

Rational and Naive Herding

Erik Eyster
Department of Economics
LSE

Matthew Rabin
Department of Economics
University of California—Berkeley

15 June 2009

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

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 Asaf Plan and 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, UCL, UCSD and Yale. 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 rational inference in social-learning settings. 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 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. Although everyone eventually chooses the same action, nobody confidently believes in its correctness because each understands that the herd, no matter how long, indicates only two signals favouring that action. 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 every person imitates others' behavior and ignores her own information, an “information cascade”. This means that the outcome is socially inefficient: despite an infinity of private signals that reveal the right action, people herd on the incorrect action with positive probability.

These fully rational models provide 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 fully rational theory of herding has implications well beyond these key insights, including some that seem unrealistic (or at least untested)

despite relating to some core behavioral and welfare issues in the literature. While most such implications have been presented in previous papers, some have been obscured by the limited set of environments widely studied, and generally have been less heralded. In this paper, we propose a simple model of naive inference in games, and show how such naive inference leads to very different implications than the fully rational model. To aid in seeing how this form of inferential naivety might contribute to the theory of herding and to help see the intuition for why it might matter so much, we now review some of the implications of full rationality in herding.

In fully rational models, information cascades occur once the information contained in a group's observed actions becomes so great that nobody's private information can ever affect his optimal action. Herds fail to eventually aggregate information if and only if an information cascade begins before the truth is revealed. This can occur in the long run only if the environment is significantly "coarse". As is widely recognized (see, e.g., Lee (1993) and Smith and Sørensen (2001)), *either* richer action spaces *or* richer signal spaces immensely reduce the probability that herds will form on the wrong action. Rich-enough action spaces ensure that actions can reveal private information. Rich-enough signal spaces allow people to add new information to their predecessors' actions. In the limit in both cases, herders converge on the efficient action. In fact, the literature contains numerous extensions adding realism to the basic model that imply that rational herds converge to certainty on the correct action.

Even in environments coarse enough to allow for non-aggregation, rational herds never *hurt* anyone in expectation: observing others's choices only helps. Rational herding is inefficient relative to pooling people's information, not relative to what they could achieve in isolation. Moreover, rational-herding models make the twin predictions that highly-likely-incorrect herds involve very unconfident beliefs and that the likelihood of beliefs being very confident

yet wrong is very low. 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. While the failure (in coarse environments) to aggregate information is an important insight, fully rational models predict that there are no observable circumstances where herds either reliably lead people astray or induce widespread confidence in false theories.

Yet apart from the informational and efficiency properties predicted by full rationality, it is worth examining more closely exactly what type of herding behavior it actually predicts. Although obscured in the canonical binary setting, rational herding generally exhibits very strong “recency” effects. Because the most recent action combines new information with the information contained in all prior actions, the subsequent mover should broadly ignore all but this action. In the continuous model developed below, in fact, each rational herder completely ignores all but his immediate predecessor. Importantly, when this ignore-all-but-the-most-recent-action principle fails due to coarseness, it can be violated in *either* direction: by imitation *or* “anti-imitation” of past actions. Depending on highly model-specific details, rational people may be more prone to play B following observed actions $AAAB$ than following $BBBB$. And consider a variant of basic herding models where $n > 1$ people move simultaneously every period, each getting independent private information and observing all previous actions. Below we show in a continuous model that, fixing behavior in period 2, the more confidence period-1 actions indicate in favour of a hypothesis, the *less* confidence people in period 3 will have in it. Since the multiple movers in period 2 *each* use the information contained in period-1 actions, to properly extract the information from period-2 actions without counting this correlated information n -fold, period-3 players must imitate

period-2 actions but subtract off period-1 actions.¹ And if period-2 agents do not sufficiently increase their confidence relative to period 1 after observing the collection of period-1 actions, this means that they each received independent evidence that the herd started in the wrong direction. When $n > 2$, if *all* $2n$ people in the first two periods indicate roughly the same confidence in one of two states, this means a rational period-3 agent will always conclude that the other state is more likely! All said, while both theoretical and experimental researchers have studied the types of environments where rational inference happens to predict imitation, in a wide range of exceedingly plausible environments it leads to either full neglect of non-recent actions or seemingly counterintuitive forms of anti-imitation.²

Especially in light of the surprisingly weak link between rational inference and the type of imitative behavior that seems the behavioral inspiration for herding models, it seems worthwhile to consider departures from full rationality that might lead to more natural imitative behavior.³ While many different departures might lead to seemingly more realistic

¹Tautologically, full information extraction means that each agent be positively influenced by all prior *signals*. The negative influence of some past *actions* is exactly because these early signals appear in magnified form in later, imitated actions.

²Another very natural class of environments where more systematic imitation may seem more likely—and recency effects impossible—is when people observe previous actions but not their order. Yet Callender and Hörner (2009) beautifully illustrate how, when prior actions’ order is unobservable, rational inference can quite readily lead (would-be) herders to follow the minority of previous actions rather than the majority. In fact, our example in the text amplifies their result about non-imitation: when everybody in the first two periods take the same actions, knowing their order is obviously irrelevant. Hence, while Callender and Hörner (2009) emphasize the role heterogeneity and order unobservability to get their “wisdom-of-minorities” result, in our alternative setting neither is necessary for this even more striking “folly-of-uniform-unanimity” result.

³While there may be environments where full rationality can be reconciled with more systematically imitative behavior, the important point is whether in the many realistic environments where it cannot, and yet imitation occurs. Moreover, many attempts to reconcile full rationality with general imitation—e.g., introducing substantial heterogeneity of preferences, so that recent actions may be aberrational and earlier

behavior, it turns out that a simple form of inferential naivety—one which seems realistic and of economic consequence in many other settings as well—predicts a more intuitive form of imitative behavior in herding.

In particular, we explore the implications of the assumption that players in games, by not attending fully to the strategic logic of the setting they are in, naively believe that each previous player’s actions reflect solely that player’s private information. This error confronts the logic of fully rational herding at its core: in the canonical binary model, a fully rational herder who observes 100 people go in sequence to restaurant A infers nothing more than two signals in favour of A. In the extreme form of naivety we model, such a person acts as if she has observed 100 signals for A. Behaviorally, this simple alternative leads directly to a propensity to imitate *all* observed previous behavior. But it also leads to very different implications for the informational and efficiency properties of herds. Naive players can herd on incorrect actions even in the many rich environments where fully rational players always converge to the correct ones.⁴ They can become extremely and wrongly confident about the actions are informative—create environments where fully rational players always learn the truth. If on-the-equilibrium-path changes in behavior pervasively induce imitation in successors, it is because the behavior is pervasively informative.

⁴In fact, we have worked out some examples not included in this paper illustrating rational and naive inference in rich economic settings. In a simple model 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 (even when signals are binary and each agent makes the binary choice between “buy” and “don’t buy” at the market-clearing price). We show that 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 lengths. 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 herders will fully learn restaurant quality by observing their predecessors’ queue/restaurant choices. Naive herders instead might become convinced the wrong restaurant

state of the world even in environments where rational players never become confident.⁵ And depending on the cost of overconfidence, inferential naivety can lead people to so over-infer from herds as to be made worse off on average by observing others actions.

In Section 2, we define for all Bayesian games a simple model of inferential naivety. In doing so, we first define 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. This severe failure of contingent thinking implies that herders simply ignore their predecessors’ actions. 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 infer from still earlier actions.⁶ While Eyster and Rabin (2005) and Eyster and Rabin (2008) define stronger forms of each of these two failures of strategic thinking, in this paper we apply only very weak versions sufficient to make unique predictions in simple herding settings.

To illustrate starkly the effect that inferential naivety has on social learning, in Section 3 we present a model where each player receives a signal from a continuum ranging from fully revealing to uninformative to (occasionally) fully misleading, and then after observing all previous actions chooses an action from a continuum that fully reveals his beliefs. In this environment, rational herders combine their private signals with the information contained in their immediate predecessors’ actions into actions that fully reveal beliefs. As a result, they converge to full confidence on the correct state of the world. By contrast, even in this

is better, generating permanent, costly queues for an inferior establishment.

⁵It also leads to identifiable circumstances where social beliefs change over time in predictable directions.

⁶Although cursed players rely too little on their predecessors’ actions, while inferentially naive players seemingly do so too much, the two concepts are not opposites, even in this environment. Inferential naivety implies not merely that herders infer too much from previous actions but also that they place far too much weight on early relative to late signals.

filthy rich environment, with positive probability naive herders converge to fully confident belief 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. Naive inference by the fourth mover, in turn, causes him to inadvertently count 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 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 early signals happen to be misleading, limit herds may so over-use them as to outweigh an infinite sequence of further signals, some of which are even in isolation arbitrarily strong, and converge to extreme actions in the wrong direction. In simulations of our main example, 11% of the time the herd converges to fully confident beliefs in the wrong state.

In Section 4 we analyze three variants of this social-learning setting, all assuming the same continuous signal and action space. First, we consider a case where each of many 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 herders 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. In this case, herding is unambiguously harmful to them in the long run. In Section 4 we also formally develop the implications discussed above for simultaneous moves by modifying the setting of Section 3 to add multiple agents each period receiving independent signals. Rational agents (per usual) always learn the true state—but only with the very exotic inferences that weight some past actions negatively rather than positively. We show (per usual) that naive herders will always weight past actions positively, and (per usual) may converge to full

confidence in the wrong state.

We then consider the case where, instead of observing all previous moves, each player observes only her immediate $k < \infty$ predecessors. Rational herders efficiently aggregate all information and converge even in this setting. We show that if $k = 1$, meaning that each player observes *solely* her immediate predecessor, naive herders (per *unusual*) behave just like rational ones and converge to the truth. Despite having the wrong theory of those beliefs' provenance, as always each naive herder correctly infers her predecessor's beliefs. But now, due to unobservability, she cannot double count earlier movers' signals; all signals are counted exactly once. When $k > 1$, however, players once more can and do overcount early signals (in the limit, when $k = 2$, by the golden ratio!), and again with positive probability they converge to the wrong limiting belief and action.

In Section 5, we return to the setting of Section 3 to 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”—leading players to under-infer from previous actions—and show in various ways that limited cursedness does not undo the core result, in that 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—and establish an interesting form of robustness. When fully informed about others' types, rational players who observe a lot always uncover the truth. Cursed players simply ignore others' signals, and inferentially naive players may again anchor with full confidence on the wrong state if the first few signals mislead.

The model of inferential naivety in this paper not only omits many other realistic types of errors people make, but is itself extreme, and hence presumably leads to some important mispredictions in its own right.⁷ As such, we doubt that the model will accommodate the

⁷One way it clearly lacks realism is by *totally* excluding strategic sophistication implied by the fully

findings of any particular experiment as well as do formal or informal theories of errors inspired by and defined in those same specific experimental contexts. Yet we hope that by sharpening the implications of the rational model and providing an equally general, portable alternative model that has ready-made, zero-degrees-of-freedom implications across different settings, this paper helps foster greater ambitions for theoretical and empirical research on herding.

This paper neither attempts to tightly fit existing evidence nor to compare our predictions systematically to other theories of departure from purely rational play. But in Section 6 we briefly discuss how our model may help to interpret existing and potential future experimental evidence on herding, arguing that there is evidence of play that is indicative of both forms of error introduced in Section 2. We also compare our model briefly to other theories of non-Bayesian play. We discuss the model’s relationships with the notion of “persuasion bias” as modeled by DeMarzo, Vayanos, and Zwiebel (2003), which similarly captures the notion of limited sophistication in social learning. Most notably, while we discuss how in broader settings the two models differ, in the settings of this paper our model corresponds exactly to the existing literature on “Level- k ” reasoning when seeded with the appropriate assumptions about Level-0 play. As such, our results can be read as providing new implications for these types of models in social-learning settings.

rational model. The necessarily far more complicated and more tenuous solution concepts defined in Eyster and Rabin (2008) permit a greater ability to study assumptions besides the two extreme ones compared in this paper. Eyster and Rabin (2008) also discuss of the interesting but problematic feature of our model that (as typical of quasi-rational models) people who make the mistakes we posit observe patterns of behavior that utterly surprise them, so that sophisticated updating might lead them to learn their mistake or otherwise change their model of the world.

2 Cursed and Naive Inference

In this section, we introduce two ways that people in strategic settings may err when inferring information—first, an extreme form of inferential neglect, and second the form of inferential naivety that is the primary focus of the paper. While Eyster and Rabin (2006, 2008) define stronger and more generally applicable variants of these solution concepts and more carefully explore their foundations, here we provide weaker and simpler versions that suffice for the herding settings that we explore in this paper. 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 more negative information about the value of the good. In a “fully cursed equilibrium”, each player correctly predicts the distribution over her opponents’ actions but entirely neglects the relationship between those opponents’ actions and private information; she plays a best response to those beliefs.

Cursed equilibrium provides a strong, general, and formulaic method of capturing the premise that players fail to fully reason through the information content in others’ actions. As it happens, however, making sharp predictions in social-learning models requires neither the full apparatus of cursed equilibrium nor the strong and tenuous equilibrium assumptions incorporated into that concept. Because players move sequentially, observe all their predecessors’ actions, and do not care directly about any other player’s actions, mistaken beliefs about one another’s actions would not affect play. Cursed equilibrium’s key implication for social learning is that players infer too little from their predecessors’ actions because they don’t fully think through how those predecessors condition their actions on information. Its most extreme and unrealistic variant has players infer nothing from previous play and simply

follow their own signals.

To formalize this idea and especially so as to characterize partial rather than full informational neglect, consider a Player l who follows Player k , and let $p(s_k|s_l)$ be Player l 's beliefs about Player k 's signal s_k conditional on her own signal s_l . In a social-learning environment, signals are correlated through the state of the world: when A is a better restaurant than B , then both Players k and l likely to receive signals that A is better. Cursed beliefs are simplest to describe when Player k uses an invertible signal-contingent strategy $a_k(s_k)$. A rational Player l who knows Player k 's strategy would infer s_k from observing a_k . By contrast, a fully cursed Player l infers nothing about s_k from a_k and forms posterior beliefs $p(s_k|a_k, s_l) = p(s_k|s_l)$. A χ -cursed Player l who observes a_k forms the posterior beliefs that $p(s_k|a_k, s_l) = (1 - \chi) + \chi p(s_k|s_l)$ and $p(s'_k|a_k, s_l) = \chi p(s'_k|s_l)$ for $s'_k \neq s_k$.⁸ When $\chi = 0$, this formula collapses to rational inference. When $\chi = 1$, this formula collapses to *fully cursed* posterior beliefs, which coincide with interim beliefs. When $\chi \in (0, 1)$, *partially cursed* successors do partially but not fully update about Player k 's signal from his action: they form posterior beliefs closer to their interim beliefs than they should. Player l plays a *cursed best response* when he plays a best response to fully cursed beliefs about his predecessors' signals. We interpret players who play cursed best responses not as best responding to erroneous beliefs that their opponents employ type-non-contingent strategies but rather as attempting to maximize their payoffs while inadvertently ignoring the relationship between others' actions and signals. Player l plays a χ -cursed *best response* when he plays a best response to χ -cursed beliefs about his predecessors' signals. Here too we interpret players who play χ -cursed best responses not as best responding to erroneous beliefs that their opponents randomize between type-contingent and type-non-contingent strategies but rather as attempting to maximize their payoffs while underappreciating the relationship between

⁸Eyster and Rabin (2005) provide a formal statement of this result as Lemma 1 and its proof.

others' actions and signals.

Eyster and Rabin (2005) show how an equilibrium version of χ -cursed best response fits a set of laboratory experiments better than Bayesian Nash equilibrium, for *any* value of $\chi \in (0, 1]$. For instance, buyers in experimental “lemons” markets succumb to the “winner’s curse” of buying when better not to buy; failing to appreciate how the seller’s willingness to sell depends upon the value of the asset leads them to overestimate the value of the asset conditional on sale. The laboratory evidence suggests that people struggle with contingent thinking, e.g., “what information would be conveyed by a seller accepting my offer for the car”.⁹

Stepping outside of formal models, there is reason to doubt people will be *severely* cursed in the social-learning settings we study in this paper. Most of Eyster and Rabin’s (2005) evidence for cursedness comes in the context of simultaneous-move games, where intuition for it is strongest: players may more severely neglect the informational content of others’ behavior 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’ behavior—e.g., when responding to others’ dropping out of a sequential auction. While people in social-learning contexts are likely to be partially cursed, the seeing-is-inferring intuition seems to accord with the psychology of contingent thinking. Insofar as cursedness as modelled by Eyster and Rabin (2005) reflects failures of contingent thinking, observing actions may mitigate cursedness while very much not overcoming the form of naivety to which we now turn—and which is the main focus of the paper.

Yet the primary reasons why we focus on naive inference are that it matters more to long-

⁹In an experiment that evokes this very type of lemons situation, but with only a single player, Charness and Levin (2008) nicely illustrate how the type of failure of contingent thinking incorporated into cursed equilibrium can extend to non-strategic settings in a way not formally captured by cursed equilibrium.

run learning and leads to striking and novel results across a broad array of social-learning environments. And as Section 5 suggests, the patterns of herding it produces in social-learning environments hold even in hybrid models where people are partially cursed at the same time as inferentially naive, as well as those that incorporate even a majority of rational and cursed players.

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 other players condition their actions on type yet neglects that other players understand the same thing. 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. We say that a player engages in *best response trailing naive inference (BRTNI) play* if he plays a best response to all his predecessors’ playing cursed best responses.¹⁰ “BRTNI players” neglect that their predecessors make informational inferences from observing their own predecessors’ actions. Analogous to cursed best response, we interpret BRTNI play as a form of limited attention: Player k simply neglects to reason through how Player i makes informational inferences from Player $j \neq i$ ’s actions.

Eyster and Rabin (2006, 2008) define solution concepts where players are partially but not fully inferentially naive. These formalizations resemble partial cursedness in modeling players who incompletely reason through how other players make informational inferences. While the extreme form of naive inference that we analyze here is surely too extreme, using it greatly simplifies our analysis without altering the qualitative results.

¹⁰“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).

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 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 , BRTNIs continue to update their beliefs that A is the better restaurant. In the extreme model we formulate, they converge to certainty that the chosen restaurant is the better one. Several experimental studies suggest that public beliefs become stronger than the rational prediction, notably Goeree, Palfrey, Rogers and McKelvey (2007) and Kübler and Weizsäcker (2005).

While existing herding experiments provide evidence for the presence of inferential naivety, most of them take place in contexts like that of the previous paragraph where the rational model predicts little more than imitative behavior and where its differences from BRTNI are least identifiable and matter least. The remainder of this paper focuses on the broader array of settings where the two models predict behavior far more nuanced than a propensity to imitate and differ dramatically from each other.

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 of rational learning's producing 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.¹¹ 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 could or does 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 $q = E[\omega|I_k] = \Pr[\omega = 1|I_k]$, Player 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 when $\omega = 0$ without any more assumptions about the model, e.g., the nature of players' information or action spaces.

Proposition 0 $\Pr[I_k|\omega = 0] \leq \frac{1-\pi}{\pi} \frac{1-q}{q}$.¹²

The maximum probability that Player k can hold information causing him to believe that $\omega = 1$ with probability q , when in fact $\omega = 0$, cannot exceed $\frac{1-\pi}{\pi} \frac{1-q}{q}$. 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. The result derives entirely from the logic of single-person decision making and holds

¹¹Public 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.

¹²This is a special case of Chamley's (2004) Proposition 2.9.

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 $n \in \mathbb{N}$ be the common set of $n + 1$ actions available to all players, and assume that each player chooses the action closest to q .¹³ Let S be the set of signals, which for simplicity we take to be denumerable, and let t be the strength of the weakest private signal, \underline{s} , in favour of $\omega = 1$ that occurs with positive probability.¹⁴ Also for simplicity, assume that each player observes all of her predecessors' actions and that actions converge.

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 $1 - t$.¹⁵ 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 three-action model ($n = 2$)

¹³As formalized below, Player k acts this way when her payoff function is $g_k(a; \omega) = -(a_k - \omega)^2$.

¹⁴Formally, $\arg \min_{s \in S} \Pr[\omega = 1 | s] = \underline{s}$ and $\Pr[\omega = 1 | \underline{s}] = t$.

¹⁵If a herd formed on $a = 1$ with public beliefs less than $1 - t$, then eventually some player would receive the signal \underline{s} and choose $a = 0$, a contradiction.

(where the middle action $a = \frac{1}{2}$ represents a safe choice for someone uncertain whether $\omega = 0$ or $\omega = 1$), the likelihood of a false herd falls to less than two percent. For $n = 3$, 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.

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 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.¹⁶

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

¹⁶By 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: 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.

and identically distributed conditional on the state.¹⁷ 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 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 $a \equiv (a_1, a_2, \dots) \in [0, 1]^N$ 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 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,

¹⁷We use S_k to denote the Player k 's signal as a random variable and s_k its realization.

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 $\ln\left(\frac{a_t}{1-a_t}\right) = \sum_{\tau < t} \ln\left(\frac{s_\tau}{1-s_\tau}\right)$. 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 feature of rational social learning—that each person fully ignores all but the most recent actions—is in fact quite striking and special to this model; Section 4.2 introduces a variant of the model where this property of rational learning fails. While rational players in models with common preferences tend to imitate their immediate predecessors, whether they ignore, imitate, or anti-imitate their predecessors' predecessors depends highly upon the context.

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. Second, by choosing actions in the continuum, players reveal their posteriors to their successors. As suggested by Lee (1992) and Smith and Sørensen (2001), either of these features suffices to guarantee 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 in fact under not-very-strong assumptions, BRTNI players may herd on wrong beliefs.

Proposition 1: Suppose that $E\left[\ln\left(\frac{S}{1-S}\right)\middle|\omega = 0\right]$ and $\text{var}\left(\left[\ln\left(\frac{S}{1-S}\right)\middle|\omega = 0\right]\right)$ are finite. 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 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 16 makes it *easier* to obtain the conclusion of Proposition 1. The finite-mean and variance assumptions ensure that the likelihood ratio has “thin tails”, ruling out, for instance, cases with a positive-probability signal 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.

The assumptions of Proposition 1 also imply that BRTNI beliefs converge almost surely to zero or one.¹⁸ BRTNI players who do not learn the true state become *fully confident* in

¹⁸The Appendix contains a proof. While there exist non-generic counterexamples described by Eyster and

the wrong state! To illustrate Proposition 1 and differentiate between BRTNI and 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 $\omega = 1$.

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.¹⁹ For each, the likelihood

Rabin (2008), convergence to certain beliefs is a generic feature of BRTNI play across learning models and constitutes another key difference from rational models.

¹⁹Since BRTNI and BNE coincide for the first two players, these should be the same; the small differences are an artifact of the simulation techniques.

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 happens 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 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 we have not formally explore this issue, the observation suggests that BRTNI play converges faster than rational play.

Although BRTNI converges fast, the next proposition establishes an interesting result in the rare event that 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 2: 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 $\omega = 1$ greater than 50%—but less than 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 50% 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.²⁰

Proposition 1 demonstrates 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 observing herds tends strongly to lead to over-confident 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 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

²⁰The result resembles the weak-beliefs-are-probably-wrong result developed in Rabin and Schrag's (1999) 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 ; 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 $n \geq 2$ 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 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 and wrong beliefs play out more generally. We begin by showing that naive long-run players who in fact each receive an infinite string of signals may be so misled by the herd that they get things wrong in the limit—doing worse under learning than they would by confining their actions to their private information. Next, we modify our original setting to allow more than one player to move simultaneously in a given period. The rational model in this setting predicts striking forms of anti-herding: rational people always increase confidence in a hypothesis when some predecessors' actions reveal *less* confidence in that hypothesis and sometimes choose actions to disagree with their immediate predecessors. BRTNI, however, behaves in more intuitive ways, herding in much the same way as our benchmark model. Next we modify the environment to limit how much information players have available to them and show how the result of Proposition 1 is robust to severe observational limitations.

4.1 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 now 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 by simply 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 the benchmark model, except now with only three players $\{A, B, C\}$ who move in sequence $A, B, C, A, B, C, A \dots$. As before, in each period t , the player on the move receives a private signal about the state and can observe all of her predecessors' actions. Maintaining our assumptions on signals, each player's private information 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. In the rational model, nothing depends upon the identities of the various players, who almost surely converge to correct limiting beliefs and actions.

Proposition 3: Suppose that three long-run BRTNI players $\{A, B, C\}$ move in sequence $A, B, C, A \dots$, that $E \left[\ln \left(\frac{S}{1-S} \right) \mid \omega = 0 \right]$ and $var \left(\left[\ln \left(\frac{S}{1-S} \right) \mid \omega = 0 \right] \right)$ are finite. Then in BRTNI play, there exists $\delta > 0$ such that $\Pr[\lim_{t \rightarrow \infty} a_t = 1 \mid \omega = 0] > \delta$.

Despite each holding a collection of signals that identifies the state, each player may end up choosing the wrong action. Play in the first three periods exactly resembles that of the baseline model, meaning that Player C overweights A 's signal s_1 in her first move. Play in the fourth period also coincides with the baseline model as A neglects that B and C

already have incorporated s_1 into their actions. The first difference emerges in period 5, where B , having chosen a_2 himself, knows that a_2 already embodies the information in a_1 and therefore does not re-count it. Letting $\ln\left(\frac{a_t}{1-a_t}\right) = \sum_{\tau \leq t} F_\tau \left(\frac{s_\tau}{1-s_\tau}\right)$ in the baseline model, BRTNI play is described by $F_\tau = 2F_{\tau+1}$, whereas here $F_\tau = F_{\tau+1} + F_{\tau+2} + F_{\tau+3}$ as shown in the Appendix. In the limit as $\tau \rightarrow \infty$, $F_\tau/F_{\tau+1}$ approaches $\psi \simeq 1.839$. 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.

4.2 Herding with Simultaneous Moves

We now turn to a very natural variant of the baseline model used above and throughout the literature. Instead of assuming that people move purely sequentially, we assume that each period two or more people might move simultaneously. In the baseline model, rational players imitate their immediate predecessors while ignoring their predecessors' predecessors. This stark recency effect does not hold in all models of social learning. We now present a model where rational players condition their actions on more than their immediate predecessors' actions—but in many cases by anti-imitating those earlier actions. BRTNI players always base their actions on as many past moves as they can observe and always imitate all predecessors rather than anti-imitate any.

We modify our baseline model to have $n \geq 1$ players move simultaneously in every period; everything else stays the same. As before, all players' signals are *iid* conditional on the state; in particular, the n players moving in period t receive different signals.

We begin by analyzing Bayesian Nash Equilibrium. Let a_t^j be the action of Player j in period t , which depends upon her private signal s_t^j as well as all previous actions.²¹ It can be shown that rational herders will choose actions to obey

²¹Note that, in the formulas that follow, we apply the convention that $0^0 = 1$ but $0^k = 0$ for all $k > 0$.

$$\ln \left(\frac{a_t^j}{1 - a_t^j} \right) = \ln \left(\frac{s_t^j}{1 - s_t^j} \right) + \sum_{i=1}^{t-1} (-1)^{i-1} (n-1)^{i-1} \left[\sum_{k=1}^n \ln \left(\frac{a_{t-i}^k}{1 - a_{t-i}^k} \right) \right]$$

To simplify exposition, let $A_t = \sum_{k=1}^n \ln \left(\frac{a_t^k}{1 - a_t^k} \right)$, the sum of period- t log odds ratios or *aggregate period- t action*, and $S_t = \sum_{k=1}^n \ln \left(\frac{s_t^k}{1 - s_t^k} \right)$, the sum of period- t signals or *aggregate period- t signal*. Using this notation and summing across j in the formula above gives

$$A_t = S_t + n \sum_{i=1}^{t-1} (-1)^{i-1} (n-1)^{i-1} A_{t-i}.$$
²²

When $n = 1$, this reduces to the familiar $A_t = S_t + A_{t-1} = \sum_{\tau \leq t} S_\tau$. When $n = 2$,

$$A_t = S_t + 2 \sum_{i=1}^{t-1} (-1)^{i-1} A_{t-i},$$

leading to $A_1 = S_1$, $A_2 = S_2 + 2A_1$, $A_3 = S_3 + 2A_2 - 2A_1$, $A_4 = S_4 + 2A_3 - 2A_2 + 2A_1$, etc.

Actions in this model depend upon past actions in unusual ways. While rational herders always imitate their immediate predecessors, they anti-imitate their predecessors' predecessors. Imagine how someone observing two parties simultaneously choosing Restaurant A would react to learning that both parties had been advised to go by the same source; she would lose confidence in A being the right choice. This drives the negative weight on predecessors' predecessors here. The core logic of rational herding differs fundamentally from pure imitation, even in a model with common preferences. Rational players carefully attend to the order of their predecessors moves, so the very same restaurant choice by Xavier yesterday that today conveys good news to Yali about Restaurant A *tomorrow* conveys negative information to Zinedine—*fixing* Yali's behavior.

Substituting for A_{t-i} recursively gives

$$A_t = S_t + n \sum_{i=1}^{t-1} S_{t-i},$$

²²Public beliefs following date t are $\sum_{i=1}^t (-1)^{i-1} (n-1)^{i-1} A_{t-i}$

where player in period t gives all signals weight one; hence, the aggregate period- t action puts weight one on s_τ^j if $\tau = t$ and weight n if $\tau < t$. Because they incorporate all past signals with equal weights, aggregate actions converge almost surely to the state. Despite wild swings in *how* rational players interpret past behavior, they learn the state eventually. Note, importantly, that the wild swings in how people use past actions won't typically show up in actions: the key is that recent actions always receive positive weight, and *typically* those actions are more extreme than earlier actions. It is when play does not converge fast enough that we would observe rational players switching.

When $n = 3$,

$$A_t = S_t + 3 \sum_{i=1}^{t-1} (-1)^{i-1} 2^{i-1} A_{t-i},$$

leading to $A_1 = S_1$, $A_2 = S_2 + 3A_1$, $A_3 = S_3 + 3A_2 - 6A_1$, $A_4 = S_4 + 3A_3 - 6A_2 + 12A_1$. The swings here are even more dramatic, amplified by exponential growth in the weights on prior actions. For instance, Player 3 strongly anti-imitates Player 1 while Player 4 even more strongly imitates Player 1. People's beliefs also move in counterintuitive ways. Consider the case where the three players in the first period all choose $a = 0.6$, each expressing 60% confidence that $\omega = 1$. If all second-period players also were to choose $a = 0.6$, then since $A_2 = S_2 + 3A_1 = A_1$, $S_2 = -2A_1 = -2S_1$, meaning that in a log likelihood sense there is twice as strong evidence for $\omega = 0$ than for $\omega = 1$. Someone who observes her six predecessors all indicate 60% confidence that $\omega = 1$ rationally concludes that there is only a 25% chance that A is better!

BRTNI players engage in far smoother play and form seemingly more intuitive beliefs, with Player j in period t choosing

$$\ln \left(\frac{a_t^j}{1 - a_t^j} \right) = \ln \left(\frac{s_t^j}{1 - s_t^j} \right) + \sum_{i=1}^{t-1} \left[\sum_{k=1}^n \ln \left(\frac{a_{t-i}^k}{1 - a_{t-i}^k} \right) \right].$$

As in the baseline model, BRTNI players interpret each past action as reflecting solely the

actor's private information and consequently weight them all equally. As before, this leads to massive over-counting of early signals:

$$\ln \left(\frac{a_t^j}{1 - a_t^j} \right) = \ln \left(\frac{s_t^j}{1 - s_t^j} \right) + n \sum_{i=1}^{t-1} (n+1)^{i-1} \sum_{k=1}^n \ln \left(\frac{s_{t-i}^k}{1 - s_{t-i}^k} \right).$$

Summing across players in period t gives the aggregate behavior

$$\begin{aligned} A_t &= S_t + n \sum_{i=1}^{t-1} A_{t-i} \\ A_t &= S_t + n \sum_{i=1}^{t-1} (n+1)^{i-1} S_{t-i}. \end{aligned}$$

When $n = 1$, we recover BRTNI play in our baseline model

$$\ln \left(\frac{a_t^j}{1 - a_t^j} \right) = \ln \left(\frac{s_t^j}{1 - s_t^j} \right) + \sum_{i=1}^{t-1} 2^{i-1} \ln \left(\frac{a_{t-i}}{1 - a_{t-i}} \right).$$

When $n = 2$,

$$A_t = S_t + 2 \sum_{i=1}^{t-1} 3^{i-1} A_{t-i}.$$

In general, players in period t give each aggregate signal $n + 1$ times as much weight as its successor (so long as that successor dates before period t). Relative to our baseline model—or $n = 1$ above—players give progressively more weight to early signals. However, since increasing n increases the total informativeness of signals in any period, this does not imply that BRTNI players are more likely to misidentify the state in the limit than when $n = 1$: in fact, a simulation gives the estimate that when $n = 2$, the probability of converging on a false herd drops from 11% to 8%, and as $n \rightarrow \infty$ BRTNI should learn the truth in period 2 by the sheer force of the Law of Large Numbers. For $n < \infty$, however, we get the now-familiar result.

Proposition 4: Suppose that $n \in \mathbb{N}$ BRTNI players move in each period, that $E \left[\ln \left(\frac{S}{1-S} \right) \mid \omega = 0 \right]$ and $\text{var} \left(\left[\ln \left(\frac{S}{1-S} \right) \mid \omega = 0 \right] \right)$ are finite. Then in BRTNI play, there exists $\delta > 0$ such that $\Pr[\lim_{t \rightarrow \infty} a_t = 1 \mid \omega = 0] > \delta$.

4.3 Learning with Limited Observation

In the baseline model, BRTNI players overweight the signals of each mover by counting them again and again through each predecessor's action. 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 5: Suppose that BRTNI players can only observe their immediate predecessors' actions. 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, 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, whereas in fact it contains all the prior signals.

While Proposition 5 establishes the ironic result that severely curtailing naive players' observation can lead them to correct inference, 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. We state formally and prove the result for the case where BRTNI players observe only their two immediate predecessors but stress that it holds when players can observe their immediate k predecessors for *any* $k > 1$. Indeed, the more predecessors they observe, the more likely BRTNI players are to herd wrongly, and as $k \rightarrow \infty$ the result converges to that of Proposition 1.

Proposition 6: Suppose that BRTNI players can only observe their two immediate pre-

decessors' actions, that $E \left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right]$ and $\text{var} \left(\left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] \right)$ are finite. Then in BRTNI play, there exists $\delta > 0$ such that $\Pr[\lim_{t \rightarrow \infty} a_t = 1 | \omega = 0] > \delta$.

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 (approximately 1.618). 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 1 essentially applies 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 $x^3 - x^2 - x - 1$.

The result that BRTNI players may converge to wrong limiting beliefs and actions also holds when players cannot observe the order of their predecessors' play.

Proposition 7: Suppose that BRTNI players can observe all their predecessors' actions (but not necessarily the order) as well as that $E \left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right]$ and $\text{var} \left(\left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] \right)$ are finite. Then in BRTNI play, there exists $\delta > 0$ such that $\Pr[\lim_{t \rightarrow \infty} a_t = 1 | \omega = 0] > \delta$.

Rational players' actions depend upon the order of their predecessor's moves. For instance, Player 3 wishes 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. 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, the result in Proposition 4 continues to hold when BRTNI players cannot observe the order of their two immediate predecessors' moves.

Together, Propositions 4 and 5 establish important robustness properties 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 as people may simply receive summary statistics of their predecessors' actions.²³

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 behavior in addition to inferential naivety. 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 behavior of a mix of BRTNI, cursed, and rational players.

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—

²³Eyster and Rabin (2008) includes a model of a changing world where a sequence of people once more observe all previous actions but must each decide whether to privately pay for cheap and up-to-date full information. Whereas rational herders periodically buy information once the informational value of the herd has decayed sufficiently (due to the changing state), naive herders believe all predecessors, including recent ones, have bought information and refrain from buying their own. Naive herders may forever mimic the initially appropriate action, despite the fact that eventually the initial conditions eventually lose all relevance to the current state, all under the mistaken belief that recent actions contain up-to-date information.

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.

We now investigate the extent to which cursedness overturns our main result. As before, we allow continuous signals but 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.

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 do converge, in this case to one-half; no player infers anything from her predecessors’ moves. More generally, Proposition 9 characterizes what happens to public beliefs in the limit. Let m_0 be the median signal when $\omega = 0$, i.e. $\Pr [s \geq m_0 | \omega = 0] = \frac{1}{2}$, and p_t be public beliefs in period t .

Proposition 8: Consider the binary-action model where $E \left[\ln \left(\frac{S}{1-S} \right) | \omega = 0 \right]$ and $var \left(\left[\ln \left(\frac{S}{1-S} \right) | \omega = 0 \right] \right)$ are finite and $\chi < 2m_0$. In χ -cursed BRTNI play, there exists some $\delta > 0$ such that $\Pr \left[\lim_{t \rightarrow \infty} p_t = 1 - \frac{\chi}{2} | \omega = 0 \right] > \delta$.

Proposition 8 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 early movers 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. Since future BRTNIs are even more likely to choose $a = 1$ after seeing an extra predecessor do so, the number of players choosing $a = 1$ less those choosing $a = 0$ is bounded from below by a random walk with positive drift. With positive probability, this process never returns to its current position, and public beliefs converge to $1 - \frac{\chi}{2}$.

Because $m_0 < \frac{1}{2}$, Proposition 9 does not apply to cases where players are very cursed. Nevertheless, the degree of cursedness compatible with wrong limiting beliefs can be substantial. In our example above where $f_0(s) = 2(1 - s)$ and $f_1(s) = 2s$, χ need only be smaller than $2m_0 = 2 - \sqrt{2} \simeq 0.59$ for there to be positive probability that public beliefs approach the wrong limit.

The analysis above showed that a limited form of our main result holds when all players are a mixed inferentially naive and (boundedly) cursed. We now explore the robustness of our predictions in a different way. In particular, we assume that 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 9: Suppose that $E \left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right]$ and $var \left(\left[\ln \left(\frac{S}{1-S} \right) \middle| \omega = 0 \right] \right)$ are finite. and that for some $n \in \mathbb{N}$, Players $n, 2n, \dots$ are BRTNI and the remainder cursed or rational. Then there exists some $\delta > 0$ for which $\Pr \left[\lim_{t \in \mathcal{B}, t \rightarrow \infty} a_t = 1 \middle| \omega = 0 \right] > \delta$.

Proposition 9 asserts that our main qualitative result—that with positive probability, BRTNI players adopt wrong limiting beliefs and actions—remains valid no matter what 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. 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.²⁴ Cursed players follow their own signals throughout. This combination of non-convergent cursed players and differently-convergent rational and BRTNI players seems an interesting implication of the underlying logic of the model, although probably unrealistic. Weakening the strong assumption that rational players know their predecessors’ types surely would slow down their convergence and should increase the plausibility of aggregate behavior.

6 Discussion and Conclusion

The purpose of this paper is to explore the implications in social-learning environments of a new theory of error in inference. Besides obvious limits to the scope of a single paper, there are two reasons why we do not concentrate on detailed analysis of experimental evidence. First, to the best of our knowledge the continuous model, especially with simultaneous moves, where our solution concept’s predictions depart most strikingly from those of the

²⁴Indeed, 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.

rational model has never been tested in the laboratory. Second, our simple solution concept omits too many other kinds of errors in strategic reasoning backed by ample intuition and evidence, and is formulated in an unrealistically extreme form itself. Yet we conclude with a brief discussion that relates our concept to existing evidence and speculate on how it might combine and compare with other theories of boundedly rational play.

Recent laboratory tests support some predictions of the rational model while also uncovering several systematic discrepancies. In a meta-study of 13 experimental datasets, Weizsäcker (2008) finds strong evidence that subjects systematically follow their predecessors far less than they should, as measured by the empirical distribution of payoffs. Researchers often attribute this pattern to “overconfidence”, whereby subjects overestimate their own private signals’ precision and therefore overuse them. Yet the psychological interpretation or foundation for such overconfidence is not always clear, and 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 others’ signals, whereas overconfidence says that subjects overuse their own signals; both lead to relative over-weighting of one’s own signal. But “overconfidence” seems to have little *a priori* psychological plausibility in these contexts, and we are unfamiliar with any direct evidence 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-interpret their private information, 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 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 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 experimental evidence on social learning fits neither overconfidence nor cursedness—but rather seems more in line with inferential naivety. Estimating a model that includes what we identify as cursed and BRTNI types among others, Kübler and Weizsäcker (2004) find evidence that most subjects behave most like BRTNI players. 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.

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. The fact that QRE leads players to correct limiting beliefs (and is also inconsistent with harmful herding), stands in marked contrast to the main results of our paper.

Another natural comparison for our model is “persuasion bias” as modelled in DeMarzo, Vayanos, and Zwiebel (2003), who study 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 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.

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). 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 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.²⁵

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 err only 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 behavior 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 uncivil. The catch is that uncivil behavior—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

²⁵Camerer, 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.

must form not only beliefs about the quality of the two restaurants but also about t 's behavior inside. To take the simplest case, suppose that uncivil behavior generates a very large positive externality—*schadenfreude*—outweighing 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 behavior 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 may or may not herd. Certainly, cursed best response does not make the unique Level-1 prediction of herding.

These different predictions seed different Level-2 and BRTNI predictions. Level-2 players assume no informational content in observed herds and therefore follow their own signals. BRTNI play is indeterminate but does not exclude herding. Eyster and Rabin's (2008) stronger INIT concept predicts herding for exactly the same reason that underlies it in the models of this paper.

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—with crowding or other externalities—the models' predictions do differ, and so it seems important to explore the frequency and implications of informational versus non-informational mistakes.

7 Appendix: Proofs

Proof of Proposition 0 From Bayes' rule,

$$\begin{aligned} \Pr[\omega = 1|I_k] &= \frac{\pi}{\pi + (1 - \pi) \frac{\Pr[I_k|\omega=0]}{\Pr[I_k|\omega=1]}} \\ \Rightarrow \frac{\Pr[I_k|\omega = 0]}{\Pr[I_k|\omega = 1]} &= \frac{1 - \pi}{\pi} \frac{1 - q}{q}, \end{aligned}$$

where $I_k = (s_k; a_1, \dots, a_{k-1})$. Because $\Pr[I_k|\omega = 1] \leq 1$, $\Pr[I_k|\omega = 0] \leq \frac{1-\pi}{\pi} \frac{1-q}{q}$. ■

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

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

or $p \geq \frac{t(2n-1)}{1-2t+2nt}$. Substituting $p = \frac{t(2n-1)}{1-2t+2nt}$ for q and $\pi = \frac{1}{2}$ into Proposition 0's bound 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 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, $\Pr\left[\ln\left(\frac{s_t}{1-s_t}\right) > -tK \mid \omega = 0\right] > \frac{t^2 K^2 - \sigma^2}{t^2 K^2}$, where $\sigma^2 = \text{var}\left[\ln\left(\frac{S}{1-S}\right) \mid \omega = 0\right] + \left(E\left[\ln\left(\frac{S}{1-S}\right) \mid \omega = 0\right]\right)^2$, which is finite by assumption. Hence,

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

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

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

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

Lemma 1: Under the assumptions of Proposition 1, BRTNI actions and beliefs converge almost surely to 0 or 1.

Proof: From above, write

$$2^{1-t}P_t = \sum_{\tau < t} 2^{-\tau} \ln \left(\frac{s_\tau}{1-s_\tau} \right). \quad (7.1)$$

Since the three series

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

converge—the inequalities following from the assumptions of Proposition 1 (using Chebyshev’s inequality to obtain the third)—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 2 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 $\omega = 1$. ■

Proof of Proposition 3 Superscripts denote player identity. BRTNI play in the first three periods does not change, from whence

$$\begin{aligned} \ln \left(\frac{a_4^A}{1 - a_4^A} \right) &= \ln \left(\frac{a_3^C}{1 - a_3^C} \right) + \ln \left(\frac{a_2^B}{1 - a_2^B} \right) + \ln \left(\frac{s_1^A}{1 - s_1^A} \right) + \ln \left(\frac{s_4^A}{1 - s_4^A} \right) \\ &= 4 \ln \left(\frac{s_1^A}{1 - s_1^A} \right) + 2 \ln \left(\frac{s_2^B}{1 - s_2^B} \right) + \ln \left(\frac{s_3^C}{1 - s_3^C} \right) + \ln \left(\frac{s_4^A}{1 - s_4^A} \right), \end{aligned}$$

and

$$\begin{aligned} \ln \left(\frac{a_5^B}{1 - a_5^B} \right) &= \ln \left(\frac{a_4^A}{1 - a_4^A} \right) + \ln \left(\frac{a_3^C}{1 - a_3^C} \right) + \ln \left(\frac{s_2^B}{1 - s_2^B} \right) + \ln \left(\frac{a_1^A}{1 - a_1^A} \right) + \ln \left(\frac{s_5^B}{1 - s_5^B} \right) \\ &= 7 \ln \left(\frac{s_1^A}{1 - s_1^A} \right) + 4 \ln \left(\frac{s_2^B}{1 - s_2^B} \right) + 2 \ln \left(\frac{s_3^C}{1 - s_3^C} \right) + \ln \left(\frac{s_4^A}{1 - s_4^A} \right) + \ln \left(\frac{s_5^B}{1 - s_5^B} \right). \end{aligned}$$

Continuing in this way, $P_t = \sum_{\tau < t} F_\tau \ln \left(\frac{s_\tau}{1 - s_\tau} \right)$, where $F_\tau = F_{\tau+1} + F_{\tau+2} + F_{\tau+3}$ and $F_{t-1} = 1$; the weight on $\ln \left(\frac{s_t}{1 - s_t} \right)$ converges to ψ times the weight on $\ln \left(\frac{s_{t+1}}{1 - s_{t+1}} \right)$, where $\psi \simeq 1.839 > \frac{9}{5}$ is the root of $x^3 - x^2 - x - 1 = 0$. Since for large T , $\sum_{\tau > T} \left| F_\tau \ln \left(\frac{s_\tau}{1 - s_\tau} \right) \right| < \sum_{\tau > T} \left| \left(\frac{9}{5} \right)^{-\tau} \ln \left(\frac{s_\tau}{1 - s_\tau} \right) \right|$, a.s. divergence of P_t follows from the argument in Lemma 1 above, replacing every instance of 2 with $\frac{9}{5}$ in (7.1) and proceeding as before. When public beliefs converge to $p \in \{0, 1\}$, actions must converge to p too; hence actions converge a.s. to zero or one. It remains to show that when $\omega = 0$ a converges with positive probability to 1. Let $K = 2\sqrt{\text{var} \left[\ln \left(\frac{S}{1-S} \right) | \omega = 0 \right] + \left| E \left[\ln \left(\frac{S}{1-S} \right) | \omega = 0 \right] \right|^2}$ and choose $T \in \mathbb{N}$ such that for each $t > T$, $F_\tau / F_{\tau+1} > \frac{9}{5}$. Since $\ln \left(\frac{s_t}{1 - s_t} \right)$ has full support, with positive probability $P_T \geq K$. For public beliefs to converge to zero, we must have $P_t < 0$ for some $t > T$, i.e.,

$\sum_{\tau>T}^{\infty} \left(\frac{9}{5}\right)^{T-\tau} \ln\left(\frac{s_{\tau}}{1-s_{\tau}}\right) < -K$. From Kolmogorov's maximal inequality for random series,

$$\begin{aligned} & \Pr \left[\max_{k \geq T} \left| \sum_{\tau>T}^k \left(\frac{9}{5}\right)^{T-\tau} \ln\left(\frac{s_{\tau}}{1-s_{\tau}}\right) \right| > K \mid \omega = 0 \right] \\ & \leq \frac{\text{var} \left[\ln\left(\frac{S}{1-S}\right) \mid \omega = 0 \right]}{\left(2\sqrt{\text{var} \left[\ln\left(\frac{S}{1-S}\right) \mid \omega = 0 \right]} + \left| E \left[\ln\left(\frac{S}{1-S}\right) \mid \omega = 0 \right] \right|\right)^2} < \frac{1}{4}. \end{aligned}$$

This implies that when $\omega = 0$ $P_t \rightarrow +\infty$ with positive probability, which implies that when $\omega = 0$ $a \rightarrow 1$ with positive probability as desired. ■

Proof of Proposition 4 The proof follows the same lines as that of Proposition 3 and hence is omitted.

Proof of Proposition 5 Clearly $\ln\left(\frac{a_1}{1-a_1}\right) = \ln\left(\frac{s_1}{1-s_1}\right)$ in BRTNI play, just as in BNE. For $k > 1$, BRTNI k deems all actions a_{k-1} possible for some s_{k-1} and hence infers that $s_{k-1} = a_{k-1}$. Because she can observe no more than her immediate predecessor's actions, $\ln\left(\frac{a_k}{1-a_k}\right) = \ln\left(\frac{a_{k-1}}{1-a_{k-1}}\right) + \ln\left(\frac{s_k}{1-s_k}\right)$. Solving the recursion gives $\ln\left(\frac{a_k}{1-a_k}\right) = \sum_{\tau=1}^k \ln\left(\frac{s_{\tau}}{1-s_{\tau}}\right)$, as in BNE. ■

Proof of Proposition 6 The proof follows the same lines as that of Proposition 3 and hence is omitted.

Proof of Proposition 7 Since BRTNI players do not heed the order of their predecessors' moves, it follows from Proposition 6.

Proof of Proposition 8 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 , $\frac{1}{2} = \Pr[s \geq m_0 \mid \omega = 0] < \Pr[s \geq \frac{\chi}{2} + \varepsilon \mid \omega = 0]$. With positive probability, the first K players get signals above $\frac{1}{2}$, with K large enough that public beliefs are $1 - \frac{\chi}{2} - \varepsilon$. Player $K + 1$ chooses $a = 1$ with

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

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

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

$\Pr[\exists \hat{t} > K : n(\hat{t}) = K]$ is less than that under random walk with $\Pr[a_t = 1] = p \forall t$. Since this is less than one, with positive probability $n(t) > K \forall t > K$. In this case, public beliefs cannot converge to $\frac{\chi}{2}$. Notice that $(1 - \chi) \cdot 1 + \chi \cdot \frac{1}{2} = 1 - \frac{\chi}{2}$ is an upper bound for χ -cursed BRTNI public beliefs. Since in this case every BRTNI believes that all her predecessors receives signals above $\frac{1}{2}$, public beliefs converge to this upper limit. Hence, when $\omega = 0$ public beliefs converge to $1 - \frac{\chi}{2}$ with positive probability. ■

Proof of Proposition 9 We express public beliefs of the k th 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 $\kappa \in \{2, \dots, k-1\}$ and $\tau \in \{1, \dots, n-1\}$, $\frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} \leq \frac{n}{n-\frac{1}{2}}$. 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 $P_{\kappa n}$ once plus the number of rational players among Players $\{(\kappa-1)n + \tau + 1, \dots, \kappa n - 1\}$ —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 $k \rightarrow \infty$,

$$\begin{aligned} \frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} &\leq \lim_{m \rightarrow \infty} \frac{(2^m - 1)(n-1) + 2^{m-1}}{(2^m - 1)(n-1) + 2^m} \\ \frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} &\leq \lim_{m \rightarrow \infty} \frac{2^m \ln 2(n-1) + 2^{m-1} \ln 2}{2^m \ln 2(n-1) + 2^m \ln 2} \\ \frac{\alpha_{(\kappa-1)n+\tau}}{\alpha_{(\kappa-2)n+\tau}} &\leq \lim_{m \rightarrow \infty} \frac{n - \frac{1}{2}}{n} = \frac{n - \frac{1}{2}}{n}, \end{aligned}$$

where the second line comes from L'Hôpital's Rule. From here, we can adapt the argument from the proof of Proposition 3, using the fixed factor $\frac{n}{n-\frac{1}{2}}$ in place of $\frac{9}{5}$, and treating the signals in blocks of n . ■

References

- 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.
- Chaness, Gary and Dan Levin.** “The Origin of the Winner’s Curse: A Laboratory Study” *AEJ: Micro* 1(1), February 2009: 207-236.
- 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. "Inferential Naivety." Mimeo, November 2006.

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

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

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 expectation". Mimeo, April 2008