The Limits to Imitation in Rational Observational Learning*

Erik Eyster and Matthew Rabin†

August 8, 2012

Abstract

An extensive literature identifies how rational people who observe the behavior of other privately-informed rational people with similar tastes may come to imitate them. While most of the literature explores when and how such imitation leads to inefficient information aggregation, this paper instead explores the *behavior* of fully rational observational learners. In virtually any setting, they imitate only some of their predecessors, and sometimes 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-imitation" 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. Anti-imitation arises from players' need to subtract off the sources of correlation in others' actions, and is mandated by rationality in settings where players observe many predecessors' actions

*We thank Min Zhang for excellent research assistance as well as Paul Heidhues and seminar participants at Berkeley, DIW Berlin, Hebrew University, ITAM, LSE, Penn, Tel Aviv University, and WZB for their comments.

†Eyster: Department of Economics, LSE, Houghton Street, London WC2A 2AE, UK. Rabin: Department of Economics, UC Berkeley, 549 Evans Hall #3880, Berkeley, CA 94720-3880, USA.

but not all recent or concurrent ones. 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 an arbitrarily large number of such episodes of extreme contrarian behavior. (JEL B49)

1 Introduction

Inference from the behavior of others is one of the ways that people guide themselves in making wise economic and social choices. Restaurant goers (to use the canonical example) decide where to go partly by the popularity of the restaurants. Investors infer good financial strategies from others' portfolios. And each new generation of society infers effective ways to live from the behavior of its elders. Despite interacting with other means of acquiring information—communication, experimentation, observation of others' success, and direct information—learning by observing the actions of others is an important facet of social and economic behavior, and has formed the basis for a massive and ongoing research program. Beginning with Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), this literature on observational learning identifies how a rational person who observes the behavior of another person with similar tastes and private information may imitate that person, even contrary to her own private information.

Yet the very logic of social inference requires that rational agents greatly limit the scope of their imitatation. When those people comprising the societies, groups, or markets whose behavior is being watched *themselves* are learning by observing the actions of people before them, the information revealed by their behavior contains a great deal of redundancy. Properly understanding this redundancy is such a formidable task that it has motivated researchers to develop models of limited-rationality learning where people do not sufficiently attend to the correlation in others' behavior. DeMarzo, Vayanos and Zwiebel (2003), for instance, model the idea that people may lack the sophistication to attribute repeated hearing of the same opinion to common information sources, and show that this "persuasion bias" generates inefficiency. And Eyster and Rabin (2010) look at the implications in simple herding contexts if people mistakingly treat the actions of all others as independent information, showing how this can lead to long-run incorrect and overconfident beliefs even in incredibly rich settings. By failing to attend to the logic that those they are imitating are themselves imitating, people end up being massively over-influnced by the early goers.

These forms of limited rationality in which people attend less than fully to redundancy lead to over-imitative behavior. This paper brings into systematic focus an obvious question: what behavior would we observe in a society of rational herders who instead attend fully to the redundancy in others' actions?

Focusing on environments where the action and observational structure is rich enough for people to extract any relevant information from all the past actions that they observe, we identify necessary and sufficient conditions for observational learning to include 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. Moreover, in conditions where people anti-imitate, there is always 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 most commonly used in the literature. We also illustrate related settings where rational herds that eventually settle down on the correct action almost surely involve many episodes of such contrarian behavior.

In Section 2, we motivate the more general results with two simple examples: unlike the familiar example of single-file herding, we suppose $n > 1$ people move simultaneously every period. Each one gets independent private information conditional upon the binary state and observes all actions in all prior periods. In the first example, players choose actions from a continuum that fully reveal their beliefs, for instance when investors with known risk attitudes allocate their portfolios across the market index fund and risk-free bonds. 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, properly extracting the information from period-2 actions without counting this correlated information multi-fold requires period-3 players to imitate period-2 actions but *subtract off* period-1 actions. This anti-imitation can take the dramatic form of players reaching conclusions contrary to all observed beliefs. Intuitively, if all period-1 and period-2 players agree on what is more likely, yet period-2 players do not sufficiently increase their confidence relative to period

1, then each of them must have received conditionally independent evidence that the herd started in the wrong direction. When $n > 2$, if *all* $2n$ people in the first two periods indicate equal confidence in either of the two states, a rational period-3 agent always concludes that the other state is more likely! Our second example also considers $n > 1$ players moving simultaneously, but modifies the canonical two-signal, two-restaurants-to-choose-from model by allowing each player to be uninformed, and to dine at home if sufficiently uncertain that she can identify the better restaurant. In this setting, a similar logic can lead a player to dine at a restaurant that her private information deems inferior even when observing only people staying home and going to the *other* restaurant. This second example illustrates how the basic logic of anti-imitation and contrarian play holds even when actions are not continuous, and a variant shows that it holds even when the order of moves is not observable.

In Section 3, we flesh out the logic of these examples more generally within a class of situations that we call "impartial inference". In contrast to partial inference, these are structures where common knowledge of rationality implies that any player who learns something about a predecessor's signal learns everything that she would wish to know from his signal. We follow Lee (1993) in working with a continuous action space that always fully reveals players' beliefs, and 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 1) j and k both observe i, 2) neither j nor k observes the other, and 3) l observes j, k, and i. 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 imitated it themselves already.[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 any game that includes a quartet of players satisfying the conditions above, there is a positive probability of a signal

[1]The result is rough because the players' observations (or lack thereof) described in parts 1 and 2 can be "indirect" in a sense made precise in Section 3.

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 l must infer that j and k both received negative information so that altogether the hypothesis is unlikely.

While in many settings such a strong form of contrarian behavior is merely a possibility, in Section 4 we illustrate a setting where it happens with near certainty. Players move multi-file and observe only those moving in the previous period; the information structure has the feature that each player can use actions in the previous period to perfectly infer public beliefs and signals in that period. We identify a set of observed actions, each of which reveals beliefs that state B is more likely than A, that lead a rational observer to conclude the opposite, namely A more likely than B. We then construct an example where such actions occur with probability arbitrarily close to one; society accomplishes complete learning through at least one (and on average many) episodes of strongly contrarian play.

We conclude in Section 5 by putting this paper in broader context, including how and why the type of anti-imitation we identify here extends quite generally outside our specific environment, including in cases where people communicate beliefs directly, and how it differs from other forms of anti-imitative behavior explored elsewhere in the literature.

2 Multi-File Movers

In this section, we illustrate the main ideas of the paper through an example that modifies the canonical models of Banerjee (1992) and Bikhchandani et al. (1992), where players move single-file, to one where players move "multi-file", namely more than a single mover per period. Like in the standard model, each player observes all players moving in previous periods but none of those moving in the current or future periods.

There are two possible states of the world, $\omega \in \{0, 1\}$, one of which is drawn from a uniform prior. In each period, three players choose actions simultaneously. Each player observes all players who moved in earlier periods but none who moves in the current or future periods. Before choosing her action, Player k receives a private signal $s_k = \Pr[\omega = 1|s_k]$;

players' signals are iid conditional on the state.

Instead of the usual binary-action setting, suppose that players can choose any action in the continuum $[0, 1]$, and that each player plays optimally by choosing the action equal to her updated probabilistic beliefs that $\omega = 1$.[2] This continuous-action-space environment was first studied in the social-learning literature by Lee (1993) and greatly facilitates the exposition because players' actions reveal their updated beliefs, which allows any successor to back out a player's private signal from her action. Lee (1993) proved that learning is complete in this setting; our focus, however, is not on efficiency but instead on behavior during the learning process.

Suppose that the three first-period players choose $(0.9, 0.9, 0.9)$, followed by the three second-period players also choosing $(0.9, 0.9, 0.9)$. Anyone observing this history of play perceives six players reveal ninety-percent confidence that $\omega = 1$. Without any private information of his own, which state would such a rational observer deem more likely?

The answer is that rational inference leads to near certain conviction that $\omega = 0$; precisely, the state $\omega = 0$ is 729 times as likely as $\omega = 1$! The first three players' actions reveal signals $9 : 1$ in favor of $\omega = 1$. Having observed all three first-round movers, for any player in period 2 to maintain $9 : 1$ beliefs in favor of $\omega = 1$ requires having a $81 : 1$ signal in favor of $\omega = 0$. (Intuitively, someone without any private information of his own who observed the first-period players would have beliefs $9 \cdot 9 \cdot 9 = 729$ in favor of $\omega = 1$ because signals are iid conditional on the state.) Because all three second-period players in the example behave identically, they all hold $81 : 1$ signals in favor of $\omega = 0$. Three $9 : 1$ signals in favor of $\omega = 1$ combined with three $81 : 1$ signals in favor of $\omega = 0$ produces a likelihood ratio of

$$\frac{\Pr[\omega = 1]}{\Pr[\omega = 0]} = \frac{9}{1} \cdot \frac{9}{1} \cdot \frac{9}{1} \cdot \frac{1}{81} \cdot \frac{1}{81} \cdot \frac{1}{81} = \frac{1}{729}.$$

If instead the three first-period players had played $(0.5, 0.5, 0.5)$—no one's private signal differed from her priors—and the second-period players again played $(0.9, 0.9, 0.9)$, then a rational observer unendowed with private information of his own would conclude that

$$\frac{\Pr[\omega = 1]}{\Pr[\omega = 0]} = \frac{1}{1} \cdot \frac{1}{1} \cdot \frac{1}{1} \cdot \frac{9}{1} \cdot \frac{9}{1} \cdot \frac{9}{1} = 729.$$

[2] The next section fleshes out preferences for which such behavior is optimal.

Rational players in the third period *anti-imitate* those in the first period. Fixing the expressed beliefs of second-period players, the more convinced first-period players are that $\omega = 1$, the more convinced third-period players are in the opposite conclusion that $\omega = 0$. The reason is the second-period actions are all positively correlated with first period actions; correct Bayesian inference requires players in the third period to account for that correlation by stripping it out. This negative partial derivative on first-period actions does not depend upon which actions are played. Suppose a first-period action goes from $x : 1$ in favor of $\omega = 1$ to $\lambda x : 1$ in favor of $\omega = 1$, for some $\lambda > 1$. To keep all second-period actions constant, a second-period signal that was $y : 1$ in favor of $\omega = 0$ must become $\lambda y : 1$ in favor of $\omega = 0$. Since that holds for all three second-period players, scaling a first-period action by the factor $\lambda > 1$, while holding second-period actions fixed, scales third-period beliefs by the factor $\frac{\lambda}{\lambda \cdot \lambda \cdot \lambda} = \frac{1}{\lambda^2} < 1$, generating anti-imitative play.

In this example, players in period three imitate those in period two and anti-imitate those in period one; they imitate exactly half of their predecessors and anti-imitate the other half. This conclusion does not change when additional periods are appended to the model: players in a fifth period would imitate predecessors in the second and fourth periods and anti-imitate those in the first and third rounds. Anti-imitation is not some fringe element of observational learning but instead one of its core components.

Third-period behavior in the example is also *contrarian*: someone who observes all of her predecessors deem $\omega = 1$ ninety percent probable—more likely than $\omega = 0$—concludes the opposite, namely $\omega = 0$ more likely than $\omega = 1$. Clearly the result is robust to not knowing the order of predecessors' moves, just as long as players know that three people move in every period.[3]

Appendix A shows that this example does not depend upon its rich action space by discretizing it to five actions and five signals. We now turn to a slight modification of the standard two-restaurant model of social learning can trigger a herd on an action before anyone sees it played. There are two restaurants in town, A and B, one of which is superior

[3]The small literature that studies social learning with unobserved order uses the analogous assumption that all players know that their predecessors moved single-file without knowing the order of those moves.

to the other; let ω_A denote the state of the world that Restaurant A is superior and ω_B denote the state that B is superior. Diners begin with uniform priors over the states and receive ternary private signals, $\{\alpha, \beta, \emptyset\}$, regarding the identity of the better restaurant. The signals are iid conditional on the state, where the α signal supports the state ω_A, the β signal supports the state ω_B, and the $\emptyset$ signal is uninformative. Specifically, for each Player k, let $\Pr[s_k = \alpha|\omega_A] = \Pr[s_k = \beta|\omega_B] = \theta(1-\eta)$ and $\Pr[\emptyset|\omega_A] = \Pr[\emptyset|\omega_B] = \eta$, where $\theta \in \left(\frac{1}{2}, 1\right), \eta \geq 0$. When $\eta = 0$, this reduces to the canonical binary-state, binary-signal information structure.

Each Player k chooses among three choices: dining in Restaurant A, dining in Restaurant B, and dining at home, represented by the action set $\{A, B, H\}$. Player k with information I_k has preferences that lead her to dine in Restaurant A iff $\Pr[\omega_A|I_k] \geq \theta$, dine in Restaurant B iff $\Pr[\omega_B|I_k] \geq \theta$, and otherwise dine at home. As in the previous example, players move triple-file.

Suppose that first-period actions are (A, H, H) and second period actions (H, H, H). A rational player in the third period infers private signals $(\alpha, \emptyset, \emptyset)$ in the first period; since anyone receiving an uninformative or α signal in the second period would have chosen Restaurant A, a rational observer concludes from the observation that all three second-period diners dined at home that they received signals (β, β, β). Hence, public beliefs after the first two periods reflect two more β signals than α. For exactly the same reason as in the standard binary-state, binary-signal model of Bikhchandani et al. (1992), this triggers a herd on B. Despite never seeing a single predecessor choose Restaurant B, the players in Period 3 initiate a herd on Restaurant B![4] Herding on a yet-to-be-played action is an extreme form of contrarian play. The fact that this behavior is also anti-imitative follows easily from symmetry: switching the one first-period A to B, while leaving all others unchanged, would launch a herd on A.

Not only does this example demonstrate that the logic of anti-imitation and contrarian play has nothing to do with rich action spaces, but it also easily extends to make the

[4]Period 3 also marks the beginning of an information cascade.

point that these phenomena do not depend upon the potentially unrealistic assumption that players know the order of the predecessors' moves used in most of the literature. Suppose that it is common knowledge that players move triple-file but no player knows the order of preceding actions. Consider once more a the third period observer who has one A and five H choices. We saw before that if the A action occurred in the first period, then two of the six players received uninformative signals. If the A action happened in the second period, then five players must have received uninformative signals (everyone except the person choosing A). The smaller η, the less likely uninformative signals, and the more likely that the one A diner moved in the first period. In the limit as $\eta \to 0$, the probability that the sole A action happened in the first period approaches one. Consequently, for small enough η, the players moving in the third period make virtually the same inferences in the unordered case as they did above. In particular, a herd forms on B after (A, H, H) and (H, H, H), even when players do not know the order of their predecessors moves.[5] Anti-imitative play is robust to unordered observations.

3 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

[5]Formally, letting S_i be the event that the sole A happened in Period $i = 1, 2$ and S be the unordered observation,

$$\Pr[\omega_A|S] = \frac{(\Pr[S_1|\omega_A] + \Pr[S_2|\omega_A])\,\frac{1}{2}}{(\Pr[S_1|\omega_A] + \Pr[S_2|\omega_A])\,\frac{1}{2} + (\Pr[S_1|\omega_B] + \Pr[S_2|\omega_B])\,\frac{1}{2}}$$

$$= \frac{3\theta(1-\eta)^4\eta^2(1-\theta)^3 + 3\eta^5\theta(1-\eta)}{3\theta(1-\eta)^4\eta^2(1-\theta)^3 + 3\eta^5\theta(1-\eta) + 3\theta^3(1-\eta)^4\eta^2(1-\theta) + 3\eta^5(1-\theta)(1-\eta)}$$

$$= \frac{\theta(1-\eta)^3(1-\theta)^3 + \eta^3\theta}{\theta(1-\eta)^3(1-\theta)^3 + \eta^3\theta + \theta^3(1-\eta)^3(1-\theta) + \eta^3(1-\theta)}$$

Because

$$lim_{\eta\to 0} \Pr[\omega_A|S] = \frac{(1-\theta)^2}{(1-\theta)^2 + \theta^2} = \Pr[\omega_A|S_1],$$

when uninformative signals are very unlikely, the inferences that Period-3 players draw from their predecessors' triple-file actions is virtually the same whether or not the players observe the order of those actions.

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.[6] Appendix B provides sufficient conditions for 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.

There are two possible states of the world, $\omega \in \{0, 1\}$, each one *ex ante* equally likely. Players in the set $\{1, 2, \ldots\}$, which can either be finite or infinite, receive private signals correlated with the state. Following Smith and Sørensen (2000), we work directly with players' updated private beliefs based upon their private signals alone: let σ_k be Player k's belief that $\omega = 1$ conditional upon her private signal by itself. 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 the case that some signal reveals the state with certainty; let $[\underline{\sigma}_k, \overline{\sigma}_k] \subseteq [0, 1]$ denote the non-degenerate convex hull of their common support. Player k's private beliefs are *unbounded* when $\underline{\sigma}_k = 0$ and $\overline{\sigma}_k = 1$ and *bounded* otherwise. To simplify exposition, we work with the log-likelihood ratio of private beliefs, $s_k := \ln\left(\frac{\sigma_k}{1-\sigma_k}\right)$; let $[\underline{s}_k, \overline{s}_k]$ be its convex hull.[7] Player k's private signal indicates that $\omega = 1$ is no less likely than $\omega = 0$ iff $s_k \geq 0$.

Let $D(k) \subset \{1, \ldots, k-1\}$ 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 lower index than another simply to mean that the former moves no later later than the latter. Let $ID(k) \subset \{1, \ldots, k-1\}$ be the subset of Player k's predecessors whom k *indirectly* observes: $l \in ID(k)$ iff there exist some path of players $k_1, k_2, \ldots, k_L$ such that $k_1 \in D(k), k_2 \in D(k_1), \ldots, k_L \in D(k_{L-1}), l \in D(k_L)$. Of course, there may be more than one path by which one player indirectly observes another,

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

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

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 if Player k indirectly observes Player j, she need not directly observe him.

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 with information I_k does this by choosing $\alpha_k = \mathbb{E}[\omega|I_k] = \Pr[\omega = 1|I_k]$, namely by choosing the action that coincides with her posteriors that $\omega = 1$. Any player who observes Player k infers Player k's updated beliefs but cannot necessarily back out 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 $\mathcal{N} = \{\{1,2,\ldots\}, \{D(1), D(2), \ldots\}\}$ as an observation structure, consisting of the players $\{1, 2\ldots\}$ and their respective sets of directly observed predecessors, which define their sets of indirectly-observed predecessors.[8] Given the observation structure $\mathcal{N}$, define its k-truncation $\mathcal{N}^k := \{\{1,2,\ldots k\}, \{D(1), D(2), \ldots D(k)\}\}$ to comprise its first k players as well as their observations sets. Throughout we assume that the observation structure is commonly known to the players.

Although a player may directly observe a large number of predecessors, many of these observations turn out to be redundant. For instance, in the classical, single-file structure with rational players choosing their actions from the continuum (Lee (1993)), no player who observes her immediate predecessor gains any useful information by observing any other predecessor, for her immediate predecessor has already correctly combined all earlier signals with his own. Similarly, suppose that each player through Player n observed all of her pre-

[8] $\mathcal{N}$ can be viewed as a directed network with players as nodes, direct observations as directed links, etc. (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.

decessors whose index number shares the parity of her own index number but no predecessor whose index number does not, namely even- and odd-numbered players move single-file in parallel. A Player $n+1$ who observed *all* of her predecessors could restrict attention to her immediate two predecessors, even and odd.

In other settings, however, players' immediate predecessors do not provide "sufficient statistics" for earlier movers indirectly observed. In our multi-file example of Section 2, players in the third period care about more than second-period actions between they wish 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* $\mathcal{N}$ *form a* diamond *if* $i \in ID(j) \cap ID(k)$, $j \notin ID(k)$, $k \notin ID(j)$, *and* $\{j,k\} \subset D(l)$.

We refer to the diamond 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.

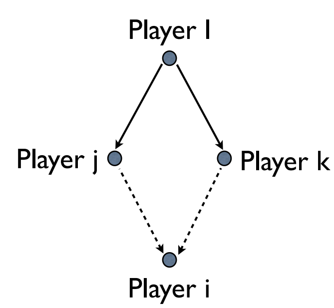

Figure 1: A Diamond
Solid Arrows Denote Direct Observation
Dashed Arrow Denote Indirect Observation

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.

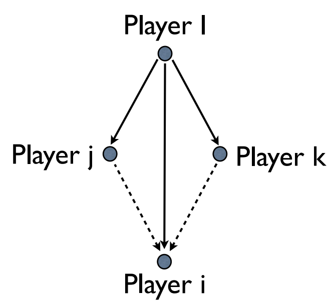

Figure 2: A Shield
Solid Arrows Denote Direct Observation
Dashed Arrows Denote Indirect Observation

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 *(II) if for each* $(s_1, \ldots, s_{k-1}) \in \times_{j<k} S_j$ *and each* $s_k \in S_k$,

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

Otherwise, Player k *achieves* partial inference *(PI).*

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. 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. 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.

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. We say that *behavior in the observation structure* $\mathcal{N}$ *is impartial* if each player in $\mathcal{N}$ achieves impartial inference.

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

Definition 4 *1. Player k* imitates *Player $j \in D(k)$ 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*

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

2. *Player* k anti-imitates *Player* $j \in D(k)$ *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. (Imitation is just the opposite.) 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 3 anti-imitates Player 1 if actions AB provide a stronger signal in favor of state B than actions BB, when it is also the case that action AA provides a stronger signal in favor of the state A than actions BA. A player who anti-imitates a predecessor always forms beliefs that tilt against that predecessor's.

Our main result is that 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.*

1. *If* $\mathcal{N}$ *contains a shield, then some player anti-imitates another.*

2. If some player in $\mathcal{N}$ anti-imitates another, then $\mathcal{N}$ contains a shield.

Clearly, in Figure 2 Player l must subtract off the action of Player i, already incorporated into both Player j and Player k's actions, in order to achieve impartial inference. The work in proving the first statement of the Proposition goes into showing any impartial-inference setting that contains a shield, regardless of the rest of the configuration, includes at least one player who anti-imitates another.

In an observation structure where all players achieve impartial inference, a player cannot imitate two predecessors who both observe an earlier, common predecessor without also anti-imitating some other predecessor.

Proposition 2 *Assume that every player in the observation structure $\mathcal{N}$ achieves impartial inference. If Player l imitates some Players i and $j \neq i$ s.t. $ID(i) \cap ID(j) \neq \emptyset$, then Player l must anti-imitate some Player k.*

Surely, it is integral to the process of social learning that different people learn from some of the same sources. Proposition 2 suggests that social learning either consists of imitating only a single predecessor or includes anti-imitation. Pluralistic social learning almost necessarily involves anti-imitation.

Not only does rational social learning often require that certain players anti-imitate others, but it 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 that are both contrary to his private signal and all the actions he observes. Two definitions help us establish some surprising results to this effect.

Definition 5 *Player k's* observational beliefs *in observation structure $\mathcal{N}$ following action profile $(a_1, a_2, \ldots a_{k-1})$ are*

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

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 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 6 *The path of play* $(a_1, a_2, \ldots, a_k)$ *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 favoring state $\omega = 0$, she attaches higher probability to state $\omega = 1$, or *vice versa*.

Proposition 3 *Assume that every player in the observation structure* $\mathcal{N}$ *achieves impartial inference.*

1. *If there is contrarian play, then* $\mathcal{N}$ *contains a shield.*

2. *If* $\mathcal{N}$ *contains a shield and all players' private beliefs are drawn from the density* $f^{(\omega)}$ *that is everywhere positive on* $[\underline{s}, \overline{s}]$, *then with positive probability there is contrarian play.*

Because contrarian play relies upon some player anti-imitating another, the first statement follows as a corollary to Proposition 1. For an intuition behind the second statement, note 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 realisation of any signal or action, then 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—Player k's observational beliefs must eventually take on the sign opposite to that of a_j.

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 not the current or future rounds. We label the n movers in round t Players $t^1, \ldots t^n$. 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.

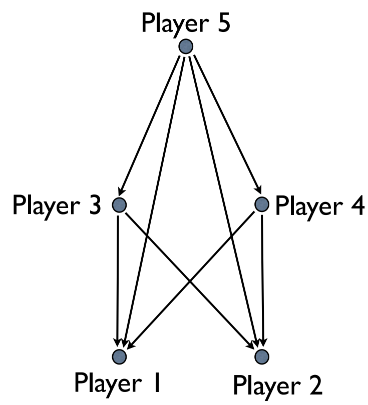

Figure 3: The First Five "Double-File" Movers
Solid Arrows Denote Direct Observation

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

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.*

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,

$$A_t = S_t + n\sum_{i=1}^{t-1}(-1)^{i-1}(n-1)^{i-1}A_{t-i}.^9$$

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. Roughly speaking, approximately half of social learning in this setting is anti-imitative!

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 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, like in the example of Section 2 where after observing three players' ninety-percent confident in $\omega = 1$, round-two players continue to assign ten-percent probability to $\omega = 0$.

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,

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

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.[10]

In some environments, people may be unable to observe their most immediate predecessors. On 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.

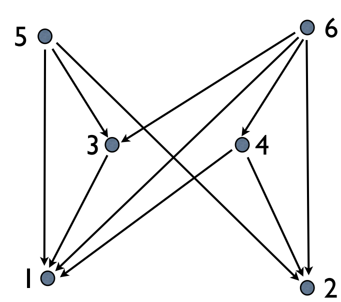

Figure 4: Single-Period Lags
Solid Arrows Denote Direct Observation

[10]This 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 first shield in this observation structure consists of Players 1, 3, 4 and 6. From the figure, it can be seen 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 $(0, 1, 1, 0, -1, -1, 0, 1, 1, \ldots)$ on actions $(a_{k-1}, a_{k-2}, \ldots)$, 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 Player 6 assigns 55% probability to $\omega = 0$, contrarian play.[11] 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 $\omega = 0$.[12]

4 Near-Certain Contrarian Play

In the last section, we showed how in an observation structure with impartial inference, shields are a necessary and sufficient sufficient condition for anti-imitation as well as for contrarian behavior to happen with positive probability. In this section, we describe an observation structure where contrarian play happens not merely with positive probability but with arbitrarily high probability.

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.

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

[12] Notice 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.

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, we consider players with bounded private beliefs. In particular, each Betty receives a draw from a distribution of binary signals $s \in \{0,1\}$ parametrized 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 $p \to \frac{1}{2}$. 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 each $\varepsilon > 0$, there exists a signal structure $p > \frac{1}{2}$ under which with probability at least $1-\varepsilon$ there exists some round t when all play is contrarian.*[13]

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 infer a net of 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 $p \to \frac{1}{2}$, contrarian play happens arbitrarily many times.

Contrarian play also 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

[13] More formally, denote the actions of the five players in round t by $(a_t^0, a_t^1, \ldots, a_t^4)$ as in the multi-file example of Section 4. Proposition 4 states that for each $\epsilon > 0$, we can find a signal structure such that with probability at least $1-\epsilon$ there exists some round t' such that the five paths of play

$$(a_1^0, a_1^1, \ldots, a_1^4; a_2^0, a_2^1, \ldots, a_2^4; \ldots; a_{t'}^0), (a_1^0, a_1^1, \ldots, a_1^4; a_2^0, a_2^1, \ldots, a_2^4; \ldots; a_{t'}^0, a_{t'}^1),$$
$$\ldots, (a_1^0, a_1^1, \ldots, a_1^4; a_2^0, a_2^1, \ldots, a_2^4; \ldots; a_{t'}^0, a_{t'}^1, \ldots, a_{t'}^4)$$

are all contrarian.

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.

5 Conclusion

Outside of the impartial-inference setting we formally analyze, anti-imitation takes 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 occurs. But the existence of the forms of anti-imitative and contrarian play that we identify do not depend upon details of our environment. (Indeed, the examples in Section 2 and Appendix A demonstrate that anti-imitation plays out in coarse-actions settings.) Many simple, natural observational structures lead rational players to anti-imitate by requiring to subtract off sources of correlation in order to fully extract information from all observed actions. If observed recent actions provide 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.

The principles we establish here are separable from the issues researchers have traditionally emphasized and debated 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 (2010) characterize necessary and sufficient conditions in general observation structures, including those we consider here, that give rise to asymptotically complete social learning. Yet neither this paper nor related ones by Banerjee and Fudenberg

(2004) and Smith and Sørensen (2008), which share the feature of Acemoglu et al. (2010) 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.

Our results hold less interest if instead we live in world where everybody learns others' *private* information instead of their "posterior" beliefs. If friends or an online community report their experiences with a restaurant separate from their beliefs about its quality, or if they divulge their private reasons for making an investment wise, independent of their overall assessment of its wisdom, then people optimally ignore others' behavior or stated beliefs in favor of their revealed private information. If we only wish to study societies where observation of others' beliefs—via actions or words—play a small role, then the days of studying observational learning should be behind us. Economists should not care much about anti-imitative behavior because they should not care about imitative behavior.

But the logic of redundancy in fact does play out as dramatically—or more dramatically—when people learning *both* from the actions of others and from their signals. 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 the café's quality. Back in her hotel room, she scours the Internet for recommended restaurants and learns that the local newspaper recently awarded Café M its maximum three stars. Assuming the tourist believes locals read the local newspaper, and everyone has common preferences, does the tourist fancy Café M more before or after reading the review?

We address this question through a simple observation structure containing only four players: Player 1 is the local paper, Players 2 and 3 locals, and Player 4 the tourist. By reading the review, Player 4 goes from being the last player in a diamond that is not a shield to the last player in a shield (i.e., moving from Player l in Figure 1 to Player l in Figure 2).

Proposition 5 *Let* $\mathcal{N} = \{\{1,2,3,4\};\{D(1)=\emptyset, D(2)=D(3)=\{1\}, D(4)=\{2,3\}\}\}$ *and* $\hat{\mathcal{N}} = \{\{1,2,3,4\};\{D(1)=\emptyset, D(2)=D(3)=\{1\}, D(4)=\{1,2,3\}\}\}$*. Then* $\hat{a}_4(a_1,a_2,a_3;s_4) \leq a_4(a_2,a_3;s_4)$ *iff* $a_1 \geq \mathbb{E}[S_1|S_1+S_2=a_2, S_1+S_3=a_3, S_4=s_4]$*.*

Following a review more positive than expected given her observations and private information ($a_1 \geq \mathbb{E}[S_1|S_1+S_2=a_2, S_1+S_3=a_3, S_4=s_4]$), the tourist updates her beliefs to form a more negative assessment of the café ($\hat{a}_4 \leq a_4$). The intuition behind the result closely matches that of Proposition 1: the more positive the newspaper's review (i.e., Player 1's action), the more negative the information content in Players 2 and 3's actions.

In this paper, we demonstrate that a simple form of anti-imitation follows from rationality due to the intrinsic logic of redundancy in social actions. Of course, anti-imitation can occur in other contexts for different reasons. It is likely to arise in many situations due to heterogeneity in tastes or in interpretation of signals. In the model of Smith and Sørensen (2000), heterogeneity of preferences, in particular opposing tastes (some people seek out spicy foods, while others avoid them), can produce anti-imitation. Park and Sabourian (2011) analyze a model with more than two states of the world, where whether a given private signal favors a high or low state depends upon prior probabilities over the states. Whenever high priors render the information content of a player's signal negative, whereas low priors render it positive, the player engages in contrarian play. Their logic that the meaning of private signals depends upon priors is unrelated to the forces at play in our model, and does not apply to any two-state model such as those in this paper. Indeed, unlike our results on contrarian play, no trader unendowed with private information in their model would engage in contrarian play.

Other papers include anti-imitation in more conventional settings like those explored here, where both tastes and interpretation of signals are shared by all involved. Callender and Hörner (2009) illustrate why one might observe anti-imitation related to the famous *overturning principle* of Smith and Sørensen (2000), which 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 predecessor. 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. When actions are coarse, players cannot extract their immediate predecessors' actions from their beliefs, and the actions of previous people may allow for fine-tuning beliefs. Whether greater belief by earlier players increases or decreases the confidence of later players depends upon details of the environment. 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 qualitative predictions cannot arise in single-file models like Callender and Hörner (2009). Most notably, these models cannot give rise to contrarian behavior. 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 why people anti-imitate, so too of course there may be many reasons why people imitate others. They may do so to conform (e.g., Bernheim (1994)), to build reputations (e.g., Scharfstein and Stein (1990)), or to benefit from safety in numbers. In all cases, imitation might prevail for reasons that are not about inference and learning, and so we may observe imitation for reasons separate from the reasons emphasized in the observational-learning literature. And, of course, the basic logic of the herding literature beginning with Banerjee (1992) and Bikhchandani et al. (1992) is always relevant: it leads to imitation of at least some players in all environments, and leads solely to imitation in certain types of single-file environments. Even so, 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 is rational, or rather arises from the many other theories of rationality. 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. As Eyster and Rabin (2010) have shown, in many settings rational and irrational models of social learning lead to very different conclusions about its efficiency. Empirical researchers could instead focus on identifying cases where rational social learning predicts forms of anti-imitation to design more powerful tests to separate rational from irrational observational learning.

References

Acemoglu, Daron, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar, "Bayesian Learning in Social Networks," Mimeo 2010.

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.

DeMarzo, Peter, Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion Bias, Social Influence, and Unidimensional Opinions," *Quarterly Journal of Economics*, 2003, *118* (3), 909–968.

Eyster, Erik and Matthew Rabin, "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.

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

Park, Andreas and Hamid Sabourian, "Herding and Contrarian Behavior in Financial Markets," *Econometrica*, 2011, *79* (4), 973–1026.

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.

6 Appendix A

The central logic of our first example, and indeed of the entire paper, does not depend at all on continuous action spaces and applies equally to discrete-action settings. A straightforward discretisation of our first example gives exactly the same inferences.

Suppose that there five possible signals: one is uninformative; two favor $\omega = 1$, with likelihood ratios $9 : 1$ and $81 : 1$; and two signals favor $\omega = 0$, again with likelihood ratios $9 : 1$ and $81 : 1$. There are five possible actions $\{0.01, 0.1, 0.5, 0.9, 0.99\}$, which are optimal for the following ranges of beliefs that $\omega = 1$: less than $\frac{1}{80}$; between $\frac{1}{80}$ and $\frac{1}{8}$; between $\frac{1}{8}$ and $\frac{7}{8}$; between $\frac{7}{8}$ and $\frac{79}{80}$; and above $\frac{79}{80}$, respectively.

Suppose that the three first-period players choose $(0.9, 0.9, 0.9)$, followed by the three second-period players also choosing $(0.9, 0.9, 0.9)$. Without any private information of his own, which state would such a rational observer deem more likely? Once again, the first three players' actions reveal signals $9:1$ in favor of $\omega = 1$. Any player in the second period receiving a $9:1$ signal in favor of $\omega = 0$, or any signal more favorable to $\omega = 1$, would form beliefs of least $81:1$ in favor of $\omega = 1$, in which case 0.99 would have been optimal, a contradiction. Hence, every second-period mover must have received an $81:1$ signal in favor of $\omega = 0$. So, once again an uninformed observer seeing six people ninety-percent convinced in $\omega = 1$ would be more than 99% convinced in $\omega = 0$. Had the first-period actions been $(0.5, 0.5, 0.5)$ instead, followed once more by second-period actions $(0.9, 0.9, 0.9)$, then yet again a rational observer un-endowed with private information would conclude that $\omega = 1$ is more that 99% likely. Hence, neither anti-imitation nor contrarian play depend upon having continuous action spaces. Nor does either depend upon whether the signals are drawn from continuous or discrete distributions. The first example in the main text assumed made no such assumption about distributions, while here we use discrete signals merely to avoid integrating out over different signals producing the same action.

7 Appendix B

We begin by introducing some concepts and notation useful for the proofs. 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, $D(1) = \overline{D}(1) = \varnothing$ and for each $k \geq 2$, $D(k) = \{1, \ldots, k-1\}$ and $\overline{D}(k) = \{k-1\}$. When two players move every round, observing (only) all players who moved in all previous rounds, $D(1) = D(2) = \varnothing$, and for $l \geq 1$, $D(2l+1) = D(2l+2) = \{1, \ldots, 2l\}$, while $\overline{D}(2l+1) = \overline{D}(2l+2) = \{2l-1, 2l\}$.

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, $ID(k) = \overline{D}(k) \sqcup \left(\cup_{j \in \overline{D}(k)} ID(j)\right)$.*

Proof.

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

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. ■

Lemma 2 *If behavior in $\mathcal{N}$ is impartial, then for each player l, there exist unique coefficients α_i^l such that $a_l = \sum_{i\in D(l)} \alpha_i^l a_i + s_l$.*

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

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

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

Lemma 3 implies that in any observation structure without diamonds where all players achieve impartial inference, all behavior is imitative.[14]

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

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

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

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

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

■

Proof of Proposition 1. We begin by proving the second statement. 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 $A = \{i : (i,j,k,l) \text{ is a diamond}, j,k \in \overline{D}(l)\}$, which is non-empty by Lemma 3. If $i \in A$, then we have a shield, so suppose otherwise. Lemma 1 implies that there exists $h \in \overline{D}(l)$ for whom $i \in ID(h)$; if $h' \in \overline{D}(l)$ 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 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 II,

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

in contradiction to l achieving II. 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 II. Hence, $\mathcal{N}$ must contain a shield.

To prove the first statement, 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 $i \in ID(\hat{l})$, II 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 $i \in ID(j) \cap ID(k)$, II implies that $\frac{da_j}{ds_i} = \frac{da_k}{ds_i} = 1$. Using $\frac{da_l}{ds_i} \in \{0,1\}$ from II 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 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

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

(by the closure of $\mathbb{Z}$ under addition and multiplication), which contradicts Player k achieving II. 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 II. Because all players achieve II, we must have that for some $k \in D(l)$, $\frac{\partial a_l}{\partial a_k} < 0$ as desired. ■

Proof of Proposition 3.

1. From Lemma 2, Player l's action can be written uniquely as $a_l = \sum_{i \in D(l)} \alpha_i^l a_i + s_l$. If for each $i \in D(l)$, $\alpha_i^l \geq 0$, and $sgn(a_i) = sgn(s_l)$, then $sgn(a_l) = sgn(s_l)$, and therefore the path of play $(a_1, \ldots, a_l)l$ 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.

2. Wlog let $\underline{s} \leq -\overline{s}$ 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 $a_i \in \left(\overline{s} - \epsilon, \overline{s} - \frac{\epsilon}{2}\right)$. Non-triviality of private beliefs permits this $\forall j$ such that $D(j) = \emptyset$; for other j, from Lemma 3,

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

Because $\underline{s} \leq -\overline{s}$, there exists a positive-measure set of signals s_j for which $a_j \in (0, \epsilon)$.

Now note that $o_{\hat{l}} \leq (\hat{l} - 2)\frac{\epsilon}{2(\hat{l}-1)} - \left(\overline{s} - \frac{\epsilon}{2}\right) < 0$ for ϵ small. ■

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 and $\Delta X(t) = X(t + 1) - X(t)$. As explained in the main text, if there exists some period τ in which $X(\tau) = +2$ and $\Delta X(\tau) = -4$, then we have strong contrarian play in that period, where every player plays not only against both her observational beliefs and private signal.

$\langle X(t) \rangle$ is a irreducible Markov chain on the countable state space of even integers. As such, all states are either persistent or transient. When $q = \frac{1}{2}$, because $\langle X(t) \rangle$ is a martingale (and non-simple random walk), all states are persistent. For $q < \frac{1}{2}$, all states are transient as the stochastic process drifts to $+\infty$. Let $p_{22}(q)$ be the probability that $\langle X(t) \rangle$ returning to the state $+2$ when starting from the state $+2$ without getting four $s = 0$ signals, when the signal structure is q. That is, define

$$\begin{aligned} p_{22}^{(n)}(q) &= \Pr[X_1 \neq 2, X_2 \neq 2, \ldots, X_{n-1} \neq 2, X_n = 2 | X(0) = 2, \Delta X(t) \neq -4], \\ p_{22}(q) &= \Pr\left[\sum_{n=1}^{\infty} p_{22}^{(n)}(q)\right] \end{aligned}$$

Because $p_{22}(q)$ is continuous in q with $p_{22}(1) = 1$ from above, for each $\epsilon > 0, \exists \overline{q}$ such that $\forall q \in \left(\frac{1}{2}, \overline{q}\right)$, $p_{22}(q) > 1 - \epsilon$. Therefore, for $q \in \left(\frac{1}{2}, \overline{q}\right)$, the probability that starting from the state $+2$, there occurs some period τ in which $X(\tau) = +2$ and $\Delta X(\tau) = -4$ is at least

$$q^4 + q^4\left(1 - q^4\right)(1-\epsilon) + q^4\left(1 - q^4\right)^2(1-\epsilon)^2 + \ldots = q^4 \sum_{k=0}^{\infty}\left(1 - q^4\right)^k (1-\epsilon)^k = \frac{q^4}{1 - (1 - q^4)(1 - \epsilon)} \to 1,$$

as $\epsilon \to 0$. The proof then follows from the fact that for q close to $\frac{1}{2}$, the probability that $X(t)$ reaches $+2$ also approaches one.

■

Proof of Proposition 5. Note that

$$
\begin{aligned}
& a_4(a_2, a_3; s_4) = a_2 + a_3 - \mathbb{E}[a_1|a_2, a_3] + s_4 \leq \hat{a}_4(a_1, a_2, a_3; s_4) = a_1 + a_2 + a_3 + s_4 \\
\iff \quad & (s_1 + s_2) + (s_1 + s_3) - \mathbb{E}[S_1|s_1 + s_2, s_1 + s_3, s_4] + s_4 \leq s_1 + s_2 + s_3 + s_4 \\
\iff \quad & a_1 = s_1 \leq E[S_1|s_1 + s_2, s_1 + s_3, s_4]
\end{aligned}
$$

■

In the main text, we assume that players achieve impartial inference. A sufficient condition for 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.

Clearly, Player 1 achieves II by choosing $a_1 = s_1$. We claim that if all Players $i \in \{1, \ldots, k-1\}$ achieve II, then so too does Player k. Define

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

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 $i \in \{1, \ldots, k-1\}$ achieving II and Lemma 1 imply that for each $j \in ID(k)$, $\beta_j^1 \geq 1$: $j \in M^1(k) \sqcup U^1(k)$. First, notice that $\forall i \in U^1(k), \forall j \in M^1(k)$, $i \notin ID(j)$; otherwise, because $\forall j \in M^1(k), \exists k_1, k_2 \in \overline{D}(k)$ 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 $i \in M^1(k)$, 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 II through a_k^1. Suppose that $\emptyset \neq M^1(k) =: \{m_1, m_2, \ldots, m_N\}$, where $m_1 < m_2 < \ldots < m_N$. Because each $m_i, i \in \{1, \ldots, N\}$, 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

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

$U^2(k) := \{j \in ID(k) : \beta_j^2 = 1\}$ and $M^2(k) := \{j \in ID(k) : \beta_j^2 \neq 1\}$. By construction, for each $l \in ID(k), l \geq m_N$, $\beta_l^2 = 1$. Because $\forall i \in U^1(k), i \notin ID(m_N)$ from above, $\frac{da_k^1}{ds_i} = \frac{da_k^2}{ds_i} = 1$ 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

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

$U^3(k) := \{j \in ID(k) : \beta_j^3 = 1\}$ 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 II by playing $a_k^{\hat{j}}$. ■