

Naive Herding in Rich-Information Settings

Erik Eyster

Department of Economics

LSE

Matthew Rabin

Department of Economics

UC Berkeley

24 March 2010

Abstract

In social-learning environments, we investigate implications of the assumption that people naively believe that each previous person's action reflects solely that person's private information, leading them to systematically imitate all predecessors even in the many circumstances where rational agents do not. Naive herders inadvertently over-weight early movers' private signals by neglecting that interim herders' actions also embed these signals. They herd with positive probability on incorrect actions across a broad array of rich-information settings where rational players never do, and—because they become fully confident even when wrong—can be harmed on average by observing others. JEL Classification: B49

1 Introduction

Beginning with Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), a theoretical literature has explored rational inference in social-learning settings. In the simplest model, a sequence of people each in turn choose 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 about the right choice. In this setting, rational agents herd: once the pattern of signals leads to two more choices of one action than the other, all subsequent people ignore their signals and take that same action. This happens because two A choices (say) on the trail of equal numbers of A and B choices reveal (given the convenient simplification that people follow their own signal when indifferent) two signals favouring A ; each subsequent mover, even with a B signal, thinks A a better bet. Generalizing this result, the rational-social-learning literature finds that when action and signal spaces are both finite, and each signal is imperfect, rational people eventually “herd” on an action because after a while an “information cascade” occurs: each person imitates others' behavior and ignores her own information. This outcome is socially inefficient: despite an infinity of private signals that reveal the right action, people herd on the incorrect action with positive probability.

This literature provides many important insights about observational social learning. Much of their basic logic about how people combine their own private information with that revealed by others' actions—and how and when herds consequently fail to aggregate information—is surely right. Yet the full-rationality theory of herding has some core implications that seem unrealistic or extreme. It assumes a level of sophistication in players' inferences that in many settings seems strained, and that predicts behavior that seems unlikely. And its predictions about the settings where herding can occur and where it *cannot* occur, as well as the nature and scope of the herding that does occur, do not invite con-

fidence that it captures the essence of observational learning. In this paper, we propose a simple model of naive inference for social learning, and show how such naive inference can lead to a qualitatively different and stronger form of herding as well as a positive probability of inefficient long-run herding in virtually any environment.

In full-rationality models, information cascades and herds occur once the information contained in observed actions becomes so great that nobody's private information can ever affect his optimal action. This can occur in the long run only if the environment is significantly "coarse". As is widely recognized (see especially Lee (1993)), rich-enough action spaces immensely reduce the probability that herds will form on the wrong action: when each action finely reflects beliefs, each new increment of information gets revealed in actions as the herd proceeds. But even when the set of actions is quite coarse, rich-enough signal spaces also severely reduce the probability of incorrect herds: Smith and Sørensen (2001) show that if there is always a positive probability of very strong signals, a would-be herd not reflecting full confidence will be overturned by a sufficiently strong signal, and subsequent movers will figure out the meaning of both overturned herds and un-overturned ones, leading actions to converge almost surely to the truth. So with a rich-enough action space *or* a rich-enough signal space, behavior converges to full efficiency. In fact, the literature contains numerous extensions adding realism to the basic model that imply such efficiency.

Even in coarse environments, however, rational herding is solely a theory of inefficiency relative to pooling people's information, not relative to what could be achieved without observing others' actions: rational herds never *hurt* anyone in expectation. Moreover, rational-herding models inherently make the prediction that herds can be frequent if and only if they are unconfident. If there is (say) a 40% chance that people's beliefs settle on the wrong guess about which of two restaurants is better, it cannot be with greater than 60% average confidence; and if people herd on being (say) 99% confident one of the restaurants is

better, then they are right 99% of the time. Full-rationality models do not in any setting predict frequent confidence in false theories. All said, it is probably fair to say that the full-rationality model predicts a relatively limited form of herding, and does so in a relatively limited set of domains. And beyond the informational and efficiency properties, Eyster and Rabin (2009) argue that even the type of *behavior* full rationality predicts has a very precarious connection with the usual image of observational learning: although obscured in the canonical binary setting, rational inference can only lead people to systematic imitation of solely the most recent actions they observe; it variously predicts that people will imitate, ignore, or “anti-imitate” earlier actions.¹

In Section 2 we define a very simple form “inferential naivety”: players realize that previous movers’ actions reflect these movers’ own signals but fail to take into account that these previous movers themselves also infer from still earlier actions.² This error confronts the logic of full-rationality herding at its core: whereas common knowledge of rationality predicts that people infer nothing more than two signals in favour of A when 100 people go in sequence to Restaurant A, the extreme form of naivety we model leads people to believe they have observed 100 signals for A. This simple alternative leads to the possibility of frequent herds

¹In the continuous model developed in this paper, for instance, each rational herder completely ignores all but his immediate predecessor. We show that when pure-recency principle fails, it can be violated in *either* direction: by imitation *or* “anti-imitation” of past actions. Indeed, Callender and Hörner (2009) beautifully illustrate how rational inference can lead would-be herders to follow the minority of previous actions rather than the majority. Eyster and Rabin (2009) illustrate yet stronger forms of anti-imitative behavior, even in the absence of knowledge-heterogeneity and order-unobservability emphasized by Callender and Hörner (2009), with simple examples where a person will take an action seemingly inconsistent with both her own signal and the beliefs of *all* previous movers!

²While Eyster and Rabin (2008) define stronger forms of this naivety, and formalize its meaning in any Bayesian game, in this paper we apply only very weak versions sufficient to make unique predictions in simple herding settings.

on incorrect actions even in rich environments where fully rational players always converge to the correct ones. And the herding that can happen is qualitatively different: naive herders become extremely confident in their wrong beliefs, and can so over-infer from herds as to be made worse off on average by observing others' actions.

We begin Section 3 by presenting a simple proposition that characterizes some severe limits on the frequency and inefficiency of herding implied by full rationality in *every* setting. But to illustrate starkly the effect of inferential naivety on social learning, Section 3 primarily analyzes an environment where each player receives a signal from a continuum ranging from fully revealing to uninformative to (very rarely) fully misleading, and then after observing all previous actions chooses an action from a continuum that fully reveals his beliefs. In this “filthy rich” environment, rational players combine their private signals with the information contained in their immediate predecessors' actions to take actions that fully reflect all signals so far, and converge in the long run to full confidence on the true state. Yet even in this environment, with positive probability naive herders converge to *fully confident* beliefs 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 inference from both predecessors leads her in fact to count the first mover's signal twice. The (naive) fourth mover, in turn, inadvertently counts the first mover's signal four-fold: once from the first action, once from the second, and twice from the third. Iterating this logic, naive herders 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 early signals happen to be misleading, limit herds may so over-use them as to outweigh the overwhelming evidence in the infinite sequence of further signals, and converge to actions reflecting extreme confidence in the wrong state. We prove that the probability of the (extremely) wrong limit beliefs is positive for any full-support distribution of signals with full support densities in both states of the world. Although it is

easy to construct examples where this happens 49% of the time, in simulations of our main example the herd converges to fully confident wrong beliefs 11% of the time.

Although the analysis of this paper is mostly limited to this one environment, the basic logic indicates that the potential for long-run misguided herding predicted by inferential naivety is quite general. To give some sense of this, in Section 4 we briefly analyze a variant of this social-learning setting that makes the naive-herding prediction even more dramatic. We consider a case where three players sequentially move in rotation an infinite number of times, each getting a fresh independent signal every period and observing all prior moves. In this setting, *each* person receives an infinite stream of private information, and of course rational players learn the truth. But because this private information is mixed with observing an infinite stream of actions whose informativeness they misread, naive herders *still* may inefficiently herd on the wrong action. Since each person gets an infinite amount of information on her own, in this setting observing others on average hurts naive herders. Reflecting the intuition from this example that naive herding might be quite general, in the concluding Section 5 we list some of the many other environments where naive herding can occur.

By the unrealistic simplification of *totally* excluding the strategic sophistication implied by the full-rationality model, and by omitting many other realistic types of errors people make, the model in this paper surely leads to some extreme and importantly wrong predictions in its own right. As such, it is meant neither to fully substitute for models incorporating other departures from full rationality, nor to tightly fit existing evidence.³ Nonetheless, in Section

³Limits to our ability to match data in experimental games also inhere in our goals and methods: by dint of being a general, formal, portable model that has ready-made, zero-degrees-of-freedom implications across different settings—rather than being a formal or informal theory of errors inspired by, defined in, and recalibrated to accommodate a specific experimental data set—our model will surely not explain specific experimental data sets as well as theories proposed in many experimental papers.

5 we briefly discuss how our model may help to interpret existing and potential experimental evidence on herding. In the process, we briefly discuss how our model contrasts with and can be combined with other types of errors people make.⁴

We conclude the paper with a brief discussion of how we feel our model might help foster greater explanatory ambition for future theoretical and empirical research on observational learning. Portable formal models such as this one might prove useful in generating the type of theoretical comparative statics that will help guide experimental work towards more powerful tests of better-identified models, and help guide researchers to a more disciplined approach to identifying the economic implications of different errors that are in evidence. Indeed, we close by reinforcing an intuition implicit in some of the formal results and discussion throughout the paper: because it is so intimately connected with the basic logic of herding, inferential naivety, more than other departures from Bayesian-Nash play, is likely to have more important long-run implications for observational learning, even if those other errors appear to explain (by some measure) more of the observed departures from pure rationality.

2 Naive Inference

In this section, we introduce the form of inferential naivety that is the primary focus of the paper. In Eyster and Rabin (2008) we formalize for general games both the concept used here, as well as stronger and more generally applicable variants of this solution concept and more carefully explore their foundations. Here we provide the weakest and simplest version

⁴We also briefly note the similarity of our model to two other theories of non-Bayesian play. Our model can be seen as a sharpening of, and portable implementation of, part of the intuition underlying DeMarzo, Vayanos, and Zwiebel's (2003) model of "persuasion bias." And, while in broader settings the two approaches differ dramatically, in the settings of this paper our model corresponds exactly to the prevailing variant of the cognitive-hierarchy approach to strategic reasoning. As such, our results can be read as providing new implications for these types of models in social-learning settings.

that suffices for the social-learning environments explored in this paper.

We say that a player engages in *best response trailing naive inference (BRTNI) play* if she plays a best response to beliefs that each of her predecessors follows his own signal, neglecting that these predecessors make informational inferences from observing their own predecessors' actions.⁵ "BRTNI" players understand that other players choose actions to best respond to their beliefs but misconstrue the provenance of those beliefs; a BRTNI player acts as if each of his predecessors' beliefs do not depend upon any observations that this predecessor made. We interpret BRTNI play as a form of limited attention and bounded rationality: Player k simply neglects to reason through how Player i makes informational inferences from Player $j \neq i$'s actions. Although the formal definition would also follow from the hypothesis that players actively believe other players are ignoring predecessors' actions, we find this to be psychologically unlikely, and it is not at all our motivation. That is, people may fail to think through that others' actions reflect these others' inference from the behavior of yet others; it is less likely that people actively believe others are too unsophisticated or choose not to make inferences. Hence, just as for Eyster and Rabin's (2005) model of cursed equilibrium, we firmly interpret our model as a formalization of a bound in players' rationality when analyzing others.

Consider the simplest of herding stories drawn from Banerjee's (1992) introduction. Each in a sequence of people chooses whether to patronize Restaurant A or Restaurant B. Diners begin with priors that A is better with 51% probability and receive *iid* private signals of which is the better restaurant; each diner observes all of her predecessors' restaurant choices. In this setting, if the first diner goes to Restaurant A , then so does everyone else. While

⁵"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).

rational diners may inefficiently “herd” on A —follow their predecessors in choosing A despite collectively possessing enough information to identify B as the better restaurant—a core intuition from the rational model is that once herding begins diners recognize it as such and stop updating their beliefs based on their predecessors’ actions. A diner who observes eight out of ten predecessors choose A —or even 98 out of 100—is no more convinced that A is the better restaurant than one who observes four out of six A choices. BRTNI players are less sophisticated, (mis)interpreting each predecessor as following his own private signal. By contrast to rational players, once a herd begins on A , BRTNI players continue to update their beliefs that A is the better restaurant. Naive players will converge to certainty that the chosen restaurant is the better one.

As formally developed and generalized in Eyster and Rabin (2009), inferential naivety builds off of a weaker form of Eyster and Rabin’s (2005) concept of “cursed equilibrium”, whereby players neglect the correlation between other players’ actions and private information. In the herding situations studied here, this severe failure of contingent thinking would imply that players simply ignore their predecessors’ actions. Although there is evidence for limited degrees of such “cursedness” in social-learning experiments, we suspect that cursedness is attenuated by sequential play where seeing someone choose A over B almost forces the observer to ask what beliefs prompted this choice. This is in contrast to simultaneous-move settings, where people condition upon observations rather than observe them. Eyster and Rabin (2009) present formal models and results combining inferential naivety with both rationality and cursedness, and Section 5 below briefly discusses the similar, contrasting, and complementary implications of these other potential errors.

Since cursed players rely too little on their predecessors’ actions, while inferentially naive ones—by taking previous actions at face-value as new information—seemingly infer too much, the two concepts might seem at first blush to be in contradiction. We show below that this is

very much not the case, in a way that calls into question the focus of much of the empirical literature: while inferential naivety does indeed push towards overweighting on prior signals, its essential property—which drives the central results of this paper—is that herders end up placing far too much weight on early relative to late signals. Indeed, it can be shown that the relative weight placed on different predecessors’ signals matters much more for limit efficiency results than does the relative weight each person places on her own versus others’ signals. Consequently, the empirical herding literature’s strong emphasis on the issue of own-vs.-other-signal weighting rather than relative weighting amongst others’ signals may be misplaced.

3 Rational and Naive Learning in a Rich Setting

Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), along with the voluminous literature that they inspired, demonstrate that rational social learning allows for the possibility of herding on a wrong alternative. While inherent and transparent in the logic of the literature, the rarity with which rational learning produces strong mistaken beliefs is perhaps insufficiently salient in economists’ conception of the rational-herding literature. Rational herders either converge to only weak public beliefs or only very infrequently herd on the wrong action.⁶ While tempting to use this literature to help understand dramatic instances of social pathology or mania where society expresses strong belief in a falsehood, this is not something that the rational-herding literature can readily deliver.

To reinforce this, consider a class of models with a binary state of the world, $\omega \in \{0, 1\}$, and priors, $\Pr[\omega = 1] = \pi$. Let I_k denote all the information available to Player k , which may include both public or private information. Let $Q_k = E[\omega | I_k] = \Pr[\omega = 1 | I_k]$, Player

⁶Public beliefs at time t are those of the player on the move after observing any public information about any predecessors’ actions but before receiving her private signal.

k 's perceived probability that $\omega = 1$ given the information set I_k . The following proposition bounds the likelihood that any Player k can form posterior beliefs $Q_k \geq q$ when $\omega = 0$ without any more assumptions about the model, e.g., the nature of players' information or action spaces.

Proposition 0 $\Pr[Q_k \geq q | \omega = 0] \leq \frac{\pi}{1-\pi} \frac{1-q}{q}$.⁷

The maximum probability that Player k can hold information causing him to believe that $\omega = 1$ with at least probability q , when in fact $\omega = 0$, cannot exceed $\frac{\pi}{1-\pi} \frac{1-q}{q}$. This bound applies to any player in any binary-state, social-learning model, including those where players have only imperfect information about their predecessors' actions, regardless of players' action spaces. It derives entirely from the logic of single-person decision making and holds in any Bayesian model of belief formation, whatever the environment. When $\pi = \frac{1}{2}$, namely equal priors, the maximum probability in any Bayesian model of social (or unsocial) learning that herders can be 99% confident in the wrong state of the world is $\frac{1}{99} \simeq 1\%$. Rational herders almost never confidently and mistakenly herd!

Extending this logic to settings where available actions and private signals might be richer shows that the probability of a confident-yet-mistaken herd is typically even more limited. To illustrate, let $A = \{0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\}$ for some positive integer n be a set of $n+1$ actions commonly available to players, and assume that each Player k chooses the action closest to her posterior q_k .⁸ Let S be the set of possible signals, which for simplicity we take to be denumerable, and let t be the strength of the weakest private signal in favor of $\omega = 1$.⁹ For simplicity, assume that each player observes all of her predecessors' actions and that actions converge.

⁷This is very close to a special case of Chamley's (2004) Proposition 2.9 in a binary-state model.

⁸As formalized below, Player k acts this way when her payoff function is $g_k(a; \omega) = -(a_k - \omega)^2$.

⁹Formally, $t = \inf \{\Pr[\omega = 1 : s] : s \in S\}$.

Corollary 0 Let $A = \{0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\}$ and $\pi = \frac{1}{2}$. Then $\Pr[\lim_{k \rightarrow \infty} a_k = 1 | \omega = 0] \leq \frac{t}{1-t} \frac{1}{2n-1}$.

For an information cascade to occur in the model with just two actions ($n = 1$), public beliefs must exceed $1 - t$.¹⁰ Combining this with Proposition 0 leads to the conclusion that in a model where the weakest positive-probability signal for $\omega = 1$ has strength t , the probability of a mistaken herd cannot exceed $\frac{t}{1-t}$. For instance, if $t = 0.05$, meaning only that once in a (possibly very, very long) while some player receives a private signal strong enough to be ninety-five percent certain of the state being $\omega = 0$, then players can wrongly herd on $\omega = 1$ no more than approximately five percent of the time. But in a four-action model ($n = 3$), mistaken herds occur with probability well under one percent. Finer action spaces reduce mistaken herding in Corollary 0 not by improving players' inference about their predecessors' information but purely by mechanically increasing the strength of public beliefs necessary for a herd. Corollary 0 shows that with $n + 1$ *ex ante* equally attractive actions, the probability of a mistaken herd approaches zero no slower than rate $\frac{1}{n}$.

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 readily 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.¹¹

¹⁰If a herd formed on $a = 1$ with public beliefs less than $1 - t$, then eventually some player would receive a signal of strength near t and choose $a = 0$, a contradiction.

¹¹By formally exploring solely the continuous-signal, continuous-action case, we not only illustrate the most striking implications of naivety but greatly simplify most of the notation and analysis. One noteworthy way that the continuous model simplifies the analysis comes in an issue arising in most models of mistaken beliefs:

There are two possible states of the world, $\omega \in \{0, 1\}$, each equally likely *ex ante*. Each player k in a countably infinite sequence receives a signal $s_k \in [0, 1]$; signals are independent and identically distributed conditional on the state.¹² When $\omega = 0$, signals have the continuous density function f_0 ; when $\omega = 1$, they have continuous 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 information structure is symmetric—for each $s \in [0, 1]$, $f_0(s) = f_1(1 - s)$ as well as 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 $a_k(a_1, \dots, a_{k-1}; s_k)$ be the action taken by Player k as a function of previous players' actions and her own private information, and let $a \equiv (a_1, a_2, \dots) \in [0, 1]^{\mathbb{N}}$ be the profile of all players' actions. This very rich action space ensures that each player's action can fully reveal her beliefs. Letting I_k be all the information available to Player k , let $E[\omega | I_k] = \Pr[\omega = 1 | I_k]$ be her probabilistic beliefs that $\omega = 1$. We assume that each Player k has a payoff function that leads her to choose $a_k = 0$ when $E[\omega | I_k] = 0$ and $a_k = 1$ when $E[\omega | I_k] = 1$, and that her optimal action a_k is an increasing function of beliefs. The precise someone with an incorrect theory of the world might observe something that she had deemed impossible. For instance, consider a modification of our model that leaves the action space intact but reduces the signal space to a finite set. Let $\bar{s} < 1$ be the strongest signal that $\omega = 1$. While BRTNI players believe no action $a > \bar{s}$ will ever be played, this proves false whenever actions converge to one. Eyster and Rabin (2008) extend the solution concept to assume that a player who observes a predecessor choosing an action too high to be consistent with naive inference believes that this predecessor received the highest possible signal. With this extension, we believe that both rational and BRTNI play in very-rich-but-finite social-learning models converge to the continuous case we explore. 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.

¹²We use S_k to denote the Player k 's signal as a random variable and s_k its realization.

shape of the payoff function affects players' actions without affecting beliefs. Purely for notational ease, we assume that every Player k has payoff function $g_k(a; \omega) = -(a_k - \omega)^2$, which is maximized by setting $a_k = E[\omega | I_k]$.

We begin by analyzing rational players. Throughout we simplify analysis by using log odds ratios, $\ln\left(\frac{a}{1-a}\right)$, the log of the ratio the player's beliefs that $\omega = 1$ versus $\omega = 0$. Given equal priors, Player 1 chooses $\ln\left(\frac{a_1}{1-a_1}\right) = \ln\left(\frac{s_1}{1-s_1}\right)$. Player 2 combines Player 1's action with his own private information to form the posterior

$$\ln\left(\frac{a_2}{1-a_2}\right) = \ln\left(\frac{a_1}{1-a_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right) = \ln\left(\frac{s_1}{1-s_1}\right) + \ln\left(\frac{s_2}{1-s_2}\right).$$

This procedure may be interpreted in two ways: Player 2 can back out Player 1's signal from her action and combine it with his own signal and the common prior. Alternatively, because the agents share a common prior, Player 2 can adopt Player 1's posterior as his own prior before incorporating his private signal. Applying this latter interpretation to Player 3 explains why Player 3 does not benefit from observing Player 1's action given that she observes Player 2's. In general then, $\ln\left(\frac{a_t}{1-a_t}\right) = \sum_{\tau \leq t} \ln\left(\frac{s_\tau}{1-s_\tau}\right)$. Behaviorally, since Player t does not observe prior movers' signals, what each Player t actually chooses is $\ln\left(\frac{a_t}{1-a_t}\right) = \ln\left(\frac{a_{t-1}}{1-a_{t-1}}\right) + \ln\left(\frac{s_t}{1-s_t}\right)$.

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 of the true state. Smith and Sørensen (2001) have shown that such "unbounded private beliefs" preclude false herds even with finite action spaces. Second, by choosing actions in the continuum, players reveal their posteriors to their successors. Lee (1992) shows that this too guarantees that rational players form beliefs and choose actions that converge almost surely to the true state.

BRTNI players depart from rational play only insofar as they neglect their predecessors' informational inferences. Clearly such error does not affect the first mover, so once more

$\ln\left(\frac{a_1}{1-a_1}\right) = \ln\left(\frac{s_1}{1-s_1}\right)$. Because the first player performs no informational inference, the second one correctly infers her 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 he would in a Bayesian Nash Equilibrium. The third player neglects how the second player incorporated the first's signal into his action. Hence, she 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. More generally, player t 's actions are described by

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

Relative to rational players, who give all signals equal weight, 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 receive 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 1 shows that BRTNI players sometimes come to hold certain yet mistaken beliefs.

Proposition 1: In BRTNI play, for each $k < 1$ there exists $\delta > 0$ such that $\Pr[a_t > k \text{ for each all } t | \omega = 0] > \delta$.

Proposition 1 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.

Our maintained assumption 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. Eyster and Rabin (2008) explain why dropping this assumption (and using the extension of BRTNI discussed in Footnote 11 makes it *easier* to obtain the conclusion of Proposition 1. Notice that the assumptions of unbounded private beliefs and signals with a continuous density representation rule out settings with positive-probability signals that reveals the state, in which case BRTNI of course eventually learns the truth.

Unlike rational beliefs, BRTNI beliefs do not form a martingale; they tend to move in a way predictable from their current level. When public beliefs $P_t > \frac{1}{2}$, beliefs tend to rise: $E[P_{t+1} | P_t] > P_t$. When public beliefs $P_t < \frac{1}{2}$, beliefs tend to fall: $E[P_{t+1} | P_t] < P_t$. Beliefs drift in this predictable way because BRTNI players in future periods reweight information already contained in current beliefs; high current beliefs indicate that future BRTNIs will

re-count stronger evidence in favour of $\omega = 1$ than $\omega = 0$, raising future beliefs. Such drift in beliefs both provides intuition for Proposition 1 as well as marking in and of itself a striking qualitative departure from a core prediction of the rational model.

Proposition 2: BRTNI actions and beliefs converge almost surely to 0 or 1.

The assumptions of Proposition 1 also imply that BRTNI beliefs converge almost surely to zero or one.¹³ BRTNI players who do not learn the true state become *fully confident* in the wrong state! To illustrate Propositions 1 and 2 as well as differentiate BRTNI from rational play, 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 $s = 0$, and when $\omega = 1$ they come from a triangular distribution with mode $s = 1$. (The two extreme signals fully reveal the state but occur with probability zero.) Table 1 reports simulations of BRTNI as well as Bayesian-Nash play for these distributions when $\omega = 1$.

¹³While Eyster and Rabin (2008) describe non-generic counterexamples, convergence to certain beliefs is a generic feature of BRTNI play across learning models and constitutes another key difference from rational models.

Player	BNE			BRTNI		
	$a \leq 0.05$	$0.05 < a \leq 0.95$	$a > 0.95$	$a \leq 0.05$	$0.05 < a \leq 0.95$	$a > 0.95$
1	0.0026	0.8998	0.0976	0.0025	0.8998	0.0977
2	0.0060	0.6905	0.3035	0.0058	0.6912	0.3030
3	0.0070	0.5059	0.4871	0.0216	0.3819	0.5965
4	0.0069	0.3684	0.6247	0.0483	0.1877	0.7640
5	0.0060	0.2708	0.7232	0.0739	0.0929	0.8332
6	0.0051	0.1995	0.7954	0.0914	0.0463	0.8623
7	0.0041	0.1482	0.8477	0.1016	0.023	0.8754
8	0.0033	0.111	0.8857	0.1068	0.0117	0.8815
9	0.0026	0.0826	0.9148	0.1098	0.0057	0.8845
10	0.0020	0.0624	0.9356	0.1115	0.0029	0.8856

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.¹⁴ 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, because they interpret Player 1 and 2's low actions as two very strong and independent pieces of evidence in favour of $\omega = 0$, only very high signals can swing actions above 0.05. Moving down Column 2 to examine later players' actions suggests that when $\omega = 1$ BRTNI players converge to $a = 0$

¹⁴Since BRTNI and BNE coincide for the first two players, these should be the same; the small differences are an artifact of the simulation techniques.

with probability approximately 11 percent. Column 3 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 its convergence. There is a 99.7% chance of BRTNI Player 10 playing an action below 0.05 or above 0.95; a rational Player 10 does so with only 93.6% chance. While Proposition 2 establishes only that BRTNI play converges—and not its speed—the simulation suggests that it converges faster than rational play.

Although BRTNI converges fast, the next proposition establishes an interesting result for the rare cases where beliefs converge slowly. In particular, when players' beliefs stabilize for a while in favour of one state over the other without converging to complete confidence in that state, they are probably wrong.

Proposition 3: For each interval $[c, d] \subset (\frac{1}{2}, 1)$ there exists $T \in \mathbb{N}$ such that if for each $t \in \{1, \dots, T\}$, $a_t \in [c, d]$ under BRTNI play, then

$$\Pr[\omega = 0 | (a_1, \dots, a_T)] > \Pr[\omega = 1 | (a_1, \dots, a_T)].$$

If for many periods BRTNI believes the likelihood of $\omega = 1$ exceeds 90%—but not 99%—then in fact it is more likely that $\omega = 0$ than $\omega = 1$. A BRTNI player at the end of a long run of high actions believes that her predecessors must all have high signals. The only reason why she would *not* conclude that $\omega = 1$ with 99% certainty is that she receives a very low signal herself. Hence, the only way that a large number of players can take actions above 90% without any single one of them reaching 99% is that if after a few pieces of evidence supporting $\omega = 1$, all subsequent signals point towards $\omega = 0$, overall indicating $\omega = 0$ more likely.¹⁵

¹⁵The result resembles the weak-beliefs-are-probably-wrong result developed in Rabin and Schrag's (1999)

Propositions 1 and 2 demonstrate that with positive probability BRTNI play culminates in the wrong limiting action. Since rational players almost surely choose the right limiting action, BRTNI players obtain strictly lower long-run average payoffs. Yet there is another sense in which they do worse than rational players: while rational players always benefit on average from observing their predecessors' actions, BRTNIs may not. Because observational learning leads to overly extreme beliefs among the inferentially naive, when the expected cost of overconfidence exceeds the added information in others' actions, BRTNI can be harmed. This turns out to be the case in a variant of our parametric example above where each Player k 's payoff function is $g_k(a_k, \omega) = -(a - \omega)^{2n}$ for some positive integer n . The higher n , the more costly it is to choose an action distant from the true state, making players reluctant to choose extremely low or high actions. We saw above that when $n = 1$, approximately 11% of the time BRTNI converges to wrong limiting beliefs and actions. This result does not depend on n because players' actions are an invertible function of beliefs regardless of n ; the precise shape of the payoff function as captured by n affects players' actions without affecting beliefs. For any n , as players become certain about the state, they converge to $a \in \{0, 1\}$. Consequently, BRTNI players obtain a long-run average payoff of approximately $-\frac{1}{9}$, the probability of settling on the wrong limiting action, $\frac{1}{9}$, times the loss from doing so, $-(1)^{2n} = -1$. A player without any opportunity to learn could not do worse than to choose $a = \frac{1}{2}$ regardless of her signal. Hence, a lower bound on her average payoff is $-\left(\frac{1}{2}\right)^{2n}$. While for our simple example used above of $n = 1$ BRTNI is better off for observing the herd, when model of "confirmatory bias", which assumes that an individual tends to misread later signals as reinforcing earlier signals. The intuition bears some resemblance to some of our results below, as well as to Callender and Hörner's (2009) result on the "wisdom of the minority" when people differ in the quality of their private information and cannot observe the order of their predecessors' moves. Someone observing one diner choose A and 3 choose B, without knowing the order, might conclude that the solitary diner is likely to have arrived last with good information.

$n \geq 2$ she is worse off, since $-\left(\frac{1}{2}\right)^{2n} \geq -\frac{1}{9}$. BRTNI does worse for having the opportunity to observe her predecessors, something impossible in any rational model.

4 Harmful Learning with Long-Run Agents

In many settings, the same people may choose actions repeatedly, learning over time both as they receive new private information as well as from observing others' choices. To take Banerjee's (1992) canonical restaurant example, most diners choose repeatedly among the same set of restaurants, learning both from their own experiences and from crowds. In this section, we show how BRTNI can lead to a striking form of harmful herding in this context: long-run BRTNI players who could learn the state almost surely simply by ignoring others and focusing on their own infinite sequences of private signals might end up taking the wrong limiting actions due to their errors in inference.

Consider a variant of the benchmark model where three players $\{A, B, C\}$ alternate moves in sequence $A, B, C, A, B, C, A, \dots$. As before, in each period t , the player moving receives a (new) private signal about the state and can observe all of her predecessors' actions. Maintaining our assumptions on signals, each player's growing collection of private signals almost surely eventually reveals the state: a player (rational or BRTNI) who simply ignored others' actions and acted solely on the basis of her private information would almost surely converge to choosing the right action. Clearly, rational players almost surely converge to correct limiting beliefs and actions. Yet:

Proposition 4: Suppose that three long-run BRTNI players $\{A, B, C\}$ move in sequence A, B, C, A, \dots . Then for each $k < 1$ there exists $\delta > 0$ such that

$$\Pr \left[\left(\frac{a_t}{1 - a_t} \right) > e^t \left(\frac{k}{1 - k} \right) \text{ for each } t \mid \omega = 0 \right] > \delta.$$

Despite holding a collection of private signals that identifies the state, each player may

end up choosing the wrong action. Proposition 4 implies that when $\omega = 0$, for any $k > \frac{1}{2}$, it sometimes happens that all long-run BRTNI players play actions above k and all converge to certain beliefs that $\omega = 1$. With long-run BRTNIs, play in the first three periods exactly resembles that of the baseline model, meaning that Player C overweights A 's signal s_1^A in her first move.¹⁶ Play in the fourth period also coincides with the baseline model as A neglects that B and C already have incorporated s_1^A into their actions. The first difference emerges in period 5, where B , having chosen a_2^B himself, knows that a_2^B already embodies the information in a_1^A and therefore does not re-count a_1^A . Letting $\ln\left(\frac{a_t}{1-a_t}\right) = \sum_{k=1}^{t-1} F_{t-k} \ln\left(\frac{s_{t-k}}{1-s_{t-k}}\right) + \ln\left(\frac{s_t}{1-s_t}\right)$, in the baseline model BRTNI play follows $F_{t-k} = 2F_{t-k+1}$, whereas here it converges to $F_{t-k} = F_{t-k+1} + F_{t-k+2}$ as shown in the Appendix. In the limit as $k \rightarrow \infty$, F_{t-k}/F_{t-k+1} approaches the Golden Ratio $\varphi \simeq 1.618$. Because later signals are geometrically discounted relative to earlier ones, a finite number of early misleading signals can lead to wrong herding, and this occurs with positive probability.

5 Discussion and Conclusion

This paper presents a very simple model of a single error that people might make, and investigates its implications in a domain chosen to highlight its consequences most starkly. We conclude by indicating various ways that our approach can be improved or extended, in the process emphasizing why we feel the qualitative insights of the model are likely to be more general.

It seems clear that our extreme form of naive inference will lead to the possibility of long-run, confident-but-wrong herding in most settings. Indeed, we analyze the full-support, continuous-action, continuous-signal setting because it is among the more difficult environments for generating herding. Results that we have derived (in both existing working papers

¹⁶Superscripts here denote the identity of long-run players.

and in personal files) in different settings reinforce this intuition.¹⁷ Other variants of the model with continuous actions and signals certainly yield similar results to this paper: when more than one player moves at a time, when the order of moves is not observed, when players don't observe recent movers, or when players *only* observe recent movers, etc., with positive probability beliefs by naive herders converge to full confidence in the wrong state.¹⁸

The basic implications also play out in models of more direct economic interest, including some that have previously been studied. In a simple variant of herding in financial markets along the lines of Glosten and Milgrom (1985) and Avery and Zemsky (1998), where prices adjust to clear the market as each new agent enters with private information, we replicate the result that full rationality implies full information aggregation. Yet naive herders may converge to the wrong beliefs once more. We have also analyzed a simple variant of the example of judging the quality of restaurants by their popularity, adding in a small negative externality imposed by queue length. Here, in fact, much more information gets revealed than when queue length doesn't matter. Indeed, there is a discontinuity—as the distaste for queuing becomes very small, rational players will fully learn restaurant quality by observing their predecessors' queuing choices. Naive herders instead might become convinced the wrong restaurant is better, generating permanent, costly queues for an inferior establishment.

¹⁷The generalization begins where the literature and our paper began. As noted in Section 2, in the canonical binary model, although BRTNI players act just as rational players do, they hold very different beliefs: in the long run BRTNI players become fully confident about which restaurant is better, even when private signals are so weak that the probability of a wrong herd is nearly 50%.

¹⁸Eyster and Rabin (2009) present what we suspect is the only notable exception: when each player observes *solely* her single immediate predecessor, naive herders behave just like rational players at each step and in the limit. Despite having the wrong theory of those beliefs' provenance, as always each naive herder correctly infers her predecessor's beliefs. In this case, however, because she does not observe earlier actions, she cannot double count earlier signals; all signals are counted exactly once. Once players observe at least two predecessors, however, naive inference can once again lead to false herds.

Generalizing our notion of inferential naivety—especially to make it less extreme—and combining it with other errors is a more difficult matter. Eyster and Rabin (2008, 2009) define various more complicated and more tenuous solution concepts that both allow less extreme forms of naivety, and (as needed for making predictions in more complicated settings) go beyond minimal dominance criteria in restricting beliefs. This includes various ways of combining inferentially naive play with both rational play and cursed play. If we allow for heterogeneity of different extreme types, fully cursed players will forever ignore others, fully rational players will figure out the truth, but neither type’s presence precludes long-run overconfident-but-wrong play by BRTNI players. We also combine within each player some inferential naivety with partial cursedness, allowing her to severely underinfer (but not totally ignore) the information content inherent in her predecessors’ actions. Because a partially-cursed player with imperfect information about the state could never observe any actions by her predecessors that would lead her to fully confident beliefs, she can never become completely convinced of either state. With that caveat, an analogue to Proposition 1 holds in this environment. However confident partially cursed BRTNI players can possibly become of the true state, they can also become that confident of the false state; for instance, if partial cursedness permits players to become up to, but no more than, 90% convinced of the true state, then with positive probability they become 90% convinced of the false state.

The model developed in this paper is ideal for use in experiments designed to reject the hypothesis that naivety is the *sole* determinant of behavior. But less extreme models of inferential naivety would aid in the more important task of investigating whether naivety exists and matters for herding. Formulations permitting both heterogeneity and other types of errors—e.g., the realistic types of more random errors (and rational expectations over those errors) captured by notions such as McKelvey and Palfrey’s (1996) notion of Quantal Response Equilibrium—would surely be essential to empirical study of what degree of in-

ferential naivety appear in experiments, and what role such naivety might play. As noted above, the existing experimental literature is generally not well-designed to differentiate among likely hypotheses about the nature of observational learning; for instance, data are generally collected in environments where extreme Bayesian-Nash play coincides with extreme BRTNI play. Despite this barrier, we now turn to a very cursory discussion of where we believe evidence from existing experiments suggests a potential role of inferential naivety.

In a meta-study of 13 experimental datasets, Weizsäcker (forthcoming) finds strong evidence that subjects systematically follow their predecessors far less than they should, as measured by the empirical distribution of payoffs. Although often attributed to “overconfidence”, by which researchers (in both the formal models and in connotation) seem to mean people’s believing their private signals are more extreme than they are and therefore over-weighting those private signals, such behaviour surely derives more from cursedness along the lines of Eyster and Rabin (2005), which implies in this context that people underweight others’ signals as a result of under-infering from their actions.¹⁹ Although such cursedness

¹⁹Despite their different predictions about beliefs, cursedness and overconfidence make similar predictions about actions in the finite-action models tested in the laboratory: cursedness implies that subjects *de facto* under-use others’ signals because they under-infer from others’ actions, whereas overconfidence says that subjects over-use their own signals, both leading to relative over-weighting of one’s own signal. But the more literal form of “overconfidence” seems to have little *a priori* psychological plausibility in these contexts, and we are unfamiliar with any attempt to directly test for it. In a typical social-learning experiment, subjects’ signals take the form of single draws from an urn. While the large psychology literature on inference identifies settings in which people over-infer from signals by magnifying those signals, we know of no 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 any evidence at all that people *per se* over-infer from a single draw from an urn in the type of symmetric-priors situations studied in the lab. In future social-learning experiments that incorporate action and signal spaces rich enough to identify the first mover’s beliefs, overconfidence in one’s own signal should show up just as strongly in the first actions or any other

may seem the opposite of naivety, as we noted above, the key feature of naive inference is not a generic “overweighting” of others’ signals, but particular patterns of *which* others’ signals influence herders more than other signals. Indeed, estimating a model that includes what we identify as cursed and BRTNI types among others, Kübler and Weizsäcker (2004) find evidence that more subjects behave in ways similar to BRTNI players than to any other type. 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; this would be the prediction of BRTNI play in their setting.²⁰

While various types of departures from pure Bayesian-rational play are clearly complementary and relatively unrelated to inferential naivety, two such departures are in this setting much more connected. DeMarzo, Vayanos, and Zwiebel (2003) formulate what they refer to as “persuasion bias” in which they provide a formal model 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. We suspect that the simple propensity of being more and more persuaded that an situations where previous actions offer no information as when choice of actions involve inferring signals from others’ actions. Cursedness predicts no systematic error in these cases.

²⁰In the traditional finite-signal, finite-action model, Goeree, Palfrey, Rogers and McKelvey (2007) show formally that in the unmodified quantal-response equilibrium—whereby players play a noisy best response to their predecessors’ actual play, with more costly actions actions played less frequently—players’ beliefs converge to certain and correct limiting beliefs. Hence, unlike inferential naivety, that there will not be harmful herding. To better fit their experimental data, they combine QRE with an *ad hoc* belief-updating rule that functions like the overconfidence described above.

action is good by seeing more and more people do it plays out independent of inferential naivety, and that our model will miss such forms of persuasion. On the other hand, our model offers sharp 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 infer incorrectly. Consequently, naive inference can be viewed as almost a refinement or elaboration of generalized persuasion bias or the propensity to be convinced by repetition. Bohren (2009) explores how various types of errors in predicting others’ information-processing capabilities can affect herding. While some type of errors lead herds to be less stable, she shows that when agents underestimate the ability of others to process observations of behavior, incorrect herds can persist in rich settings for much the same intuition as we establish here. Guarino and Jehiel (2009) develop a rich-action-space model of “coarse inference” in social learning that assumes players understand the relation between predecessors’ actions and the state, but not between their actions and private signals. They replicate our prediction that early signals will be overweighted relative to later ones; unlike BRTNI, however, their model does not lead to mistaken beliefs and actions in the limit.

A leading non-rational model of behavior, 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)) is less similar to BRTNI in motivation than is persuasion bias, but more similar in terms of formalism. In the simplest form of such models, 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 than 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.²¹ Although Eyster and Rabin (2009) discuss examples where the two models will in fact differ, in the settings of this paper, where the two coincide, our results provide a new set of implications for Level- k models.

We conclude by emphasizing two of the key motivations for introducing this notion of inferential naivety: its prevalence across settings, and its likely economic impact. The psychology of underattentiveness to how others extract information from behavior seems to us a key aspect people's (mis)learning across a variety of settings. Our emphasis on a formal and portable model that aspires to capturing general tendencies, rather than fitting data in particular contexts, can give substance to these claims: to judge the merits of an 'explanation' of the degree of herding present or absent in one setting we obviously would like to know what predictions this explanation makes in other settings. The pure-rationality theory of observational learning makes some strikingly different, and strikingly unexplored, predictions across

²¹Camerer, 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 $\{1, \dots, k-1\}$, with Level k and $k-1$ sharing beliefs about the relative frequencies of levels $k-2$ and below. 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.

various settings. We've shown that inferential naivety changes some of these conclusions. Although not formalized here, the strength and robustness of misinference from inferential naivety gives every reason to believe that it may have more profound economic implications in observational-learning settings than other departures from pure rationality. Essentially all other types of mistakes (that we know of) explored in the theoretical or experimental literatures seem more likely to undermine than to exacerbate the possibility that society will tend to overconfident beliefs because of herding. By contrast, inferential naivety—by its core logic of suggesting that people might neglect how little additional information there in the imitative behavior of others who are themselves inferring from society—is likely to be the type of error that might truly lead groups of people systematically astray.

6 Appendix: Proofs

Proof of Proposition 0 Let $\bar{I}_k = \{I_k = (s_k; a_1, \dots, a_{k-1}) : Q_k \geq q\}$. From Bayes' Rule,

$$\Pr[\omega = 1 | \bar{I}_k] = \frac{\pi}{\pi + (1 - \pi) \frac{\Pr[\bar{I}_k | \omega=0]}{\Pr[\bar{I}_k | \omega=1]}} \geq q$$

$$\Rightarrow \frac{\Pr[\bar{I}_k | \omega = 0]}{\Pr[\bar{I}_k | \omega = 1]} \leq \frac{\pi}{1 - \pi} \frac{1 - q}{q},$$

Because $\Pr[\bar{I}_k | \omega = 1] \leq 1$, $\Pr[\bar{I}_k | \omega = 0] \leq \frac{\pi}{1 - \pi} \frac{1 - q}{q}$. ■

Proof of Corollary 0 When public beliefs are that $\Pr[\omega = 1 | (a_1, \dots, a_{k-1})] = p$, Player k with private beliefs t herds on $a = 1$ if and only if

$$\Pr[\omega = 1 | I_k] = \frac{pt}{pt + (1 - p)(1 - t)} \geq \frac{2n - 1}{2n},$$

or $p \geq \frac{1}{1 + \frac{t}{1-t} \frac{1}{2n-1}}$. Using $p = \frac{1}{1 + \frac{t}{1-t} \frac{1}{2n-1}}$ for q and $\pi = \frac{1}{2}$ in Proposition 0 gives the result. ■

Proof of Proposition 1 Choose $k \in (\frac{1}{2}, 1)$ and let $K = \ln\left(\frac{k}{1-k}\right) > 0$. 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 $P_4 = 2P_3 + \ln\left(\frac{S_3}{1-S_3}\right) > 2 \cdot 4K - 3K = 5K$, etc. In general, $P_t > (t+1)K$, and so $\ln\left(\frac{a_t}{1-a_t}\right) = P_t + \ln\left(\frac{S_t}{1-S_t}\right) > (t+1)K - tK = K$ as desired. Now,

$$\begin{aligned} \Pr\left[\ln\left(\frac{S_t}{1-S_t}\right) < -tK \mid \omega = 0\right] &< \Pr\left[\left|\ln\left(\frac{S_t}{1-S_t}\right)\right| > tK \mid \omega = 0\right] \\ &< \frac{1}{(tK)^2} E\left[\left|\ln\left(\frac{S}{1-S}\right)\right|^2 \mid \omega = 0\right], \end{aligned}$$

where the final inequality comes from Markov's Inequality. Also,

$$\begin{aligned} Q &:= E\left[\left(\ln\left(\frac{S}{1-S}\right)\right)^2 \mid \omega = 0\right] = \int_0^1 \left(\ln\left(\frac{S}{1-S}\right)\right)^2 f_0(s) ds \\ &\leq M \int_0^1 \left(\ln\left(\frac{S}{1-S}\right)\right)^2 ds = M \left(\int_0^1 (\ln(S))^2 + (\ln(1-S))^2 - 2\ln(S)\ln(1-S) ds\right) \\ &< 2M \int_0^1 (\ln(S))^2 ds = 4M, \end{aligned}$$

for $M := \sup\{f_0(s) : s \in [0, 1]\}$, which is finite by the continuity of f_0 . Define $\tau = \min\{t \in \mathbb{N} : Q < t^2 K^2\}$ so that for each $t \geq \tau$, $\frac{t^2 K^2 - Q}{t^2 K^2} \in (0, 1)$, and let $C(K) := \prod_{t=1}^{\tau-1} (1 - F_0(-tK)) > 0$.

Then

$$\begin{aligned} \Pr\left[\left(\frac{S_t}{1-S_t}\right) > e^{-tK}, \forall t \mid \omega = 0\right] &> C(K) \prod_{t \geq \tau} \frac{t^2 K^2 - Q}{t^2 K^2} = C(K) \exp\left\{\sum_{t \geq \tau} \ln\left(\frac{t^2 K^2 - Q}{t^2 K^2}\right)\right\} \\ &= C(K) \exp\left\{\sum_{t \geq \tau} -\frac{Q}{z_t}\right\}, \end{aligned}$$

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

$$\begin{aligned} \Pr\left[\left(\frac{S_t}{1-S_t}\right) > e^{-tK}, \forall t \mid \omega = 0\right] &> C(K) \exp\left\{\sum_{t \geq \tau} -\frac{Q}{t^2 K^2}\right\} \\ &> C(K) \exp\left\{\sum_{t \geq 1} -\frac{Q}{t^2 K^2}\right\} = C(K) \exp\left\{-\frac{Q\pi}{6K^2}\right\} > 0. \end{aligned}$$

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

Proof of Proposition 2: From above, write

$$2^{1-t}P_t = \sum_{\tau < t} 2^{-\tau} \ln \left(\frac{s_\tau}{1-s_\tau} \right). \quad (6.1)$$

Since the three series

$$\begin{aligned} \sum_{\tau=1}^{\infty} E \left[2^{-\tau} \ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] &= 2E \left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] \\ \sum_{\tau=1}^{\infty} \text{var} \left[2^{-\tau} \ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] &= \frac{1}{3} \text{var} \left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] \\ \sum_{\tau=1}^{\infty} \Pr \left[2^{-\tau} \left| \ln \left(\frac{S}{1-S} \right) \right| \geq 1 \right] &\leq \sum_{\tau=1}^{\infty} 4^{-\tau} \text{var} \left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] \end{aligned}$$

converge—the first two follow from finiteness of the second moment (and therefore first moment) established in the proof of Proposition 1 and the third from Chebyshev’s inequality—Kolmogorov’s Three-Series Theorem implies that $2^{1-t}P_t$ converges a.s. Since $2^{1-t}P_t = 0$ iff $\ln \left(\frac{s_t}{1-s_t} \right) = -2P_{t-1}$ and $\ln \left(\frac{s_t}{1-s_t} \right)$ is atomless with negative mean when $\omega = 0$, this can happen for only finitely many t ; hence, $2^{1-t}P_t$ converges a.s. to something other than zero. This implies that P_t diverges a.s., and so a_t converges a.s. to 0 or 1. ■

Proof of Proposition 3 Let $[u, v] \subset \mathbb{R}_{++}$. Define $T_1 = \lfloor \frac{v}{u} + 1 \rfloor$, so that $(T_1 - 1)u \leq v < T_1 u$. Choose $\delta \in (0, T_1 u - v)$. Suppose that for each $t < T_1$, $\ln \left(\frac{a_t}{1-a_t} \right) \in [u, v]$ (equivalent to $a_t \in [c, d] \subset (\frac{1}{2}, 1)$ for $c = \frac{e^u}{1+e^u}$ and $d = \frac{e^v}{1+e^v}$). For Player $T_1 + 1$,

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

If $\ln \left(\frac{a_{T_1+1}}{1-a_{T_1+1}} \right) \leq v$, 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 $T_1 v - \delta T_2 < 0$ and set $T = T_1 + T_2$. We claim that if $\ln \left(\frac{a_t}{1-a_t} \right) \in [u, v]$ for each $t \in \{1, \dots, T\}$, then $\Pr[\omega = 0 | (a_1, \dots, a_T)] > \Pr[\omega = 1 | (a_1, \dots, a_T)]$. To see that, note that the first T_1 players have signals with log likelihoods no larger than v (otherwise one would choose an action with log odds above v), and the next T_2 have signals with log

likelihoods no larger than $-\delta$. Since $T_1 v - \delta T_2 < 0$, Bayesian beliefs after T periods ascribe higher probability to $\omega = 0$ than to $\omega = 1$. Finally, note that since in each state the distribution of signals has full support, this event has positive probability. ■

Proof of Proposition 4 Choose $k \in (\frac{1}{2}, 1)$ and let $K = \ln\left(\frac{k}{1-k}\right) > 0$. When $\omega = 0$, with positive probability $\ln\left(\frac{s_1}{1-s_1}\right) \geq K$, $\ln\left(\frac{s_2}{1-s_2}\right) \geq K$, and $\ln\left(\frac{s_3}{1-s_3}\right) \geq 0$, in which case $\ln\left(\frac{a_1^A}{1-a_1^A}\right) \geq K$, $\ln\left(\frac{a_2^B}{1-a_2^B}\right) \geq 2K$, and $\ln\left(\frac{a_3^C}{1-a_3^C}\right) \geq 3K$. We claim that if for each $t \geq 4$, $\ln\left(\frac{S_t}{1-S_t}\right) > -(t-3)K$, then for each t , $\ln\left(\frac{a_t}{1-a_t}\right) > tK$ as desired. We prove by induction the following statements about BRTNI A , B and C 's actions: for $t = 3\tau$, $\ln\left(\frac{a_t}{1-a_t}\right) > \left(\frac{t^2}{6} + \frac{t}{2}\right)K$; for $t = 3\tau - 1$ for some positive integer τ , $\ln\left(\frac{a_t}{1-a_t}\right) > \left(\frac{(t+1)^2}{6} + \frac{t+1}{6}\right)K$, and for $t = 3\tau - 2$, $\ln\left(\frac{a_t}{1-a_t}\right) > \left(\frac{(t+2)^2}{6} - \frac{t+2}{6}\right)K$. Indeed, for $\tau = 1$, $\ln\left(\frac{a_3}{1-a_3}\right) > \left(\frac{3^2}{6} + \frac{3}{2}\right)K = 3K$, $\ln\left(\frac{a_2}{1-a_2}\right) > \left(\frac{3^2}{6} + \frac{3}{6}\right)K = 2K$, and $\ln\left(\frac{a_1}{1-a_1}\right) = \left(\frac{3^2}{6} - \frac{3}{6}\right)K = K$ as desired. It is equally straightforward to verify the claim for $\tau = 2$. Suppose that these inequalities hold for each $\tau \leq k$. Consider BRTNI A who moves in period $3k+1$ and whose accumulated signals are, for $k \geq 3$, $S_{3k+1} = \sum_{t \leq 3k+1} \ln\left(\frac{s_t^A}{1-s_t^A}\right) > -K + K - 4K - 7K - \dots - (3k+1-3)K = -K - 3K\left(\frac{(k-1)k}{2}\right)$. Using the fact that BRTNI $3k+1$ chooses $\ln\left(\frac{a_{3k+1}}{1-a_{3k+1}}\right) = \ln\left(\frac{a_{3k}}{1-a_{3k}}\right) + \ln\left(\frac{a_{3k-1}}{1-a_{3k-1}}\right) + S_{3k+1}$ and the induction hypothesis,

$$\begin{aligned} \ln\left(\frac{a_{3k+1}}{1-a_{3k+1}}\right) &> \left(\frac{(3k)^2}{6} + \frac{3k}{2}\right)K + \left(\frac{3k^2}{6} + \frac{3k}{6}\right)K - K - 3K\left(\frac{(k-1)k}{2}\right) \\ &> \left(\frac{3}{2}k^2 + \frac{5}{2}k + 1\right)K = \left(\frac{(t+2)^2}{6} - \frac{t+2}{6}\right)K > tK \end{aligned}$$

as desired. The proofs for BRTNI B moving in period $3k+2$ and BRTNI C moving in $3k+3 = 3(k+1)$ follow exactly the same lines and are therefore omitted. Together, these three formulas establish that when the claim holds for each $\tau \leq k$, it holds for $\tau = k+1$ too. Finally, because the proof that when $\Pr\left[\ln\left(\frac{S_t}{1-S_t}\right) > -(t-3)K, \forall t \geq 3 \mid \omega = 0\right]$ follows identical lines to the proof that $\Pr\left[\ln\left(\frac{S_t}{1-S_t}\right) > -tK, \forall t \geq 3 \mid \omega = 0\right]$ contained in the Proof of Proposition 1 and is therefore omitted. ■

References

- Avery, Christopher and Peter Zemsky.** “Multidimensional Uncertainty and Herd Behavior in Financial Markets.” *American Economic Review*, 88(4), Sept. 1998: 724-748
- Banerjee, Abhijit.** “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics* 107(3), August 1992: 797-817.
- 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.
- Bohren, J. Aislinn.** “Information-Processing Uncertainty in Social Learning” Mimeo, April 2009.
- Callender, Steven and Johannes Hörner.** “The Wisdom of the Minority”, *Journal of Economic Theory* (2009).
- 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
- Çelen, Bogaçhan and Shachar Kariv.** “An Experimental Test of Observational Learning under Imperfect Information”, *Economic Theory* 26: 677-699.
- Chamley, Christophe P.** *Rational Herds: Economic Models of Social Learning*. Cambridge University Press, 2004.
- 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.
- 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
- Eyster, Erik and Matthew Rabin.** “Cursed Equilibrium.” *Econometrica* 73(5), Septem-

ber 2005: 1623-1672.

Eyster, Erik and Matthew Rabin. "Naive Herding" Mimeo, August 2008.

Eyster, Erik and Matthew Rabin. "Rational and Naive Herding". CEPR Discussion Paper DP7351, July 2009.

Glosten, LR and Paul Milgrom. "Bid, ask and transaction prices in a market-maker market with heterogeneously informed traders". *Journal of Financial Economics*, 1985

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

Guarino, Antonio and Philippe Jehiel. "Social Learning with Coarse Inference". World Economy & Finance Research Programme Discussion Paper 0050, July 2009.

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

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

Lee, In-Ho. "'On the Convergence of Informational Cascades,'" *Journal of Economic Theory*, Vol.61, 1993, 395-411.

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

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

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

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

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.

Weizsäcker, Georg "Do we follow others when we should? A simple test of rational expectations". *American Economic Review*, forthcoming.