(l) \leq \frac{1}{2}\right\}$ and $b = \min\left\{l \in L^* : s(l) > \frac{1}{2}\right\}$. Define

$$K_b = \min\left\{n \in \mathbb{N} : \left(\frac{s(b)}{s(b)}\right)^n \underline{s} \geq 1\right\}$$

and

$$K_a = \min\left\{n \in \mathbb{N} : \left(\frac{s(a)}{s(a)}\right)^n \overline{s} \leq 1\right\},$$

When players play b and receive K_b $\overline{s}$ signals in a row, then they can never again play a. Likewise, when players play a and receive K_a $\underline{s}$ signals in a row, then they can never again play b. But the probability that players never receive K_b $\overline{s}$ signals in a row when playing b nor K_a $\underline{s}$ signals in a row when playing a is zero, which implies that with probability one BRTNI play converges to a unique limit point that is either a_1 or a_K . **Q.E.D.**

Proof of Proposition 2: 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 $\ln\left(\frac{a_t}{1-a_t}\right) = P_t + \ln\left(\frac{s_t}{1-s_t}\right) > (t+1)K - tK = K$ as desired. From Chebyshev's Inequality, $\Pr\left[\ln\left(\frac{s_t}{1-s_t}\right) > -tK \middle| \omega = 0\right] > \frac{t^2K^2 - \sigma^2}{t^2K^2}$, and hence

$$\begin{aligned}\Pr\left[\left(\frac{s_t}{1-s_t}\right) > e^{-tK}, \forall t \middle| \omega = 0\right] &> \prod_t \frac{t^2K^2 - \sigma^2}{t^2K^2} = \exp\left\{\sum_t \ln\left(\frac{t^2K^2 - \sigma^2}{t^2K^2}\right)\right\} \\ &= \exp\left\{\sum_t -\frac{\sigma^2}{z_t}\right\},\end{aligned}$$

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

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

Finally, note that the result holds for $k \leq \frac{1}{2}$ because it holds for any $k > \frac{1}{2}$. **Q.E.D.**

Proof of Proposition 3: 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 4: The argument is stated in the text.

Proof of Proposition 5: Similar to Proof of Proposition 2 and hence omitted.

Proof of Proposition 6: Since BRTNI players pay no heed to the order of their predecessors' moves, it follows from Proposition 2.

Proof of Proposition 7: Assume that $\chi < 2m_0$ and choose $\varepsilon > 0$ such that $\chi + 2\varepsilon < 2m_0$, or $\frac{\chi}{2} + \varepsilon < m_0$. By definition of m_0, $\frac{1}{2} = \Pr[s \geq m_0 | \omega = 0] < \Pr\left[s \geq \frac{\chi}{2} + \varepsilon \middle| \omega = 0\right]$.

Suppose that the first K players get signals above $\frac{1}{2}$, with K large enough that public beliefs are $1 - \frac{\chi}{2} - \varepsilon$; this occurs with positive probability. Player $K+1$ chooses $a = 1$ with

$$\begin{aligned}
& \Pr\left[s : \frac{s\left(1-\frac{\chi}{2}-\varepsilon\right)}{s\left(1-\frac{\chi}{2}-\varepsilon\right)+(1-s)\left(\frac{\chi}{2}+\varepsilon\right)} \geq \frac{1}{2} \middle| \omega = 0 \right] \\
= & \Pr\left[s \geq \frac{\chi}{2}+\varepsilon \middle| \omega = 0 \right] \equiv p > \frac{1}{2}.
\end{aligned}$$

Public beliefs at time t depend only on the Markov process

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

and since

$$\Pr[a_t = 1 | n(t) > K, \omega = 0] > \Pr[a_t = 1 | n(t) = K, \omega = 0] = p > \frac{1}{2},$$

$\Pr\left[\exists \widehat{t} > K : n\left(\widehat{t}\right) = K\right]$ is less than that under random walk with $\Pr[a_t = 1] = p \ \forall t$. Since this is less than one, 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. **Q.E.D.**

Proof of Proposition 8: As in the proof of Proposition 2, we write the log likelihood of public beliefs of the kth BRTNI Player, Player kn, as a function of previous signals:

$$P_{kn} = \sum_{\kappa=0}^{k-1} \sum_{\tau=1}^{n-1} \alpha_{\kappa n+\tau} \ln\left(\frac{s_{\kappa n+\tau}}{1-s_{\kappa n+\tau}}\right).$$

We claim that for each $\kappa \in \{2, \ldots, k-1\}$ and $\tau \in \{1, \ldots, n-1\}$, $\frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} \leq \frac{n}{n-\frac{1}{2}}$. To see this, notice that the fewest times $s_{(\kappa-2)n+\tau}$ can appear in $P_{(\kappa-1)n}$ is once, which happens when $\tau = n-1$ or all of the players between Player $(\kappa-2)n+\tau$ and Player $(\kappa-1)n$ are fully cursed. The signal $s_{(\kappa-1)n+\tau}$ appears in $P_{\kappa n}$ is once plus the number of rational players among Players $\{(\kappa-1)n+\tau+1, \ldots, \kappa n-1\}$—call the cardinality of this set R—for a total of $R+1$ times. The signal $s_{(\kappa-2)n+\tau}$ appears in $P_{\kappa n}$ at least $R+2$ times. If $\kappa = k-1$, then

$$\frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} = \frac{R+1}{R+2} \leq \frac{n}{n+1} \leq \frac{n-\frac{1}{2}}{n},$$

as desired, where the first inequality follows from $R \leq n-1$. As $k \to \infty$,

$$\frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} \leq \lim_{m\to\infty} \frac{(2^m-1)(n-1)+2^{m-1}}{(2^m-1)(n-1)+2^m}$$
$$\frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} \leq \lim_{m\to\infty} \frac{2^m \ln 2(n-1)+2^{m-1}\ln 2}{2^m \ln 2(n-1)+2^m \ln 2}$$
$$\frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} \leq \lim_{m\to\infty} \frac{n}{n-\frac{1}{2}} = \frac{n-\frac{1}{2}}{n},$$

where the second line comes from L'Hôpital's Rule. Since the limit is approached from below, this gives the desired result for any k. From here, we can adapt the argument from the proof of Proposition 2, using the fixed factor $\frac{n-\frac{1}{2}}{n}$ in place of 2, and treating the signals in blocks of n. **Q.E.D.**

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 Behaviour". *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] Camerer, Colin F., Ho, Teck-Hua, and Juin-Kuan Chong. "A Cognitive Hierarchy Model of Games", *Quarterly Journal of Economics*, 119(3), August 2004: 861-898

[5] Çelen, Bogaçhan and Shachar Kariv. "An Experimental Test of Observational Learning under Imperfect Information", *Economic Theory* 26: 677-699.

[6] Crawford, Vincent P. and Nagore Iriberri. "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica* 75(6), November 2007: 1721–1770.

[7] DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. "Persuasion Bias, Social Influence, and Unidimensional Opinions" *Quarterly Journal of Economics* 118(3), August 2003: 909-968

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

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

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

[11] 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.

[12] Jehiel, Philippe. "Analogy-Based Expectation Equilibrium." *Journal of Economic Theory* 123, 2005: 81–104.

[13] Jehiel, Philippe and Frederic Koessler. "Revisiting Games of Incomplete Information with Analogy-Based Expectations," *Games and Economic Behaviour*, 62 (2), March 2008: 533-557.

[14] 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

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

[16] McKelvey, Richard D. and Thomas R. Palfrey. "A Statistical Theory of Equilibrium in Games." *Japanese Economic Review* 47(2), June 1996: 186-209.

[17] Nagel, Rosemarie. "Unraveling in Guessing Games: An Experimental Study". *American Economic Review*, 85(5), 1995: 1313-1326

[18] Nöth, Markus and Martin Weber. "Information Aggregation with Random Ordering: Cascades and Overconfidence", *Economic Journal* 113(484), 2003: 166-189.

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

[20] Rabin, M. and Schrag, J., "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics* 114, 1999: 37-82.

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

[22] Stahl, Dale O. II and Paul W. Wilson. "Experimental evidence on players' models of other players". *Journal of Economic Behaviour & Organization* 25(3), 1994: 309-327.