

Rational Observational Learning

Erik Eyster and Matthew Rabin*

February 19, 2011

Abstract

An extensive literature identifies how privately-informed rational people who observe the behavior of other privately-informed rational people with similar tastes may come to imitate those people, emphasizing when and how such imitation leads to inefficiency. This paper investigates not the efficiency but instead the *behavior* of fully rational observational learners. In virtually any setting apart from that most commonly studied in the literature, rational observational learners imitate only some of their predecessors and, in fact, frequently contradict both their private information and the prevailing beliefs that they observe. In settings that allow players to extract all relevant information about others' private signals from their actions, we identify necessary and sufficient conditions for rational observational learning to include “anti-imitative” behavior, where, fixing other observed actions, a person regards a state of the world as less likely the more a predecessor’s action indicates belief in that state.

*Eyster: Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE United Kingdom, (email: e.eyster@lse.ac.uk); Rabin: Department of Economics, University of California, Berkeley, 549 Evans Hall #3880, Berkeley, CA 94720-3880 USA, (email: rabin@econ.berkeley.edu). The final examples of Section 2 originally appeared as a component of the paper entitled “Rational and Naïve Herding”, CEPR Discussion Paper DP7351. For valuable research assistance on this and the earlier paper, we thank Asaf Plan, Zack Grossman, and Xiaoyu Xia. We thank seminar participants at Berkeley, Columbia, Harvard, LSE, NYU, Pompeu Fabra, Sabanci, Toulouse, UBC and Yale for helpful comments. Rabin thanks the National Science Foundation (grants SES-0518758 and SES-0648659) for financial support.

Such anti-imitation follows from players’ need to subtract off sources of correlation to interpret information (here, other players’ actions) correctly, and is mandated by rationality in settings where players can observe many predecessors’ actions but cannot observe all recent or concurrent actions. Moreover, in these settings, there is always a positive probability that some player plays contrary to both her private information and the beliefs of *every* single person whose action she observes. We illustrate a setting where a society of fully rational players nearly always converges to the truth via at least one such episode of extreme contrarian behavior. (JEL B49)

Keywords: social networks, observational learning, rationality

1 Introduction

An extensive literature—beginning with Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992)—identifies how a rational person who learns by observing the behavior of others with similar tastes and private information may be inclined to imitate those parties’ behavior, even in contradiction to her own private information. Yet the literature’s special informational and observational structure combined with its focus on information aggregation have obscured the precarious connection between rational play and imitation. Some recent papers have illustrated departures from the prediction that rational learning leads simply to imitation: in many natural settings, rational players imitate some previous players, but “anti-imitate” others. These papers derive anti-imitation from players’ imperfect ability to extract others’ beliefs from their actions. This paper complements those by investigating a conceptually distinct reason for rational anti-imitation in a class of rich-information settings where each player *is* perfectly able to extract the beliefs of any player whose action she observes. We identify necessary and sufficient conditions for observational learning to involve some instances of “anti-imitation”—where, fixing others’ actions, some player revealing a greater belief in a hypothesis causes a later player to believe less in it. In these same conditions, there is a positive probability of contrarian behavior, where at least one player contradicts both her private information and the revealed beliefs of every single person she

observes. These conditions hold in most natural settings outside of the single-file, full-observation structure previously emphasized in the literature. We also illustrate related settings where rational herds almost surely involve at least one episode of such contrarian behavior.

In the canonical illustrative example of the literature, a sequence of people choose in turn one of two options, A or B , each observing all predecessors' choices. Players receive conditionally independent and equally strong private binary signals about which option is better. The rationality of imitation is easy to see in this setting: early movers essentially reveal their own signals, and should be imitated. Once the pattern of signals leads to, say, two more choices of A than B , subsequent rational agents will imitate the majority of their predecessors rather than follow their own signals because they infer more signals favoring A than B . Yet this canonical binary setting obscures that such "single-file" rational herding does not predict global imitation. It predicts something far more specific: it is ubiquitously rational only to imitate the single most recent action, which combines new private information with all the information contained in prior actions. To a first approximation, prior actions should be ignored. The symmetric binary-signal, binary-action model obscures this prediction because the most recent player's action is never in the minority. The ordered sequence AAB can never occur, for instance, since Player 3 would ignore her signal following AA . However, in any common-preference situation with a private-information structure that allows AAB to occur with positive probability, Player 4 will interpret it to indicate that B is the better option.

Not only should players be more influenced by the most recent actions than prior ones, but for many signal structures prior actions in fact should count *negatively*, so that Player 4 believe more strongly in B 's optimality given the observation AAB than BBB . When only a small random proportion of players are informed, for instance, Player 4 believes the probability that B is the better option given BBB is little more than 50%, since the first person was probably just guessing, and followers probably just imitating. Following AAB , following AAB , on the other hand, there is a much stronger reason to believe B to be the

better option, since only an informed person would overturn a herd. It can be shown, in fact, that when signals are weak and very few people are informed, over 63% of eventual herds involve a very extreme form of anti-majority play—at least one uninformed player will follow the most recent person’s action, despite it contradicting *all* prior actions. Another very natural class of environments where systematic imitation may seem more likely—and recency effects obviously impossible—is when people observe previous actions without seeing their order. Yet Callender and Hörner (2009) show that with similar heterogeneity in the quality of people’s private signals, rational inference quite readily can lead people to follow the *minority* of previous actions. With some people 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.¹

The logic underlying all of these examples of anti-imitation relates to the “coarseness” of available actions. When a player’s action does not perfectly reveal his beliefs, earlier actions shed additional light on his beliefs by providing clues as to the strength of the signal he needed to take his action. Yet a second, conceptually distinct form of anti-imitative behavior highlighted in Eyster and Rabin (2009) can occur even in much richer informational environments. Consider a simple alternative to the single-file herding models that pervade the literature. Suppose $n > 1$ people move simultaneously every period, each getting independent private information and observing all previous continuous actions that fully reveal people’s beliefs. Fixing behavior in period 2, the more confidence period-1 actions indicate in favor of a hypothesis, the *less* confidence period-3 actors will have in it. The logic is simple: 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

¹The models of Smith and Sørensen (2008), Banerjee and Fudenberg (2004), and Acemoglu, Dahleh, Lobel and Ozdaglar (2010) all encompass settings where players observe a random subset of predecessors. Although not the subject of their work, in these models rational social learning also leads to anti-imitation for the same reason as in Callender and Hörner (2009): players can only partially infer their observed predecessors’ beliefs from these predecessors’ actions.

information n -fold, period-3 players must imitate period-2 actions but *subtract off* period-1 actions. In turn, period-4 players will imitate period-3 players, anti-imitate period-2 players, and imitate period-1 players. Indeed, *every single* player in the infinite sequence outside periods 1 and 2 will anti-imitate almost half her predecessors. Moreover, this anti-imitation can take a dramatic form: 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!

In Section 2 we model general observation structures that allow us to flesh out this logic more generally within the class of situations we call “impartial inference”. We say that a situation is one of *impartial inference* whenever common knowledge of rationality implies that any player who learns something from a previous set of players’ signals in fact learns everything that she would wish to know from those signals. (The “impartial” here means not partial—either information is fully extracted, or not at all.) This immediately rules out “coarse” actions, so that we focus solely on the case where actions fully reveal beliefs. Our first proposition provides necessary and sufficient conditions on the observation structure for players to achieve impartial inference. We define Player k to “indirectly observe” Player j if Player k observes some player who observes some player who \dots observes Player j . Roughly speaking, then, a rich-action setting generates impartial inference if and only if whenever a Player l indirectly observes the actions of Player j and k —neither of whom indirectly observes the other and both of whom indirectly observe Player i —then Player l also observes Player i .²

²The statement is only rough because it suffices for Player l to indirectly observe some Player m who satisfies the above desiderata or for Player l to indirectly observe some Player m who in turn indirectly observes i and satisfies the statement expressed in the text for Player i . The canonical single-file herding models, the “multi-file” model from above, and for instance a single-file model where each player observes the actions and order of only the players before her, are all games of impartial inference. Multi-file models

Focusing on games of impartial inference allows for surprisingly simple necessary and sufficient conditions for anti-imitation. Essentially, anti-imitation occurs in an observational environment if and only if it contains a foursome of players i, j, k, l where 1) j and k both indirectly observe i , 2) neither j nor k indirectly observes the other, and 3) l indirectly observes both j and k and observes i .³ Intuitively, as in the n -file herding example above, Player l must weight Players j and k positively to extract their signals, but then must weight Player i 's action negatively because both j and k have weighted it themselves already. A more striking conclusion emerges in these settings when signals are rich: there is a positive probability of a sequence of signals such that at least one player will have 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 l must infer that j and k both received negative information so that altogether the hypothesis is unlikely.⁴

While in most natural settings such a strong form of contrarian behavior is merely a possibility, in Section 3 we illustrate a setting where it happens with near certainty. To keep within the framework of impartial inference, we use the following contrived set-up: in each round, an identifiable player receives no signal, while four others each receive conditionally independent and identically distributed binary signals of which of the two states obtain; each player observes only the actions (that fully reveal beliefs) of the five players in the previous round. Despite only observing the previous round's actions, all players in each round t can

where each player gets a private signal but players only observe the previous round of actions (and not the full history) are not games of impartial inference.

³When players have unbounded private beliefs, then this condition is sufficient for some player to anti-imitate another, though not necessarily Player i, j, k nor l . Proposition 2 both weakens the sufficient condition along the lines of the previous footnote and provides a necessary condition for anti-imitation that does not include unbounded private beliefs.

⁴Such "contrarian" behavior—believing the opposite of all your observed predecessors—cannot occur in single-file models with partial inference like that of Callender and Hörner (2009), where anti-imitation derives from players' using the overall distribution of actions to refine their interpretation of individual actions. Clearly, if all players have a coarse belief favoring A over B, then no inference by any observer about the identity of the most recent mover could lead him to believe B more likely.

infer precisely how many of each signal have occurred through round $t - 1$: the no-signal person in round $t - 1$ reveals the total information through round $t - 2$, while the four other movers in round $t - 1$ reveal their signals through the differences in their beliefs from those of the no-signal person. In this case, a round- t player observing the no-signal person in round $t - 1$ revealing beliefs equivalent to two signals favoring (say) option B, but all four signalled people in round $t - 1$ revealing beliefs of only a single signal in favor of B would know each of them received an A signal, making the total number of signals through round $t - 1$ two in favor of A. In this case, each player in round t —even one holding a B signal—believes A more likely than B, despite having only seen predecessors who believed B more likely than A. We prove that in the limit as signals become very weak the probability that such an episode occurs at least once approaches certainty. In fact, when signals tend to their un-informative limit, it will happen arbitrarily many times.

The class of formal models we examine in this paper is clearly quite stylized. But the forms of anti-imitative and contrarian play that we identify do not depend upon details of our environment such as the richness of signal or action spaces. Many simple, natural observational structures would lead players to rationally anti-imitate because they require players to subtract sources of correlation in order to rationally extract information from different actions. If observed recent actions provide some independent information, then they should all be imitated. But if all those recent players are themselves imitating earlier actions, those earlier actions should be subtracted.

We speculate that the strong forms of anti-imitation and contrarian play predicted by the full-rationality model will not be common in practice. Whether this speculation turns out to be right or wrong, this paper provides an abundance of guidance that can be used to help determine whether approximate rationality truly governs observational learning. Settings like ours provide much more powerful tests of whether and why people imitate than current experiments, typically set in the very rare setting where rationality and virtually all other theories predict imitation.

The realism and abundance of settings where full rationality predicts anti-imitation is

important for a second reason: models of departures from rational play that robustly predict imitation give rise to very different informational and efficiency properties than models of rational social learning. Eyster and Rabin (2010) propose a simple alternative model of naïve inference where players, by dint of under-appreciating the fact that their predecessors are also making informational inferences from behavior, take all actors' behavior as face-value indications of their private information. This naïve inference directly leads to universal imitation. And its inference leads far more robustly to a much stronger and far less efficient form of herding than predicted by any rational model. We conclude the paper in Section 4 with a brief discussion of the relationship between imitation and inefficiency, speculating that in fact many theories of bounded rationality that do lead to long-run efficiency and prevent the form of overconfident mislearning that Eyster and Rabin (2010) predicts can happen will also involve anti-imitation! As such, if anti-imitative behavior should turn out to be rare in practice, then we will learn not just the limits of rationality in observational-learning settings but also that societal beliefs might frequently converge to highly confident but incorrect beliefs.

2 Impartial Inference and Anti-Imitative Behavior

In this section, we consider 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. We focus on environments where rational players completely extract all of the payoff-relevant information to which they have access. 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.

There are two possible states of the world, $\omega \in \{0, 1\}$, with each state *ex ante* equally likely. Each Player k in the set of players $\{1, 2, \dots\}$ receives a private signal σ_k whose density conditional upon the state being ω is $f(\sigma_k|\omega)$. We assume that for each $\omega \in \{0, 1\}$, $f(\sigma_k|\omega)$ is everywhere positive and continuous on the support $[0, 1]$. We also assume that $\frac{f(\sigma_k|\omega=1)}{f(\sigma_k|\omega=0)}$

strictly increases in σ_k , which allows us to normalize signals such that $Pr[\omega = 1|\sigma_k] = \sigma_k$, where σ_k has support $\Sigma_k \subset [0, 1]$. Players' signals are independent conditional upon the state. Following (Smith and Sørensen 2000), we say that players have *unbounded private beliefs* when $\Sigma_k = [0, 1]$ for each Player k and *bounded private beliefs* otherwise. To simplify exposition, we work with signals' log-likelihood ratios, $s_k := \ln \left(\frac{f(\sigma_k|\omega=1)}{f(\sigma_k|\omega=0)} \right)$; let S_k be the support of s_k . Player k 's private signal indicates that $\omega = 1$ is no less likely than $\omega = 0$ iff $s_k \geq 0$.

Let $D(k) \subset \{1, \dots, k-1\}$ be the subset of Player k 's predecessors whose actions k observes; to distinguish this direct form of observation from less direct forms, we refer to it as “direct observation” henceforth. When $k-1 \notin D(k)$, we can interpret Players $k-1$ and k as moving simultaneously; we interpret one player's having a higher number than another simply as indicating that the former moves weakly later than the latter. Let $ID(k) \subset \{1, \dots, k-1\}$ be the subset of Player k 's predecessors whom k *indirectly* observes: $l \in ID(k)$ iff there exist some path of players k_1, k_2, \dots, k_L such that $k_1 \in D(k), k_2 \in D(k_1), \dots, k_L \in D(k_{L-1}), l \in D(k_L)$. Of course, there may be more than one path by which one player indirectly observes another, a possibility that plays a crucial role in our analysis below. If Player k directly observes Player j , then she must also indirectly observe him, but not necessarily *vice versa*. The indirect observation (ID) relation defines a strict partial order on the set of players.

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 her expectation of $-(\alpha - \omega)^2$ given all her information, I_k . Players do this by choosing $\alpha_k = E[\omega|I_k]$, namely by choosing actions that coincide with their posteriors that $\omega = 1$. Any player who observes Player k can therefore back out Player k 's beliefs from her action but cannot necessarily infer Player k 's private signal. For simplicity, as with signals, we identify actions by their log-likelihoods, $a_k := \ln \left(\frac{\alpha_k}{1-\alpha_k} \right)$. Player k optimally chooses $a_k \geq 0$ iff she believes $\omega = 1$ at least as likely as $\omega = 0$.

We refer to $\mathcal{N} = \{\{1, 2, \dots\}, \{D(1), D(2), \dots\}\}$ as an observation structure or a network, consisting of the players $\{1, 2, \dots\}$ and their respective sets of directly observed predecessors,

which define their sets of indirectly-observed predecessors.⁵ When it causes no ambiguity, we abuse notation by referring to \mathcal{N} as the set of players in the network \mathcal{N} . For any set A of action profiles $a := (a_1, a_2, \dots)$, we say that \mathcal{N} *admits* A if for each open set of actions B that contains A , $\Pr[B] > 0$. Given the network \mathcal{N} , its k -truncation $\mathcal{N}^k := \{\{1, 2, \dots, k\}, \{D(1), D(2), \dots, D(k)\}\}$ comprises its first k players as well as their observations sets.

Player k may observe a predecessor both directly and indirectly. Define $\overline{D}(k) = \{j \in D(k) : \forall i \in D(k), j \notin ID(i)\}$, the set of players whom Player k indirectly observes only by directly observing. In the classical single-file model, for example, $D(1) = \overline{D}(1) = \emptyset$ and for each $k \geq 2$, $D(k) = \{1, \dots, k-1\}$ and $\overline{D}(k) = \{k-1\}$. 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 $\overline{D}(2l+1) = \overline{D}(2l+2) = \{2l-1, 2l\}$. The “only-observe-directly” set $\overline{D}(k)$ plays an important role in our analysis and is non-empty whenever $D(k)$ is non-empty.⁶

Although a player may directly observe a large number of predecessors, many of those observations turn out to be redundant. For instance, in the classical, single-file structure with rational players, no player who observes her immediate predecessor gains any useful information by observing any other predecessor. 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 , $ID(k) = \overline{D}(k) \cup \left(\bigcup_{j \in \overline{D}(k)} ID(j)\right)$.*

⁵Formally, \mathcal{N} is a directed network with players as nodes, direct observations as directed links, etc. (See, e.g., Jackson (2008).) Because network-theoretic language does not clarify or simplify our results, we prefer to use 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.


⁶Because $\max D(k) = \max \overline{D}(k)$ one set is empty iff the other one is too.

Those predecessors who belong to $\overline{D}(k)$ collectively have access to all the information that Player k could ever hope to incorporate into her own decision. In most papers of the social-learning literature, each Player k 's only-observe-directly set $\overline{D}(k)$ is a singleton or has the property that $\cap_{j \in \overline{D}(k)} ID(j) = \emptyset$ —no two predecessors in $\overline{D}(k)$ share a common action observation. Either assumption frees Player k from concern that two distinct, non-redundant observations incorporate the same information.

As described in the introduction, we are particularly interested in networks that contain what we call “diamonds”: two players j and k both observe a common predecessor i but not each other, while some fourth player l observes both j and k .

Definition 1 *The distinct players i, j, k, l in the network \mathcal{N} form a diamond if $i \in ID(j) \cap ID(k)$, $j \notin ID(k)$, $k \notin ID(j)$, and $\{j, k\} \subset ID(l)$.*


We refer to the diamond by the ordered quadruple (i, j, k, l) —where $i < j < k < l$ —and say that the network has a diamond if it contains four players who form a diamond.



An important subset of a network's diamonds are those diamonds in which Player l also directly observes Player i .



Definition 2 *The distinct players i, j, k, l in the network \mathcal{N} form a shield if $i \in ID(j) \cap ID(k)$, $j \notin ID(k)$, $k \notin ID(j)$, $\{j, k\} \subset ID(l)$ and $i \in D(l)$.*

The diamond (i, j, k, l) is a shield if and only if $i \in D(l)$. Every shield must be a diamond, but not *vice versa*.



We say that the network has a shield if it contains players who form a shield. Finally, we are interested in diamonds (i, j, k, m) that if not shields—if Player m does not observe Player i —closely resemble shields in one of two ways: either Player m indirectly observes some Player l who belongs to the shield (i, j, k, l) , or Player m forms the shield (l, j, k, m) with some Player l who indirectly observes Player i .

Definition 3 *The diamond (i, j, k, m) in the network \mathcal{N} circumscribes a shield if (i) it is a shield, (ii) there exists some player $l < m$ such that (i, j, k, l) is a shield and $l \in ID(m)$, or (iii) there exists some player $l < m$ such that (l, j, k, m) is a shield and $i \in ID(l)$.*



Whenever the diamond (i, j, k, m) circumscribes a shield, Player m has access to the three signals s_i, s_j and s_k through three different channels, indirectly observes someone who has access to the three signals through three different channels, or has access to $s_j, s_k, s_i + s_l$ for some Player l through three different channels (in which case m need not disentangle s_i from s_l).

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 that is not a shield of Figure 1, for instance, the final observer l cannot discern the correlation in j and k 's beliefs through some 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. Although we do not formally analyze networks of partial inference, the form of anti-imitative behavior that we describe appears equally in these settings. Hence we define:

Definition 4 *Player k achieves impartial inference (II) if for each $(s_1, \dots, s_{k-1}) \in \times_{j < k} S_j$ and each $s_k \in S_k$,*

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

Otherwise, Player k achieves partial inference (PI).

A player who achieves impartial inference can never 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. As is clear from the example in the Introduction, a typical setting with discrete actions is unlikely to involve impartial inference when the signal structure is richer than the action structure, for even the second mover cannot fully recover the first mover’s signal from her action. But because our formal results concern rich actions where each person’s beliefs are fully revealed to all observers, in our setting the possibility of partial inference stems entirely from inability to disentangle the signals which generate the constellation of observed beliefs.

Note that 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 extracted enough of the information combined in their 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. We say that *behavior in the network \mathcal{N} is impartial* if each player in \mathcal{N} achieves impartial inference.

When each player's only-observe-directly set lacks two predecessors sharing a common observation, then all players achieve impartial inference by combining their private information solely with observations from predecessors in their only-observe-directly sets. In this case—which covers most of the social-learning literature— $\overline{D}(k)$ is “sufficient” for $D(k)$.

Lemma 2 *If for each Player k , $\overline{D}(k)$ is a singleton set or has the property that $\cap_{j \in \overline{D}(k)} ID(j) = \emptyset$, then each Player l achieves impartial inference by choosing $a_l = \sum_{j \in \overline{D}(l)} a_j + s_l$.*

In the single-file model, Lemma 2 implies that players achieve impartial inference by combining their private information with their immediate predecessor's action. Moreover, it implies that when the actions in $\overline{D}(l)$ are all independent conditional on the state, Player l rationally imitates the observed actions of everyone she indirectly observes only by directly observing, and ignores any other actions that she observes.

Of course, many networks have architectures that preclude rational players from achieving impartial inference. We wish to distinguish those networks that allow rational players to achieve impartial inference from those that do not. Especially in the case where actions reveal beliefs, it turns out that the impartiality of inference is intimately related to the role of diamonds and shields in the observational structure.

Proposition 1 *If every diamond in \mathcal{N} circumscribes a shield, then behavior in \mathcal{N} is impartial. If all players have unbounded private beliefs, then behavior in the network \mathcal{N} is impartial only if every diamond in the network circumscribes a shield.*

When (i, j, k) form the “base” of a diamond, the first Player l to complete the diamond (i, j, k, l) must directly observe i to achieve impartial inference with unbounded private beliefs. Roughly speaking, because a_j and a_k both weight s_i , Player l must learn s_i to uncover the correlation between a_j and a_k due to s_i to avoid double counting it in her own action. Accomplishing this requires directly observing a_i , for merely indirectly observing it via some other action a_m would not permit her to disentangle s_i from s_m .

The gap between necessary and sufficient conditions in Proposition 1 derives from the fact that certain discrete signal structures allow a Player k who observes Player j to disentangle j 's signal from j 's priors without further information. Consider for instance the diamond (i, j, k, l) in Figure 1 that is not a shield— $i \notin D(l)$ —but where $s_j = 0$ with certainty, namely j lacks an informative signal. Because $a_j = a_i + s_j = s_i$, $a_k = a_i + s_k = s_i + s_k$, Player l can use a_j and a_k to infer s_i, s_j and s_k and thereby achieve impartial inference.

We now turn our attention to the behavioral rules that players use in such shields to achieve impartial inference. To begin, we define anti-imitation more precisely:

Definition 5 *Player k anti-imitates Player j if for each $a_{-j} \in \mathbb{R}^{|\{i < k, i \neq j\}|}$ and each $s_k \in S_k$, (i) for each $a_j, a'_j \in \mathbb{R}$ such that $a_j < a'_j$*

$$a_k(a_j, a_{-j}; s_k) \geq a_k(a'_j, a_{-j}; s_k)$$

and (ii) there exist some $a_j, a'_j \in \mathbb{R}$ such that $a_j < a'_j$ and

$$a_k(a_j, a_{-j}; s_k) > a_k(a'_j, a_{-j}; s_k)$$

Player k anti-imitates Player j if k 's confidence in a state of the world never moves in the same direction as j 's—holding everyone else's action fixed—and sometimes moves in the opposite direction. Note that this formal definition of anti-imitation is stronger than the one described in the Introduction because it insists that the effect on belief of changing a player's action is weakly negative for every combination of others' actions. In the context of coarse action spaces, we do not say that Player 4 anti-imitates Player 1 if actions AAB provide a stronger signal in favor of state B than actions BBB , when it is also the case that action ABA provides a stronger signal in favor of the state A than actions BBA . A player who anti-imitates a predecessor always forms beliefs that tilt against that predecessor's.

Our main result is that in rich networks where players' actions reveal their beliefs, private beliefs are unbounded, and behavior is impartial, anti-imitation occurs if and only if the network contains a diamond. Roughly speaking, in settings where players observe some predecessors but not all of their most recent ones, certain players become less confident in

a state the more confident they observe certain of their predecessors becoming. That is, rational social learning requires anti-imitation whenever there are diamonds:

Proposition 2 *Suppose that behavior in \mathcal{N} is impartial. If some player in \mathcal{N} anti-imitates another, then \mathcal{N} contains a diamond. If all players have unbounded private beliefs and \mathcal{N} contains a diamond, then some player anti-imitates another.*

Not only does rational social learning in general observation structures often require that certain players anti-imitate others, but it also may lead to some players' forming beliefs that go against all of their information. That is, a player may form beliefs that are both contrary to his private signal and all the predecessors whose beliefs he observes. Two definitions will help us establish some surprising results to this effect.

Definition 6 *Player k 's observational beliefs in network \mathcal{N} following action profile $(a_1, a_2, \dots, a_{k-1})$ are*

$$o_k(a_1, \dots, a_{k-1}) := \ln \left(\frac{\Pr[\omega = 1 | \mathcal{N}^k; (a_1, \dots, a_{k-1})]}{\Pr[\omega = 0 | \mathcal{N}^k; (a_1, \dots, a_{k-1})]} \right)$$

A player's observational beliefs are those (in log-likelihood form) that she would arrive at after observing any actions visible to her but before learning her own private signal. In models where all players observe all of their own predecessors, observational beliefs are often called "public beliefs". In our setting, because the subset of Player k 's predecessors observed by Player $l \geq k$ may differ from those observed by Player $m \geq k$, observational beliefs are neither common nor public. In any case, a rational Player k chooses $a_k = o_k + s_k$, which optimally combines her own private information with that gleaned from her predecessors.


Definition 7 *The path of play (a_1, a_2, \dots, a_k) is contrarian if either (i) $\forall j \in D(k), a_j < 0$ and $o_k > 0$ or (ii) $\forall j \in D(k), a_j > 0$ and $o_k < 0$.*

A contrarian path of play arises for a Player k when despite all her observations favouring state $\omega = 0$, she attaches higher probability to state $\omega = 1$, or *vice versa*.

Proposition 3 *If players have unbounded private beliefs and the network \mathcal{N} contains a diamond and has impartial inference, then \mathcal{N} admits contrarian play.*

Impartial inference and the existence of a diamond imply that each player’s action is a linear combination of the actions of her predecessors she observes plus her signal. Because the weights in this linear combination do not depend upon the realisation of any signal or action, if Player k attaches a negative weight to Player j ’s action, as the magnitude of a_j becomes large—and all other actions are held fixed, something possible because private beliefs are unbounded— a_k must eventually take on the opposite sign as a_j .

We conclude the section by illustrating the features above with a simple example drawn from Eyster and Rabin (2009). Our general network encompasses simple, natural variants of the standard model, such as the one where rather than move “single-file” like in the standard model, n players here move “multi-file” in each round, each player observing all players moving in prior rounds but not the current or future rounds. When $n \geq 2$, this network includes diamonds that are shields and admits contrarian play. Figure 5 illustrates the first five movers in a double-file setting.



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

To succinctly describe behavior in this model, let $A_t = \sum_{k=1}^n a_t$, the sum of round- t actions or *aggregate round- t action*, and $S_t = \sum_{k=1}^n s_t$, the sum of round- t signals or *aggregate*

round- t signal.

Clearly $A_1 = S_1$, so for a player in round 2 with signal s_2 , $a_2 = s_2 + A_1$, in which case $A_2 = S_2 + nA_1$. Likewise, a player in round three wishes to choose $a_3 = s_3 + S_2 + S_1$. 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 imitate each round-2 player and know that *each* of those players is using all round-1 actions. Since they do not want to count those n -fold, they subtract off $n-1$ of the round-1 aggregate actions. In general,

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

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

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

For $t \geq 3$, Player t anti-imitates approximately half of her predecessors. This implies that nearly all social learners in the infinite sequence engage in substantial anti-imitation.

Whatever n , substituting for A_{t-i} recursively gives

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

where players in round t give all signals unit weight; hence, the aggregate round- 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 do learn the state eventually. Note, importantly, that the wild swings in how people use past actions typically do not find their way into actions: recent actions always receive positive weight, and *typically* they are more extreme than earlier actions. It is when play does not converge fast enough that we would observe rational players switching. Roughly speaking, approximately half of social learning in this setting is anti-imitative!

⁷Observational following round $t-1$ are $\sum_{i=1}^t (-1)^{i-1} (n-1)^{i-1} A_{t-i}$

To see the crispest form of contrarian play, note that when there are three players, we will observe the following pattern in the first three rounds. When $n = 3$,

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

leading to $A_1 = S_1$, $A_2 = S_2 + 3A_1$, and $A_3 = S_3 + 3A_2 - 6A_1$. 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 $\alpha = 0.6$, each expressing 60% confidence that $\omega = 1$. If all second-period players also were to choose $\alpha = 0.6$, 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! In general, in *odd* periods, complete agreement by predecessors always leads players to contradictory beliefs.⁸

3 Guaranteed Contrarian Behavior

In the last section, we showed how shields are a necessary and (essentially) sufficient condition to produce guaranteed anti-imitation and possible contrarian behavior. In this section, we give an example of a network architecture that guarantees with arbitrarily high probability at least one episode of contrarian play.

In every round, five players move simultaneously, four of them named Betty. Each Betty observes the actions of all five players in the previous round plus her own private signal.

⁸This 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.

The fifth player, Gus, observes all players in the previous round but is known to have no private signal of his own. No player observes any player who does not move in the round immediately before his or her round.

For simplicity, in this section we consider players with bounded private beliefs. In particular, each Betty receives a draw from a distribution of binary signals $s \in \{0, 1\}$ parameterized by $p := \Pr[s = 1 | \omega = 1] = \Pr[s = 0 | \omega = 0]$. We refer to p as the signal structure, which decreases in informativeness as $p \rightarrow \frac{1}{2}$. Our result is that eventually a round will occur in which Gus takes an action that reveals an aggregate of two more $s = 1$ than $s = 0$ signals, while all four Bettys choose actions indicating an aggregate of one more $s = 1$ than $s = 0$ signals. In the following round, Gus will choose an action that reveals an aggregate of two more $s = 0$ than $s = 1$ signals, while all Bettys then will choose actions indicating at least one more $s = 0$ than $s = 1$, contrarian behavior!

Proposition 4 *For any $\varepsilon > 0$, there exists a signal structure $p > \frac{1}{2}$ under which with probability at least $1 - \varepsilon$ that there exists some round t when all players' play is contrarian.*

In any round where the net number of $s = 1$ signals is two and all four current signals are $s = 0$, players in the following round have one, two, or three $s = 0$ signals—depending on their private information—and therefore exhibit contrarian behavior. When signals are weak, the net number of $s = 1$ signals approaches a random walk, in which case such a round would occur with certainty. Short of this limit, such an episode happens with near certainty. Indeed, as $p \rightarrow \frac{1}{2}$, contrarian play happens arbitrarily many times.

While we have not proven it, we suspect that contrarian play would arise in a setting with a large number of symmetric, privately informed players acting in every period, i.e., without Gus there to allow impartial inference by keeping track of observational beliefs. As the number of players in every round grows arbitrarily large, the distribution of their actions almost perfectly reveals observational beliefs, essentially replicating Gus. Moreover, signals every period will resemble their parent distribution and, like in our example, come to resemble a random walk with drift. Once more, as signals become less and less informative, the conditions that guarantee contrarian play will be satisfied with near certainty.

In fact, as signals become very weak in this setting, first-round movers are almost as likely to be right as wrong. Yet eventually a majority of signals will come to favor the true state. So long as the net number of signals favouring one state goes from -2 or more in one round to +2 or more in the next—so long as the switch does not happen following a round where there the number of signals favouring the two states differs by no more than one—then at this time all players, irrespective of their signals, take actions favoring a state deemed most likely by only a minority of first-round movers and no player since. As the number of players becomes large, the probability that the switch takes place in such a round approaches one. Thus, in this setting the probability that all players in some round take actions contrary to every one of their predecessors except a minority of first-round movers approaches one-half.

4 Conclusion

In settings rich enough to allow all players to perfectly infer the relevant private information of all the predecessors whose actions they observe, we show that when players do not observe all of their most recent predecessors or contemporaneous movers, rational observational learning must involve some anti-imitation. We limit our analysis to rich-information settings in order both to crisply articulate the observational conditions for which anti-imitation occurs and to sharply differentiate two sources of anti-imitation; yet intuition and many examples suggest that anti-imitation is a quite general feature of social learning across domains. Because we doubt its empirical plausibility, we close by speculating on the implications of social-learning without anti-imitation. If people neither anti-imitate, nor virtually ignore their predecessors' actions—an even more dubious empirical prediction—then it seems likely that in many contexts observational learning will entail inevitable overconfident and inefficient herding. That is, because behavioral anti-imitation is so fundamentally necessary for avoiding the sort of over-counting of early signals that leads to mis-inference, we suspect that anti-imitation is not merely a feature of rational herding but also of *any* theory of observational learning that does not lead to frequent instances of overconfident and wrong

herds.

If people do not realize that they must “subtract off” the earlier actions they observe from their interpretation of later actions, then the effect of early actions on long-run beliefs will become immense. This “social confirmation bias” leads to both false herds in the many situations where rational observational learning rules them out, *and* to extreme confidence in those wrong herds. This immense over-counting can only be mitigated without anti-imitation if late movers are barely influenced by recent observations. This too seems very unlikely.⁹

In fact, Eyster and Rabin (2010) explore the implications for social learning of one simple theory of inference that incorporates non-anti-imitation. We assume there that players fail to attend to the strategic logic of the setting they inhabit by naively believing that each predecessor’s actions reflects solely her private information. This simple alternative leads people to imitate *all* observed previous behavior, regardless of the observation structure. It also leads to very different implications for the informational and efficiency properties of herds. Naïve players can herd on incorrect actions even in the many rich environments where fully rational players always converge to the correct ones. They become extremely, and often wrongly, confident about the state of the world in nearly all environments, including ones where rational players never become confident. Moreover, depending upon 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 classical herding settings, even minimal departures of this sort from the rational calculus of recency can make long-run inefficiency very likely. Consider, for instance, a situation where people move single-file and take binary actions, and very few of them receive any private information—yet those who do are very well informed. Formally, a tiny fraction ϵ^2 receive binary private signals that match the state of the world with the very high probability $1 - \epsilon$. As $\epsilon \rightarrow 0$, rational players eventually learn the true state with near certainty. Frequently this involves anti-majority play, because informed players know that the first herd

⁹Even if behaviorally plausible, such stubborn beliefs would likely lead to the (more conventional) form of long-run inefficiency whereby society’s beliefs never converge even in the presence of an unboundedly large number of signals.

most likely begins with an uninformed Player 1 choosing randomly and a string of uninformed successors following suit. Only an informed player bucks the herd, and so his uninformed immediate successors will act against the majority of observed actions by following him. What happens when players are not *so* willing as a fully rational person to follow a minority? Consider, for instance, a population of people almost all of whom when uninformed choose one action over the other whenever it has been chosen at least $K > 1$ times than the other. When $K = 10$, for instance, no uninformed player would choose A over B when the number of her predecessors choosing B is more than ten more than the number choosing A . Such rules allow players to anti-imitate, just not as dramatically as they do in the rational model. As $\epsilon \rightarrow 0$, then, with only about 50% chance does the herd converge to the correct action: if the herd starts out wrong, then nobody is likely to get a signal suggesting that they should overturn it until too late. An informed player may follow her own signal and buck the herd, but the likelihood of having enough of them to ever bring the majority in favor of Player 1's action to within K is negligible. Similarly, even if a minority of players is fully rational, they would not overturn the herd. Hence, even players fairly willing to follow a minority—merely not quite *so* willing as rationality prescribes—fall dramatically short of efficient information aggregation. Similarly, in the case studied by Callender and Hörner (2009) where the order of moves is not observed, and where the rational anti-imitation seems an even less likely logic to prevail, an unwillingness to follow the minority is likely to lead to observational learners to converge to the wrong action with very high probability.

More than a willingness to follow a minority over majority, examples like the setting examined in Section 2 suggest that “enough” rational anti-imitation to prevent inefficiency relative to the rational model may be unrealistic indeed. In that setting, when each person's signal is relatively weak, there is close to 50% chance that by Round 2 people will unanimously be playing the wrong action. If people in this situation merely combined their own signals with the beliefs of 2 of the 10 observed (rather than all 10), it is guaranteed that beliefs will become stronger and stronger in the wrong direction once it starts. If people neither massively ignore the beliefs of those they observe nor contradict every single person's

beliefs they observe, overconfident herding will occur.

More generally, in rich-action settings like that of this paper, in all but knife-edge cases, behavioral rules that do not anti-imitate, like the naivety of the last paragraph, lead to one of two undesirable consequences. Like naïve play of the previous paragraph, behavioral rules that put substantial weight on predecessors' actions converge with positive probability to the worst possible action. Behavioral rules that put insubstantial weight on predecessors' actions fail to converge to actions corresponding to certainty in one state or the other. In sum, without anti-imitation, observational learning is unlikely to produce convergence to the correct beliefs even in the richest of environments.

References

- Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar,** “Bayesian Learning in Social Networks,” Mimeo 2010.
- 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.
- 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.
- Callender, Steven and Johannes Hörner,** “The Wisdom of the Minority,” *Journal of Economic Theory*, 2009, *144* (4), 1421–1439.

Eyster, Erik and Matthew Rabin, “Rational and Naïve Herding,” CEPR Discussion Paper DP7351 2009.

— and —, “Naïve Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2010, 2 (4), 221–243.

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.

Jackson, Matthew, *Social and Economic Networks*, Princeton University Press, 2008.

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.

5 Appendix

Proof of Lemma 1.

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

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

Proof of Lemma 2. We prove this by induction. Clearly Player 1 achieves impartial inference by choosing $a_1 = s_1$ as per the claim. Suppose that Players $i \in \{1, \dots, k-1\}$ choose a_i as claimed. If $\overline{D}(k) = \{j\}$ for some $j < k$, then since j does II by assumption, so too does k through $a_k = a_j + s_k$. Suppose then that $\overline{D}(k)$ has more than one element. Then

$$\begin{aligned} a_k &= \sum_{j \in \overline{D}(k)} \left(\sum_{i \in ID(j)} s_i + s_j \right) + s_k = \sum_{j \in \overline{D}(k)} s_j + \sum_{j \in \overline{D}(k)} \sum_{i \in ID(j)} s_i + s_k \\ &=: \sum_{l \in ID(k)} \beta_l s_l + s_k, \end{aligned}$$

where the first equality comes from the induction hypothesis that j achieves II. By Lemma 1, for each $l \in ID(k)$, $\beta_l \geq 1$; by the assumption that $\cap_{j \in \overline{D}(k)} ID(j) = \emptyset$, for each $l \in ID(k)$, $\beta_l \leq 1$; together, these yield the result. ■

Proof of Proposition 1. First statement Suppose that all diamonds circumscribe shields. Again we prove the result by induction. Once more Player 1 achieves II by choosing $a_1 = s_1$. We prove that if all Players $i \in \{1, \dots, k-1\}$ achieve II, then so too does Player k . Define

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

Players $i \in \{1, \dots, k-1\}$ achieving II and Lemma 1 imply that for each $j \in ID(k)$, $\beta_j^1 \geq 1$. Define $U(k) := \{j \in ID(k) : \beta_j^1 = 1\}$ and $M(k) := \{j \in ID(k) : \beta_j^1 > 1\}$. First, notice that $\forall i \in U(k), \forall j \in M(k), i \notin ID(j)$, for, otherwise, because $\forall j \in M(k), \exists k_1, k_2 \in \overline{D}(k)$ 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 $i \in M(k)$, a contradiction. This implies that if through adding some linear combination of actions from players in $M(k)$ to a_k^1 , Player k manages to set the weights on all signals in $M(k)$ equal to one, then Player k achieves II. If $M(k) = \emptyset$, then Player k achieves II through a_k^1 . If $M(k) \neq \emptyset, M(k) \subset D(k)$, then Player k can achieve II as follows. Define $\hat{j}_1 := \max M(k)$, and

$$a_k^2 := a_k^1 - (\beta_{\hat{j}_1}^1 - 1) a_{\hat{j}_1} =: \sum_{j \in ID(k)} \beta_j^2 s_j + s_k,$$

where by construction now for each $l \in ID(k), l \geq \hat{j}_1$, $\beta_l^2 = 1$. Define $\hat{j}_2 := \max M(k) \setminus \{\hat{j}_1\}$ and

$$a_k^3 := a_k^2 - (\beta_{\hat{j}_2}^2 - 1) a_{\hat{j}_2} =: \sum_{j \in ID(k)} \beta_j^3 s_j + s_k,$$

where by construction now for each $l \in ID(k), l \geq \hat{j}_2$, $\beta_l^3 = 1$. Continuing in this way gives II for the case where $M(k) \subset D(k)$. Suppose, however, that some $i \in M(k) \cap (D(k))^c$. Since i belongs to a diamond with k and two members of $\overline{D}(k)$, yet $i \notin D(k)$, then because by assumption this diamond circumscribes a shield, there exists some player $l < k$ and $k_1, k_2 \in \overline{D}(k)$ such that (l, k_1, k_2, k) is a shield and $i \in ID(l)$. Wlog let l be the lowest-indexed such player. (Case (ii) of the definition of circumscribing a shield cannot apply

here because then $k_1, k_2 \in ID(l)$, $l \in ID(k)$ implies $k_1, k_2 \notin \overline{D}(k)$, a contradiction.) Since $l \in D(k) \cap M(k)$, the iterative procedure described above eventually must reach some step p such that a_k^p puts weight one on s_l . Because $l \in ID(k_1) \cap ID(k_2)$, k_1, k_2 achieving II must both put weight one on s_l and s_i , which implies that s_l has weight one in a_k^p iff s_i has weight one.

Second statement Assume that players have unbounded private beliefs and achieve II. Towards a contradiction, suppose also that some diamond does not circumscribe a shield. Let (i, j, k, m) be such a diamond with the lowest possible index of m (the first diamond not circumscribing a shield as ordered by the highest-indexed player, the “tip”, of the diamond). Because condition (ii) of Definition 3 fails, we can assume wlog that $j, k \in \overline{D}(m)$. As per above in the proof of the first statement, define

$$a_m^1 := \sum_{j \in \overline{D}(m)} a_j + s_m =: \sum_{j \in ID(m)} \beta_j^1 s_j + s_m$$

and $M(m) := \{j \in ID(m) : \beta_j^1 > 1\}$. Clearly, $i \in M(m)$. For Player m to achieve II, she must subtract off (i) a_i or (ii) some a_l for which $i \in ID(l)$ and $l \in M(m)$. Because m cannot observe i if (i, j, k, m) is not a shield, we focus on case (ii). But this implies that (l, j, k, m) is a shield, and $i \in ID(l)$, which in turn implies that (i, j, k, m) circumscribes a shield, a contradiction. ■

Proof of Proposition 2. The if direction follows directly from the proof of Proposition 1, where when i is the first player to form the base of a diamond, and k the first player to form the tip of a diamond with i , then Player k anti-imitates Player i . For the other direction, if \mathcal{N} contains no diamond, then consider

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

By Lemma 1 and impartial inference, for each $j \in ID(k)$, $\beta_j \geq 1$. Because \mathcal{N} contains no diamonds, $\beta_j \leq 1$. ■

Proof of Proposition 3. By Proposition 2, the network has negative weighting. Take the first player to do negative weighting, at the head of a shield—Player k —and let Player j be

the last player whom Player k weights negatively. Suppose that all players in $ID(k)$ other than j play actions in $(0, \epsilon)$ for $\epsilon > 0$ small and j plays action in $(0, \frac{1}{\epsilon})$. With unbounded private beliefs, for each $\epsilon > 0$, these actions happen with strictly positive probability. Note that

$$a_k < |\overline{D}(k)| < k \cdot \epsilon - \frac{1}{\epsilon},$$

using the expression for a_k from the proof of Proposition XXX. Hence, for ϵ sufficiently small, $a_k < 0$ despite all of k 's action observations being positive, a positive-probability instance of anti-unanimity. ■

Proof of Proposition 4.

Suppose wlog $\omega = 1$ so $q := 1 - p < \frac{1}{2}$ is probability of receiving $s = 0$ signal. Let $X(t)$ be the number of $s = 1$ signals through round $t - 1$ minus the number of $s = 0$ signals. Let $E(t, q)$ be the event that $X(t) = 2$ and $|E(t, q)| = \sum_{t=1}^{\infty} 1_{E(t, q)}$. As $q \rightarrow \frac{1}{2}$, $X(t)$ approaches a recurrent random walk, so that $\forall N \in \mathbb{N}, \forall \epsilon > 0, \exists \delta > 0$ such that $\forall q \in (\frac{1}{2} - \delta, \frac{1}{2})$, $Pr[|E(t, q)| > N] > 1 - \epsilon$.

Because signals are conditionally iid, $Pr[E(t, q)](1 - q)^4$ is the probability that $X(t) = 2$ and all signals in round t are $s = 0$, leading to contrarian play in round $t + 1$. Hence, the probability of contrarian play is at least

$$Pr[|E(t, q)| > N] (1 - (1 - (1 - q)^4)^N) > (1 - \epsilon) (1 - (1 - (1 - q)^4)^N).$$

Choosing ϵ sufficiently small and N sufficiently large yields the result. ■