

Extensive Imitation is Irrational and Harmful*

Erik Eyster and Matthew Rabin[†]

September 7, 2013

Abstract

Rationality leads people to imitate those with similar tastes but different information. But people who imitate common sources develop correlated beliefs, and rationality demands that later social learners take this redundancy into account. This implies severe limits to rational imitation. We show that (i) in most natural observation structures besides the canonical single-file case, full rationality dictates that people must “anti-imitate” some of those they observe; and (ii) in *every* observation structure full rationality dictates that people who do not anti-imitate can, in essence, imitate at most one person among predecessors who share common information. We also show that in a very broad class of settings, virtually any learning rule in which people regularly *do* imitate more than one person without anti-imitating others will lead to a positive (and, in some environments, arbitrarily high) probability of people converging to confident and wrong long-run beliefs. When testing either the rationality or the efficiency of social learning, researchers should not focus on whether people follow others’ behavior—but instead whether they follow it too much. (JEL B49)

*We thank Paul Heidhues, seminar participants at Berkeley, DIW Berlin, Hebrew University, ITAM, LSE, Oxford, Penn, Tel Aviv University, and WZB, as well as the editor and referees, for their comments. We are grateful to the Russell Sage Foundation for its hospitality, and Katie Winograd and especially Claire Gabriel for their help with the historical literature on treating syphilis and other medical practices, and Tristan Gagnon-Bartsch and Min Zhang for excellent research assistance.

[†]Eyster: Department of Economics, LSE, Houghton Street, London WC2A 2AE, UK. Rabin: Department of Economics, UC Berkeley, 530 Evans Hall #3880, Berkeley, CA 94720-3880, USA.

1 Introduction

Inference from the behavior of others is one of the ways that people guide themselves in making wise social, economic and medical choices. Diners (to use the canonical example) decide where to eat partly by the popularity of restaurants. Investors infer good financial strategies from others' portfolios. And each new generation of doctors and individuals infers effective medical practice from the behavior and beliefs of their elders. Despite interacting with other means of information acquisition—communication, experimentation, observation of others' success and failure—learning by observing other people's actions and beliefs is an important facet of social and economic behavior. It has formed the basis for an extensive and ongoing research program: beginning with Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), the literature on observational learning identifies how a rational person who observes the behavior of another person with similar tastes and private information may follow that person, even contrary to her own private information. *Imitation is rational.*

In this paper we show some striking ways that the logic of social inference requires that rational agents greatly limit their imitation. When people comprising the societies, groups, or markets whose behavior is being watched *themselves* are learning by observing the actions of a common set of people before them, the information in their behavior contains a great deal of redundancy. This intrinsic redundancy implies severe limits to rational imitation. In most natural environments besides the canonical single-file case, full rationality dictates that people must “anti-imitate” some of those they observe. And in the remaining environments, it dictates that each individual must imitate a group of predecessors who share common information by treating them, in a way we formalize, as if they were only one person. *Abundant imitation is irrational.*

We also show that in a very broad class of settings, learning rules in which people regularly *do* imitate more than one person without anti-imitating others will lead to a positive probability of people converging to confident but wrong long-run beliefs. Although there are ways that individuals can behave irrationally that lead to socially optimal long-run information aggregation, any irrational form of social learning where a sequence of individuals each

violate the imitate-as-if-only-one-person principle will lead to long-run inefficiency—in the form of confident but wrong shared social beliefs. In many environments, this probability is very high. *Abundant imitation is socially harmful.*

To illustrate many principles of our more abstract general analysis in a context closer to the type of settings that inspires our analysis, we present in Section 2 an extended example of medical (mis)learning. In it, doctors treating an illness experiment with an infinity of medications, each of which can be good or bad. Good medications lead to recovery more frequently than bad medications, but patients treated with either stochastically recover and succumb. Doctors learn about medicine quality from two sources: their record of success with their own patients, and the medicines prescribed by their predecessors (but not their predecessors’ success). In this setting, rational doctors are guaranteed to eventually settle on a good medicine. But their path there exhibits several counter-intuitive properties. First, they do very limited imitation. A medicine used with success by an early doctor may then be adopted by an enormous number of subsequent doctors without conveying any additional information; rational doctors understand this, and abandon the medicine with only moderate evidence of its badness. Second, doctors frequently “anti-imitate” their predecessors: fixing a medicine’s popularity last year, the more popular it was two years ago, the less effective a rational doctor judges it today. These intrinsic implications of rationality not only seem counterfactual, but are crucial to the efficiency of long-run medical learning. We show that doctors who instead fail to appreciate the redundancy in their predecessors actions can be quite likely to permanently adopt an ineffective drug.

In Section 3, we flesh out the logic of observational learning more generally within a class of situations satisfying a condition that we call “impartial inference”. In contrast to partial inference, these are structures where common knowledge of rationality implies that any player who learns anything about a predecessor’s signal learns *everything* relevant about that signal. Impartial inference places a joint restriction on the players’ action spaces and observations. We follow Lee (1993) in working with a continuous action space that fully reveals players’ posterior beliefs, and describe in the Appendix a condition on observations

that guarantees impartial inference. Although it does not drive our qualitative results, working in settings with impartial inference allows us to crisply articulate of our results. Indeed, in this framework we identify surprisingly simple necessary and sufficient conditions for anti-imitation. Roughly speaking, an observational structure generates anti-imitation if and only if it contains a quartet of players i, j, k, l where (i) j and k both observe i , (ii) neither j nor k observes the other, and (iii) l observes j, k , and i . Intuitively, Player l must imitate Players j and k to use their signals, but then must anti-imitate Player i , whom j and k have both already imitated, in order to avoid “double-counting” i ’s private information.¹ Many natural settings include such configurations, for instance those where in every period $n > 1$ people move simultaneously, each observing all predecessors who moved in prior periods, as well as those where every player observes all but her most recent $m \geq 1$ predecessors, perhaps due to a time-lag in actions being reported. In fact, the canonical single-file model is one of the very few structures that does not generate anti-imitation.

Our main result is that rational players who do not anti-imitate perform very limited imitation. Specifically, absent anti-imitation, no player can imitate any two players who commonly observe a shared predecessor. Because in most social-learning environments, any two people beyond the first two movers whom a person observes share some common observation, rational social learners who do not anti-imitate do not imitate more than one person. In addition, we show that in any game that includes a quartet of players satisfying the conditions above, there is a positive probability of a signal sequence which leads at least one player to form beliefs opposite to both the observed beliefs of everybody she observes and her own signal. Intuitively, if Player i is observed to believe strongly in a hypothesis and Players j and k only weakly, then Player l must infer that j and k both received negative information so that altogether the hypothesis is unlikely.

In Section 4, we explore what happens when social learners anti-imitate too little and imitate too much because they fail to account for the redundancy in their predecessors’ actions. Someone who neglects redundancy and observes many people 60% confident in the

¹The result is rough because the players’ observations (or lack thereof) described in parts (i) and (ii) can be “indirect” in a sense made precise in Section 3.

effectiveness of a new hybrid seed will become more than 60% confident herself. We formalize such “redundancy neglect” by assuming that there is some number of predecessors, N , and overweighting factor, $c > 0$, such that whenever a player’s N immediate predecessors all raise their actions by some amount x , the player raises her own action by at least $(1 + c)x$. For example, if $N = 10$ and $c = 1$, then every player treats her 10 immediate predecessors as if they embody two pieces of conditionally independent information. People who act this way exponentially overweight early actions and signals, which allows early evidence to inefficiently trump all later evidence. We show that any society of redundancy neglecters converges with positive probability to complete confidence in the wrong state of the world.

We conclude in Section 5 by putting this paper in broader context, including how and why the limits to imitation we identify here extends quite generally outside our specific environment, including in cases where people communicate beliefs directly. When testing *either* the rationality of social learning *or* for whether observed (rational or irrational) social learning leads to correct versus overconfident-but-wrong social beliefs, researchers should investigate not whether people follow others’ behavior or listen to others’ advice—but whether they do so too much.

2 Mercurial Beliefs

From the late 15th century, until the discovery and adoption of penicillin in the mid-20th century, syphilis wreaked havoc on the world. Tens of millions of people suffered from the disease, and millions died from it.² Hyman (1941) estimates that as late as 1941, 4.4% of Americans were infected, causing 40,000 deaths per year.³ From the 16th to early 20th century, the leading treatment for syphilis was mercury, despite its high cost and horrible and

²Even today, millions of people are infected every year, with much less harm in wealthy countries, and harm in developing countries due to economic rather than medical reasons.

³This infection rate resembles estimates for soldiers of the 19th century. Acton (1846) estimates that 4% of troops quartered in Britain were infected with syphilis at any time, while Parascandola (2008) estimates that during the American Civil War, 12% of Union Army soldiers were treated for syphilis.

lethal side effects. Although it was understood at least as far back as the beginning of the 19th century that mercury led to poisoning and renal disease (Wells (1812), cited by George (2011)), mercury was used for two reasons: to alleviate symptoms, and to rid the patient of the *treponema pallidum* bacterium responsible for the disease.⁴ Although we have not found convincing evidence that mercury truly did relieve symptoms (much less do so more than the symptoms it created), when applied topically its anti-inflammatory properties may have soothed sores, and its toxic properties may have inhibited cell growth, including new sores (O’Shea (1990)). Most of the debate, however, centered on the curative properties of mercury, and much of the mercury taken for syphilis was taken orally. Despite evidence from medical experiments that mercury could kill the *treponema pallidum* bacterium when both were injected into the human body simultaneously (Tongue (1801)), no strong evidence shows that mercury as employed by doctors killed syphilis in patients already infected.⁵ Many doctors, notably those working in military hospitals who kept the best statistics of their day, observed that patients who were untreated or unaggressively treated fared better than those treated by mercury (Ferguson (1813), Devergie (1831)). Some, merely suggestive, evidence comes from the famously unethical Tuskegee trials. From 1932 to 1972 the U.S. army studied 600 African American men, withholding treatment from the two-thirds of the

⁴The 19th century medical literature contains over 400 reports of mercury intoxication (Goldwater (1972)).

⁵According to O’Shea (1990) p. 392,

“Mercury is a potent diuretic and in toxic doses it induces salivation. It was thought that by inducing diuresis and salivation that the syphilitic ‘virus’ would be excreted, aborting the illness. Th[is] premise is a fallacious one and resulted in grievous clinical errors.”

Doctors could have been right that mercury worked but wrong about why. Despite the seeming consensus in its ineffectiveness, it is surprisingly difficult to judge mercury’s efficacy. The long-term and stochastic nature of both the effects of syphilis and the good and bad effects of mercury makes the lack of randomized controlled trials or epidemiological studies that convincingly address selection effects in who received mercury especially unfortunate. (Perhaps surprisingly, some evidence suggests positive selection in patients treated with mercury because doctors believed that sicker patients could not tolerate its side effects. See Colles (1881).) The voluminous scientific literature on syphilis and mercury (3000 articles on syphilis were published in nineteenth century medical journals, of which 1000 discussed mercury (O’Shea (1990)) has been inconclusive.

sample infected with syphilis. The mortality rate was 40% higher for the syphilitics in the study than for the non-syphilitics—but about 10% lower than matched syphilitics outside the study, who received standard pre-penicillin treatment, primarily mercury (Lasagna (1975)). Despite its dubious efficacy, mercury continued to be widely used until the discovery of penicillin.⁶ Why? Although it is remarkably hard to prove even to this day whether it was, on net, a bad idea, it seems very unlikely that the benefits outweighed the costs. If so, how could doctors get it wrong for over 400 years? And, if not, what can we make of the extreme faith in its curative powers without being able to discern its efficacy?

Doctors who have access to clinical trials—and *exclusively use that information*—can (if good at statistics) evaluate a medicine’s effectiveness by simply analyzing the evidence. But those who rely on the beliefs and practices of their predecessors must carefully attend to the logic of redundancy. We illustrate here in intuitive terms how rationality places severe limits on imitation and often mandates that people “anti-imitate” some predecessors.

Consider a society that experiments with an infinite number of medicines to potentially treat a disease, where each medicine is good with probability $k \in (0, 1)$. A diseased person treated with a good medicine recovers with probability x ; one treated with a bad medicine recovers with probability $y < x$. An untreated convalescent recovers with probability $u < y$.⁷ To illustrate the possibility that long-term mislearning can be very likely, we focus on the case where k , x , and y are all relatively low, so that doctors consider lots of potential medicines and get noisy feedback of effectiveness so that most cases of recovery with a medicine are “false positives”. Because we assume no statistical irrationality *per se*, we assume that *all doctors understand false positives*. Doctors observe other doctors’ prescriptions but not their success. This set-up guarantees an infinite number of maximally effective drugs. Long-run

⁶Mercury was partially superseded by the arsenic derivative salvarsan (compound 606) after its discovery in 1909. However, the medical profession did not believe salvarsan rendered mercury obsolete: standard practice was to combine salvarsan with mercury, and later to other heavy metals, such as the less-toxic bismuth. (See, for instance, the recommended best practice of the League of Nations (Martenstein (1935).)

⁷Mislearning is more dramatic when $u > y$, which may better fit the case of mercury and syphilis, but we concentrate on the case where $u < y$ for analytic simplicity.

harm from mislearning occurs whenever society settles upon an ineffective medicine.

In each period $t = 1, 2, \dots$, each of D doctors active for a single period treats some number of patients, where the size of each doctor’s clientele is an independent draw from a Poisson process: the doctor sees Z patients, where $\Pr(Z = z) = (1 - q)q^z$ for some parameter $q \in (0, 1)$. Doctors do not know how many patients they will treat until after treating the last. All doctors in period $t + 1$ observe the treatment history of all those in periods 1 to t but not the effectiveness of their treatments, which is the crux of observational learning. Each doctor attempts to maximize recovery amongst her own patients.⁸

Consider the behavior of a rational doctor in period 1 when $y \geq u$. Faced with a first patient, she tries some medicine at random. If her patient succumbs, then she switches to a random new medicine for any second patient and continues switching until observing her first recovery, at which point she continues with the “successful” medicine. Letting f be the number of failures observed after her first success, the doctor switches as soon as $\frac{x(1-x)^f}{y(1-y)^f} < 1$, namely whenever the odds ratio that the medicine is good relative to bad falls below 1. This implies that she switches whenever $f \geq f^*(x, y) := \frac{\log(x) - \log(y)}{\log(1-y) - \log(1-x)}$. When $x = 0.2$ and $y = 0.1$, for instance, $f^*(x, y) = 5.8$, meaning that the first doctor to successfully use a medicine abandons it if the next six patients fail to recover.

What will later doctors do, and where will long-run medical beliefs settle? Doctors continue to experiment with new medicines or those tried by earlier doctors; should they temporarily settle on a bad medicine, eventually some doctor sees enough patients to reject it. In essence, each doctor’s unboundedly large potential clientele enables arbitrarily strong signals, preventing herding on a bad medicine. Consequently, common knowledge of rationality implies that for each $1 \geq x > y \geq u \geq 0$, $D \geq 1$, $k > 0$, and $q > 0$, society eventually settles on a good medicine.

To succinctly describe behavior, we define a few categories of experience. Medicine m is *tried* by Doctor j if he treats one or more patients with it. It is *accepted* by Doctor j if he treats two or more patients with it, without switching to another medicine. It is *rejected*

⁸Many lessons developed here hold or are even strengthened when doctors try to maximize total survival, but the analysis is more complicated.

by Doctor j if he tries it before switching to another medicine. It receives *qualified initial endorsement* if the first period in which it is accepted by some doctor, it is accepted by only one, and used fewer than $1 + f^*(x, y)$ times.

Some features of behavior in Bayesian Nash equilibrium are intuitive. First, doctors use their private information: a doctor who tries a medicine successfully does not reject it immediately. Second, doctors imitate one another: a medicine rejected by at least one doctor in period t , and accepted by no other in that same period, is never tried by any subsequent doctor.⁹ And if, entering period t , some medicines have been accepted by at least one doctor without being rejected, then all doctors in period t will try one of these. These forms of imitation are, however, relatively weak.

Any doctor who has never seen a medicine receiving qualified initial endorsement used on more than $1 + f^*(x, y)$ patients by any previous doctor rejects it if her first $f^*(x, y)$ patients succumb. This reflects her understanding that doctors after the first treat multiple patients not due to success but rather imitation. When $q = \frac{1}{2}$, $x = \frac{1}{5}$, and $y = \frac{1}{10}$, even though the first accepted medicine will ultimately be rejected if bad, this usually takes quite a while, even in the extremely unlikely event that it never again yields a false positive. An average of 63 doctors try an accepted, bad medicine before its ultimate rejection, 31 of whom accept it, on average. Rationality requires that a doctor who observes an initial doctor and 31 followers all use a medicine twice or more without abandoning it should herself reject it after a mere six failures, despite a failure rate for the good medicine of 80%. When instead $q = \frac{1}{3}$, rational doctors reject the medicine upon the same six failures even after seeing, on average, 2100 doctors try the medicine, 730 of whom accept it, on average.

Full rationality requires doctors to be acutely aware that thousands of their predecessors' treatments convey no information whatsoever. We regard this intense attention to the extreme redundancy in treatments as most likely to break down in ways that matter. But below we also establish results about how rational social learning demands even more dramatic departures from imitation.

⁹This echoes Smith and Sørensen's (2000) "overturning principle": since the prior doctor observed everyone whom you do, plus his private information, without private information it is optimal to follow him.

Compare the following two situations, where again $x = \frac{1}{5}$ and $y = \frac{1}{10}$. Situation A comprises the following history: in period 1, 10 doctors saw patients; one tried the drug *Deppisol* and continued to prescribe it through his fourth and final patient; no other medicine tried in period 1 was accepted; in period 2, all 11 doctors with patients tried *Deppisol*, with two having only one patient, seven having two to five patients each without rejecting the drug, and two switching away from *Deppisol* after six applications. A rational doctor in period 3 knows that the patient in period 1 receiving *Deppisol* recovered—the reason his doctor accepted it—but infers no other positive evidence. The 11 doctors who tried it in period 2 were merely imitating, and the nine who used it more than once gave no indication of success—since it is rational to continue with *Deppisol* even after five patients succumb—but instead only failure, for the two doctors rejecting *Deppisol* after six patients reveal that all succumbed. Hence, the period 3 doctor would rationally switch drugs.

Situation B comprises the following history: in period 1, no doctor tried *Deppisol*, and no other medicine was accepted; in period 2, four doctors, each with two patients, tried *Deppisol*; three of them switched after the first patient, while one accepted *Deppisol*; no other medicine was accepted. In Situation B, a rational doctor in period 3 would use *Deppisol*. Although the three doctors rejecting it provide three bad signals, the one good signal from the one doctor who accepts *Deppisol* more than makes up for them, given the low recovery rate of a good medicine. In contrast to Situation A, where rationality demands very limited imitation, in Situation B it demands some degree of enthusiasm despite the high rejection rate. By comparing across situations, we find that rational behavior is anti-imitative. In each period, more doctors accepted and fewer doctors rejected *Deppisol* in Situation A than in Situation B. Nevertheless, a rational observer concludes that *Deppisol* is more likely to be effective in Situation B than Situation A. Early popularity of a medicine renders later popularity much less meaningful. Section 3 illustrates more striking forms of anti-imitation in other settings.

What happens when the logic of the redundancy in medical beliefs largely eludes doctors? Doctors who imitate too much because they neglect the redundancy in their predecessors' actions may fail to reject bad medicines after the very few failures amongst their own patients

that trigger rational rejection. The extreme form of naive inference modeled in Eyster and Rabin (2010), which corresponds to doctors’ neglecting that their predecessors’ actions depend upon more than their private information, virtually guarantees herds on false medicines for many parameter values. A doctor seeing (say) 700 doctors accept a medicine would judge it as nearly certainly good, even if hundreds of his own patients succumbed on it. As $q \rightarrow 0$, a society of such inferentially naive doctors almost certainly herds on the first accepted medicine. But as $k \rightarrow 0$, the first accepted medicine is nearly certainly bad, in which case this form of naivety nearly guarantees that society adopts a bad medicine.

Just as we do for more abstract environments in Section 4, we can show here that much weaker forms of redundancy neglect also lead to positive probability of long-run error; moreover, in this setting, social mislearning happens with very high probability. Suppose that all doctors use learning rules that satisfy three properties. First, a doctor who has observed no medicine either accepted or rejected continues to use a medicine that has not failed her. Second, no doctor uses a medicine that has been rejected but never accepted. Third, each doctor is “humble” in the sense of according comparable weight to her predecessors’ treatments and her own experience: a doctor who sees N doctors accept a given medicine, and none accept any other, prescribes that medicine so long as the number of her own patients who succumb on it, d , satisfies $\left(\frac{x}{y}\right)^N \left(\frac{1-x}{1-y}\right)^d > H$. Inferential naivety as modeled by Eyster and Rabin (2010) implies $H = 1$. Larger values of H allow for doctors to heavily downplay observed beliefs relative to their own experience. *If* each predecessor’s acceptance really did signal at least one additional recovery, then “humble” doctors would discount their predecessors’ treatments by neglecting multiple recoveries by any single doctor. However, in settings where q is low (doctors see few patients), this error pales in comparison to that of reading independent evidence into each predecessor’s acceptance of the medicine. Indeed, for each $1 > x > y > 0$, $D \geq 1$, $k > 0$, and $H > 0$, as $q \rightarrow 0$, the first medicine to accompany a recovery gets used forever. This happens even when most medicines are bad and, hence, doctors regard recovery as most likely a false positive. Since the probability that the first success comes from a bad medicine equals $\frac{1-k}{1+k}$, as $q \rightarrow 0$ and $k \rightarrow 0$, the probability that

doctors forever prescribe a bad medicine approaches 1. Section 4 shows how far weaker forms of redundancy neglect than the one here generate social mislearning.

Although the example above more closely parallels the structure of the next two sections, variations on it yield even more dramatic implications of rationality. Suppose that doctors move single-file but cannot observe the order of their predecessors' moves. In this case, the success-failure asymmetry of the signal structure yields a particularly simple and counterintuitive result: a doctor should try a previously tried medicine if and only if it no doctor has rejected it. Intuitively, no matter how many doctors have accepted the medicine, the doctor who initially tried before rationally rejecting the medicine signals stronger information. Even in this case, no bad drug will ever be used in the long run. A “classical herd” can form against a drug—society ceases learning with weak beliefs in its ineffectiveness—but never in favor.

A second variant comes when doctors observe only a random subset of those acting in the previous period. Suppose that 100 doctors act per period, where each one observes each doctor in the previous period with independent chance 0.9, but no doctors in earlier periods. Consider a doctor in period 13, say, whose period 12 observations comprise nine doctors prescribing Deppisol and 76 doctors prescribing different, and mutually distinct, medicines. When doctors who randomize over medicines do so independently, Deppisol is *extremely unlikely* to have been randomly chosen by nine. This makes the most likely history one where Deppisol was the medicine of choice going into period 11, when it was rejected by a single doctor, unbeknownst to the nine doctors using Deppisol in period 12. Now suppose that some of these nine accept Deppisol, while none reject it, and no doctor accepts any other medicine. Then our doctor in period 13 should not try Deppisol despite not seeing it rejected and instead seeing it (uniquely) accepted! In other scenarios of this world of “stochastic recent observation”, a bad medicine is most likely to be rejected *because* it is reasonably popular. This not only resembles some of the counterintuitive (and counterfactual) predictions of Section 3, but also points to the likelihood of the forms of inefficiency illustrated in Section 4: if doctors imitate rather than reject the most popular medicine after seeing it accepted by some and rejected by none, then society may very likely ultimately adopt a bad medicine.

3 Impartial Inference and the Limits of Imitation

In this section, we consider a (seemingly) very different environment, observation structures more general than those used in the classical models by Banerjee (1992) and Bikhchandani et al. (1992), where players move single-file after observing all of their predecessors' actions. For analytic tractability, like Lee (1993) we focus on environments where rational players completely extract all of the payoff-relevant information to which they have access.¹⁰ Proposition 6 in the Appendix provides sufficient conditions for common knowledge of rationality to imply such full extraction. In these settings, we provide necessary and sufficient conditions for rational social learning to include anti-imitation, meaning that some player's action decreases in some predecessor's observed action, holding everyone else's action fixed. We also show that absent anti-imitation, people cannot rationally imitate very much: no one can imitate two predecessors who share a common observation, unless she also anti-imitates.

There are two possible states of the world, $\omega \in \{0, 1\}$, each one *ex ante* equally likely. Players in the set $\{1, 2, \dots\}$ receive private information correlated with the state. Following Smith and Sørensen (2000), we work directly with players' updated private beliefs based upon their private information alone: let σ_k be Player k 's belief that $\omega = 1$ conditional upon her private information. In state ω , Player k 's private beliefs are drawn from the distribution $F_k^{(\omega)}$; players' beliefs are independent conditional upon the state. We assume that $F_k^{(0)}$ and $F_k^{(1)}$ are mutually absolutely continuous to rule out cases where some positive-probability private information reveals the state with certainty; let $[\underline{\sigma}_k, \bar{\sigma}_k] \subseteq [0, 1]$ denote the non-degenerate convex hull of their common support. Player k 's private beliefs are *unbounded* when $\underline{\sigma}_k = 0$ and $\bar{\sigma}_k = 1$ and *bounded* otherwise. To simplify exposition, we work with the log-odds ratio of private beliefs, $s_k := \ln\left(\frac{\sigma_k}{1-\sigma_k}\right)$, which we refer to as Player k 's (*private*) *signal*; let $[\underline{s}_k, \bar{s}_k]$ be the convex hull of its range.¹¹ Player k 's private signal indicates that $\omega = 1$ is at least as likely than $\omega = 0$ iff $s_k \geq 0$. Let $s^k := (s_1, s_2, \dots, s_k)$.

¹⁰A special case of Avery and Zemsky (1998) with binary signals and absent noise traders shares this feature.

¹¹When Player k 's private beliefs are unbounded, we abuse notation by writing $[\underline{s}_k, \bar{s}_k] = [-\infty, \infty]$.

We denote by $D(k) \subset \{1, \dots, k-1\}$ the subset of Player k 's predecessors whose actions k observes. When $k-1 \notin D(k)$, we can interpret Players $k-1$ and k as moving simultaneously; one player's having a lower index than another means simply that the former moves no later than the latter. Let $ID(k) \subset \{1, \dots, k-1\}$ be the subset of Player k 's predecessors whom k observes *indirectly*: $l \in ID(k)$ if and only if there exist some path of players $(k, k_1, k_2, \dots, k_L, l)$ such that each observes the next. A player can indirectly observe another through many such paths, a possibility that plays a crucial role in our analysis below. We refer to $\mathcal{N} = \{\{1, 2, \dots\}, \{D(1), D(2), \dots\}\}$ as an observation structure, consisting of the players $\{1, 2, \dots\}$ and their respective sets of observed predecessors, which define their sets of indirectly-observed predecessors.¹²

After observing any predecessors visible to her as well as learning her own private signal, Player k chooses the action $\alpha_k \in [0, 1]$ to maximize the expectation of $-(\alpha_k - \omega)^2$ given all her information, I_k . Player k does this by choosing $\alpha_k = \mathbb{E}[\omega | I_k] = \Pr[\omega = 1 | I_k]$, namely by choosing the action equal to her posteriors that $\omega = 1$. Any player who observes Player k can infer Player k 's updated beliefs but not necessarily k 's private signal. For simplicity, as with signals, we identify actions by their log-odds ratios, $a_k := \ln\left(\frac{\alpha_k}{1-\alpha_k}\right)$. Player k optimally chooses $a_k \geq 0$ iff she believes $\omega = 1$ at least as likely as $\omega = 0$ and would choose $a_k = s_k$ whenever observing no predecessor. Let $a^k = (a_1, \dots, a_k)$ and $a_{-j}^k = (a_1, \dots, a_{j-1}, a_{j+1}, \dots, a_k)$. Throughout this section, we assume that it is common knowledge that players are rational.

Although a player may observe a large number of predecessors, many of these observations turn out to be redundant. For instance, in the customary single-file structure, when the range of actions spans the range of posteriors (e.g., Lee (1993)), no player who observes

¹² \mathcal{N} can be viewed as a directed network, with players as nodes and observations as edges or links, that must be acyclic to respect the order of actions. (See, e.g., Jackson (2008).) Because network-theoretic language neither clarifies nor simplifies our results, we prefer the game-theoretic term "observation". A small literature in network economics (notably Bala and Goyal (1998) and Golub and Jackson (2010)) differs from our work in three important ways: networks are undirected; players take actions infinitely often, learning from one another's past actions; and players are myopic.


her immediate predecessor gains any useful information from observing any more distant predecessor, whose signal has already been used by the immediate predecessor.

In other settings, however, players’ immediate predecessors do not provide “sufficient statistics” for earlier movers indirectly observed. In Section 2, a rational doctor in period 3 attends to more than second-period actions because she wishes to strip out the common correlation from second-period actions.

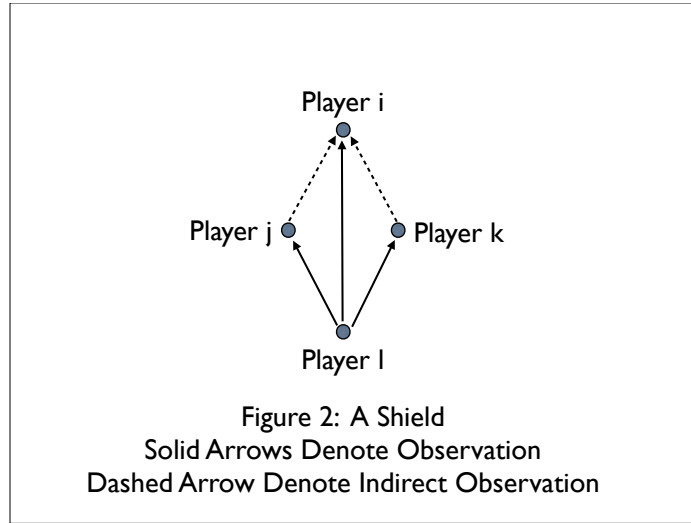
It turns out that observation structures with common correlation that lack sufficient statistics must include configurations that we call “diamonds”: two players j and k both indirectly observe a common predecessor i but not each other, while some fourth player l directly observes both j and k .

Definition 1 *The quadruple of distinct players (i, j, k, l) in the observation structure \mathcal{N} forms a diamond if $i \in ID(j) \cap ID(k)$, $j \notin ID(k)$, $k \notin ID(j)$, and $\{j, k\} \subset D(l)$.*

Figure 1 illustrates a diamond, to which we refer by the ordered quadruple (i, j, k, l) —where $i < j < k < l$ —and say that the observation structure contains a diamond if it includes four players who form a diamond. The observation structures of the canonical models of Banerjee (1992) and Bikhchandani et al. (1992), where players move single-file after observing all predecessors, do not include diamonds because for any pair of distinct players i, j , either $i \in D(j) = ID(j)$ or $j \in D(i) = ID(i)$.



In this paper, we wish to abstract from difficulties that arise when players can partially but not fully infer their predecessors' signals. In the diamond of Figure 1, for instance, the final observer l cannot discern the correlation in j and k 's beliefs attributable to common observation of i . Rational inference therefore requires l to use her priors on the distribution of the different signals that i, j and k might receive. To avoid becoming mired in these complications, we concentrate on situations of “impartial inference” in which the full informational content of all signals that influence a player's beliefs can be extracted. Figure 2 contains a special kind of diamond that we call a “shield” and differs from Figure 1 by allowing Player l to back out all of her predecessors' signals because she also observes Player i .



Definition 2 *The quadruple of distinct players (i, j, k, l) in the observation structure \mathcal{N} forms a shield iff it is a diamond and $i \in D(l)$.*

In shields like that of Figure 2, players do not suffer the problem of being able to partially but not fully infer certain predecessors' signals from observed actions.

Definition 3 *Player k achieves impartial inference if for each $s^{k-1} \in \times_{j < k} [\underline{s}_j, \bar{s}_j]$ and each $s_k \in [\underline{s}_k, \bar{s}_k]$,*

$$\alpha_k = \arg \max_{\alpha} \mathbb{E} \left[-(\alpha - \omega)^2 \mid \cup_{j \in ID(k)} \{s_j\} \cup \{s_k\} \right].$$

Otherwise, Player k achieves partial inference.

A player who achieves impartial inference cannot improve her expected payoff by learning the signal of anyone whom she indirectly observes. In the classical binary-action-binary-signal herding model, making the natural and usual assumption that a player indifferent between the two actions follows her signal, prior to formation of a herd, each player can infer all of her predecessors' signals exactly; once the herd begins, however, players can infer nothing about herders' signals, so inference is partial. Like the example in Section 2, a typical setting where actions do not reveal posteriors is unlikely to involve impartial inference, for even the second mover cannot fully recover the first mover's signal from her action. Because we work in a rich-action space where each person's beliefs are fully revealed to all observers, partial inference in our setting stems entirely from inability to disentangle the signals that generate the constellation of observed beliefs. Nevertheless, we have already seen in Section 2 an example of unordered moves demonstrating that impartial inference is not necessary for the form of anti-imitative behavior studied in this paper.¹³

Impartial inference does not imply that a player can identify the signals of all those players whom she indirectly observes—but merely that she has gleaned sufficient information from those signals that any deficit does not lower her payoff. For instance, when each player observes only her immediate predecessor, she achieves impartial inference despite an inability to separate her immediate predecessor's signal from his own predecessors' signals. Proposition 6 in the Appendix demonstrates that if every diamond is a shield, then common knowledge of rationality implies that every player achieves impartial inference.

We now turn our attention to the behavioral rules that players use to achieve impartial inference. Player k 's strategy $a_k(a^{k-1}; s_k)$, maps her observations and private signal into actions. We begin with precise definitions of imitation and anti-imitation:

Definition 4 *Player k imitates Player j if $a_k(a_j, a_{-j}^k; s_k)$ is weakly increasing (but not constant) in a_j . Player k anti-imitates Player j if $a_k(a_j, a_{-j}^k; s_k)$ is weakly decreasing (but not constant) in a_j .*

¹³The variant of the classic single-file model where players' actions span the continuum (Lee (1993)) satisfies impartial inference.

Player k anti-imitates Player j if k 's action never moves in the same direction as j 's—holding everyone else's action fixed—and sometimes moves strictly in the opposite direction. (Imitation is just the opposite.) A player who anti-imitates a predecessor always forms beliefs that tilt against that predecessor's.

The first limit to imitation that we explore is the occurrence of anti-imitation. In rich-observation structures where every player achieves impartial inference, rational social learning includes anti-imitation if and only if the observation structure contains a shield. Roughly speaking, in settings where players observe some predecessors without observing all of their most recent ones, certain players become less confident in a state the more confident they observe certain of their predecessors becoming.

Proposition 1 *Assume that every player in the observation structure \mathcal{N} achieves impartial inference. Then \mathcal{N} contains a shield if and only if some player anti-imitates another.*

In Figure 2, Player l must anti-imitate (or subtract off in log-odds form) the action of Player i , which Player j and Player k have both incorporated into their actions, in order not to double-count the signal of Player i . Proposition 1 shows that any impartial-inference setting that contains a shield, regardless of the rest of the configuration, includes at least one player who anti-imitates another.

The logic of redundancy plays out as dramatically—or more dramatically—when people learn from *both* other people's actions and public information. In Figure 2, because the first player chooses $a_i = s_i$, and all succeeding players observe a_i , we can interpret s_i simply as public information that is exogenously revealed. Under this interpretation, Proposition 1 implies that Player l anti-imitates the public information.

Anti-imitation gives rise to counter-intuitive comparative statics across observation structures. Consider a rational tourist who one night after dinner notices many locals queueing outside of Café M. From this, she forms some initial beliefs about its quality. Back in her hotel, she scours the Internet for good restaurants and learns that the local newspaper recently awarded Café M its maximum three stars. Assuming the tourist believes that locals

read the local newspaper, and everyone has common preferences, does the tourist fancy Café M more before or after reading the positive review?

We address this question in a simple observation structure with only four players: Player i is the local paper, Players j and k locals, and Player l the tourist. The observation structures before and after the tourist reads the review are described by Figures 1 and 2, respectively. By reading the review, the tourist goes from being the last player in a diamond that is not a shield (Player l in Figure 1), to the last player in a shield (Player l in Figure 2).

Proposition 2 *Let \mathcal{N} be the observation structure depicted in Figure 1 and $\widehat{\mathcal{N}}$ that of Figure 2. Then $\widehat{a}_l(a_i, a_j, a_k; s_l) \leq a_l(a_j, a_k; s_l)$ iff $a_i \geq \mathbb{E}[s_i | s_i + s_j = a_j, s_i + s_k = a_k, s_l]$.*

Following a better review than expected given her observations and private information ($a_i \geq \mathbb{E}[s_i | s_i + s_j = a_j, s_i + s_k = a_k, s_l]$), the tourist updates her beliefs to form a more negative assessment of the café ($\widehat{a}_l \leq a_l$). The intuition behind the result closely matches that of Proposition 1: the more positive the newspaper’s review, the more negative the information conveyed by Players j and k ’s actions.

Even in settings where rational people do not anti-imitate, they still do not do very much imitation. For instance, in single-file settings where everyone observes all predecessors, each person imitates only her immediate predecessor. More generally, in an observation structure where all players achieve impartial inference and do not anti-imitate, no player imitates two predecessors who both observe an earlier, common predecessor.

Proposition 3 *Assume that every player in the observation structure \mathcal{N} achieves impartial inference. If no player anti-imitates any other player, then Player k imitates the distinct Players i and j only if they share no common indirect observation (i.e., $ID(i) \cap ID(j) = \emptyset$).*

Absent anti-imitation, a player imitates more than one predecessor only if those predecessors share no common information. But “sharing no common observation” excludes virtually all social learners over time in almost all settings—except at the beginning of learning, it is unlikely in any setting of interest to economists that two people observed by later movers

(are known to) share no information. Rational social learning does not predict very much imitation without anti-imitation *in any setting*.

Rational social learning also may lead some players to form beliefs on the opposite side of their priors than all their information. That is, a player may form beliefs contrary to both his private signal and all the actions he observes. A definition helps us to establish a surprising result to this effect.

Definition 5 *Player k 's action a_k is contrarian given signal s_k and history a^{k-1} iff $a_k \neq 0$ and $\text{sgn}(a_k) = -\text{sgn}(s_k) = -\text{sgn}(a_j)$ for every $j \in D(k)$.*

Player k 's action is contrarian whenever it indicates state ω most likely, whereas his private signal and the expressed beliefs of everyone he observes suggest state $1 - \omega$ more likely.

Proposition 4 *Assume that every player in the observation structure \mathcal{N} achieves impartial inference.*


1. *If some player's action is contrarian, then \mathcal{N} contains a shield.*
2. *If \mathcal{N} contains a shield and players' private signals are drawn from the density $f^{(\omega)}$ that is everywhere positive on $[\underline{s}, \bar{s}]$, then with positive probability some player's action is contrarian.*

Because contrarian play relies upon some player anti-imitating another, the first statement follows as a corollary to Proposition 1. An intuition for the second statement comes from the fact that Bayes' Rule and impartial inference imply that each player's action is a linear combination of the actions she observes as well as her private signal. Because the weights in this linear combination do not depend upon the realization of any signal or action, if Player k attaches a negative weight to Player j 's action, then as the magnitude of a_j becomes large—fixing all other actions—Player k 's beliefs, given a weak private signal, must eventually take on the sign opposite to that of a_j .¹⁴

¹⁴The previous draft of this paper described an environment where contrarian play occurs with probability arbitrarily close to one.

3.1 Examples

To illustrate our model, we now consider two simple yet natural variants. In the first, n players move “multi-file” in each *round*, each player observing all players moving in prior rounds but none of those in the current or future rounds. When $n \geq 2$, this observation structure includes shields and admits contrarian play. Figure 3 illustrates the first five movers in a double-file setting.



In Figure 3, the quartets $(1, 3, 4, 5)$ and $(2, 3, 4, 5)$ both constitute shields.

We denote by s_t^1, \dots, s_t^n and a_t^1, \dots, a_t^n the signals and actions, respectively, of the n players who move simultaneously in round t . To succinctly describe behavior in this model, let $A_t = \sum_{k=1}^n a_t^k$, the sum of round- t actions or *aggregate round- t action*, and $S_t = \sum_{k=1}^n s_t^k$, the sum of round- t signals or *aggregate round- t signal*.

Because for each Player k moving in round 1, $a_1^k = a_1^k$, we have $A_1 = S_1$, so for a player in round 2 with signal s_2 , $a_2 = s_2 + A_1$, in which case $A_2 = S_2 + nA_1$. A player in round 3 wishes to set $a_3 = s_3 + S_2 + S_1$. Because she observes only A_2 and A_1 and knows that $A_2 = S_2 + nA_1$ as well as that $A_1 = S_1$, she chooses $a_3 = s_3 + A_2 - nA_1 + A_1$, so that $A_3 = S_3 + nA_2 - n(n-1)A_1$. Players in round 3 anti-imitate those in round 1 because they know that *each* of the n players from round 2 whom they imitate uses all round-1 actions; in order to avoid counting round-1 actions n -fold, players in round 3 must subtract off $n-1$

times the round-1 aggregate action. In general,

$$A_t = S_t + n \sum_{i=1}^{t-1} (-1)^{i-1} (n-1)^{i-1} A_{t-i}.^{15}$$

Whenever $n \geq 2$, this formula attaches negative coefficients to approximately half of past actions. In contrast to the canonical single-file model, in the multi-file model there is a sense in which approximately half of social learning is anti-imitative.

When $n = 1$, $A_t = S_t + A_{t-1} = \sum_{\tau \leq t} S_\tau$: players only imitate their immediate predecessors and do not anti-imitate. When $n = 2$, players imitate and anti-imitate alternating rounds with coefficients of constant magnitude:

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

When $n = 3$, the coefficients on past actions grow exponentially:

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

leading to $A_1 = S_1$, $A_2 = S_2 + 3A_1$, and $A_3 = S_3 + 3A_2 - 6A_1$. Players in round 3 strongly anti-imitates those in round 1, while those in round 4 even more strongly imitates those in round 1. People's beliefs also move in counterintuitive ways. Consider the case where the three players in round 1 all choose $\alpha = 0.6$, each expressing 60% confidence that $\omega = 1$.¹⁶ If all round 2 players also choose $\alpha = 0.6$, then since $A_2 = S_2 + 3A_1 = A_1$, $S_2 = -2A_1 = -2S_1$, meaning that the evidence for $\omega = 0$ is twice as strong as that for $\omega = 1$, in log-odds terms. Hence, someone who observes her six predecessors all indicate 60% confidence that $\omega = 1$ rationally concludes that there is only a 25% chance that $\omega = 1$! In general, in *odd* periods, complete agreement by predecessors always leads players to contradictory beliefs.¹⁷

¹⁵Observational beliefs following round $t - 1$ are $\sum_{i=1}^{t-1} (-1)^{i-1} (n-1)^{i-1} A_{t-i}$.


¹⁶Alternatively, they choose $a = \ln\left(\frac{0.6}{0.4}\right)$. We partially switch our exposition away from log-odds ratios here to express beliefs in the most natural way.

¹⁷This cannot happen with two players per round, where a player who chooses α after seeing the two previous rounds choose α has signal $1 - \alpha$. With three players, the same pattern can emerge even if actions increase over rounds: by continuity, nothing qualitative would change when actions $(0.6, 0.6, 0.6)$ are followed by actions $(0.61, 0.61, 0.61)$. Hence, it is not the presence or absence of trends that matters but instead how trends compare to what they would be if later signals supported earlier signals.

This example also demonstrates that observation of the order of moves does not drive anti-imitation or contrarian play; as long as it is known that players move three at a time, a unique path of play gives rise to all six actions coinciding, which leads an observer of ordered moves to the same inference as one of unordered moves.

Whatever n , actions converge almost surely to the state, despite wild swings in how rational players interpret past behavior. Note, importantly, that the wild swings in interpretation typically do not make their way into actions: recent actions always receive positive weight and *typically* are more extreme than earlier actions. It is precisely when play does not converge fast enough that we observe contrarian play.

Our second example depicts environments where people cannot observe their most immediate predecessors. In financial markets, for example, a trader may not see a trade placed momentarily before her own. Now we examine the case where players move single-file, and each player observes all of her predecessors except for the most recent one. Figure 4 shows the first six such players.



The first shield in this observation structure consists of Players 1, 3, 4 and 6. From the figure, it can be deduced that $a_3 = a_1 + s_3$ and $a_4 = a_2 + a_1 + s_4$; consequently, Player 6 achieves impartial inference by choosing $a_6 = a_4 + a_3 - a_1 + s_6$. In general, Player k puts

weights $(0, 1, 1, 0, -1, -1, 0, 1, 1, \dots)$ on actions $(a_{k-1}, a_{k-2}, \dots)$, or

$$a_k = \sum_{i=0} (a_{k-2-6i} + a_{k-3-6i} - a_{k-5-6i} - a_{k-6-6i}) + s_k.$$

Roughly speaking, players ignore one-third of their predecessors, imitate one-third, and anti-imitate one-third.

Unlike in the multi-file example, unanimity does not produce contrarian play. In order for Player 6 to be contrarian, Player 1 must have stronger beliefs in favor of a state, say $\omega = 1$, than either Players 3 or 4. For example, if Player 1 assigns $\omega = 1$ probability $\frac{5}{6}$, while Players 3 and 4 assign it only probability $\frac{2}{3}$, then a Player 6 with a neutral or no signal assigns $\frac{5}{9}$ probability to $\omega = 0$, yielding contrarian play.¹⁸ Intuitively, seeing Players 3 and 4 independently revise their confidence in $\omega = 1$ down from the level Player 1's confidence provides strong evidence for Player 6 that $\omega = 0$.¹⁹

4 Redundancy Neglect

Section 3 has established that rational players anti-imitate in most settings and do very little imitation. What happens if people do not anti-imitate and imitate more more broadly than predicted by full rationality? In this section, we consider the implications of not anti-imitating and “redundacy neglect”.

Throughout this section, we assume that players have unbounded private beliefs described by densities that are uniformly bounded. We first make two assumptions about how the behavioral rules that describe people's actions (again, in log-odds form) depend upon their private beliefs (again, in log-odds form).

Definition 6 *Players use social-learning rules that are strictly and boundedly increasing in private signals if*

¹⁸The formula for a_6 gives log odds ratio $\ln \left(\frac{\Pr[\omega=1]}{\Pr[\omega=0]} \right) = \ln \left(\frac{2}{1} \cdot \frac{2}{1} \cdot \left(\frac{5}{1} \right)^{-1} \right) = \ln \left(\frac{4}{5} \right)$

¹⁹Notice that it does not matter to Player 6 that Player 4's action embodies Player 2's signal because that does not produce a diamond; if we modified the observation structure to eliminate Player 2 and gave Player 4 the signal $s_2 + s_4$, only Player 5's action would change.

1. (strictly increasing) for each Player t , and each $a^{t-1} \in \mathbb{R}^{t-1}$,

$$\hat{s}^t > s^t \Rightarrow a_t(a^{t-1}, \hat{s}_t) > a_t(a^{t-1}, s_t)$$

2. (boundedly increasing) there exists $K \in \mathbb{R}_{++}$ such that for each Player t , each $a^{t-1} \in \mathbb{R}^{t-1}$, and each $s^t, \hat{s}^t \in \mathbb{R}$,

$$|a_t(a^{t-1}, \hat{s}_t) - a_t(a^{t-1}, s_t)| \leq K |\hat{s}_t - s_t|$$

Strictly increasing social-learning rules exclude cases where people ignore their private signals. Boundedly increasing social-learning rules exclude cases where people put arbitrarily high weight on their private signals. Bayesian belief updating implies that $a_t(a^{t-1}, s_t) = g_t(a^{t-1}) + s_t$, for some function g_t , which satisfies boundedly increasing for $K = 1$. Boundedness is weak enough to encompass many forms of non-Bayesian belief updating, including cases where people overconfidently overweight their own signal.

Our main behavioral assumption is that people neglect the redundancy in their predecessors' actions: someone who observes many people 60% convinced in the effectiveness of a new hybrid seed will become more than 60% convinced herself.

Definition 7 *Players use social-learning rules that neglect redundancy if there exist an integer N and a constant $c > 0$ with the property that for each Player $t \geq N + 1$, each $a^{t-N-1} \in \mathbb{R}^{t-N-1}$, each $s_t \in \mathbb{R}$, and each $z' > z \geq 0$,*

$$a_t(a^{t-N-1}, \underbrace{z', z', \dots, z'}_{N \text{ times}}, s_t) - a_t(a^{t-N-1}, \underbrace{z, z, \dots, z}_{N \text{ times}}, s_t) \geq (1 + c)(z' - z).$$

Suppose that Player t 's two immediate predecessors choose actions based solely on their private signals ($a_{t-1} = s_{t-1}, a_{t-2} = s_{t-2}$); if both raised their action from some z to z' , then Player t , if observing both, would increase her action by $2(z' - z)$. Hence, social learners who satisfy redundancy neglect for $c = 1$ and $N = 2$ treat their two immediate predecessors as if their actions conveyed only private information. Since generally this will not be the case—and cannot be the case when these predecessors themselves satisfy redundancy neglect for $c =$

1 and $N = 2$ —redundancy neglect embodies the error of reading more than one conditionally independent piece of information into recent predecessors’ actions.²⁰ Redundancy neglect is a joint assumption about observation structure and imitation and encompasses all sorts of combinations of assumptions about whom people observe and whom they imitate.²¹ By virtue of assuming that players essentially imitate more than one person, redundancy neglect is the antithesis of the limited imitation we explore in Section 3.

When players’ actions coincide with their beliefs, society converges to certain beliefs that $\omega = 1$ whenever $\lim_{t \rightarrow \infty} a_t = +\infty$ and to certain beliefs that $\omega = 0$ whenever $\lim_{t \rightarrow \infty} a_t = -\infty$.

Proposition 5 *Suppose that players use social-learning rules that are strictly and boundedly increasing in private signals as well as neglect redundancy, and that no player anti-imitates any other. Then, with positive probability, society converges to the action that corresponds to certain beliefs in the wrong state.*

Redundancy neglect and the absence of anti-imitation do not merely prevent society from learning but instead cause it to mislearn, as in Section 2 where redundancy-neglecting doctors can converge to using a bad medicine with near certainty. Consider again redundancy neglect with $c = 1$ and $N = 2$. Loosely speaking, Players 3 and 4 double-count Players 1 and 2; Players 5 and 6 double-count Players 3 and 4 and therefore quadruple-count Players 1 and 2, given the assumption that they do not anti-imitate anyone; etc.

When players move single-file and observe all predecessors, BRTNI players (Eyster and Rabin (2010)), who interpret each predecessor’s action as her private signal, satisfy redun-

²⁰Redundancy neglect is an error in an environment where everybody does it. But someone whose predecessors do not themselves overuse actions in this way—if, for instance, predecessors are uninfluenced by anybody before them (either due to irrationality or lack of observations)—may rationally neglect redundancy. Our formal result about inefficient learning applies only to settings where everyone neglects redundancy, but we conjecture that mislearning occurs when people neglect redundancy in an average sense.

²¹It also allows people to under-infer from their predecessors, as in partially-cursed equilibrium (Eyster and Rabin (2005)), so long as they neglect redundancy: someone can treat all predecessors’ actions as half as informative as they are at the same time as she mistakenly imitates many predecessors instead of just one. The condition is, intuitively, that the sum total of influence from underweighting individuals and overcounting predecessors is greater than the influence of one person, correctly interpreted.

dancy neglect with $c = N - 1$ for each N and mislearn with positive probability. BRTNI players give the past 100 players 100 times the weight that they should get. Proposition 5 establishes that far milder over-imitation leads society astray.²² For instance, if everyone treats their predecessors' actions as embodying just two conditionally independent signals, instead of one, then society sometimes converges to complete confidence in the wrong state. One example of this is where two people move in every round, as in the first example of Section 3, and everyone imitates only the two people moving in the previous round. Proposition 5 demonstrates the fragility of efficient learning: even mild failure by people to understand redundancy of others' beliefs leads society astray. Moreover, it implies that many other forms of error do not eliminate the possibility of wrong herding. For instance, even people who massively overweight their private signals mislearn with positive probability. The effect of redundancy neglect trumps that of all other errors in the context of social learning.

5 Conclusion

Outside of the impartial-inference setting we formally analyze, the limits to imitation and anti-imitation take on more complicated patterns. We limit our analysis to rich-information settings in order to crisply articulate the observational conditions under which our form of anti-imitation and simple limit on imitation occur. But their existence does not depend upon details of our environment. (Indeed, the example in Section 2 demonstrates that limited imitation and anti-imitation play out in settings without impartial inference.) Many simple, natural observational structures lead rational players to anti-imitate by requiring them to subtract off sources of correlation in order to fully extract information from all observed actions. When observed recent actors provide independent information, they should all be imitated. But when those recent players themselves imitate earlier actions, those earlier

²²Nevertheless, when players move single-file and observe only their immediate predecessors, BRTNI players do not neglect redundancy and have correct long-run beliefs. This discrepancy illustrates how redundancy neglect is a joint assumption on behavior and observation structure or, alternatively, a reduced form for various inferential errors in various contexts.

actions should be subtracted. The only circumstances when a rational person should imitate two or more predecessors are when she knows that these predecessors share no common observation, or when she anti-imitates.

The principles we establish here are distinct from the issues researchers have traditionally emphasized in the observational-learning literature. Different papers reach very different conclusions about the prevalence of herding and asymptotic efficiency depending on assumptions about the environment—richness and range of signals, richness and range of actions, heterogeneity in tastes, congestion costs, small errors, observability of some versus all of the population, observability of order of the moves, etc. Acemoglu, Dahleh, Lobel and Ozdaglar (2011) characterize necessary and sufficient conditions in general observation structures, including those we consider here, that give rise to asymptotically complete social learning. Yet neither that paper, nor related ones by Banerjee and Fudenberg (2004) and Smith and Sørensen (2008) that share the feature of Acemoglu et al. (2011) that players only observe subsets of their predecessors, focus on how players’ learning manifests itself in behavior. Not focusing on behavior short of the limit, these authors do not identify conditions determining whether rationality leads to anti-imitation.

A concern sometimes expressed about the relevance of observational learning is that in many settings information can be conveyed by other means—such as communication or even direct access to predecessors’ signals—to such a degree that people learn little from observing others’ actions. If people simply communicate with each other rather than take observable actions, is the entire literature rendered irrelevant? In fact, one reason to study the “rich-action” model is that it encompasses the case of communication. If everybody reports truthfully their beliefs about the state of the world, then the rich-action case is perhaps the most relevant, and, of course, our results hold even here.

Callender and Hörner (2009) illustrate a form of anti-imitation different from ours, related to the famous *overturning principle* of Smith and Sørensen (2000). The overturning principle states that in a single-file model where each player observes all of her predecessors, any player un-endowed with private information optimally imitates her immediate predeces-

Callender and Hörner (2009) analyze behavior in a binary-action model where no player observes the order of her predecessors' moves—yet observes all such moves—and each player is either perfectly uninformed or perfectly informed about the state. The overturning principle implies that any uninformed player who knew the identity of her immediate predecessor would imitate that player. Callender and Hörner (2009) show, when the order of previous moves is not observed, that the over-turning principle can lead to a dramatic form of anti-imitation. After certain histories, uninformed players optimally follow the *minority* of previous actions. When some people are much better informed than others, the most likely interpretation of seeing (say) four people sitting in Restaurant A and only one in Restaurant B is that the loner is a well-informed local bucking the trend rather than an ignorant tourist. That is, when the order of play is unobserved, it can be inferred that the minority choice is the most recent.

The forms of anti-imitative and contrarian behavior studied in this paper follow exclusively from players' need to subtract off correlations when imitating more than one predecessor. Some of our these predictions, most notably contrarian play, cannot arise in single-file models like Callender and Hörner (2009). More generally, single-file models (whether players' observations are ordered or un-ordered) satisfy a monotonicity property violated in the presence of shields. In any single-file, un-ordered model like that of Callender and Hörner (2009), if the lowest action in a first history is strictly higher than the highest action in a second history, then observed beliefs after the first history must exceed those after the second. This follows almost as a corollary of Smith and Sørensen's (2000) overturning principle: a player un-endowed with private information wishes to mimic her immediate predecessor. This intrinsically means that the inferred information in a single-file environment must lie within the range of beliefs consistent with at least one of the observed actions. Plainly, contrarian play in our model violates such monotonicity.

Just as there might be other reasons for anti-imitation, so too there may be other reasons for imitation. People may imitate to conform (e.g., Bernheim (1994)), to build reputations (e.g., Scharfstein and Stein (1990)), or to benefit from safety in numbers (penguins). We

think one of the uses of the findings in this paper is to help discern with greater power than has been done so far whether observed imitation derives from rational inference, irrational inference, or some other cause. For instance, Cai, Chen and Fang (2009) conduct a field experiment to distinguish imitative social learning from salience. In it, they inform diners at a chain of Beijing restaurants of the most popular dishes from the previous week or of “featured dishes”. By comparing the same dish as popular or featured, Cai, Chen and Fang (2009) show that diners react more strongly to popularity than to salience. Although this convincingly establishes that diners imitate, it does not uncover whether such imitation is rational or irrational. We have seen that full rationality and redundancy neglect lead to very different conclusions about the long-run efficiency of social learning. Rather than test whether people imitate their predecessors, empirical researchers could design more powerful tests to separate rational from irrational observational learning by evaluating whether people imitate their predecessors too much.

References

- Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar,** “Bayesian Learning in Social Networks,” *Review of Economic Studies*, 2011, 78 (4), 1201–1236.
- Acton, William,** “Observations on Venereal Diseases in the United Kingdom,” Statistical Reports in the Army, Navy and Merchant Service, with Remarks on the Mortality from Syphilis in the Metropolis 1846.
- Avery, Christopher and Peter Zemsky,** “Multidimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review*, 1998, 88 (4), 724–748.
- Bala, Venkatesh and Sanjeev Goyal,** “Learning from Neighbours,” *Review of Economic Studies*, 1998, 65 (3), 595–621.

- Banerjee, Abhijit**, “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.
- and **Drew Fudenberg**, “Word-of-Mouth Learning,” *Games and Economic Behavior*, 2004, *46* (1), 1–22.
- Bernheim, B Douglas**, “A Theory of Conformity,” *Journal of Political Economy*, 1994, *102* (5), 841–877.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch**, “A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades,” *Journal of Political Economy*, 1992, *100* (5), 992–1026.
- Cai, Hongbin, Yuyu Chen, and Hanming Fang**, “Observational Learning: Evidence from a Randomized Natural Field Experiment,” *American Economic Review*, 2009, *99* (3), 864–882.
- Callender, Steven and Johannes Hörner**, “The Wisdom of the Minority,” *Journal of Economic Theory*, 2009, *144* (4), 1421–1439.
- Colles, Abraham**, *Selections from the Works of Abraham Colles: Consisting Chiefly of His Practical Observations on the Venereal Disease, and on the Use of Mercury*, New Sydenham Society, 1881.
- Devergie, Marie Nicolas**, *Clinique de la Maladie Syphilitique*, Paris, 1831.
- Eyster, Erik and Matthew Rabin**, “Naïve Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2010, *2* (4), 221–243.
- Ferguson, William**, “Observations on the Venereal Disease in Portugal, as affecting the Constitutions of the British Soldiery and Natives,” *Medico-Chirurgical Transactions (today Journal of the Royal Society of Medicine)*, 1813, *4*, 1–16.
- George, Charles**, “Mercury and the Kidney,” *Journal of Nephrology*, 2011, *24*, S126–S132.

- Goldwater, Leonard**, *Mercury: A History of Quicksilver*, York Press, 1972.
- Golub, Benjamin and Matthew O. Jackson**, “Naïve Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2010, 2 (1), 112–149.
- Hyman, Harold Thomas**, “A Critical Evaluation of the Results of Routine Conservative Treatment of Syphilis,” *Bulletin New York Academy of Medicine*, 1941, 16 (6), 467–482.
- Jackson, Matthew**, *Social and Economic Networks*, Princeton University Press, 2008.
- Lasagna, Louis**, *The V. D. Epidemic*, Temple University Press, 1975.
- Lee, In-Ho**, “On the Convergence of Informational Cascades,” *Journal of Economic Theory*, 1993, 61, 395–411.
- Martenstein, Hans**, “The Treatment of Syphilis: An Inquiry in Five Countries,” *British Medical Journal*, 1935, 1 (3880), 1038–1039.
- O’Shea, J.G.**, “Two Minutes with Venus, Two Years with Mercury: Mercury as an Antisyphilitic Chemotherapeutic Agent,” *Journal of the Royal Society of Medicine*, 1990, 83 (6), 392–395.
- Parascandola, John**, *Sex, Sin and Science: A History of Syphilis in America*, Praeger, 2008.
- Scharfstein, David and Jeremy Stein**, “Herd Behavior and Investment,” *American Economic Review*, 1990, 80 (3), 465–479.
- Smith, Lones and Peter Sørensen**, “Pathological Outcomes of Observational Learning,” *Econometrica*, 2000, 68 (2), 371–398.
- and — , “Rational Social Learning by Random Sampling,” Mimeo 2008.

Tongue, James, “An Inaugural Dissertation Upon the Three Following Subjects: I. An Attempt to Prove, that the Lues Venerea was Not Introduced Into Europe from America; II. An Experimental Inquiry Into the Modus Operandi of Mercury in Curing the Lues Venerea; III. Experimental Proofs that the Fues Venerea and Gonorrhoea are Two Distinct Forms of Disease.” PhD dissertation, University of Pennsylvania 1801.

Wells, William Charles, “On the presence of the red matter and serum of blood in the urine of dropsy, which has not originated from scarlet fever,” *Transactions of the Society to Improve Medic-Chirurgical Knowledge*, 1812, 3, 194–240.

6 Appendix: Proofs

We begin by introducing some concepts and notation useful for the proofs. Define $\bar{D}(k) := \{j \in D(k) : \forall i \in D(k), j \notin ID(i)\}$, the set of players whom Player k indirectly observes only by directly observing. In the classical single-file model, for example, $D(1) = \bar{D}(1) = \emptyset$ and for each $k \geq 2$, $D(k) = \{1, \dots, k-1\}$ and $\bar{D}(k) = \{k-1\}$. When two players move every round, observing (only) all players who moved in all previous rounds, $D(1) = D(2) = \emptyset$, and for $l \geq 1$, $D(2l+1) = D(2l+2) = \{1, \dots, 2l\}$, while $\bar{D}(2l+1) = \bar{D}(2l+2) = \{2l-1, 2l\}$.

Lemma 1 states that any predecessor whom Player k indirectly observes she indirectly observes through someone in her only-observe-directly set.

Lemma 1 *For each Player k , $ID(k) = \bar{D}(k) \cup \left(\cup_{j \in \bar{D}(k)} ID(j)\right)$.*

Proof.

$$ID(k) = D(k) \cup \left(\cup_{j \in D(k)} ID(j)\right) = \bar{D}(k) \cup \left(\cup_{j \in D(k)} ID(j)\right) = \bar{D}(k) \cup \left(\cup_{j \in \bar{D}(k)} ID(j)\right),$$

where the first equality follows from the definition of ID , the second by the definition of $\bar{D}(k)$, and the third once more by the definition of $\bar{D}(k)$ together with transitivity of the ID relation. ■

Lemma 2 *If every player in the observation structure \mathcal{N} achieves impartial inference, then for each player l , there exist unique coefficients γ_i^l such that $a_l = \sum_{i \in D(l)} \gamma_i^l a_i + s_l$.*

Proof. Let $\{\gamma_i^l\}$ and $\{\hat{\gamma}_i^l\}$ be two such sets of coefficients. Towards a contradiction, suppose that these differ, and define $\hat{i} = \max\{i : \gamma_i^l \neq \hat{\gamma}_i^l\}$. Then

$$\frac{da_l}{ds_{\hat{i}}} = \frac{\partial a_l}{\partial a_{\hat{i}}} \frac{da_{\hat{i}}}{ds_{\hat{i}}} + \sum_{i \in D(l), i > \hat{i}} \frac{\partial a_l}{\partial a_i} \frac{da_i}{ds_{\hat{i}}} = \gamma_{\hat{i}}^l \cdot 1 + \sum_{i \in D(l), i > \hat{i}} \gamma_i^l \cdot 1 = \hat{\gamma}_{\hat{i}}^l \cdot 1 + \sum_{i \in D(l), i > \hat{i}} \hat{\gamma}_i^l \cdot 1.$$

Because γ_i^l and $\hat{\gamma}_i^l$ coincide for any $i > \hat{i}$, by definition, $\gamma_{\hat{i}}^l = \hat{\gamma}_{\hat{i}}^l$, a contradiction. ■

Lemma 3 implies that in any observation structure without diamonds where all players achieve impartial inference, all behavior is imitative.²³

Lemma 3 *Suppose every player in the observation structure \mathcal{N} achieves impartial inference. If for some player l , $\forall j, k \in \bar{D}(l)$, $ID(j) \cap ID(k) = \emptyset$, then $a_l = \sum_{k \in \bar{D}(l)} a_k + s_l$.*

Proof of Lemma. Write $a_l = \sum_{k \in \bar{D}(l)} a_k + s_l =: \sum_{i \in ID(l)} \beta_i s_i + s_l$. Lemma 1 implies that $\beta_i \geq 1$ for each $i \in ID(l)$. The assumption that $\forall j, k \in \bar{D}(l)$, $ID(j) \cap ID(k) = \emptyset$ implies that $\beta_i \leq 1$. Hence, the formula for a_l gives impartial inference. By Lemma 2, it is unique.

■

Lemma 4 *If $j \in \bar{D}(l)$ and all players achieve impartial inference, then $\frac{\partial a_l}{\partial a_j} = 1$.*

Proof of Lemma. If $\frac{\partial a_l}{\partial a_j} \neq 1$, then since $j \in \bar{D}(l)$, $\frac{da_l}{ds_j} = \frac{\partial a_l}{\partial a_j} \frac{\partial a_j}{\partial s_j} = \frac{\partial a_l}{\partial a_j} \neq 1$, in contradiction to Player l achieving impartial inference. ■

Proof of Proposition 1. We begin by proving the “if” direction. Let l be the first player to anti-imitate in an observation structure with impartial inference. Let i be a player whom l anti-imitates; plainly, $i \in D(l)$. Define $B := \{i' : (i', j, k, l) \text{ is a diamond, } j, k \in \bar{D}(l)\}$, which is non-empty by Lemma 3. If $i \in B$, then we have a shield, so suppose otherwise. Lemma 1 implies that there exists $h \in \bar{D}(l)$ for whom $i \in ID(h)$; if $h' \in \bar{D}(l)$ such that $i \in ID(h')$, then if $h' \neq h$, (i, h, h', l) is a shield, so we assume that h is unique. Since l anti-imitates

²³In fact, Proposition 6 below implies that in any observation structure without diamonds, all players do achieve impartial inference.

i , there must exist some $k \in D(l)$, $k \neq h$ such that $i \in ID(k)$ and l imitates k ; otherwise, using Lemma 4 and the assumption that h achieves impartial inference,

$$\frac{da_l}{ds_i} = \frac{\partial a_l}{\partial a_i} \frac{da_i}{ds_i} + \frac{\partial a_l}{\partial a_h} \frac{da_h}{ds_i} < \frac{\partial a_l}{\partial a_h} \frac{da_h}{ds_i} = 1,$$

in contradiction to l achieving impartial inference. Notice that for $k \neq h, l$, if $i \in ID(k)$, $k \in D(l)$, then if $k \notin ID(h)$, (i, k, h, l) is a shield, so we suppose that $k \in ID(h)$. Note that

$$\frac{da_l}{ds_k} \geq \frac{\partial a_l}{\partial a_k} \frac{da_k}{ds_k} + \frac{\partial a_l}{\partial a_h} \frac{da_h}{ds_k} = \frac{\partial a_l}{\partial a_k} \cdot 1 + 1 \cdot 1 > 1$$

because l is the first player to anti-imitate, and using Lemma 4 and impartial inference for h . But $\frac{da_l}{ds_k} > 1$ contradicts l achieving impartial inference. Hence, \mathcal{N} must contain a shield.

To prove the “only if” direction, let $\hat{l} = \min\{l : (i, j, k, l) \text{ is a diamond}\}$, which exists because every shield is a diamond. We claim that we can take $j, k \in \overline{D}(\hat{l})$. To see why, define $\hat{j} = \max\{j\} \cup \{h \in \overline{D}(\hat{l}) : j \in D(h)\}$ and $\hat{k} = \max\{k\} \cup \{h \in \overline{D}(\hat{l}) : k \in D(h)\}$ and note that $(i, \hat{j}, \hat{k}, \hat{l})$ is a diamond since otherwise, if $\hat{j} = \hat{k}$, then (i, j, k, \hat{j}) is a diamond, in contradiction to the definition of \hat{l} . Since $i \in ID(\hat{l})$, impartial inference implies that

$$\frac{da_{\hat{l}}}{ds_i} = \sum_{l \in D(\hat{l}), l \neq i} \frac{\partial a_{\hat{l}}}{\partial a_l} \frac{da_l}{ds_i} + \frac{\partial a_{\hat{l}}}{\partial a_i} \frac{da_i}{ds_i} = 1$$

Assume $\frac{\partial a_{\hat{l}}}{\partial a_l} \geq 0 \forall l \in D(\hat{l}), l \neq i$, otherwise we’re done. Since $j, k \in \overline{D}(\hat{l})$, $\frac{\partial a_{\hat{l}}}{\partial a_j} = \frac{\partial a_{\hat{l}}}{\partial a_k} = 1$ by Lemma 4. Because $i \in ID(j) \cap ID(k)$, impartial inference implies that $\frac{da_j}{ds_i} = \frac{da_k}{ds_i} = 1$. Using $\frac{da_l}{ds_i} \in \{0, 1\}$ from impartial inference as well as $\frac{da_i}{ds_i} = 1$ gives

$$\frac{da_{\hat{l}}}{ds_i} \geq 2 + \frac{\partial a_{\hat{l}}}{\partial a_i} = 1,$$

from whence $\frac{\partial a_{\hat{l}}}{\partial a_i} \leq -1$, as desired. ■

Proof of Proposition 2. We have

$$\begin{aligned} a_l(a_j, a_k; s_l) &= a_j + a_k - \mathbb{E}[a_i | a_j, a_k, s_l] + s_l \geq \hat{a}_l(a_i, a_j, a_k; s_l) = a_j + a_k - a_i + s_l \\ \iff (s_i + s_j) + (s_i + s_k) - \mathbb{E}[s_i | a_j = s_i + s_j, a_k = s_i + s_k, s_l] + s_l &\geq s_1 + s_j + s_k + s_l \\ \iff a_i = s_i \geq E[s_i | s_i + s_j, s_i + s_k, s_l]. \end{aligned}$$

■

Proof of Proposition 3. We first claim that $\forall j, k \in \mathcal{N}$, $\frac{\partial a_k}{\partial a_j} \in \mathbb{Z}$. Towards a contradiction, suppose that for some j, k , $\frac{\partial a_k}{\partial a_j} \notin \mathbb{Z}$; wlog take k to be the first such k and $j(k) = \max \left\{ j : \frac{\partial a_k}{\partial a_j} \notin \mathbb{Z} \right\}$ and notice that

$$\frac{da_k}{ds_{j(k)}} = \frac{\partial a_k}{\partial a_{j(k)}} \frac{da_{j(k)}}{ds_{j(k)}} + \sum_{j \in D(k), j > j(k)} \frac{\partial a_k}{\partial a_j} \frac{da_j}{ds_{j(k)}} = \frac{\partial a_k}{\partial a_{j(k)}} + \sum_{j \in D(k), j > j(k)} \frac{\partial a_k}{\partial a_j} \frac{da_j}{ds_{j(k)}} \notin \mathbb{Z},$$

which contradicts Player k achieving impartial inference. Now suppose Player l imitates some Players i and $j \neq i$ s.t. $h \in ID(i) \cap ID(j) \subset ID(l)$. Then

$$\begin{aligned} \frac{da_l}{ds_h} &= \frac{\partial a_l}{\partial a_i} \frac{da_i}{ds_h} + \frac{\partial a_l}{\partial a_j} \frac{da_j}{ds_h} + \sum_{k \in D(l), k \neq i, j} \frac{\partial a_l}{\partial a_k} \frac{da_k}{ds_h} \\ &\geq 2 + \sum_{k \in D(l), k \neq i, j} \frac{\partial a_l}{\partial a_k} \frac{da_k}{ds_h}, \end{aligned}$$

using the previous result as well as the assumption that Players i, j achieve impartial inference. Because all players achieve impartial inference, we must have that for some $k \in D(l)$, $\frac{\partial a_l}{\partial a_k} < 0$, which contradicts our assumption of no anti-imitation. ■

Proof of Proposition 4. Part 1: From Lemma 2, Player l 's action can be written uniquely as $a_l = \sum_{i \in D(l)} \gamma_i^l a_i + s_l$. If for each $i \in D(l)$, $\alpha_i^l \geq 0$, and $\text{sgn}(a_i) = \text{sgn}(s_l)$, then $\text{sgn}(a_l) = \text{sgn}(s_l)$, and therefore the path of play (a_1, \dots, a_l) cannot be contrarian. Hence, contrarian play requires that $\alpha_i^l < 0$ for some i, l . Proposition 1 then implies that \mathcal{N} contains a shield.

Part 2: Wlog let $\underline{s} \leq -\bar{s}$ and take $\epsilon > 0$ small. The existence of a shield implies the existence of a diamond; let $\hat{l} = \min\{l : (i, j, k, l) \text{ is a diamond}\}$, and let i be the last Player whom \hat{l} anti-imitates.

For each $j < \hat{l}$, $j \neq i$, let $a_j \in \left(0, \frac{\epsilon}{2(\hat{l}-1)}\right)$ and $a_i \in \left(\bar{s} - \epsilon, \bar{s} - \frac{\epsilon}{2}\right)$. Non-triviality of private beliefs permits this $\forall j$ such that $D(j) = \emptyset$; for other j , from Lemma 3,

$$a_j = \sum_{k \in \overline{D}(j)} a_k + s_j \leq \bar{s} - \frac{\epsilon}{2} + (\hat{l} - 2) \frac{\epsilon}{2(\hat{l} - 1)} + s_j < \bar{s} + s_j$$

Because $\underline{s} \leq -\bar{s}$, there exists a positive-measure set of signals s_j for which $a_j \in \left(0, \frac{\epsilon}{2(\hat{l}-1)}\right)$.

Now note that $a_i \leq (\hat{l} - 2)\frac{\epsilon}{2(\hat{l}-1)} - (\bar{s} - \frac{\epsilon}{2}) + s_i < \epsilon + s_i - \bar{s} < 0$ for $s_i < \bar{s}$ and ϵ small enough, which implies positive-probability contrarian play. ■

Proof of Proposition 5. Wlog assume $\omega = 0$. Choose c and N such that actions satisfy redundancy neglect and K to satisfy boundedness. Choose $\epsilon > 0$ such that $\epsilon \in (0, \bar{s}_k)$ for each $k = 1, \dots, 2N$. With positive probability $s_k \geq \epsilon$ for each $k = 1, \dots, N$. Choose $\varepsilon > 0$ such that $\varepsilon < a_k(0, \epsilon)$ for each $k = 1, \dots, N$, which exists by strict monotonicity in private signals. For each $k = 1, \dots, N$,

$$a_k = a_k(a^{k-1}, s_k) \geq a_k(0, s_k) \geq \varepsilon,$$

where the first inequality follows from the assumption of no anti-imitation. With positive probability, $s_k \geq \epsilon$ for each $k = N+1, \dots, 2N$, which implies that for each $k = N+1, \dots, 2N$,

$$a_k = a_k(a^{k-1}, s_k) \geq a_k(\underbrace{\varepsilon, \varepsilon, \dots, \varepsilon}_{N \text{ times}}, s_k) \geq a_k(\underbrace{\varepsilon, \varepsilon, \dots, \varepsilon}_{N \text{ times}}, \epsilon) \geq (1+c)\varepsilon,$$

where the first inequality follows from the assumption of no anti-imitation, the second from strict increasingness in private signals, and the third from redundancy neglect.

We claim that if $s_{iN+j} > -\frac{(i-1)c^2}{K}\varepsilon$ for $i = 2, 3, \dots$ and $j = 1, \dots, N-1$ that $a_{iN+j} \geq (1+ic)\varepsilon$. For $i = 2$, note

$$a_{2N+j} \geq (1+c)^2\varepsilon - K\frac{(2-1)c^2}{K}\varepsilon = (1+2c)\varepsilon,$$

where the first term after the inequality follows from redundancy neglect and no anti-imitation (given each of the last $j-1$ players chose $a > (1+2c)\varepsilon > (1+c)\varepsilon$), and each of the preceding $N-(j-1)$ players chose $a > (1+c)\varepsilon$) and the second term after the inequality follows from the assumption about signals and the assumption that social-learning rules are boundedly increasing. If the claim holds for $i = 2, \dots, k$, then

$$a_{(k+1)N+j} \geq (1+c)(1+kc)\varepsilon - K\frac{(k+1-1)c^2}{K}\varepsilon = (1+(k+1)c)\varepsilon,$$

as desired (where, again, the first term after the inequality follows from redundancy neglect combined with the result for $i = k$, and the second by the assumption about signals and boundedly increasing social-learning rules).

Thus, $\lim_{t \rightarrow \infty} a_t = +\infty$, namely actions converge to the one corresponding to full confidence in the wrong state. We can adapt the argument in the proof of Proposition 3 of Eyster and Rabin (2010) to prove that such signal paths have positive probability.²⁴ ■

In the main text, we assume that players achieve impartial inference. A sufficient condition for common knowledge of rationality to imply impartial inference in our model is that every diamond be a shield.

Proposition 6 *If every diamond in the observation structure \mathcal{N} is also a shield, then common knowledge of rationality implies that every player achieves impartial inference.*

Proof of Proposition 6. Clearly, Player 1 achieves impartial inference by choosing $a_1 = s_1$. We claim that if all Players $i \in \{1, \dots, k-1\}$ achieve impartial inference, then so too does Player k . Define

$$a_k^1 := \sum_{j \in \bar{D}(k)} a_j + s_k =: \sum_{j \in ID(k)} \beta_j^1 s_j + s_k.$$

Define $U^1(k) := \{j \in ID(k) : \beta_j^1 = 1\}$ and $M^1(k) := \{j \in ID(k) : \beta_j^1 > 1\}$. Players $i \in \{1, \dots, k-1\}$ achieving impartial inference and Lemma 1 imply that for each $j \in ID(k)$, $\beta_j^1 \geq 1$: $j \in M^1(k) \cup U^1(k)$. First, notice that $\forall i \in U^1(k), \forall j \in M^1(k), i \notin ID(j)$; otherwise, because $\forall j \in M^1(k), \exists k_1, k_2 \in \bar{D}(k)$ s.t. $j \in ID(k_1) \cap ID(k_2), i \in ID(j)$ implies $i \in ID(k_1) \cap ID(k_2)$ and therefore $i \in M^1(k)$, a contradiction.

Common knowledge of rationality implies that Player k knows the β_j^1 's. If $M^1(k) = \emptyset$, alternatively $U^1(k) = ID(k)$, then Player k achieves impartial inference through a_k^1 . Suppose that $\emptyset \neq M^1(k) =: \{m_1, m_2, \dots, m_N\}$, where $m_1 < m_2 < \dots < m_N$. Because each $m_i, i \in$

²⁴That paper showed that the event $s_t > -kt \forall t$ occurs with positive probability for any constant $k > 0$, which is our task here for the special case where $N = 1$. For $N > 1$, the result follows as a corollary: we know that for a fixed $j = 1, \dots, N$, the event that $s_{tN+j} > -kt \forall t$ has positive probability, call it P_j ; since signals are conditionally i.i.d., this implies that $s_{tN+j} > -kt \forall j, t$ has probability $\prod_{j=1}^N P_j > 0$.

Although it in no way affects the argument of that proof nor its corollary here, we note for completeness that the last displayed line of the proof of Proposition 3 in Eyster and Rabin (2010) contains a typo through omission of a squared sign on the π term, as $\sum_{t \geq 1} \frac{1}{t^2} = \frac{\pi^2}{6}$ rather than the $\frac{\pi}{6}$ implied by that proof.

$\{1, \dots, N\}$, belongs to a diamond (m_i, g, h, k) for $g, h \in \overline{D}(k)$, and that diamond is also a shield by assumption, $m_i \in D(k)$. Define

$$a_k^2 := a_k^1 - (\beta_{m_N}^1 - 1) a_{m_N} =: \sum_{j \in ID(k)} \beta_j^2 s_j + s_k,$$

$U^2(k) := \{j \in ID(k) : \beta_j^2 = 1\}$ and $M^2(k) := \{j \in ID(k) : \beta_j^2 \neq 1\}$. By construction, for each $l \in ID(k), l \geq m_N, \beta_l^2 = 1$. From above, $\forall i \in U^1(k), i \notin ID(m_N), \frac{da_k^1}{ds_i} = \frac{da_k^2}{ds_i} = 1$ so $U^1(k) \subseteq U^2(k)$. Since $m_N \in (U^1(k))^c \cap U^2(k)$ by construction, $U^1(k) \subsetneq U^2(k)$ and, hence, $M^2(k) \subsetneq M^1(k)$.

Define

$$a_k^3 := a_k^2 - (\beta_{m_{N-1}}^2 - 1) a_{m_{N-1}} =: \sum_{j \in ID(k)} \beta_j^3 s_j + s_k,$$

$U^3(k) := \{j \in ID(k) : \beta_j^3 = 1\}$ and $M^3(k) := \{j \in ID(k) : \beta_j^3 \neq 1\}$. By the same argument as before, $U^2(k) \subsetneq U^3(k)$ and, hence, $M^3(k) \subsetneq M^2(k)$. Iterating produces two strictly nested sequences of sets, the $(M^j(k))_j$ decreasing and $(U^j(k))_j$ increasing. In at most $k - 1$ steps, this terminates with $(U^{\hat{j}}(k)) = ID(k)$, allowing Player k to achieve impartial inference by playing $a_k^{\hat{j}}$. ■