Herding with Collective Preferences*

S. Nageeb Ali† Navin Kartik‡

February 21, 2011

Abstract

This paper studies a simple model of observational learning where agents care not only about the information of others but also about their actions. We show that despite complex strategic considerations that arise from forward-looking incentives, herd behavior can arise in equilibrium. The model encompasses applications such as sequential elections, public good contributions, and leadership charitable giving.

Keywords: Social learning, observational learning, herd behavior, payoff interdependence, sequential voting, momentum.

JEL Classication: D7, D8

*This paper supersedes our working paper "Social Learning in Elections." We thank Daron Acemoglu, Susan Athey, Doug Bernheim, Vince Crawford, Eddie Dekel, Jeff Ely, Nir Jaimovich, Fahad Khalil, Jon Levin, Mark Machina, Marc Meredith, Roger Myerson, Marco Ottaviani, Joel Sobel, Lin Zhang, anonymous referees, and a number of seminar and conference audiences for helpful comments and discussions. Ali thanks the Institute for Humane Studies, the John M. Olin Foundation, and the National Science Foundation for support. Kartik thanks the National Science Foundation for support, and the Institute for Advanced Study at Princeton for hospitality and support. We appreciate proof-reading by Uliana Loginova.

†Department of Economics, University of California, San Diego. Email: snali@ucsd.edu.

‡Department of Economics, Columbia University. Email: nkartik@gmail.com.

Contents

> "...when New Yorkers go to vote next Tuesday, they cannot help but be influenced by Kerry's victories in Wisconsin last week. Surely those Wisconsinites knew something, and if so many of them voted for Kerry, then he must be a decent candidate."
>
> — Duncan Watts in *Slate* Magazine, February 24, 2004

1 Introduction

Many collective choice mechanisms take place over time, revealing the choices of some agents to others. A prominent example is a sequential election such as the U.S. presidential primary system and roll-call voting procedures used by city councils, Congressional bodies, and many organizations. The U.S. Primaries sequentially aggregate a number of different individual choices, not only those by voters but also campaign donors and political leaders who decide over time which candidate to endorse. Similarly, contributions for public goods are often raised gradually; for example, leaders typically decide how to direct their seed money while anticipating the choices of followers (List and Lucking-Reiley, 2002; Andreoni, 2006). Some settings that are not about collective choice *per se* also share the feature that agents who make choices over time may benefit from other agents making good decisions: for example, when adopting new technologies, farmers in developing countries may care about others' choices because of risk-sharing arrangements in their community.

In such environments, agents are often only partially informed about the merits of each alternative and benefit from aggregating their information. When individual decisions are made sequentially rather than simultaneously, there is the potential for signaling or information revelation by early actors and, in turn, learning or herding by later actors. The vast literature on observational learning, initiated by Banerjee (1992) and Bikhchandani et al. (1992), has offered a useful framework and a number of results about information aggregation when agents learn from the choices of others. However, there is a challenge in porting these insights to the kinds of problems mentioned above: with few exceptions (which we discuss below), most existing models assume that an agent's payoff is independent of the action of others, and hence do not apply to any setting with payoff interdependence.

This paper identifies a class of environments with payoff interdependence in which herd behavior can arise. As in the canonical framework, we study a sequence of agents who choose between two actions in a fixed sequential order, observing the choices of all prior agents. Our departure is to allow an agent's payoff to depend on the profile of actions of any subset of all agents—as opposed to only her own action—and, as usual, a binary state of the world. We require that in the high (resp., low) state, each agent's payoff weakly increase (resp. decrease) in the entire action profile, but agents may differ in how they care about the choices of others. This monotonicity assumption corresponds to each agent having a (weak) preference for others to make the "right" decision, which is why we refer to the setting as one of *collective preferences*. In addition to the payoff interdependence, there is a standard information externality because each agent privately observes a binary signal that is only partially informative about the state.

When an agent cares about the actions of future individuals, she faces potentially complex forward-looking incentives. For instance, in a sequential electoral mechanism, an instrumentally-motivated voter recognizes that her decision affects her payoff only when it changes the electoral outcome. Accordingly, she must account for both the informational content of being pivotal and that her vote could affect the way in which subsequent voters behave. Crucial to these considerations are her beliefs about the strategies followed by subsequent voters. Similarly, a technology adopter who values all individuals making the correct choice may be tempted to be contrarian so as to enhance informational efficiency, for instance if she believes that choosing what currently appears the better technology is very likely to trigger a herd and thus prevent aggregating future agents' information.

Our main result is that herding can emerge as the outcome of strategic behavior even when individuals have such forward-looking incentives. The equilibrium we characterize, *Sincere Behavior*, takes a surprisingly simple and tractable form: each agent uses all currently available information—the prior, observed history of actions, and private signal—to form an expectation about the state and then selects her optimal action *as if* she only cares about her own action. In other words, sincere behavior is identical to how agents would behave if their payoffs were independent as in the standard framework. Consequently, as in the standard setting, this equilibrium generates a herd once the informational content of agents' choices swamp the private information of any individual. Unlike the standard setting, however, optimality of sincere behavior in our model is genuinely a strategic equilibrium phenomenon: if an agent does not expect future agents to act sincerely, then it typically would no longer be optimal for her to act sincerely either.

It is unexpected that sincere behavior is an equilibrium despite payoff interdependence, so let us highlight the assumptions that are important for the result. First is the nature of payoff interdependence we study. Our model is not one of (positive) network externalities, in which conformity is intrinsically valued, nor of congestion effects, in which conformity is intrinsically disliked. Rather, we have collective preferences in the sense mentioned earlier. This captures the essence of some applications, including collective choice problems, and also allows us to focus on the signaling or information revelation motive that intuitively makes herd behavior difficult to sustain with payoff interdependence. Second, we assume that there exists a belief about the relative likelihood of each state such that when any agent holds this belief, she is indifferent between all action profiles. While this is clearly restrictive, we discuss examples that fit the framework in Section 2. Finally, our model has a commonly known precision of information.

Besides its intrinsic interest, the equilibrium characterization shows that standard insights from the literature on observational learning, such as fragility of mass behavior, can be relevant even with payoff interdependence and forward-looking incentives. It should be noted, however, that we only characterize one equilibrium of the model. Nevertheless, there are at least three reasons this finding is of interest. First, the sincere-behavior equilibrium exists for the entire range of payoff specifications that fit our model, whereas little is known generally about other equilibria. Second, even in the special cases where more informationally-efficient equilibria are known to exist,[1] the possibility of rational herding presents a cautionary note on the scope for successful information aggregation. Third, the equilibrium is compatible with naive agents who don't condition on complicated "pivotal" considerations.

Although the payoff interdependence in our model encompasses multiple applications, we were motivated initially by sequential voting, which we examined in more detail in a prior version of this paper. Sequential elections, such as the U.S. Presidential Primaries, are believed to have *momentum* effects because later voters tend to follow the choices of earlier voters (see Knight and Schiff 2010 for a recent empirical study). One explanation that has been suggested for momentum is that of *cue-taking*, where voters learn the relative merits of each alternative by observing earlier votes and then vote accordingly (Bartels, 1988). This idea mirrors the sincere-behavior equilibrium characterized here. An implication is that even though a large electorate collectively has almost full information, the votes of early voters sets the course for the election and suppresses later voters' information. This provides some support for the view that sequential elections give early voters too much influence (Palmer, 1997) and can be senstive to shocks that directly affect only relatively few voters. Indeed, in our model, an early voter whose vote counts for less than those of other individual voters—in the extreme, even when her vote is a straw vote—can nevertheless play a powerful role through her influence on subsequent voters.

There are a few other papers that also study observational learning with payoff interdependence, some specifically in the context of sequential voting. In unpublished work, Wit (1997) and Fey (2000) analyze a version of sincere behavior in simple majority rule elections with instrumen-

[1] We discuss this in more detail later, but see Wit (1997), Fey (2000), and Dekel and Piccione (2000) for instrumental-voting elections and Smith and Sorensen (2008) for a model of altruistic agents.

tal voting, although their focus is largely on how non-sincere equilibria can be selected through belief-based refinements. Since we deal with a much broader class of environments—not only elections with different voting rules and voter motivations, but also non-electoral applications—our approach is entirely different. Callander (2007) studies herding in an infinite voter model in which each voter intrinsically likes to vote for the winner, in addition to having the usual common-value component of preferences. Dekel and Piccione (2000) show that symmetric binary-agenda sequential elections with instrumentally-motivated voters have equilibria in which voting behavior is independent of history. Such equilibria replicate the outcomes of simultaneous voting games and thereby attain informational efficiency in large elections (Feddersen and Pesendorfer, 1997). Outside of sequential elections, observational learning has also been studied in coordination problems (Dasgupta, 2000), common-value auctions (Neeman and Orosel, 1999), settings with network externalities (Choi, 1997), and when agents partially internalize the welfare of future agents (Smith and Sorensen, 2008).

2 Preliminaries

2.1 Model

There is a finite population of agents indexed $1, \ldots, n$, who take decisions in a roll-call sequence one after another. When it is agent i's turn to act, she chooses an action $a_i \in \{0, 1\}$, having observed the history of prior actions, $a_1, \ldots, a_{i-1}$. There is an unknown state of the world, $\omega \in \{0, 1\}$, drawn from a common prior distribution with $\pi := \Pr(\omega = 1) \geq \frac{1}{2}$. Before choosing her action, each agent i receives a private signal, $s_i \in \{0, 1\}$. Conditional on ω, signals are drawn independently from a Bernoulli distribution with precision $\gamma_\omega \in (0, 1)$, i.e., for each ω, $\Pr(s_i = \omega|\omega) = \gamma_\omega$. Note that signal precisions could differ across states. We assume the strict monotone likelihood ratio property, which translates here as $\gamma_0 + \gamma_1 > 1$. In addition to her private signal about the state, each agent also has a preference type $t_i \in T$, where T is any non-empty set. The vector of preference types $(t_1, \ldots, t_n)$ is drawn from some distribution τ on T^n that is independent from the state or signals but otherwise arbitrary (in particular, τ may be correlated across players).[2]

A player i's von-Neumann Morgenstern utility depends on her preference type, t_i, the vector of

[2] We do not specify the observability structure about preference types because it plays no role in our analysis, hence any structure will do. For concreteness, one may like to fix ideas on preference types being privately observed.

actions, $\mathbf{a} = (a_1, \ldots, a_n)$, and the state ω. It is represented by a function $u_i(\mathbf{a}, t_i, \omega)$. To formulate assumptions on preferences, define the following notation: given any $\mathbf{a}_{-i} := (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$, let $\mathbf{a}^+_{-i} := (a_1, \ldots, a_{i-1}, 1, a_{i+1}, \ldots, a_n)$ and $\mathbf{a}^-_{-i} := (a_1, \ldots, a_{i-1}, 0, a_{i+1}, \ldots, a_n)$. We maintain the following two assumptions throughout:

(A1) For all i and t, $u_i(\cdot, t, 1)$ is non-decreasing and $u_i(\cdot, t, 0)$ is non-increasing.[3]

(A2) There exists $c \in (0,1)$ such that for any i, j, $\mathbf{a}_{-j}$, and t,

$$c[u_i(\mathbf{a}^+_{-j}, t, 1) - u_i(\mathbf{a}^-_{-j}, t, 1)] = (1-c)\left[u_i(\mathbf{a}^-_{-j}, t, 0) - u_i(\mathbf{a}^+_{-j}, t, 0)\right].$$

Assumption (A1) is a monotonicity property that captures *collective preferences*: each agent weakly prefers higher action profiles if the state is 1 and weakly prefers lower actions profiles if the state is 0.[4] Generally then, agents have *state-dependent* preferences over action profiles, which is what makes this a problem of social learning. (A1) is very different from and indeed precludes network externalities, because (A1) implies that for any fixed profile of actions $\mathbf{a}_{-j}$, agent i's preference over j's action (where j can be the same as i) does not depend on $\mathbf{a}_{-j}$, rather it only depends on the state. Note also that because of collective preferences, there is a sense in which each agent would like to convey information about the state to subsequent agents.

Assumption (A2) says that for any player i, given any profile of actions for all players excluding j (where j can be the same as or different from i), the ratio between states of the gain in utility from j changing her action in the preferred direction is a constant. Given (A1), (A2) is necessary and sufficient for all types to share a common belief threshold, c, such that for any fixed profile of actions of any $n-1$ players, they would like the remaining player to take action $a = 1$ if and only if the probability of state 1 is at least c. To see this, note that if agent i of type t ascribes probability μ to state 1, then for any profile of actions $\mathbf{a}_{-j}$, the difference in i's expected utility between $a_j = 1$ and $a_j = 0$ is

$$\begin{aligned} & \mu[u_i(\mathbf{a}^+_{-j}, t, 1) - u_i(\mathbf{a}^-_{-j}, t, 1)] + (1-\mu)\left[u_i(\mathbf{a}^+_{-j}, t, 0) - u_i(\mathbf{a}^-_{-j}, t, 0)\right] \\ = \; & [u_i(\mathbf{a}^+_{-j}, t, 1) - u_i(\mathbf{a}^-_{-j}, t, 1)]\left(\frac{\mu - c}{(1-\mu)(1-c)}\right), \end{aligned} \tag{1}$$

where the equality follows from Assumption (A2). As the term in square brackets in (1) is non-

[3]The order on the action profile space, $\{0,1\}^n$, is standard vector order: $\mathbf{a} > \mathbf{a}'$ if $a_i \geq a'_i$ for all i with strict inequality for some i.

[4]Such a monotonicity assumption is common in the observational learning literature without payoff interdependence, even in models that allow considerable heterogeneity in agents' preferences (e.g. Goeree et al., 2006). However, see the discussion of confounded learning in Smith and Sorensen (2000) for an exception.

negative by Assumption (A1), agent i prefers that agent j choose action 1 if and only if μ exceeds c (with indifference if $\mu = c$).[5] Note, however, that we are holding fixed the profile of actions $\mathbf{a}_{-j}$ here, just as in the statement of (A2). Since the game is dynamic, an agent's action may affect the actions taken by subsequent players, which can matter significantly for the agent. This is the heart of the strategic issues that arise in the model. Plainly, (A2) is a restrictive assumption, but it is satisfied in common-value elections and some other applications, as we explain in the following sub-section.

All aspects of the model except the realization of private signals and possibly the preference types (cf. fn. 2) are common knowledge.

2.2 Examples

Four examples help illustrate how the model can be applied to a variety of settings.

Example 1: Only information externality. The canonical example from Bikhchandani et al. (1992) obtains when there is a single preference type, t, and for all i, $u_i(\mathbf{a}, t, \omega) = \mathbf{1}_{\{a_i=\omega\}}$. Here, no agent cares directly about the action of any other agent.

Example 2: Sequential voting with symmetric payoff interdependence. For a voting application, suppose the n agents are voting over a binary agenda, $\{0, 1\}$. The winner or outcome of the election, $W(\mathbf{a})$, is determined by a Q-rule: $W(\mathbf{a}) = 1$ if $|\{i : a_i = 1\}| \geq Qn$, and $W(\mathbf{a}) = 0$ otherwise; the threshold $Q \in [0, 1]$ can be arbitrary. Voters want to elect the "right" candidate: there is a single type t, and for every voter i, $u_i(\mathbf{a}, t, \omega) = \mathbf{1}_{\{W(\mathbf{a})=\omega\}}$. In this baseline voting model, voting is purely instrumental because voters care only about who gets elected.[6]

The current framework can also incorporate richer motivations for voters. For instance, some voters may care about whether their own vote is for the right candidate or not, in the spirit of *expressive voting*. This is accommodated by adding a type t' such that for all i, $u_i(\mathbf{a}, t', \omega) = \mathbf{1}_{\{a_i=\omega\}}$. It is also straightforward to include types whose utility is a combination of type t and

[5] Readers familiar with the information aggregation in voting literature may find it helpful to drawn an analogy with a "common threshold of reasonable doubt" in those models (e.g. Feddersen and Pesendorfer, 1998). If payoffs are symmetric across states for all action profiles (i.e., $u_i(\mathbf{a}, t, 1) + u_i(\mathbf{a}, t, 0)$ is constant for all $\mathbf{a}$, i, and t) then c takes the familiar value of 1/2.

[6] Fey (2000) and Wit (1997) studied this case with simple majority rule ($Q = \frac{1}{2}$) and assuming symmetric signal precisions across states. Callander (2007) further assumes that each voter also intrinsically enjoys voting with the majority. This is akin to positive network externalities, and as previously discussed, violates our Assumption (A1).

t', so that they care about both the collective outcome and their own vote. More interestingly, some voters may also care about *margins of victory*, for example if this influences the policy chosen by the elected candidate (cf. Razin, 2003). This can be captured by adding a type t'' with payoff function, say, $u_i(\mathbf{a}, t'', \omega) = (\omega - \frac{1}{2}) \sum_{j=1}^n a_j$. So type t'' wants the better candidate to get as many votes as possible, even if the winner as specified by the Q-rule is unaffected.

Example 3: Resource allocation with common but asymmetric weights. Suppose that each agent i has an amount of indivisible resources $x_i \geq 0$ (time, money, etc.) that must be given to one of two projects, 0 and 1, and the endowment of resources is commonly known.[7] Normalize $\sum_i^n x_i \leq 1$. Each project can be a success or a failure; the success probability depends on intrinsic quality (the state) and the total amount of resources it receives. In particular, suppose that for any project $p \in \{0, 1\}$, $\Pr(p \text{ succeeds}|\mathbf{a}, \omega) = \sum_{i=1}^n x_i \mathbf{1}_{\{a_i = p = \omega\}}$.[8] Then, if individuals' payoffs are determined by project success, we have $T = \{t\}$ and for all i, $u_i(\mathbf{a}, t, \omega) = \sum_{j=1}^n x_j \mathbf{1}_{\{a_j = \omega\}}$.

The important feature here is that agents are asymmetric because of their resource endowment. The setup captures some essential features of applications like leadership charitable giving, technology adoption, and volunteering for grassroots organizations. An additional application is to campaign contributions. While some contributors may be motivated by purely collective preferences of maximizing the probability that the right candidate from a party gets elected, others may only care about this conditional on them having contributed to the winning person, perhaps because their contribution is intended to buy favors. This can be captured by adding a type t' with, for example, $u_i(\mathbf{a}, t', \omega) = \sum_{j=1}^n x_j \mathbf{1}_{\{a_j = a_i = \omega\}}$.

Example 4: Directed altruism with heterogenous weighting. Suppose that agents are linked according to a directed graph, $g = \{ij, \ldots, i'j'\}$, where $ij \in g$ means that there is a directed link from i to j, capturing that agent i is altruistic toward agent j. The "selfish" component of payoffs for any player i are those in the canonical non-interdependent model, say $v(a_i, \omega) = \mathbf{1}_{\{a_i = \omega\}}$. There is a single preference type, t, and the net payoff for any player i is given by $u_i(\mathbf{a}, t, \omega) := v(a_i, \omega) + \delta \sum_{j \neq i: ij \in g} v(a_j, \omega)$, where $\delta > 0$. The important feature here is that the graph g—which is mapped into the u_i's—allows for each player to care about the actions of a different set of other players, and for this to be common knowledge.[9] In this application,

[7] By assuming that x_i must be allocated to one of two projects, we are abstracting here from free-riding issues; the indivisibility is to maintain the binary action structure.

[8] While this formulation assumes no minimum threshold of resources that must be surpassed by the high-quality project, it would be straightforward to incorporate such an aspect. For instance, if all that matters is whether a project is of high quality and receives more resources than some threshold, $\bar{x} \in (0, 1)$, then $\Pr(p \text{ succeeds}|\mathbf{a}, \omega) = \mathbf{1}_{\{\bar{x} \leq \sum_{i=1}^n x_i \mathbf{1}_{\{a_i = p = \omega\}}\}}$. With obvious changes to the players' payoff specifications described next, this fits into our framework.

[9] If g is the complete graph and each agent i internalizes any future agent j's utility but discounted by δ^{j-i} for

no matter the network structure, agents still act in sequence and observe the entire history of play; the network specifies the structure of payoff interdependence.[10] Adding preference types in a straightforward way would also allow the model to capture lack of common knowledge about whether agents are altruistic or not, how altruistic they are, or exactly which other agents each agent cares about.

2.3 Strategies and Equilibrium

Denote by $G\left(\pi, \gamma_0, \gamma_1, \tau(\cdot); n, \{u_i(\cdot)\}_{i=1}^{n}\right)$ the dynamic game defined above with prior π, signal precisions γ_0 and γ_1, preference type distribution $\tau(\cdot)$, and n agents with payoff functions $\{u_i(\cdot)\}$. Throughout the subsequent analysis, we use the term *equilibrium* to mean a (weak) Perfect Bayesian equilibrium of this game (Fudenberg and Tirole, 1991).[11] Let $h^i \in \{0,1\}^{i-1}$ be the history of actions observed by agent i prior to his choice, with $h^1 := \emptyset$. A pure strategy for player i is a map $\alpha_i : T \times \{0,1\}^{i-1} \times \{0,1\} \to \{0,1\}$, where $\alpha_i(t_i, h^i, s_i)$ is i's action when she has preference type t_i, observes history h^i, and has received signal s_i. We say that agent i acts or chooses *informatively* following a history h^i if for all $t_i \in T$, $\alpha_i\left(t_i, h^i, s_i\right) = s_i$. An agent acts *uninformatively* if her action does not depend on her signal. There is a *herd* on action $a \in \{0,1\}$ at some history h^i if every agent $j \geq i$ chooses action a uninformatively.

Given a strategy profile, α, and any history h^j that can occur on the path of play of α, let $\mu_i(h^j, s_i; \alpha)$ denote the posterior probability that Bayes rule generates for player i on state 1 after observing history h^j and given her private signal s_i. For any history h^j that is off the path of play given strategy profile α, define $\mu_i\left(h^j, s_i; \alpha\right)$ as the posterior probability that Bayes rule generates given signal s_i and the maximal sub-history of h^i that is consistent with on-path behavior, i.e., ignoring all actions that are off path.

some common $\delta \in (0,1)$, then this setting becomes similar to the altruism model of Smith and Sorensen (2008), who study the planner's problem or socially optimal equilibrium when there are an infinite number of agents.

[10] Acemoglu et al. (2010) analyze a complementary set of issues about observational learning in networks by maintaining payoff independence and instead using a network to represent the structure of observability of actions.

[11] The sincere-behavior equilibrium we characterize is also a sequential equilibrium (Kreps and Wilson, 1982).

3 Sincere Behavior

3.1 Definition and Characterization

The following strategy profile is the focus of this paper.

Definition 1. *A strategy profile, α^{SB}, has* ***sincere behavior*** *if every agent i with type t_i and signal s_i who faces history h^i chooses actions as follows:*

1. $\alpha_i^{SB}(t_i, h^i, s_i) = 0$ *if* $\mu_i\left(h^i, s_i; \alpha^{SB}\right) < c$.
2. $\alpha_i^{SB}(t_i, h^i, s_i) = 1$ *if* $\mu_i\left(h^i, s_i; \alpha^{SB}\right) > c$.
3. $\alpha_i^{SB}(t_i, h^i, s_i) = s_i$ *if* $\mu_i\left(h^i, s_i; \alpha^{SB}\right) = c$.

An agent i is ***sincere*** *if her strategy is α_i^{SB}.*

The first two parts of the definition specify that given the history of actions and her private signal, an agent chooses action 0 if her posterior belief that the state is 1 is strictly less than c, and chooses action 1 if the belief is strictly greater than c. The third part is a tie-breaking rule which stipulates that if the posterior belief is exactly c, the agent follows her signal. This choice of tie-breaking rule is inessential because ties arise only for a non-generic constellation of parameters.[12]

Definition 1 is indirect in that it specifies players' behavior as a function of their posterior beliefs rather then as a function of history. However, by proceeding recursively from the first player, it is routine to check that it produces a unique and well-defined strategy profile; Proposition 1 below verifies this. To interpret sincere behavior, recall that c is the common threshold of doubt, so that if a player only cared about her own action, it would be optimal to choose action 1 (resp. 0) if she believes the probability of state 1 is larger (resp. smaller) than c. An agent behaves sincerely if she acts in a myopically optimal fashion after using Bayes rule to update on the state given her private signal and the history of observed actions, with a conjecture that prior agents have behaved sincerely. Note that sincere behavior posits ignoring any unexpected or off-path actions because of how we have defined μ_i.

In the traditional observational learning environment without payoff interdependence (Example 1), it is well known and straightforward that any equilibrium must have sincere behavior

[12] Consider any threshold c, and fix any strategy profile α satisfying the first two conditions of Definition 1. It is straightforward to show that $\{(\pi, \gamma_0, \gamma_1) : \mu_i(h^i, s_i; \alpha) = c \text{ for some } i, h^i, s_i\}$ is of (Lebesgue) measure zero in the space of priors and signal precisions.

modulo the tie-breaking rule. In settings with payoff interdependence, this is far from obvious because of the myopic nature of sincerity. For example, sincere behavior does not explicitly account for pivot considerations crucial to a sequential voting context (Example 2) or strategically influencing future players in a resource allocation or network altruism setting (Examples 3 and 4).

Studying the strategic incentives for agents requires a more direct characterization of how action choices are affected by history in a sincere behavior profile. These dynamics are similar to those in the standard observational learning environment and can be summarized by two state variables. In any history, h^i, the *action lead* for action 1 over action 0, $\Delta(h^i)$, is defined as

$$\Delta(h^i) := \sum_{j=1}^{i-1} \left(\mathbf{1}_{\{a_j=1\}} - \mathbf{1}_{\{a_j=0\}}\right). \tag{2}$$

The second state variable, called the *phase*, summarizes whether voters are continuing to learn about the state (denoted phase L), or learning has terminated because of a herd on one of the actions (denoted phase 1 or 0 in the natural way). To define it, we need *threshold action leads* for each agent, $n_1(i)$ and $n_0(i)$, such that a herd on the respective action begins at i's turn only if the action lead (for action 1 over action 0) has reached the respective threshold. Thus, for any $i > 1$, $n_1(i)$ is the smallest action lead such that for any history h^i with $\Delta(h^i) = n_1(i)$, sincere behavior dictates that agent i choose action 1 even if her private signal is 0.[13] Similarly, the threshold $n_0(i)$ is the smallest action lead such that a sincere agent i would choose action 0 with signal 1, assuming all prior agents have acted informatively. Appendix A.1 provides an explicit construction of these thresholds.[14]

[13]Two facts about sincere behavior are implicit here: first, so long as prior agents are acting informatively, a sufficient statistic for how the history affects an agent's beliefs is the history's action lead (the exact sequence of actions does not matter); second, a larger action lead for 1 can only make an agent choose action 1 when she otherwise would not, but not choose 0 when she otherwise would have chosen 1. Also, sincere behavior may require that agent i choose action 0 when obtaining signal 0 even if *every* preceding agent has chosen action 1; in this case, we set $n_1(i) := i$. Finally, set $n_1(1) := 1$. These points apply analogously to the subsequent discussion about $n_0(\cdot)$.

[14]The range of $n_1(\cdot)$ and $n_0(\cdot)$ need not respectively be positive and negative; see the discussion following Proposition 1.

The phase map $\Psi : h^i \to \{L, 0, 1\}$ is now defined by

$$\forall i > 1, \Psi\left(h^i\right) := \begin{cases} \Psi\left(h^{i-1}\right) & \text{if } \Psi\left(h^{i-1}\right) \in \{0,1\} \\ 1 & \text{if } \Psi\left(h^{i-1}\right) = L \text{ and } \Delta\left(h^i\right) = n_1(i) \\ 0 & \text{if } \Psi\left(h^{i-1}\right) = L \text{ and } \Delta\left(h^i\right) = n_0(i) \\ L & \text{otherwise,} \end{cases} \tag{3}$$

with the initial condition $\Psi(h^1) := x \in \{0, 1\}$ if the first agent's action when sincere is $a_1 = x$ independent of her signal, and $\Psi(h^1) := L$ otherwise.[15]

The following result provides the alternative characterization of sincere behavior, avoiding any reference to beliefs.

Proposition 1. *Every game $G(\pi, \gamma_0, \gamma_1, \tau(\cdot); n, \{u_i\}_{i=1}^n)$ has a unique sincere behavior strategy profile. For each $i \leq n$, the threshold functions $n_1(i)$ and $n_0(i)$ and the map $\Psi(h^i)$ defined above are such that in this strategy profile, each agent acts*

1. *informatively if $\Psi(h^i) = L$;*
2. *uninformatively if $\Psi(h^i) \in \{0, 1\}$, choosing $a = \Psi(h^i)$.*

Moreover, the thresholds $n_1(i)$ and $n_0(i)$ are independent of the population size, n.

The proof of this result, and all subsequent ones not in the text, is the Appendix. The Proposition says that sincere behavior is essentially characterized by the agent-specific herding thresholds on the action lead. If signal precisions are identical in both states ($\gamma_1 = \gamma_0$), the thresholds are invariant to an agent's index and take on a familiar form; for example, if $\pi = c = 1/2$ and $\gamma_1 = \gamma_0 > 1/2$, then $n_1(i) = 2$ and $n_0(i) = -2$ for every $i > 2$. When signal precisions are asymmetric ($\gamma_1 \neq \gamma_0$), the herding thresholds typically vary across agents and hence sincere behavior can be somewhat subtle. For example, even if $c = 1/2$, a herd could begin on an action at some history in which only a minority of prior players have chosen that action, i.e. for some i we may have $n_1(i) < 0$ or $n_0(i) > 0$.

3.2 Rationality of Sincerity

Our main result is:

[15] By Bayes rule, $\Psi(h^1) = 0$ if $\frac{\gamma_1 \pi}{\gamma_1 \pi + (1-\gamma_0)(1-\pi)} < c$, $\Psi(h^1) = 1$ if $\frac{(1-\gamma_1)\pi}{(1-\gamma_1)\pi + \gamma_0(1-\pi)} > c$, and $\Psi(h^1) = L$ otherwise.

Theorem 1. *The sincere behavior strategy profile is an equilibrium for any payoff interdependence structure that satisfies Assumptions (A1) and (A2).*

Following the discussion in Section 3.1, it is clear that the path of play in the sincere-behavior equilibrium is genuinely history-dependent: there is generally no outcome-equivalent equilibrium in a simultaneous counterpart of our model.[16]

In proving Theorem 1, we establish that generically, incentives in both the learning and herding phases are strictly satisfied at every history in which an agent may be "pivotal" in the sense that with positive probability her action will affect her payoff, given the history and strategies of other agents. This implies that if an agent's action has some (possibly small) impact on her own payoff for every fixed profile of others' actions, then her incentives to play sincerely are strict at every history she may encounter. This is the case in Examples 1, 3, and 4 of Section 2.2.[17] The sequential voting setting of Example 2 would also fit so long as all voters have arbitrarily small but positive expressive-voting preferences. Even in settings where an agent's own action does not affect her payoff for some profile of others' actions, versions of strict incentives will often hold. For instance, in the purely-instrumental voting version of Example 2, incentives are generically strict in the learning phase so long as the winner of the election is not already determined.

3.3 Proof

This section outlines the proof for Theorem 1. As a piece of notation, let $(h^j, a_j, ..., a_k)$ denote the history h^{k+1} that obtains when actions $a_j, \ldots, a_k$ follow history h^j. Also, let s_{-i} denote a profile of signal realizations for all players excluding i.

We begin by observing a monotonicity property in how one player's action affects subsequent actions. Since sincere behavior implies that no player is mixing and actions do not depend on preference types, any realized signal profile s_{-i} combined with i's own action a_i determines a unique action profile that is played; denote this action profile by $\sigma_i(a_i, s_{-i})$.

Lemma 1 (Monotonicity Lemma)**.** *Under sincere behavior, $\sigma_i(1, s_{-i}) > \sigma_i(0, s_{-i})$ for any agent i and signal profile s_{-i}.*

[16] In particular, in a sequential election application with purely instrumental voting, the uninformative behavior that occurs once a herd has been triggered in a sincere-behavior equilibrium is for entirely different reasons than the uninformative voting in the asymmetric history-independent equilibria identified by Dekel and Piccione (2000, Theorem 2), where voters who vote uninformatively do so in a pre-ordained way.

[17] In Example 3, so long as $x_i > 0$ for all i.

Proof. Fix i and s_{-i}. The signals $s_1, \ldots, s_{i-1}$ determine the actions prior to i, so we can treat these actions as fixed, say h^i, and focus on the actions of players after i. The result is trivially true for $i = n$, so assume $i < n$. Notice that under sincere behavior, if any agent $j > i$ with signal s_j chooses $a_j = 1$ given any history h^j, she would also do so following a history $h'_j > h^i$. By induction from $j = i + 1$ to $j = n$, this implies that given $s_{i+1}, \ldots, s_n$, if the action profile following history $h^{i+1} = (h^i, 0)$ is $(a_{i+1}, \ldots, a_n)$, then it is $(a'_{i+1}, \ldots, a'_n) \geq (a_{i+1}, \ldots, a_n)$ following history $h^{i+1} = (h^i, 1)$. *Q.E.D.*

Next, we introduce a useful way to think about when a player is "pivotal" in the sense of her action affecting her payoffs. Define the event in which a player i of type t_i is pivotal (for her own payoff) as

$$Piv_i(t_i) := \{s_{-i} : u_i(\sigma_i(1, s_{-i}), t_i, \omega) \neq u_i(\sigma_i(0, s_{-i}), t_i, \omega) \text{ for some } \omega\}.$$

By the Monotonicity Lemma (Lemma 1), for any player i and type t_i,

$$Piv_i(t_i) = \{s_{-i} : u_i(\sigma_i(\omega, s_{-i}), t_i, \omega) > u_i(\sigma_i(1-\omega, s_{-i}), t_i, \omega) \text{ for some } \omega\}. \tag{4}$$

Let $U_i(a_i, t_i | h^i, s_i)$ denote player i's expected utility from action a_i when she has type t_i, faces a history h^i, and has a private signal s_i. If $\Pr(Piv_i(t_i) | h^i, s_i) = 0$, then both actions are optimal for i. Agent i's action changes her expected utility if and only if $s_{-i} \in Piv_i(t_i)$, and so it suffices to consider only those histories in which $\Pr(Piv_i(t_i) | h^i, s_i) > 0$.[18] Therefore, in such cases,

$$U_i\left(a_i, t_i | h^i, s_i\right) > U_i\left(\tilde{a}_i, t_i | h^i, s_i\right) \Leftrightarrow U_i\left(a_i, t_i | h^i, s_i, Piv_i(t_i)\right) > U_i\left(\tilde{a}_i, t_i | h^i, s_i, Piv_i(t_i)\right).$$

When a player conditions on being pivotal, the Monotonicity Lemma (Lemma 1) implies that her strategic decision concerns whether she wants to induce a stochastic increase or decrease in the action profile. By Assumption (A2), the optimal choice hinges on the belief regarding the state: if, taking all relevant information including the pivotal event into consideration, her belief that the state is 1 exceeds (resp., is below) c, then she prefers stochastically increasing (resp., decreasing) actions. The following lemma formalizes this point.

Lemma 2 (Belief Threshold Lemma)**.** *For any agent i and type t_i: if for some $\mathbf{a}$ and $\mathbf{a}' > \mathbf{a}$,*

[18]Notice that $\Pr\left(Piv_i(t_i) | h^i, 1\right)$ need not equal $\Pr\left(Piv_i(t_i) | h^i, 0\right)$, but because signal profiles have full support under the two states, if either is strictly positive, then so is the other.

$u_i(\mathbf{a}', t_i, \omega) \neq u_i(\mathbf{a}, t_i, \omega)$ *for some* ω*, then*

$$\text{sign}\left(\sum_{\omega} u_i\left(\mathbf{a}', t_i, \omega\right) \Pr\left(\omega\right) - \sum_{\omega} u_i\left(\mathbf{a}, t_i, \omega\right) \Pr\left(\omega\right)\right) = \text{sign}\left(\Pr(\omega = 1) - c\right).$$

Proof. Pick any i, t_i, and $\mathbf{a}' > \mathbf{a}$. For any $j \in \{1, \ldots, n\}$, define an operator $\varphi_j\left(\mathbf{a}', \mathbf{a}\right) := \left(a_1', \ldots, a_{j-1}', a_j, \ldots, a_n\right)$, and let $\varphi_0\left(\mathbf{a}', \mathbf{a}\right) := \mathbf{a}$. It follows that[19]

$$\begin{aligned}
& \sum_{\omega} \Pr\left(\omega\right)\left[u_i\left(\mathbf{a}', t_i, \omega\right) - u_i\left(\mathbf{a}, t_i, \omega\right)\right] \\
= \ & \sum_{\omega} \Pr\left(\omega\right) \sum_{j=1}^{n} \mathbf{1}_{\left\{a_j' > a_j\right\}} \left[u_i\left(\left(\varphi_j\left(\mathbf{a}', \mathbf{a}\right)\right)_{-j}^{+}, t_i, \omega\right) - \left(u_i\left(\varphi_j\left(\mathbf{a}', \mathbf{a}\right)\right)_{-j}^{-}, t_i, \omega\right)\right] \\
= \ & \sum_{j=1}^{n} \mathbf{1}_{\left\{a_j' > a_j\right\}} \sum_{\omega} \Pr\left(\omega\right) \left[u_i\left(\left(\varphi_j\left(\mathbf{a}', \mathbf{a},\right)\right)_{-j}^{+}, t_i, \omega\right) - \left(u_i\left(\varphi_j\left(\mathbf{a}', \mathbf{a}\right)\right)_{-j}^{-}, t_i, \omega\right)\right]. \qquad (5)
\end{aligned}$$

Assumption (A2) implies that for each j, the interior summation in (5) is either zero or has the same sign as $\Pr(\omega = 1) - c$. Finally, the hypothesis that $u_i(\mathbf{a}', t_i, \omega) \neq u_i(\mathbf{a}, t_i, \omega)$ for some ω implies that for some j with $a_j' > a_j$, the interior summation has the same sign as $\Pr(\omega = 1) - c$, which must be different from zero. *Q.E.D.*

Building on the preceding lemmas, the following result provides sufficient conditions on an agent's pivotal event that ensures an unambiguous optimal choice for her action.

Lemma 3 (Pivotal Lemma)**.** *Fix an agent* i*, type* t_i*, signal* s_i*, and a history* h^i *such that* $\Pr\left(Piv_i\left(t_i\right) | h^i, s_i\right) > 0$*. Given sincere behavior by all other agents:*

1. *If for every* $(a_{i+1}, \ldots, a_n)$ *that is possible given* h^i*,* $Piv_i(t_i)$*, and* $a_i = 0$*,*

$$\Pr(\omega = 1 | (a_{i+1}, \ldots, a_n), a_i = 0, h^i, s_i, Piv_i(t_i)) \geq c, \qquad (6)$$

   then it is optimal for i *to choose* $a_i = 1$*. Moreover, if the inequality is strict for some* $(a_{i+1}, \ldots, a_n)$*, then it is strictly optimal for* i *to choose* $a_i = 1$*.*

2. *If for every* $(a_{i+1}, \ldots, a_n)$ *that is possible given* h^i*,* $Piv_i(t_i)$*, and* $a_i = 1$*,*

$$\Pr(\omega = 1 | (a_{i+1}, \ldots, a_n), a_i = 1, h^i, s_i, Piv_i(t_i)) \leq c, \qquad (7)$$

[19] Recall that for any $\tilde{\mathbf{a}}_{-j}$, $\tilde{\mathbf{a}}_{-j}^{+} := (\tilde{a}_1, \ldots, \tilde{a}_{j-1}, 1, \tilde{a}_{j+1}, \ldots, \tilde{a}_n)$ and $\tilde{\mathbf{a}}_{-j}^{-} := (\tilde{a}_1, \ldots, \tilde{a}_{j-1}, 0, a_{j+1}, \ldots, a_n)$.

then it is optimal for i to choose $a_i = 0$. Moreover, if the inequality is strict for some $(a_{i+1}, \ldots, a_n)$, then is strictly optimal for i to choose $a_i = 0$.

The first part of the Lemma examines agent i's belief about the state conditioning on all observable information (the public history and her private signal), an action choice of $a_i = 0$, and a vector of subsequent action choices; it says that if this belief attributes c or more weight to $\omega = 1$ for *every* vector of subsequent action choices, then it is optimal for agent i to instead choose $a_i = 1$. The second part of the lemma is analogous but considers behavior after a choice of $a_i = 1$.

A consequence of Lemma 3 is that once the herding phase begins, each agent strategically chooses to herd.

Lemma 4 (Herding Phase Lemma)**.** *Assume all other agents are behaving sincerely. It is strictly optimal for an agent i to choose $a_i = x$ at any history h^i such that $\Psi(h^i) = x \in \{0, 1\}$ and $\Pr\left(Piv_i\left(t_i\right)|h^i\right) > 0$.*

Proof. Assume $\Psi\left(h^i\right) = 1$. Since all future players act uninformatively,

$$\Pr(\omega = 1|(a_{i+1}, \ldots, a_n), h^i, s_i, Piv_i(t_i)) = \Pr(\omega = 1|h^i, s_i)$$

for all $(a_{i+1}, \ldots, a_n)$; moreover, this probability is strictly larger than c for all s_i since $\Psi(h^i) = 1$. By the Pivotal Lemma (Lemma 3), agent i strictly prefers to choose $a_i = 1$. An analogous argument applies when $\Psi(h^i) = 0$. *Q.E.D.*

It remains to check whether an agent will find it optimal to act informatively in the learning phase, assuming that all other agents behave sincerely. Several issues emerge when considering these incentive constraints. First, given the payoff interdependence, informative behavior may not even be optimal in a simultaneous game. For example, it is well-known in the instrumental-voting electoral context that informative voting is not generally an equilibrium for arbitrary voting rules (e.g. Austen-Smith and Banks, 1996). Second, there are forward-looking incentives that are arise due to the combination of sequential decision making and payoff interdependence. For example, an agent who has a signal that favors the tide could behave in a contrarian manner to stem the onset of a herd and induce others to utilize their signals, even if such behavior is myopically suboptimal.

The following result establishes that even though these considerations are paramount, equilibrium incentives in the learning phase incentives do hold.

Lemma 5 (Learning Phase Lemma)**.** *Assume all other agents are behaving sincerely. Consider a history h^i such that $\Psi(h^i) = L$ and $\Pr\left(Piv_i\left(t_i\right)|h^i\right) > 0$. It is optimal for agent i to choose $a_i = s_i$; for generic parameters, it is strictly optimal.*

Proof. Fix an agent $i < n$ with a type t_i and a history h^i such that $\Psi(h^i) = L$ and $\Pr\left(Piv_i\left(t_i\right)|h^i\right) > 0$. Suppose that $s_i = 1$; the analysis is analogous for the other signal. Assume the sincere behavior profile. If $i = n$, the conclusions follow directly from the fact that $\Psi(h^i) = L$ implies $\mu_i(t_i, h^i, s_i; \alpha^{SB}) \geq c$ (generically strictly). So assume $i < n$. It suffices to show that for every $(a_{i+1}, ..., a_n)$ that is possible following $a_i = 0$, (6) is satisfied, because then the Pivotal Lemma (Lemma 3) implies that it is optimal for i to choose $a_i = 1$. Note that optimality must be generically strict in the parameter space because the left-hand side of (6) does not depend on c.

We proceed by partitioning the possible action profiles $(a_i = 0, a_{i+1}, \ldots, a_n)$ into three sets, $\mathcal{A}_0$, $\mathcal{A}_1$, and $\mathcal{A}_L$, as follows: for $x \in \{L, 0, 1\}$, $\mathcal{A}_x \subseteq \{0, 1\}^{n-i+1}$ is the set of those action profiles $(0, a_{i+1}, ..., a_n)$ such that $\Psi\left(\left(h^i, 0, a_{i+1}, ..., a_n\right)\right) = x$. In other words, for $x \in \{0, 1\}$, $\mathcal{A}_x$ is the set of action profiles in which $a_i = 0$ and an x-herd occurs at some point before the end of the game; $\mathcal{A}_L$ is the set of action profiles in which $a_i = 0$ that do not result in a herd. We analyze action profiles in each of these three sets separately, showing that (6) is satisfied for any profile. Intuitively, profiles in $\mathcal{A}_0$ are the most likely to lead to a failure of (6), because under sincere behavior, such profiles are most suggestive that the state is in fact 0. Accordingly, we provide the argument for this case here, and relegate the other two cases ($\mathcal{A}_1$ and $\mathcal{A}_L$) to the Appendix.

<u>Case 1</u>: (Triggering a 0-herd.) Fix any action profile $(a_i = 0, a_{i+1}, ..., a_n) \in \mathcal{A}_0$, and for all $j > i$, let $h^j := (h^i, a_i = 0, a_{i+1}, \ldots, a_{j-1})$. Let $k := \min\{m : \Psi\left(h^m\right)\} = 0$. It must be that $k > i$ because $\Psi\left(h^i\right) = L$; moreover, since $\Psi\left(h^{k-1}\right) = L$, Lemma A.3 implies that $a_{k-1} = 0$. Since all agents with index of k or greater choose action 0 regardless of their signals, only the action choices $(a_{i+1}, ..., a_{k-1})$ are informative among $(a_{i+1}, \ldots, a_n)$. Thus,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right) = \Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_{k-1}\right), a_i = 0, h^i, s_i = 1\right).$$

Since all actions before a herd begins are equally informative, and $a_{k-1} = a_i = 0$, it follows that the RHS of the above expression equals player $k-1$'s belief if he had signal $s_{k-1} = 1$ and observed history h^{k-1}. Therefore,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right) = \mu_{k-1}(h^{k-1}, 1; \alpha^{SB}) \geq c, \tag{8}$$

where the weak inequality follows from $\Psi(h^{k-1}) = L$.

To show that (8) implies the desired (6), we need to add conditioning on the pivotal event, $Piv_i(t_i)$: pick any $s_{-i} \in Piv_i(t_i)$. By the Monotonicity Lemma (Lemma 1), $\sigma_i(1, s_{-i}) > \sigma_i(0, s_{-i})$. Therefore, when conditioning on being pivotal, agent i gets the additional information that some subset of agents (possibly empty) with index k or higher choose action 1 if $a_i = 1$ and action 0 if $a_i = 0$. Since higher actions suggest higher signals, it is then intuitive that i's belief about state 1 (weakly) increases when also conditioning on $Piv_i(t_i)$, which yields (6). A rigorous argument for this step is quite involved, however, and deferred to the Appendix.

This completes the analysis of Case 1. As already noted, Cases 2 and 3—respectively about triggering a herd on action 1 and not triggering a herd at all—are also dealt with in the Appendix.

Q.E.D.

Theorem 1 follows directly from Lemmas 4 and 5, recalling that any action is optimal for agent i at a history h^i such that $\Pr(Piv_i(t_i)|h_i) = 0$.

4 Concluding Remarks

Given the ubiquity of institutions and mechanisms that aggregate information over time and in which the behavior of some agents is observable to others, it is important to understand whether insights of the observational learning literature can be applied to settings where agents' preferences have payoff interdependence. We have shown that these insights do carry over in a simple model in which players have collective preferences in the sense of Assumptions (A1) and (A2).

In the sincere-behavior equilibrium characterized here, agents reveal their information through their actions until the history is more informative than the private signal obtained by any individual agent; subsequent agents then herd. As is well known, in a sufficiently large population, such a herd will occur with probability close to one,[20] and with positive probability has an arbitrarily large fraction of agents choosing action $a \neq \omega$ (where ω is true state). On the other hand, by assumption (A1), the action profile $(\omega, \ldots, \omega)$ is Pareto-optimal in state ω, and with a large number of agents there is collectively close to full-information about the state. Thus, information is asymptotically inefficiently aggregated in the sincere-behavior equilibrium. In the context of a purely instrumental-voting electoral application of our model, this message contrasts with the

[20]Standard arguments imply that for every $\varepsilon > 0$, there exists $\overline{n} < \infty$ such that for all $n > \overline{n}$, if agents play the sincere-behavior equilibrium in the game $G(\pi, \gamma_0\gamma_1, \tau(\cdot); n, \{u_i\}_{i=1}^n)$, then $\Pr[\Psi(h^n) \in \{0,1\}] > 1 - \varepsilon$.

positive result of Dekel and Piccione (2000), who show the existence of history-independent equilibria in this setting that do asymptotically achieve full information aggregation. Since sincere behavior is in undominated strategies, our results at least suggest that information aggregation may be more difficult in a sequential election than a simultaneous counterpart.[21]

Moving beyond the simplest purely-instrumental voting setting—for example, with even a small possibility of expressive-voting preferences or preferences about margins of victory (see Example 2 in Section 2.2)—little is known about whether there are any other equilibria besides sincere behavior in our model. This is also generally the case for non-electoral applications of the model (but see fn. 9 for one exception). We hope this will be addressed in future research.

More generally, there are numerous ways to extend the simple framework studied here. We have restricted attention to binary actions, but in some contexts, agents may choose from more than two alternatives and face an issue of coordination. Our analysis also abstracts from the possibility of multiple signals of different precisions. While one would not expect sincere behavior to always be an equilibrium with payoff interdependence and heterogeneous qualities of information, it would be useful to investigate whether some form of herding can yet persist.

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2010): "Bayesian Learning in Social Networks," Unpublished.

Andreoni, J. (2006): "Leadership Giving in Charitable Fund-Raising," *Journal of Public Economic Theory*, 8, 1.

Austen-Smith, D. and J. S. Banks (1996): "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *The American Political Science Review*, 90, 34–45.

Banerjee, A. (1992): "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107, 797–817.

Bartels, L. M. (1988): *Presidential Primaries and the Dynamics of Public Choice*, Princeton, NJ: Princeton University Press.

[21] In the corresponding simultaneous election environment, Feddersen and Pesendorfer (1997) establish that *any* sequence of undominated (and symmetric) equilibria lead to asymptotic full information aggregation. See Bhattacharya (2008) for conditions on preferences under which aggregation fails even in a simultaneous election.

Bhattacharya, S. (2008): "Preference Monotonicity and Information Aggregation in Elections," Mimeo, University of Pitsburgh.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, 100, 992–1026.

Callander, S. (2007): "Bandwagons and Momentum in Sequential Voting," *Review of Economic Studies*, 74, 653–684.

Choi, J. (1997): "Herd Behavior, the "Penguin Effect," and the Suppression of Informational Diffusion: an Analysis of Informational Externalities and Payoff Interdependency," *Rand Journal of Economics*, 28, 407–425.

Dasgupta, A. (2000): "Social Learning with Payoff Complementarities," Mimeo, London School of Economics.

Dekel, E. and M. Piccione (2000): "Sequential Voting Procedures in Symmetric Binary Elections," *Journal of Political Economy*, 108, 34–55.

Feddersen, T. and W. Pesendorfer (1997): "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica*, 65, 1029–1058.

——— (1998): "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting," *The American Political Science Review*, 92, 23–35.

Fey, M. (2000): "Informational Cascades and Sequential Voting," University of Rochester Working Paper.

Fudenberg, D. and J. Tirole (1991): *Game Theory*, Cambridge, MA: MIT Press.

Goeree, J., T. Palfrey, and B. Rogers (2006): "Social learning with private and common values," *Economic Theory*, 28, 245–264.

Knight, B. and N. Schiff (2010): "Momentum and Social Learning in Presidential Primaries," Forthcoming in the *Journal of Political Economy*.

Kreps, D. and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.

List, J. and D. Lucking-Reiley (2002): "The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign," *Journal of Political Economy*, 110, 215–233.

Neeman, Z. and G. Orosel (1999): "Herding and the winner's curse in markets with sequential bids," *Journal of Economic Theory*, 85, 91–121.

Palmer, N. A. (1997): *The New Hampshire Primary and the American Electoral Process*, Westport: Praeger Publishers.

Razin, R. (2003): "Signaling and Election Motivations in a Voting Model with Common Values and Responsive Candidates," *Econometrica*, 71, 1083–1119.

Smith, L. and P. Sorensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

——— (2008): "Informational Herding and Optimal Experimentation," Unpublished.

Wit, J. (1997): "Herding Behavior in a Roll-Call Voting Game," Department of Economics, University of Amsterdam.

Appendix

A.1 Herding Thresholds

We begin with preliminaries that formally construct the action lead thresholds $n_0(i)$ and $n_1(i)$ for each constellation of parameters $(\pi, \gamma_0, \gamma_1, \tau(\cdot); n, \{u_i\}_{i=1}^n)$.

For any history h^i, define the *public likelihood ratio*, $\lambda(h^i) := \frac{\Pr(\omega=1|h^i)}{\Pr(\omega=0|h^i)}$, which captures how informative the history h^i is. Denote $\underline{\lambda} := \left(\frac{c}{1-c}\right)\left(\frac{1-\gamma_0}{\gamma_1}\right)$ and $\bar{\lambda} := \left(\frac{c}{1-c}\right)\left(\frac{\gamma_0}{1-\gamma_1}\right)$. Under sincere behavior, agent i acts informatively so long as $\lambda(h^i) \in [\underline{\lambda}, \bar{\lambda}]$, chooses 0 regardless of her signal if $\lambda(h^i) < \underline{\lambda}$, and chooses 1 regardless of her signal if $\lambda(h^i) > \bar{\lambda}$. To characterize this behavior in terms of the lead for action 1 over action 0, $\Delta(h^i)$, define the function

$$f(\gamma_\omega, \gamma_{\omega'}) := \frac{\gamma_\omega}{(1-\gamma_{\omega'})}$$

on the domain $\{(x, y) : x, y \in (0,1), \text{ and } x + y > 1\}$. The function f yields likelihood ratios for each signal realization, i.e. $f(\gamma_0, \gamma_1) = \frac{\Pr(s=0|\omega=0)}{\Pr(s=0|\omega=1)}$ while $f(\gamma_1, \gamma_0) = \frac{\Pr(s=1|\omega=1)}{\Pr(s=1|\omega=0)}$. Note that f strictly exceeds 1 over its domain because of the strict monotone likelihood ratio property, $\gamma_0 + \gamma_1 > 1$.

For each integer $i > 0$ and any integer k where $|k| < i$, define the function $g_i(k) := (f(\gamma_1, \gamma_0))^k \left(\frac{f(\gamma_1,\gamma_0)}{f(\gamma_0,\gamma_1)}\right)^{\frac{i-k-1}{2}}$. For a history h^i where the lead for action 1 is k (i.e. $\Delta(h^i) = k$) and all prior agents have acted informatively, $g_i(k) = \frac{\Pr(h^i|\omega=1)}{\Pr(h^i|\omega=0)}$; thus $g_i(k)$ returns the likelihood ratio for such a history, and by Bayes' rule equals $\left(\frac{1-\pi}{\pi}\right)\lambda(h^i)$. For any i, $g_i(k)$ is strictly increasing in k. Define $\underline{g} := \left(\frac{1-\pi}{\pi}\right)\underline{\lambda}$ and $\bar{g} := \left(\frac{1-\pi}{\pi}\right)\bar{\lambda}$.

Given π, γ_0, and γ_1, we next explicitly construct the herding thresholds $\{n_1(i)\}_{i=1}^{\infty}$ and $\{n_0(i)\}_{i=1}^{\infty}$. Begin with the former: for all i such that $g_i(i-1) \leq \bar{g}$, set $n_1(i) = i$. These are cases where even if all agents before i have chosen action 1, $s_i = 0$ still induces a posterior weakly less than c. If $g_i(i-1) > \bar{g}$, set $n_1(i)$ to be the unique integer such that $i - n_1(i)$ is odd, $g_i(n_1(i)) > \bar{g}$, and $g_i(\Delta) \leq \bar{g}$ for any $\Delta \in \{k \in \{-i+1, ..., n_1(i)-2\} : i-k \text{ is odd}\}$. Since $g_i(\cdot)$ is strictly increasing, $n_1(i)$ is uniquely defined. Note that the reason we require $i - n_1(i)$ to be odd is because the feasible action leads at time i are $\{-i+1, -i+3, \ldots, i-1\}$.

Similarly, define $\{n_0(i)\}_{i=1}^{\infty}$ as follows. For all i such that $g_i(-(i-1)) \geq \underline{g}$, set $n_0(i) = -i$. These are cases where even if all agents before i have chosen action 0, $s_i = 1$ still induces a posterior weakly greater than c. If $g_i(-(i-1)) < \underline{g}$, set $n_0(i)$ to be the unique integer such that $i - n_0(i)$ is odd, $g_i(n_0(i)) < \underline{g}$, and $g_i(\Delta) \geq \underline{g}$ for any $\Delta \in \{k \in \{n_0(i)+2, ..., i-1\} : i-k \text{ is odd}\}$. As before, since $g_i(\cdot)$ is strictly increasing, $n_0(i)$ is uniquely defined.

These values of $n_0(i)$ and $n_1(i)$ define $\Psi(\cdot)$ as in equation (3) from the text.

We now record some useful facts about how these thresholds vary across agents and implications for the relationship between phases, action leads, and the actions that trigger herds.

Lemma A.1. *For any $i > 1$, $n_1(i-1) \leq n_1(i) + 1$ and $n_0(i-1) \geq n_0(i) - 1$.*

Proof. The proof uses the observation that for any $k \in \{-i+1, \ldots, i-1\}$,

$$
\begin{aligned}
g_i(k-1) &= \frac{(f(\gamma_1,\gamma_0))^{\frac{i+(k-1)-1}{2}}}{(f(\gamma_0,\gamma_1))^{\frac{i-(k-1)-1}{2}}} \\
&= \left(\frac{1}{f(\gamma_0,\gamma_1)}\right)\left(\frac{(f(\gamma_1,\gamma_0))^{\frac{(i-1)+k-1}{2}}}{(f(\gamma_0,\gamma_1))^{\frac{(i-1)-k-1}{2}}}\right) \\
&= \left(\frac{1}{f(\gamma_0,\gamma_1)}\right) g_{i-1}(k) \\
&< g_{i-1}(k),
\end{aligned}
$$

because $f(\gamma_0, \gamma_1) > 1$. We prove the result for $n_1(\cdot)$; it is analogous for $n_0(\cdot)$. Suppose, to contradiction, that for some $i > 1$, $n_1(i-1) > n_1(i) + 1$. This immediately implies that $n_1(i) < i - 1$. Since $n_1(i)$ is then a feasible action lead at time i, we further have $n_1(i) \leq i - 3$, and therefore, $g_i(i-3) > \bar{g}$ because g_i is strictly increasing. Consider the following two mutually exclusive and exhaustive cases.

Case 1: Suppose $n_1(i-1) = i - 1$. Then $g_{i-1}(i-2) \leq \bar{g}$, and therefore by the observation above, $g_i(i-3) < \bar{g}$, leading to a contradiction.

Case 2: Suppose $n_1(i-1) < i - 1$. Then the feasibility of action lead $n_1(i-1)$ at time $i-1$ combined with $n_1(i-1) > n_1(i) + 1$ implies that $n_1(i-1) \geq n_1(i) + 3$. We have

$$g_{i-1}(n_1(i) + 1) \leq g_{i-1}(n_1(i-1) - 2) \leq \bar{g} < g_i(n_1(i)), \tag{9}$$

where the first inequality is by the monotonicity of $g_{i-1}(\cdot)$, while the latter two inequalities are by the definition of $n_1(\cdot)$. But (9) contradicts the observation above. *Q.E.D.*

Lemma A.2. *If $\Psi(h^i) = L$ then $\Delta(h^i) \in \{n_0(i) + 1, \ldots, n_1(i) - 1\}$.*

Proof. The proof is by induction. The claim is obviously true for $i = 1$, so fix any $i > 1$ and suppose, to contradiction, that the claim is true for all $j < i$, yet $\Psi(h^i) = L$ and $\Delta(h^i) > n_1(i) - 1$ (an analogous argument applies for $\Delta(h^i) < n_0(i) + 1$). By definition of the phase map, $\Delta(h^i) \neq n_1(i)$, hence $\Delta(h^i) \geq n_1(i) + 2$ because the action lead at any given period only takes values in increments of two. Thus $\Delta(h^{i-1}) \geq n_1(i) + 1$ and since $\Psi(h^{i-1}) = L$, it follows from the induction hypothesis that $n_1(i-1) \geq n_1(i) + 2$. This contradicts Lemma A.1. *Q.E.D.*

Lemma A.3. *For any i and $x \in \{0, 1\}$, $\Psi(h^i) = L$ and $\Psi((h^i, a_i)) = x$ imply $a_i = x$.*

Proof. Assume the hypotheses. We provide the argument for $x = 1$; it is analogous for the other case. Assume, to contradiction, that $a_i = 0$ and let $h^{i+1} := (h^i, a_i = 0)$. Since $\Psi(h^i) = L$, Lemma A.2 implies that $\Delta(h^i) \leq n_1(i) - 1$. This implies $\Delta(h^{i+1}) = \Delta(h^i) - 1 \leq n_1(i) - 2 < n_1(i+1)$, where the strict inequality is by Lemma A.1. But then, by the definition of the phase map, $\Psi(h^{i+1}) \neq 1$, a contradiction with the hypotheses. *Q.E.D.*

A.2 Proposition 1

Proof. The proposition is true for agent 1: if $n_1(1) = 0$, then $\Psi(h^1) = 1$ and sincere behavior involves agent 1 choosing 1 regardless of his signal. Similarly, if $n_0(1) = 0$, then $\Psi(h^1) = 0$, and sincere behavior involves agent 1 choosing 0 regardless of his signal. If $n_1(1) = 1$ and $n_0(1) = -1$, then $\Psi(h^1) = L$ and Sincere Behavior involves agent 1 acting informatively. To proceed by induction, fix some $i > 1$ and assume that the proposition's claim about behavior is true for all agents $j < i$.

Case 1: $\Psi(h^i) = L$: All preceding agents have acted informatively. Since the signals possessed by agents are exchangeable, only the number of actions that each alternative has received matters and not the actual sequence of actions. Thus, we can define a function $\tilde{\mu}_i(\Delta, s_i) = \mu_i\left(h^i, s_i; \alpha^{SB}\right)$ where $\Delta = \Delta(h^i)$. By Bayes' rule,

$$\frac{\tilde{\mu}_i(\Delta, 1)}{1 - \tilde{\mu}_i(\Delta, 1)} = \left(\frac{\pi}{1-\pi}\right)\left(\frac{\gamma_1}{1-\gamma_0}\right) g_i(\Delta).$$

Notice that $\Psi(h^i) = L$ implies that $\Delta \geq n_0(i) + 1$ by Lemma A.2, and therefore, $g_i(\Delta) \geq \underline{g}$. It follows that $\tilde{\mu}_i(\Delta, 1) \geq c$, and therefore, sincere behavior requires that agent i choose action 1, following her signal. Similarly, using Bayes' rule,

$$\frac{\tilde{\mu}_i(\Delta, 0)}{1 - \tilde{\mu}_i(\Delta, 0)} = \left(\frac{\pi}{1-\pi}\right)\left(\frac{1-\gamma_1}{\gamma_0}\right) g_i(\Delta).$$

Notice that $\Psi(h^i) = L$ implies that $\Delta \leq n_1(i) - 1$ by Lemma A.2, and therefore, $g_i(\Delta) \leq \bar{g}$. It follows that $\tilde{\mu}_i(\Delta, 0) \leq c$, and therefore, sincere behavior requires that agent i choose action 0 following her signal.

Case 2: $\Psi(h^i) = 0$. Then all agents who chose prior to the first time Ψ took on the value 0 chose informatively, whereas no agent acted informatively thereafter. Let $j \leq i$ be such that $\Psi(h^j) = 0$ and $\Psi(h^{j-1}) = L$; hence, $\Delta(h^j) = n_0(j)$. Then, $\mu_j\left(h^j, s_j; \alpha^{SB}\right) = \tilde{\mu}_j(n_0(j), s_j)$. Since all choices after that of agent $j-1$ are uninformative, $\mu_i\left(h^i, s_i; \alpha^{SB}\right) = \mu_j\left(h^j, s_i; \alpha^{SB}\right) = \tilde{\mu}_j(n_0(j), s_i)$. Suppose $s_i = 1$. Since $g_j(n_0(j)) < \underline{g}$, it follows that $\tilde{\mu}_j(n_0(j), 1) < c$, and therefore sincere behavior requires that agent i choose 0 following $s_i = 1$. A fortiori, since $\tilde{\mu}_j(n_0(j), 0) < \tilde{\mu}_j(n_0(j), 1)$, she must also choose 0 following $s_i = 0$.

Case 3: $\Psi(h^i) = 1$. This argument is omitted since it is entirely analogous to Case 2 above. *Q.E.D.*

A.3 Lemma 3

Proof. We prove the first item; the argument is analogous for the second. To ease notation in the proof, let $\mathcal{P}_i$ be shorthand for the triple $\langle h^i, s_i, Piv_i(t_i)\rangle$, and let $\mathbf{a}'$ denote a profile of actions for players $i+1,\ldots,n$. We have

$$\begin{aligned} U_i(0,t_i|\mathcal{P}_i) &= \sum_{\mathbf{a}'}\sum_{\omega} u_i\left(\left(h^i,0,\mathbf{a}'\right),t_i,\omega\right)\Pr(\omega,\mathbf{a}'|a_i=0,\mathcal{P}_i) \\ &= \sum_{\mathbf{a}'}\sum_{\omega} u_i\left(\left(h^i,0,\mathbf{a}'\right),t_i,\omega\right)\Pr(\mathbf{a}'|a_i=0,\mathcal{P}_i)\Pr(\omega|\mathbf{a}',\mathcal{P}_i) \\ &= \sum_{\mathbf{a}'}\Pr(\mathbf{a}'|a_i=0,\mathcal{P}_i)\sum_{\omega} u_i\left(\left(h^i,0,\mathbf{a}'\right),t_i,\omega\right)\Pr(\omega|\mathbf{a}',\mathcal{P}_i), \end{aligned} \tag{10}$$

where the 2nd equality uses the fact that $\Pr(\omega|a_i,\mathbf{a}',\mathcal{P}_i)=\Pr(\omega|\mathbf{a}',\mathcal{P}_i)$ because i's action contains no information about the state. We next perform a change of variables to replace each $\mathbf{a}'$ by an equivalent set of signal profiles. For this purpose, observe that any history h^i, signal profile $s:=(s_{i+1},\ldots,s_n)$, and action a_i deterministically map into some action profile $\mathbf{a}'$; with abuse of notation, we denote this function by $\mathbf{a}'(s,a_i,h^i)$. Let $S(h^i)$ be a set of equivalence classes of signal profiles of players $i+1,\ldots,n$ such that for any $z\in S(h^i)$, $s\in z$, and $\tilde{s}\in z$, $\mathbf{a}'(s,0,h^i)=\mathbf{a}'(\tilde{s},0,h^i)$. With more abuse of notation, we also write $\mathbf{a}'(z,0,h^i)$ for any $z\in S(h^i)$. It then follows from (10) that

$$\begin{aligned} U_i(0,t_i|\mathcal{P}_i) &= \sum_{z\in S(h^i)}\Pr(z|a_i=0,\mathcal{P}_i)\sum_{\omega} u_i\left(\left(h^i,0,\mathbf{a}'\left(z,0,h^i\right)\right),t_i,\omega\right)\Pr(\omega|z,\mathcal{P}_i) \\ &= \sum_{z\in S(h^i)}\Pr(z|\mathcal{P}_i)\sum_{\omega} u_i\left(\left(h^i,0,\mathbf{a}'\left(z,0,h^i\right)\right),t_i,\omega\right)\Pr(\omega|z,\mathcal{P}_i), \end{aligned} \tag{11}$$

where the 2nd equality uses the fact that signal profile realizations of other players are independent of i's action. Now consider what happens if i plays $a_i=1$ instead of $a_i=0$. Each signal profile of subsequent players, s, induces an action profile $\mathbf{a}'(s,1,h^i)$, which may be different from $\mathbf{a}'(s,0,h^i)$. Thus,

$$\begin{aligned} &U_i(1,t_i|\mathcal{P}_i) \\ =& \sum_{z\in S(h^i)}\Pr(z|\mathcal{P}_i)\sum_{\omega}\Pr(\omega|z,\mathcal{P}_i)\sum_{s\in z} u_i\left(\left(h^i,1,\mathbf{a}'\left(s,1,h^i\right)\right),t_i,\omega\right)\Pr(s|\omega,z,\mathcal{P}_i) \\ =& \sum_{z\in S(h^i)}\Pr(z|\mathcal{P}_i)\left\{\begin{array}{c} \sum_{\omega}\Pr(\omega|z,\mathcal{P}_i)\sum_{s\in z} u_i\left(\left(h^i,1,\mathbf{a}'\left(s,1,h^i\right)\right),t_i,\omega\right)\Pr(s|\omega=0,z,\mathcal{P}_i) \\ +\Pr(\omega=1|z,\mathcal{P}_i)\sum_{s\in z} u_i\left(\left(h^i,1,\mathbf{a}'\left(s,1,h^i\right)\right),t_i,1\right)\Pr(s|\omega=1,z,\mathcal{P}_i) \\ -\Pr(\omega=1|z,\mathcal{P}_i)\sum_{s\in z} u_i\left(\left(h^i,1,\mathbf{a}'\left(s,1,h^i\right)\right),t_i,1\right)\Pr(s|\omega=0,z,\mathcal{P}_i) \end{array}\right\}, \end{aligned} \tag{12}$$

where it is important to keep in mind that the equivalence classes of signal profiles are those corresponding to $a_i = 0$. The remainder of the proof shows that (12) is larger (weakly or strictly, as appropriate) than (11), which proves the desired result.

First note that for any $(a_{i+1} \ldots, a_n) = \mathbf{a}'(z, 0, h^i)$, (6) is equivalent to $\Pr(\omega = 1|z, \mathcal{P}_i) \geq c$. Also, for any $z \in S(h^i)$ and $s \in z$, $\mathbf{a}'(s, 1, h^i) \geq \mathbf{a}'(z, 0, h^i)$ by the argument used in the Monotonicity Lemma (Lemma 1). We claim that for any $z \in S(h^i)$ such that $\Pr(\omega = 1|z, \mathcal{P}_i) \geq c$,

$$\begin{aligned} & \sum_{\omega} \Pr(\omega|z, \mathcal{P}_i) \sum_{s \in z} u_i\left(\left(h^i, 1, \mathbf{a}'\left(s, 1, h^i\right)\right), t_i, \omega\right) \Pr(s|\omega = 0, z, \mathcal{P}_i) \\ \geq \; & \sum_{\omega} u_i\left(\left(h^i, 0, \mathbf{a}'\left(z, 0, h^i\right)\right), t_i, \omega\right) \Pr(\omega|z, \mathcal{P}_i). \end{aligned} \tag{13}$$

To see this, observe that for any $z \in S(h^i)$, (13) is equivalent to

$$\begin{aligned} & \sum_{\omega} \Pr(\omega|z, \mathcal{P}_i) \sum_{s \in z} u_i\left(\left(h^i, 1, \mathbf{a}'\left(s, 1, h^i\right)\right), t_i, \omega\right) \Pr(s|\omega = 0, z, \mathcal{P}_i) \\ \geq \; & \sum_{\omega} \Pr(\omega|z, \mathcal{P}_i) \sum_{s \in z} u_i\left(\left(h^i, 0, \mathbf{a}'\left(s, 0, h^i\right)\right), t_i, \omega\right) \Pr(s|\omega = 0, z, \mathcal{P}_i), \end{aligned}$$

which in turn is equivalent to

$$\begin{aligned} & \sum_{s \in z} \left(\sum_{\omega} u_i\left(\left(h^i, 1, \mathbf{a}'\left(s, 1, h^i\right)\right), t_i, \omega\right) \Pr(\omega|z, \mathcal{P}_i) \right) \Pr(s|\omega = 0, z, \mathcal{P}_i) \\ \geq \; & \sum_{s \in z} \left(\sum_{\omega} u_i\left(\left(h^i, 0, \mathbf{a}'\left(z, 0, h^i\right)\right), t_i, \omega\right) \Pr(\omega|z, \mathcal{P}_i) \right) \Pr(s|\omega = 0, z, \mathcal{P}_i). \end{aligned} \tag{14}$$

(13) follows from (14) and the observation that for any $s \in z$,

$$\sum_{\omega} u_i\left(\left(h^i, 1, \mathbf{a}'\left(s, 1, h^i\right)\right), t_i, \omega\right) \Pr(\omega|z, \mathcal{P}_i) \geq \sum_{\omega} u_i\left(\left(h^i, 0, \mathbf{a}'\left(z, 0, h^i\right)\right), t_i, \omega\right) \Pr(\omega|z, \mathcal{P}_i)$$

because $\mathbf{a}'(s, 1, h^i) \geq \mathbf{a}'(z, 0, h^i)$ and Lemma 2 applies.

Similarly, (13) holds with a strict inequality if $\Pr(\omega = 1|z, \mathcal{P}_i) > c$. Finally, since $\omega = 1$ makes higher signals more likely for any player than $\omega = 0$, it also follows from the monotonicity of $\mathbf{a}'(\cdot, 1, h^i)$ and (A1) that for any $z \in S(h^i)$,

$$\begin{aligned}&\sum_{s\in z} u_i\left(\left(h^i, 1, \mathbf{a}'\left(s, 1, h^i\right)\right), t_i, 1\right) \Pr\left(s|\omega = 1, z, \mathcal{P}_i\right) \\ \geq \ &\sum_{s\in z} u_i\left(\left(h^i, 1, \mathbf{a}'\left(s, 1, h^i\right)\right), t_i, 1\right) \Pr\left(s|\omega = 0, z, \mathcal{P}_i\right). \end{aligned} \tag{15}$$

Using these facts about (13) and (15), we conclude that (12) is weakly larger than (11), and strictly so if (6) holds strictly for some $(a_{i+1}, \ldots, a_n)$. *Q.E.D.*

A.4 Lemma 5

Proof. We first complete the analysis of Case 1 that was begun in the main text, and then indicate how Cases 2 and 3 follow in a similar vein.

Case 1: We must show that (8) implies (6). To reduce notational burden, we drop the dependence of $Piv_i(t_i)$ on t_i for the remainder of the analysis of Case 1. It suffices to establish that

$$\Pr\left(\omega = 1|\left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, Piv_i\right) \geq \Pr\left(\omega = 1|\left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right).$$

Since $\Psi(h^j) = L$ for all $j < k$, the conditioning event on the right-hand side above reveals that $s_j = a_j$ for all $j < k$ $(j \neq i)$. Thus, denoting

$$Piv_i' := Piv_i \cap \{s_{-i} : s_j = a_j \text{ for } j < k, j \neq i\},$$

it follows that

$$\Pr\left(\omega = 1|\left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, Piv_i\right) = \Pr\left(\omega = 1|\left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, Piv_i'\right). \tag{16}$$

Consider the space $\{0, 1\}^{n-k+1}$ and let $s^{k:n}$ be a generic element, which denotes a vector of signal for players k through n. We write $s^{k:n}_m$ to denote the signal of player $m \in \{k, \ldots, n\}$ (as opposed to the m^{th} coordinate of the vector). Create a partition as follows:

$$\zeta_j = \left\{ s^{k:n} : \sum_{m\in\{k,..,n\}} s^{k:n}_m = j \right\}, \text{ where } j \in \{0, ..., n-k+1\}.$$

In words ζ_j consists of all signal profiles for players $k, \dots, n$ that contain exactly j one signals.

By the law of iterated expectations,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right) = \mathbb{E}_F\left[\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, \zeta_j\right)\right], \quad (17)$$

where $F(j)$ is the cumulative distribution whose density function is

$$f(j) := \Pr\left(\zeta_j | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right).$$

Next, note that for any j,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, \zeta_j, Piv_i'\right) = \Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, \zeta_j\right)$$

because once we condition on a ζ_j, Piv_i' can only add information about which subset of players $k, \dots, n$ have the j one signals and which have the $n-k+1-j$ zero signals, but this does not affect the posterior about the state. Therefore, by iterated expectations again,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, Piv_i\right) = \mathbb{E}_G\left[\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, \zeta_j\right)\right], \quad (18)$$

where G is the cumulative distribution whose density function is

$$g(j) := \Pr\left(\zeta_j | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, Piv_i'\right).$$

Since ζ_j consists of signal profiles with j signals of one,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1, \zeta_j\right)$$

is strictly increasing in j. Therefore, by (17) and (18), it suffices to prove that G first-order stochastically dominates (FOSD) F, which is the content of the remainder of the proof.

Let $N_j := |\zeta_j|$ and $P_j := |\zeta_j \cap Piv_i'|$.[22] Since each element of ζ_j is equally likely (because signals are iid conditional on the state), the definitions of $f(\cdot)$ and $g(\cdot)$ yield that for any $j \in \{0, \dots, n-k+1\}$ and $s^{k:n} \in \zeta_j$,

$$f(j) = N_j \Pr\left(s^{k:n} | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right) \quad (19)$$

[22]As usual, $|X|$ denotes the cardinality of a set X.

and

$$\begin{aligned}
g(j) &= \frac{\Pr\left(\zeta_j \cap Piv_i' | \left(a_{i+1},...,a_n\right), a_i = 0, h^i, s_i = 1\right)}{\sum_{j' \in \{0,...,n-k+1\}} \Pr\left(\zeta_{j'} \cap Piv_i' | \left(a_{i+1},...,a_n\right), a_i = 0, h^i, s_i = 1\right)} \\
&= \frac{P_j \Pr\left(s^{k:n} | \left(a_{i+1},...,a_n\right), a_i = 0, h^i, s_i = 1\right)}{\sum_{j' \in \{0,...,n-k+1\}} \Pr\left(\zeta_{j'} \cap Piv_i' | \left(a_{i+1},...,a_n\right), a_i = 0, h^i, s_i = 1\right)}. \qquad (20)
\end{aligned}$$

It is well known that G FOSD F if the density g likelihood-ratio dominates f, i.e. for any $j \in \{0, \ldots, n-k\}$, $g(j)/f(j) \leq g(j+1)/f(j+1)$. Substituting from (19) and (20), this is equivalent to showing that for any $j \in \{0, \ldots, n-k\}$,

$$\frac{P_{j+1}}{N_{j+1}} \geq \frac{P_j}{N_j}. \qquad (21)$$

Now we claim that $s^{k:n} \in Piv_i'$ and $\tilde{s}^{k:n} > s^{k:n}$ imply $\tilde{s}^{k:n} \in Piv_i'$. To see this, note that since we are holding fixed the signals of players $1 \ldots, k-1$, $\tilde{s}^{k:n}$ must lead to a (weakly) higher action profile than $s^{k:n}$ following $a_i = 1$. On the other hand, $a_i = 0$ leads to the same action profile under both $s^{k:n}$ and $\tilde{s}^{k:n}$ since a $0-$herd is triggered at person k's turn. The two previous statements imply that if i is pivotal given $s^{k:n}$, she is also pivotal given $\tilde{s}^{k:n}$.

The above fact can be seen to imply that for every $j \in \{0, ..., n-k\}$,

$$\frac{P_{j+1}}{P_j} \geq \frac{n-k+1-j}{j+1}. \qquad (22)$$

Since for any j, $N_j = |\zeta_j| = \frac{(n-k+1)!}{j!(n-k+1-j)!}$ (the first equality is by definition; the second by the binomial formula), it follows that for any $j \in \{0, \ldots, n-k+1\}$,

$$\frac{N_{j+1}}{N_j} = \frac{n-k+1-j}{j+1},$$

which combines with (22) to imply the desired inequality (21).

<u>Case 2</u>: (Triggering a 1-herd.) Fix any action profile $(a_i = 0, a_{i+1}, ..., a_n) \in \mathcal{A}_1$, and for all $j > i$, let $h^j := (h^i, a_i = 0, a_{i+1}, \ldots, a_j)$. Let $k := \min\{m : \Psi\left(h^m\right)\} = 1$. It must be that $k > i$ because $\Psi\left(h^i\right) = L$; moreover, since $\Psi\left(h^{k-1}\right) = L$, Lemma A.3 implies that $a_{k-1} = 1$. Since all agents with index of k or greater choose action 1 regardless of their signals, only the action choices

$(a_{i+1}, ..., a_{k-1})$ are informative among $(a_{i+1}, \ldots, a_n)$. Thus,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right) = \Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_{k-1}\right), a_i = 0, h^i, s_i = 1\right).$$

Since all actions before a herd reveal signals and are equally informative, and $a_i = 0$ while $s_i = 1$, it follows that the right-hand-side above is strictly larger than player k's belief at h^k with $s_k = 1$ (since h^k "wrongly" features $a_i = 0$), i.e. $\mu_k(h^k, 1; \alpha^{SB})$, which is strictly larger than c because $\Psi(h^k) = 1$. Therefore,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right) > c. \tag{23}$$

The remaining step to conclude that (6) holds is to introduce conditioning on $Piv_i(t_i)$; this is entirely analogous to Case 1, hence omitted.

<u>Case 3</u>: (No herd is triggered.) Fix any action profile $(a_i = 0, a_{i+1}, ..., a_n) \in \mathcal{A}_L$, and for all $j > i$, let $h^j := (h^i, a_i = 0, a_{i+1}, \ldots, a_j)$. Since no herd is triggered, all actions are informative. It follows $\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_{n-1}\right), a_i = 0, h^i, s_i = 1\right)$ is strictly larger than player n's belief at h^n given $s_n = 1$, i.e. $\mu_n(h^n, 1; \alpha^{SB})$, which is weakly larger than c because $\Psi(h^n) = L$. Therefore,

$$\Pr\left(\omega = 1 | \left(a_{i+1}, ..., a_n\right), a_i = 0, h^i, s_i = 1\right) > c. \tag{24}$$

Finally, note that since all players signals are known to i given the conditioning above, further conditioning on $Piv_i(t_i)$ adds no information, so (6) directly follows. *Q.E.D.*