OPTIMAL EVIDENCE ACCUMULATION ON SOCIAL NETWORKS

BHARGAV KARAMCHED*, SIMON STOLARCZYK†, ZACHARY P. KILPATRICK‡, AND KREŠIMIR JOSIĆ§

Abstract. A fundamental question in biology is how organisms integrate sensory and social evidence to make decisions. However, few models describe how both these streams of information can be combined to optimize choices. Here we develop a normative model for collective decision making in a network of agents performing a two-alternative forced choice task. We assume that rational (Bayesian) agents in this network make private measurements, and observe the decisions of their neighbors until they accumulate sufficient evidence to make an irreversible choice. As each agent communicates its decision to those observing it, the flow of social information is described by a directed graph. The decision-making process in this setting is intuitive, but can be complex. We describe when and how the *absence* of a decision of a neighboring agent communicates social information, and how an agent must marginalize over all unobserved decisions. We also show how decision thresholds and network connectivity affect group evidence accumulation, and describe the dynamics of decision making in social cliques. Our model provides a bridge between the abstractions used in the economics literature and the evidence accumulator models used widely in neuroscience and psychology.

Key words. decision-making; probabilistic inference; social networks

1. Introduction. Understanding how organisms use sensory and social information to make decisions is of fundamental interest in biology, sociology, and economics [13, 16, 20]. Psychologists and neuroscientists have developed a variety of experimental approaches to probe how humans and other animals make choices. Particularly popular are variants of the two-alternative forced choice task (2AFC) where an observer is asked to decide between two options based on information obtained from one or more noisy observations [7, 24, 47, 48]. The 2AFC task has motivated several mathematical models that successfully explain how humans use sequentially-presented evidence to make decisions [9, 15, 55].

Most evidence accumulation models take the form of drift-diffusion stochastic processes that describe the information gathered by lone observers [7]. However, humans and many other animals do not live in isolation and make decisions based on more than their own private observations. Animals watch each other as they forage [29, 34]. Stock traders, while not privy to all of their competitor’s information, can still observe each other’s decisions. To make the best choices many biological agents thus also take into account the observed choices of others [37, 42, 45].

Here we address the question of how idealized agents in a social network should combine a sequence of private measurements with observed decisions of other individuals to choose between two options. We refer to the information an agent receives from its neighbors as *social information*, and information available only to the agent as *private information*. As agents do not share their private information with others directly, they only reveal their beliefs through their choices. These choices are based on private and social information, and thus reveal something about the total evidence an agent has collected.

We assume that private measurements, and, as a consequence, observed choices can improve the odds of making the right decision. However neither type of information affords certainty about which choice is correct. We take a probabilistic (Bayesian) approach to describe the behavior of rational agents who make *immutable* decisions once they have accrued sufficient evidence. Their choices thus provide information about their beliefs at a single point in time.

There are two reasons for assuming that agents only share their decision: First, many organisms communicate their decisions, but not the evidence they used to reach them. For example, herding animals in motion can only communicate their chosen direction of movement [14, 40]. Animals that forage in groups may communicate their preferred feeding locations [19, 49], but not the evidence, or the process which they used to decide. Human traders can see their competitor’s choice to buy or sell a stock, but may not have access to the evidence that lead to these actions. Second, if agents communicate all information they gathered to their neighbors, the problem is mathematically trivial as every agent obtains all evidence provided to the entire network.

*Department of Mathematics, University of Houston, Houston TX 77204 (bhargav@math.uh.edu),

†Department of Mathematics, University of Houston, Houston TX 77204 (spstolar@math.uh.edu),

‡Department of Applied Mathematics, University of Colorado, Boulder CO 80309 (zpkilpat@colorado.edu),

§Departments of Mathematics, and Biology and Biochemistry, University of Houston, Houston TX 77204, Department of BioSciences, Rice University, Houston TX 77005 (josic@math.uh.edu); Bhargav Karamched and Simon Stolarczyk are co-first authors; Zachary Kilpatrick and Krešimir Josić are co-last authors.

The behavior and performance of rational agents in a social network can depend sensitively on how information is shared [1, 35]. Surprisingly, in some cases rational agents perform better in isolation, on average. This can happen when social information from a small subset of agents dominates the decisions of the collective, as in classical examples of *herding behavior* [5, 21]. On the other hand, if after a single observation agents share their belief and continue to do so repeatedly in a recurrent network, then all agents can asymptotically come to agree on the most probable choice given the totality of evidence obtained by the network [2, 22, 23, 35]. In this case of *asymptotic learning*, all agents are eventually able to make use of all the private information in the network by observing the preferences of their neighbors.

In contrast the decisions of agents in our model are immutable. As a consequence, asymptotic learning typically does not occur. We also show how a rational agent who only observes a portion of the network needs to marginalize over the decisions of all unobserved agents to correctly integrate an observed decision. Interestingly when decision thresholds are asymmetric, even the absence of a decision can communicate information deterministically until a decision is reached. In recurrent networks observing the absence of a decision by a neighbor who has observed your own indecision can lead to additional social information, and this can happen repeatedly. This is akin to situations addressed in the literature on *common knowledge* [2, 22, 33].

Our evidence accumulation model shows how to best combine streams of private, and social evidence to make decisions. Such evidence accumulation models have been shown to provide excellent description of decision making in various tasks, and there is evidence that they can explain the formation of social decisions [29]. They may thus point to common mechanisms for human decision-making.

2. Definitions and setup. We consider a set of agents who accumulate evidence to decide between two states, H^+ or H^- . Each agent is rational (Bayesian): They compute and compare the probability of each state (or hypothesis), based on all evidence they accrue. We assume that agents make a decision once the conditional probability of one of the states, given all the accumulated observations, crosses a predetermined threshold [7, 59].

To make a choice, agents gather both private (Priv) and social (Soc) information. Private information comes from a sequence of noisy observations (measurements) of the state, $H \in \{H^+, H^-\}$. Agents also gather social information by observing each other's decisions, as in previous models of foraging animal groups [34], consumer networks [1], and opinion exchange on social networks [60]. Thus agents do not share their private measurements directly, but only observe whether their neighbors have made a choice, and, if so, what that choice was.

Evidence accumulation for a single agent. The problem of a single agent accumulating private evidence to decide between two options has been thoroughly studied [7, 24, 46, 47, 55, 56, 59]. Assume that an agent makes a sequence of noisy observations, $\xi_{1:t}$ with $\xi_i \in \Xi$, $i \in \{1, \dots, t\}$, and $\Xi \subset \mathbb{R}$ finite. The observations ξ_i are independent and identically distributed, conditioned on the true state $H \in \{H^+, H^-\}$,

$$P(\xi_{1:t}|H^\pm) = \prod_{i=1}^t P(\xi_i|H^\pm) = \prod_{i=1}^t f_\pm(\xi_i),$$

where $f_\pm(\xi) := P(\xi|H^\pm)$ are the measurement distributions, which are probability mass functions since Ξ is finite. Observations ξ are drawn from the same set Ξ in either state $H^\pm$, and states are distinguished by the different measurement distributions $f_\pm(\xi)$.

To compute $P(H|\xi_{1:t})$ the agent uses Bayes' Rule: For simplicity, we assume that the agent knows the measurement distributions $f_\pm(\xi)$, uses a flat prior, $P(H^+) = P(H^-) = 1/2$, and that observations are conditionally independent. Thus, the log likelihood ratio (LLR) of the two states at time t is

$$y_t := \log \left(\frac{P(H^+|\xi_{1:t})}{P(H^-|\xi_{1:t})} \right) = \sum_{s=1}^t \log \left(\frac{P(\xi_s|H^+)}{P(\xi_s|H^-)} \right) = y_{t-1} + \log \left(\frac{P(\xi_t|H^+)}{P(\xi_t|H^-)} \right), \quad (2.1)$$

where $y_0 = 0$, since both states are equally likely *a priori*, $\log(P(H^+)/P(H^-)) = 0$. We also refer to the LLR as the *belief* of the agent.

An ideal agent continues making observations while $\theta_- < y_t < \theta_+$, and makes a decision after acquiring sufficient evidence, choosing H^+ (H^-) once $y_t \geq \theta_+$ ($y_t \leq \theta_-$). We assume $\theta_- < 0 < \theta_+$.

TABLE 2.1
Notation guide.

Notation guide	
Notation	Meaning
$H^\pm$	state of the environment (correct hypothesis)
LLR	log-likelihood ratio
$\xi_t^{(i)}$	observation made at time t by agent i
$\xi_{1:t}^{(i)}$	vector of observations made by agent i at times 1 through t
Ξ	set of all possible observations, ξ
$f_\pm(\xi)$	measurement (observation) distributions, $P(\xi H^\pm)$
$\text{Priv}_{1:t}^{(i)}$	private evidence accumulated by agent i through time t
$\text{Soc}_t^{(i)}$	social evidence given by agent i to its neighbors at time t
$I_t^{(i)}$	total information available to agent i at time t .
$y_t^{(i)}$	belief (LLR) of agent i based on private and social information at time t
$\theta_\pm$	decision threshold associated with $H^\pm$. When the thresholds are symmetric about 0, <i>i.e.</i> when $\theta_+ = -\theta_-$, we use the notation $\theta = \theta_+$ and $-\theta = \theta_-$.
Θ	interval of possible agent LLRs (beliefs) prior to a decision, <i>i.e.</i> $\Theta = (\theta_-, \theta_+)$
$d_t^{(i)}$	decision state of agent i at time t
$y_{t,k}^{(i)}, d_{t,k}^{(i)}, \text{Soc}_{t,k}^{(i)}$	LLR, decision state, and social information available to agent i at the k th substep of equilibration after a private observation at time t
$S_\pm(t)$	survival probability for an independent agent under $H^\pm$.
N	total number of agents
Π	all possible LLR increments due to a single private observation
$N^{(i)}$	set of agents that agent i observes (neighbors)
$U^{(i)}$	set of agents that agent i cannot see (unobserved)
$R(t, \mathcal{I})$	LLR from observing agent i has belief in interval $\mathcal{I}$ at time t , $\log \left(\frac{P(y_t^{(i)} \in \mathcal{I} H^+)}{P(y_t^{(i)} \in \mathcal{I} H^-)} \right)$

3. Multiple agents. Limited bandwidth or physical distances between agents can limit their interactions [30, 41, 61]. We identify agents with a set of vertices, $V = \{1, \dots, N\}$, and communication between agents with a set of directed edges, E , between these vertices [25]. The agent at the tail of each directed edge communicates information to the agent at the edge’s head.

As in the case of a single observer, we assume that each agent, i , makes a sequence of noisy, identically-distributed measurements, $\xi_{1:t}^{(i)}$, of the state, H , from a state-dependent distribution, $f_\pm(\xi_j^{(i)}) = P(\xi_j^{(i)} | H^\pm)$. We assume that the observations are independent in time and between agents, conditioned on the state, $P(\xi_{1:t}^{(i_1)}, \dots, \xi_{1:t}^{(i_k)} | H) = \prod_{l=1}^k \prod_{j=1}^t P(\xi_j^{(i_l)} | H)$ for any sequence of measurements, $\xi_{1:t}^{(i)}$, and set of agents, $i_1, \dots, i_k$, in the network. This conditional independence of the measurements is a strong assumption, and is unlikely to hold in practice. We return to this point in the Discussion.

An agent gathers social evidence by observing whether its neighbors have made a decision, and, once they have, what that decision is. Each agent thus gathers private and social evidence, and makes a decision when its belief (LLR) about the two states crosses one of the thresholds, $\theta_- < 0 < \theta_+$. Importantly, *once an agent has made a decision, it cannot change it*. The absence of a decision thus communicates that an agent has not gathered sufficient evidence to make a choice, and hence that this agent’s belief (LLR) is still in the interval (θ_-, θ_+) .

For simplicity, we assume that the distributions and thresholds, $\theta_\pm$, are identical across agents. The theory is similar if agents have different, but known, measurement distributions. The assumption that the measurement distributions, $f_\pm(\xi)$, are discrete, simplifies some convergence arguments. In subsequent work,

FIG. 3.1. A pair of unidirectionally coupled agents deciding between the states H^+ and H^- on the basis of private and social evidence. (a) Schematic of the information flow in the network. Agent 1 accumulates their own observations, $\xi_{1:t}^{(1)}$, which result in a sequence of decision states, $d_{1:t}^{(1)}$, that is observed by agent 2. In addition, agent 2 gathers its own observations, $\xi_{1:t}^{(2)}$, to make a decision. (b) Sample trajectories for the beliefs (LLRs) of the agents. Decisions are made when an agent's belief crosses a threshold, $\theta_\pm = \pm\theta$ in this case. A decision of agent 1 leads to a jump in the belief of agent 2.

we will extend our analysis to the case of continuous measurement distributions and time.

Evidence accumulation with two agents. To illustrate how an agent integrates private and social information to reach a decision, we use the example network shown in Fig. 3.1a. We will show that even the absence of a decision can provide social information.

Let $I_t^{(i)}$ be the total information available to agent i at time t . The information available to agent 1, $I_t^{(1)}$, consists only of private observations. However, agent 2 makes private observations, and obtains social information from agent 1, both of which constitute $I_t^{(2)}$. Agents base their decisions on the computed LLR, or belief, $y_t^{(i)} = \log \left[\frac{P(H^+ | I_t^{(i)})}{P(H^- | I_t^{(i)})} \right]$. Each agent i makes a choice at the time $T^{(i)}$ when $y_t^{(i)}$ crosses one of the thresholds, θ_+ or θ_- .

Since $I_t^{(1)} = \xi_{1:t}^{(1)}$, the belief of agent 1 is described by Eq. (2.1). At each time t , agent 2 observes the resulting decision state, $d_t^{(1)}$, of agent 1 (Fig. 3.1a), where

$$d_t^{(1)} = \begin{cases} -1, & y_t^{(1)} \leq \theta_-, \\ 0, & y_t^{(1)} \in (\theta_-, \theta_+), \\ 1, & y_t^{(1)} \geq \theta_+. \end{cases} \quad (3.1)$$

The decision state captures whether agent 1 made a decision by time t , and, if so, what that decision was.

Agent 2 also makes private observations, $\xi_t^{(2)} \in \Xi$. Assuming that the two agents make private observations synchronously, the information available to agent 2 at time t is $I_t^{(2)} = (\xi_{1:t}^{(2)}, d_{1:t-1}^{(1)})$: Decision $d_t^{(1)}$ is thus not known to agent 2 until time $t + 1$. The observed decision state of agent 1 and the private observation, $\xi_t^{(2)}$, determine the belief of agent 2 at time t ,

$$y_t^{(2)} = \log \frac{P(H^+ | \xi_{1:t}^{(2)}, d_{1:t-1}^{(1)})}{P(H^- | \xi_{1:t}^{(2)}, d_{1:t-1}^{(1)})} = \underbrace{\sum_{l=1}^t \log \frac{P(\xi_l^{(2)} | H^+)}{P(\xi_l^{(2)} | H^-)}}_{\text{Priv}_{1:t}^{(2)}} + \underbrace{\log \frac{P(d_{1:t-1}^{(1)} | H^+)}{P(d_{1:t-1}^{(1)} | H^-)}}_{\text{Soc}_t^{(1)}}. \quad (3.2)$$

The last equality follows from our assumption that the observations are conditionally independent (See Proposition A.1 in Appendix A). The belief of agent 2 is thus a sum of the LLR corresponding to private observations ($\text{Priv}_{1:t}^{(2)}$), and the LLR corresponding to social information ($\text{Soc}_t^{(1)}$). The stochastic process $\text{Priv}_{1:t}^{(2)}$ is defined in the same way for both agents, and is equivalent to that of a lone observer given in Eq. (2.1).

4. Social Information. We next ask how much evidence is provided by social information, that is, by observing the decision state of other agents in the directed network in Fig. 3.1a. We show that even observations of non-decisions, $d_t^{(1)} = 0$, can be informative.

Decision Evidence As decisions are immutable, once an agent makes a choice, the social information it provides to its neighbors is fixed. Assume agent 1 chooses H^+ at time $T^{(1)}$. Then the belief of agent 2 at time $t > T^{(1)}$ is:

$$y_t^{(2)} = \text{Priv}_{1:t}^{(2)} + \text{Soc}_t^{(1)} = \log \left(\frac{P(H^+ | \xi_{1:t}^{(1)})}{P(H^- | \xi_{1:t}^{(1)})} \right) + \log \left(\frac{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^+)}{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^-)} \right).$$

Intuitively, if, for instance, $d_{T^{(1)}}^{(1)} = 1$, then agent 2 knows that $y_{T^{(1)}}^{(1)} \geq \theta_+$. Agent 1 has reached its decision based on private observations, and hence none of the social information obtained by agent 2 is redundant. The belief of agent 1 at the time of the decision, $T^{(1)}$ could exceed threshold. The evidence it has obtained may thus exceed θ_+ . However, this excess is small if the evidence obtained from each measurement is small.

The following proposition claims that, after observing a choice, agent 2 updates its belief by estimating the amount of evidence gathered by agent 1, *i.e.* by an amount close to θ_+ (Fig. 3.1b).

PROPOSITION 4.1. *If $d_{T^{(1)}}^{(1)} = \pm 1$, then*

$$\theta_{\pm} \leq \pm \log \left[\frac{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^+)}{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^-)} \right] < \theta_{\pm} + \epsilon_{T^{(1)}}^{\pm},$$

where

$$|\epsilon_{T^{(1)}}^{\pm}| \leq \sup_{\xi \in \Xi_{T^{(1)}}^{\pm}} \left(\log \frac{f_+(\xi)}{f_-(\xi)} \right),$$

and $\Xi_{T^{(1)}}^{\pm}$ is the set of all observations $\xi_{T^{(1)}}$ that trigger a $H^{\pm}$ decision exactly at time $T^{(1)}$.

We prove Proposition 4.1 in Appendix B.

Social information prior to a decision. To understand why the absence of a decision can provide information, consider the case when $\theta_+ \neq -\theta_-$. As one of the thresholds is closer to naught, in general each agent is more likely to choose one option over another by some time t . The absence of the more likely choice therefore reveals to an observer that the agent has gathered evidence favoring the alternative.

We will show that the lack of a decision is informative only in such asymmetric cases below.

DEFINITION 4.2. *The measurement distributions $P(\xi | H^+) = f_+(\xi)$ and $P(\xi | H^-) = f_-(\xi)$ are symmetric if $f_+(-\xi) = f_-(\xi)$ for every $\xi \in \Xi$. When $\theta_+ = -\theta_-$ we say that the thresholds are symmetric.*

It is frequently assumed, and experiments are frequently designed, so that threshold and measurement distributions are symmetric [50]. However, there are a number of interesting consequences when introducing asymmetries into the reward or measurement distributions [4], which suggest subjects adapt their priors or decision thresholds [27]. In examples we will assume that agents use asymmetric decisions thresholds ($\theta_+ \neq -\theta_-$) due to a known asymmetry in the 2AFC task.

We call the social information agent 2 gathers before observing a choice by agent 1, the *non-decision evidence*. Using the decomposition in Eq. (3.2), we have

$$\text{Soc}_t^{(1)} = \log \frac{P(d_t^{(1)} = 0 | H^+)}{P(d_t^{(1)} = 0 | H^-)}, \quad (4.1)$$

before agent 1 makes a decision. Let $\Theta := (\theta_-, \theta_+)$, and define the *survival probabilities* that the stochastically evolving belief remains within Θ as

$$S_{\pm}(t) = P(d_t^{(1)} = 0 | H^{\pm}) = P(y_s^{(1)} \in \Theta, 0 \leq s \leq t | H^{\pm}).$$

Then the social information provided by the absence of the decision by time t is,

$$\text{Soc}_t^{(1)} = \log \frac{S_+(t)}{S_-(t)}. \quad (4.2)$$

Note that if, for example, $S_-(t) \leq S_+(t)$ for all $t \geq 0$, then H^- decisions more often occur sooner than H^+ decisions, and $\log[S_+(t)/S_-(t)] \geq 0$. Thus, observing that an agent has not made a choice by time t provides evidence in favor of the choice that requires more time to make.

Note also that social information depends only on the thresholds and measurement distributions of both agents, and not on the specific measurements of agent 2. Until agent 1 makes a decision, *the social information is deterministic*, i.e. independent of realization. We will see that a related result holds in more complex networks.

The next proposition shows that the symmetry of the thresholds and evidence distributions implies the absence of social information until a decision is made.

PROPOSITION 4.3. *If the measurement distributions and thresholds are symmetric then $\text{Soc}_t^{(1)} = 0$ when $d_t^{(1)} = 0$.*

Proof. Let $\mathcal{S}^{(1)}(t)$ be the set of all observation sequences of agent 1, $\xi_{1:t}^{(1)}$, not leading to a decision. Then for all $0 \leq s \leq t$,

$$\sum_{l=1}^s \log \frac{f_+(\xi_l^{(1)})}{f_-(\xi_l^{(1)})} = \sum_{l=1}^s \log \frac{f_-(-\xi_l^{(1)})}{f_+(-\xi_l^{(1)})} = - \sum_{l=1}^s \log \frac{f_+(-\xi_l^{(1)})}{f_-(-\xi_l^{(1)})}.$$

Here we have used the symmetry of measurement distributions in the first equality. Since the thresholds, and hence the interval Θ , are symmetric about 0, it follows that $\sum_{l=1}^s \log \frac{f_+(\xi_l^{(1)})}{f_-(\xi_l^{(1)})} \in \Theta$, if and only if $\sum_{l=1}^s \log \frac{f_+(-\xi_l^{(1)})}{f_-(-\xi_l^{(1)})} \in \Theta$. Therefore, $\xi_{1:t}^{(1)} \in \mathcal{S}^{(1)}(t)$, if and only if $-\xi_{1:t}^{(1)} \in \mathcal{S}^{(1)}(t)$. Moreover,

$$S_+(t) = \int_{\mathcal{S}^{(1)}(t)} f_+(\xi_{1:t}^{(1)}) d\xi_{1:t}^{(1)} = \int_{\mathcal{S}^{(1)}(t)} f_-(-\xi_{1:t}^{(1)}) d\xi_{1:t}^{(1)} = \int_{\mathcal{S}^{(1)}(t)} f_-(\xi_{1:t}^{(1)}) d\xi_{1:t}^{(1)} = S_-(t),$$

and thus observing the absence of a decision of agent 1 is uninformative, as

$$\text{Soc}_t^{(1)} = \log \frac{P(d_t^{(1)} = 0 | H^+)}{P(d_t^{(1)} = 0 | H^-)} = \log \frac{S_+(t)}{S_-(t)} = 0.$$

□

We show below that when measurement distributions and thresholds are symmetric, non-decisions are uninformative in any social network.

When thresholds and measurement distributions are not symmetric then typically $S_-(t) \neq S_+(t)$ and $\log[S_+(t)/S_-(t)] \neq 0$. Thus, the absence of a decision of agent 1 can provide evidence for one of the choices. We provide a concrete example next.

Example: Belief as a Biased Random Walk. In the following example we chose the measurement distributions, $f_{\pm}(\xi)$, so that the belief increments due to private measurements satisfy $\log \frac{f_+(\xi)}{f_-(\xi)} \in \{+1, 0, -1\}$ with respective probabilities $\{p, s, q\}$. This ensures that a decision based solely on private evidence results in a belief (LLR) exactly at threshold¹. Beliefs then evolve as biased random walks on the integer lattice. Details of the model and the analysis are in Appendix C.

Agent 2 makes private observations and observes the decisions of agent 1, which are made only using private information. Agent 1 and 2's beliefs are described by Eq. (2.1) and Eq. (3.2), respectively. Realizations of these processes are shown in Fig. 4.1. The social evidence obtained by agent 2 is independent of the particular realization, until the decision of agent 1, whereas private information is realization-dependent. When thresholds are small, an expression for social evidence can be obtained explicitly (See Appendix C).

FIG. 4.1. An example of the belief evolution in a two-agent, unidirectional network shown in Fig. 3.1. (a) The belief of agent 1 is a random walk on the integers. For boundaries at $\theta_- = -1$ and $\theta_+ = 2$, initially two observations favoring H^+ are required for decision $d_2^{(1)} = 1$. The belief of agent 2 is the sum of stochastically evolving private evidence, and deterministic social evidence until observing a choice. If agent 2 observes agent 1 choosing $H^\pm$, its belief jumps by $\theta_\pm$. (b) The same processes with $\theta_- = -2$ and $\theta_+ = 40$. The parameters used, and an analysis of the belief dynamics is given in Appendix C. Code to produce all the figures in the manuscript is available at <https://github.com/Bargo727/NetworkDecisions.git>.

FIG. 4.2. Statistics of decisions in a two-agent network with (coupled), and without (uncoupled) communication from agent 1 to 2. The lower threshold is fixed at $\theta_- = -2$ throughout. (a) The fraction of times both agents selected the correct choice as a function of asymmetry in the system as measured by θ_+/θ_- . The dashed lines are asymptotic probabilities of the correct choice in the limit $\theta_+ \rightarrow \infty$ for uncoupled agents. (b) First passage time distributions for the LLR in the case $\theta_+ = 40$. (c) Relative percent difference in decision times for agent 2 in coupled versus uncoupled trajectories as a function of asymmetry in the system shown for different values of p .

First-step analysis shows that social evidence is not accrued significantly beyond the first few observations when decision thresholds are small. We show this for $\theta_- = -1$ and $\theta_+ = 2$ in Fig. 4.1a, with $\text{Soc}_t^{(1)} \rightarrow 1/2$ as $t \rightarrow \infty$ for our choice of parameter values (See Appendix C).

¹In the remainder of this work we use the same measurement distributions in all examples. We always choose p, q, s so that decisions based solely on private evidence result in a LLR exactly at threshold.

In general, social information, $\text{Soc}_t^{(1)}$, will converge in the direction of the larger threshold, in the absence of a decision. Intuitively, when $|\theta_+| > |\theta_-|$ then more observations and a longer time are required for an agent's belief to reach θ_+ , on average. In this case, observing that agent 1 has not chosen H^- after a small initial period suggests that this agent has evidence favoring H^+ .

To illustrate the impact of asymmetry in the measurement distributions, we varied the probability, p , of an observation favoring H^+ , while keeping the increments in the belief fixed. When agent 2 observes the decisions of agent 1, the probability that both reach the correct decision is larger than when they both gather information independently (See Fig. 4.2a). In particular, as p/q decreases so that private observations provide less information, the impact of social information on accuracy is strengthened. With θ_- small, little evidence is needed for an agent to choose H^- . However, as θ_+ is increased, social information pulls the belief of agent 2 in the direction of θ_+ over time, but more evidence is required to reach the upper threshold θ_+ .

Social information also affects decision times, particularly in the case of strongly asymmetric thresholds (Fig. 4.2b). An early peak in the distributions represents decisions corresponding to the smaller threshold, θ_- , while the latter peak corresponds to the opposite decision when the belief crosses $\theta_+ \gg -\theta_-$. Furthermore, as p/q increases, the difference in decision times between the agents decreases (See Fig. 4.2c), as social information speeds up agent 2's decisions more when private measurements are unreliable. Also, social information plays a more significant role for thresholds $\theta_{\pm}$ with stronger asymmetry.

Remark: The impact of social information in this example is small, unless the difference in thresholds is very large. However, this impact can be magnified in larger networks: Consider a star network in which an agent observes the decision of $N > 1$ other agents. If these are the only social interactions, the independence of private measurements implies that social information obtained by the central agent is additive. Until a decision is observed, this social information equals $N\text{Soc}_t^{(1)}$, where $\text{Soc}_t^{(1)}$ is as defined in Eq. (4.2). Thus the impact of non-decision information can be magnified in larger networks. However, as these cases are computationally more difficult to deal with, we do not discuss them in detail here.

5. Two-agent recurrent network. We next assume that the two agents can observe, and react to, each other's choices. As in the previous section, we assume that at each time the agents synchronously make independent, private observations, $\xi_t^{(i)} \in \Xi$, and update their beliefs. The agents then observe each other's decision state $d_t^{(j)}$ ($j \neq i$), and use this social information to update their belief again. Knowing that a belief has been updated, and observing the resulting decision state can provide new information about an agent's belief. Hence, social information can continue to be exchanged, even before the next private observation is made (See Fig. 5.1 for illustration).

We describe several properties of this evidence exchange process: As in the case of a unidirectional network, social information is additive. It evolves deterministically up to the time of a decision of the observed agent. Once a decision is made, the social information that is communicated approximately equals the belief threshold (θ_+ or θ_-) crossed by the LLR of the deciding agent. We also show that the exchange of social information after an observation terminates in a finite number of steps either when indecisions provide

FIG. 5.1. In a recurrent network of two agents the LLRs, $y_{t,n}^{(1)}$ and $y_{t,n}^{(2)}$, of the two observers are updated recursively. Agents update their belief after private observations, $\xi_t^{(i)}$, as well as observations of the subsequent series of decision states, $d_{t,n}^{(j)}$, of their neighbor ($j \neq i$). After a finite number of steps, N_t , no further information can be obtained by observing each other's decision, and the two agents make their next private observation, $\xi_{t+1}^{(i)}$, synchronously. The process continues until one of the agents makes a decision.

no further social information, or when one of the agents makes a choice.

Such information exchange has been discussed in the literature on *common knowledge* and rational learning in social networks [33]. This body of work shows that rational agents that repeatedly announce their preferred choice must eventually reach agreement [2, 18, 22, 23, 38]. Typically, it is assumed that information is shared by announcing a preference that can be changed as more information is received. Our assumption that decisions are immutable means that agreement is not guaranteed.

We show how social information is exchanged using induction. We describe the basic case following an observation at $t = 1$ in some detail. Following exchanges are similar, as the belief is updated recursively.

Social information exchange after the first observation. After the first private observation, $\xi_1^{(i)} \in \Xi$, at $t = 1$, the beliefs of the agents are $y_{1,0}^{(1)}, y_{1,0}^{(2)}$, where $y_{1,0}^{(i)} = \text{Priv}_1^{(i)}$. Let

$$\Pi = \left\{ \log \frac{P(\xi|H^+)}{P(\xi|H^-)} \mid \xi \in \Xi \right\} \quad (5.1)$$

be the set of all possible increments due to a private observation, so that $\text{Priv}_j^{(i)} \in \Pi, i = 1, 2$. As the set of possible observations, Ξ , is finite, so is Π . We will describe how the two agents exchange social information with their neighbor until observing a decision, or until no further information can be exchanged. The second subscript, n in $y_{1,n}^{(j)}$, denotes the steps in this subprocess of social information exchange preceding a decision, or next private measurement.

We again associate a belief, $y_{1,n}^{(i)}$, with a corresponding decision state, $d_{1,n}^{(i)}$, as in Eq. (3.1). If neither of the first two private observations leads to a decision, then $d_{1,0}^{(i)} = 0$, and $y_{1,0}^{(1)}, y_{1,0}^{(2)} \in \Theta$, where, $\Theta \equiv (\theta_-, \theta_+)$ for $i = 1, 2$. Importantly, the fact that agent i observed that its counterpart did not decide means that they know $y_{1,0}^{(j)} \in \Theta$, for $i \neq j$.

To update their belief, agents compare the probability of all available evidence under the two hypotheses, $P(\xi_1^{(i)}, d_{1,0}^{(j)}|H^+)$, and $P(\xi_1^{(i)}, d_{1,0}^{(j)}|H^-)$. As $d_{1,0}^{(j)}$ is independent of $\xi_1^{(i)}$ for $i \neq j$, their updated beliefs are

$$\begin{aligned} y_{1,1}^{(i)} &= \log \frac{P(\xi_1^{(i)}, d_{1,0}^{(j)}|H^+)}{P(\xi_1^{(i)}, d_{1,0}^{(j)}|H^-)} = \log \frac{P(\xi_1^{(i)}|H^+)}{P(\xi_1^{(i)}|H^-)} + \log \frac{P(d_{1,0}^{(j)} = 0|H^+)}{P(d_{1,0}^{(j)} = 0|H^-)} \\ &= y_{1,0}^{(i)} + \log \frac{P(d_{1,0}^{(j)} = 0|H^+)}{P(d_{1,0}^{(j)} = 0|H^-)} = \text{Priv}_1^{(i)} + \underbrace{\log \frac{P(y_{1,0}^{(j)} \in \Theta|H^+)}{P(y_{1,0}^{(j)} \in \Theta|H^-)}}_{\text{Soc}_{1,1}} \\ &= \text{Priv}_1^{(i)} + \text{Soc}_{1,1}. \end{aligned}$$

We omit the superscripts on the social information, since $\text{Soc}_{1,1}^{(1)} = \text{Soc}_{1,1}^{(2)}$. Non-decisions provide the same social evidence to both agents because we assumed the agents are identical. Since the agents know the measurement distributions, $f_{\pm}(\xi)$, the survival probabilities, $P(y_{1,0}^{(i)} \in \Theta|H^+)$, can be computed as in the previous section. Depending on the thresholds and measurement distributions, $\text{Soc}_{1,1}$ may be positive, negative, or zero.

If $y_{1,1}^{(1)}, y_{1,1}^{(2)} \in \Theta$, no decision is made after the first exchange of social information, and $d_{1,1}^{(i)} = 0$ for $i = 1, 2$. In this case, agent i knows that $y_{1,1}^{(j)} \in \Theta$, for $j \neq i$, and so

$$\theta_- < y_{1,0}^{(j)} + \text{Soc}_{1,1} < \theta_+.$$

Thus observing $d_{1,1}^{(j)} = 0$ informs agent i that $y_{1,0}^{(j)} \in \Theta \cap (\Theta - \text{Soc}_{1,1}) \equiv \Theta_{1,1}$, for $i \neq j$. More precisely,

$$P(d_{1,1}^{(j)} = 0|d_{1,0}^{(j)} = 0, H^+) = P(y_{1,0}^{(j)} \in \Theta_{1,1}|H^+). \quad (5.2)$$

Any initial measurement $\xi_1^{(j)}$ that would lead to a belief $y_{1,0}^{(j)} \notin \Theta_{1,1}$ would have lead to a decision at this point. This would end further evidence accumulation. Thus the other agent either observes a decision, or knows that $y_{1,0}^{(j)} \in \Theta_{1,1}$.

We will deal with the consequence of observing a decision subsequently, and first consider only observations that do not lead to a decision. For such observations, we can compute the probability of available evidence under either hypothesis by noting that some sequences of decisions have zero probability:

$$\begin{aligned}
 & P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(i)} = 0, d_{1,0}^{(j)} = 0, d_{1,1}^{(j)} = 0 | H^\pm) \\
 &= \sum_{D=0,1,-1} P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(i)} = D, d_{1,0}^{(j)} = 0, d_{1,1}^{(j)} = 0 | H^\pm) \\
 &= P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(i)} = 0, d_{1,1}^{(j)} = 0 | H^\pm) \\
 &= P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(j)} = 0 | H^\pm).
 \end{aligned} \tag{5.3}$$

The second equality follows from the fact that an observation $\xi_1^{(i)}$ leads to $\text{Priv}_1^{(i)} \in \Theta_{1,1}$, which is inconsistent with $d_{1,1}^{(i)} = \pm 1$. Therefore, $P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(i)} = \pm 1, d_{1,0}^{(j)} = 0, d_{1,1}^{(j)} = 0 | H^\pm) = 0$. The last equality follows from immutability of decisions.

We also observe that

$$\begin{aligned}
 P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(j)} = 0 | H^\pm) &= P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0 | H^\pm) P(d_{1,1}^{(j)} = 0 | \xi_1^{(i)}, d_{1,0}^{(i)} = 0, H^\pm) \\
 &= P(\xi_1^{(i)} | H^\pm) P(d_{1,1}^{(j)} = 0 | d_{1,0}^{(i)} = 0, H^\pm),
 \end{aligned} \tag{5.4}$$

where we used

$$P(d_{1,1}^{(j)} = 0 | \xi_1^{(i)}, d_{1,0}^{(i)} = 0, H^\pm) = P(d_{1,1}^{(j)} = 0 | d_{1,0}^{(i)} = 0, H^\pm),$$

as $d_{1,1}^{(j)}$ depends only on the observed decision of agent i , and $P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0 | H^+) = P(\xi_1^{(i)} | H^+)$, because $d_{1,0}^{(i)} = 0$ is implied by $\xi_1^{(i)}$ generating $\text{Priv}_1^{(i)} \in \Theta_{1,1}$.

Combining Eqs. (5.2) and (5.4) gives the updated beliefs,

$$\begin{aligned}
 y_{1,2}^{(i)} &= \log \frac{P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(i)} = 0, d_{1,0}^{(j)} = 0, d_{1,1}^{(j)} = 0 | H^+)}{P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, d_{1,1}^{(i)} = 0, d_{1,0}^{(j)} = 0, d_{1,1}^{(j)} = 0 | H^-)} \\
 &= \text{Priv}_1^{(i)} + \log \frac{P(y_{1,0}^{(j)} \in \Theta_{1,1} | H^+)}{P(y_{1,0}^{(j)} \in \Theta_{1,1} | H^-)} \\
 &:= \text{Priv}_1^{(i)} + \text{Soc}_{1,2}.
 \end{aligned}$$

Again, if $y_{1,2}^{(i)} \in \Theta$, neither agent makes a decision, and both will observe $d_{1,2}^{(i)} = 0$.

The same argument used to obtain Eq. (5.3) and Eq. (5.4) shows that for any $l > 0$,

$$\begin{aligned}
 & P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, \dots, d_{1,l}^{(i)} = 0, d_{1,0}^{(j)} = 0, \dots, d_{1,l}^{(j)} = 0 | H^\pm) \\
 &= P(\xi_1^{(i)}, d_{1,l-1}^{(i)} = 0, d_{1,l}^{(j)} = 0 | H^\pm) = P(\xi_1^{(i)} | H^+) P(d_{1,l}^{(j)} = 0 | d_{1,l-1}^{(i)} = 0, H^+),
 \end{aligned}$$

when $\xi_1^{(i)}$ generates $\text{Priv}_1^{(i)} \in \Theta_{1,l-1}$, where $\Theta_{1,l-1} \equiv \bigcap_{n=0}^{l-1} (\Theta - \text{Soc}_{1,n})$. Thus if the l -th exchange of social information results in no decision, each agent updates its belief recursively as

$$y_{1,l}^{(i)} = \text{Priv}_1^{(i)} + \log \frac{P(y_{1,0}^{(j)} \in \Theta_{1,l-1} | H^+)}{P(y_{1,0}^{(j)} \in \Theta_{1,l-1} | H^-)} \equiv \text{Priv}_1^{(i)} + \text{Soc}_{1,l}, \tag{5.5}$$

where $\text{Soc}_{1,0} = 0$.

This exchange of social information continues until one of the agents makes a choice, or when the absence of a decision does not lead to new social information [23]. The second case occurs at a step, N_1 , at which the absence of a decision would not provide any further information about the belief of either agent, that is when

$$\Pi \cap \Theta_{1,N_1} = \Pi \cap \Theta_{1,N_1+1}, \tag{5.6}$$

where Π is defined in Eq. (5.1). In this case, $\text{Soc}_{1,N_1} = \text{Soc}_{1,N_1+1}$, and, if neither agents decide, their beliefs would not change at step $N_1 + 1$. As both agents know that nothing new is to be learned from observing their neighbor, they then make the next private observation.

We denote the total social information gained after the exchange is complete by Soc_1 , and the belief at the end of this social information exchange by $y_1^{(i)} = y_{1,N_1}^{(i)} = \text{Priv}_1^{(i)} + \text{Soc}_1$. If no decision is made at this point, then agent j knows that $y_{1,0}^{(j)} = \text{Priv}_1^{(i)} \in \left(\bigcap_{n=0}^{N_1} (\Theta - \text{Soc}_{1,n})\right) \equiv \Theta_1$.

The process can be explained simply: The absence of a decision provides sequentially tighter bounds on the neighbor's belief. When the agents can conclude that these bounds do not change from one step to the next, the absence of a decision provides no new information, the exchange of social information ends, and both agents make the next private observation.

Importantly, this process is *deterministic*: Until a decision is made, the social information gathered on each step is the same across realizations, *i.e.* independent of the private observations of the agent.

Social information exchange after an observation at time $t > 1$. The integration of private information from each individual measurement, $\xi_t^{(i)}$, is again followed by an equilibration process. The two agents observe each others' decision states until nothing further can be learned. To describe this process, we proceed inductively, and extend the definitions introduced in the previous section.

Let $\text{Priv}_t^{(i)} = \log \left[P(\xi_t^{(i)} | H^+) / P(\xi_t^{(i)} | H^-) \right]$ be the private information obtained from an observation at time t . For the inductive step assume that each equilibration process after the observation at time t ends either in a decision, or allows each agent i to conclude that the accumulated private evidence satisfies $\sum_{k=1}^t \text{Priv}_k^{(j)} \in \Theta_t$ for some set $\Theta_t \subset \Theta$. Note that Θ_1 was defined above, and we will define the other sets in the sequence recursively. We have proved the base step above, since we have shown that following equilibration after observation $\xi_1^{(i)}$, either one of the agents makes a decision, or each agent i knows that the private information of its counterpart satisfies $\text{Priv}_1^{(j)} \in \Theta_1$ for $j \neq i$. We then define

$$\Theta_{1:t} = \left\{ \text{Priv}_{1:t}^{(j)} \mid \sum_{k=1}^l \text{Priv}_k^{(j)} \in \Theta_l \text{ for all } 1 \leq l \leq t \right\}.$$

Thus, agent $i \neq j$ knows that any sequence $\text{Priv}_{1:t}^{(j)}$ that did not lead to a decision by agent j must lie in $\Theta_{1:t}$, and this is the largest such set. Note that the condition in this definition is given in terms of a sum of LLRs given by private information, so that $\Theta_{1:t}$ is not the same as $\{\Theta_1, \dots, \Theta_t\}$.

To define the social information gathered during equilibration following the observation at time t , let $\Theta_{t,0} = \Theta_{t-1}$, $\text{Soc}_{t,0} = \text{Soc}_{t-1}$, and set

$$\text{Soc}_{t,l} := \log \frac{P(\text{Priv}_{1:t-1}^{(i)} \in \Theta_{1:t-1}, \text{ and } \text{Priv}_{1:t}^{(i)} \in \Theta_{t,l-1} | H^+)}{P(\text{Priv}_{1:t-1}^{(i)} \in \Theta_{1:t-1}, \text{ and } \text{Priv}_{1:t}^{(i)} \in \Theta_{t,l-1} | H^-)}, \quad (5.7)$$

for $l \geq 1$, where $\Theta_{t,l-1} \equiv \bigcap_{n=0}^{l-1} (\Theta - \text{Soc}_{t,n})$.

THEOREM 5.1. *Assume that, in a recurrent network, two agents have made a sequence of private observations, $\xi_{1:t}$, followed by l observations of the subsequent decision states of each other. If neither agent has made a decision, then the belief of each is given by*

$$y_{t,l}^{(i)} = \sum_{k=1}^t \text{Priv}_k^{(i)} + \text{Soc}_{t,l}, \quad (5.8)$$

for $1 \leq t, i = 1, 2$. The exchange of social information terminates in a finite number of steps after an observation, either when a decision is made, or after no further social information is available at some step $l = N_t$. The private evidence in Eq. (5.8) is a random variable (depends on realization), while the social evidence is independent of realization, and equal for both agents.

Proof. The proof is inductive, with the basic case proved in the previous section. The inductive step follows similarly. By a slight abuse of notation, let $d_{1:t,1:l}^{(i)}$ denote the sequence of decision states up to the l -th step in the equilibration process. If no decision has been made at this step, we write $d_{t,l}^{(i)} = 0$. This implies that no previous decision has been made, and we denote this by writing $d_{1:t,1:l}^{(i)} = 0$.

In the induction step, we assume that equilibration following private observation $\xi_t^{(i)}$ terminates on step N_t . Conditional independence of the observations implies that $P(\xi_{1:t+1}^{(i)}, d_{1:t,1:N_t}^{(i)} = 0, d_{1:t,1:N_t}^{(j)} = 0 | H^\pm) = P(\xi_{t+1}^{(i)} | H^\pm) P(\xi_{1:t}^{(i)}, d_{1:t,1:N_t}^{(i)} = 0, d_{1:t,1:N_t}^{(j)} = 0 | H^\pm)$, so that

$$y_{t,0}^{(i)} = \log \frac{P(\xi_{1:t+1}^{(i)}, d_{1:t,1:N_t}^{(i)} = 0, d_{1:t,1:N_t}^{(j)} = 0 | H^+)}{P(\xi_{1:t+1}^{(i)}, d_{1:t,1:N_t}^{(i)} = 0, d_{1:t,1:N_t}^{(j)} = 0 | H^-)} = \sum_{k=1}^{t+1} \text{Priv}_k^{(i)} + \text{Soc}_{t+1,0},$$

where we used $\text{Soc}_{t+1,0} = \text{Soc}_t$.

Suppose that no decision has been made in the following $l \geq 0$ equilibration steps, so that $d_{1:t+1,1:l}^{(i)} = 0$ for $i = 1, 2$. For all sequences of measurements $\xi_{1:t+1}^{(i)}$ that are consistent with this absence of a decision, we have

$$P(\xi_{1:t+1}^{(i)}, d_{1:t+1,1:l}^{(i)} = 0, d_{1:t+1,1:l}^{(j)} = 0 | H^\pm) = P(\xi_{1:t+1}^{(i)}, d_{t+1,l-1}^{(i)} = 0, d_{t+1,l}^{(j)} = 0 | H^\pm)$$

by marginalizing over all intervening decisions, and the final decision of agent i , as in Eq. (5.3). As in Eq. (5.4) we have

$$P(\xi_{t+1}^{(i)}, d_{t+1,l-1}^{(i)} = 0, d_{t+1,l}^{(j)} = 0 | H^\pm) = P(\xi_{1:t+1}^{(i)} | H^\pm) P(d_{t+1,l}^{(j)} = 0 | d_{t+1,l-1}^{(i)} = 0, H^\pm),$$

and therefore,

$$y_{t,l+1}^{(i)} = \log \frac{P(\xi_{1:t+1}^{(i)} | H^+) P(d_{t+1,l}^{(j)} = 0 | d_{t+1,l-1}^{(i)} = 0, H^+)}{P(\xi_{1:t+1}^{(i)} | H^-) P(d_{t+1,l}^{(j)} = 0 | d_{t+1,l-1}^{(i)} = 0, H^-)} = \sum_{k=1}^{t+1} \text{Priv}_k^{(i)} + \text{Soc}_{t+1,l+1},$$

where $\text{Soc}_{t+1,l+1}$ is defined in Eq. (5.7). This exchange of social information stops at step $l = N_{t+1}$ at which point $\Pi \cap \Theta_{t+1,N_t} = \Pi \cap \Theta_{t,N_t+1}$, and neither agent learns anything further from the absence of a decision by its counterpart.

□

Belief update after a decision. Previously we described social information exchange in the absence of a decision. The following proposition shows what happens when the belief of one of the agents crosses a threshold.

PROPOSITION 5.2. *Suppose that in a recurrent two-agent network, agent i makes a decision after a private observation at time $T^{(i)}$, during the n th step of the subsequent social information exchange process, $d_{T^{(i)},n}^{(i)} = \pm 1$. Then agent $j \neq i$, updates its belief as*

$$y_{T^{(i)},n+1}^{(j)} = \text{Priv}_{1:T^{(i)}}^{(j)} + \text{Soc}_{T^{(i)},n+1},$$

where

$$\theta_\pm \leq \pm \text{Soc}_{T^{(i)},n+1} < \theta_\pm + \epsilon_{T^{(i)},n}^\pm \quad \text{if } d_{T^{(i)},n}^{(i)} = \pm 1.$$

In the case $n = 0$,

$$\pm \epsilon_{T^{(i)},0}^\pm \leq \sup_{\xi \in \Xi_{T^{(i)}}^\pm} \left(\log \frac{f_+(\xi)}{f_-(\xi)} \right),$$

and $\Xi_{T^{(i)}}^{\pm}$ is the set of all observations leading to a choice of $H^{\pm}$ at timestep $(T^{(i)}, 0)$. If $n > 0$,

$$\pm \epsilon_{T^{(i)}, n}^{\pm} \leq \pm (\text{Soc}_{T^{(i)}, n} - \text{Soc}_{T^{(i)}, n-1}).$$

We prove Proposition 5.2 in Appendix D.

As we discussed above, agent j may obtain social information before observing the decision of its counterpart. However, this earlier social information is subsumed by the information obtained from observing a decision. If the two agents have different decision thresholds, the expression for the post-decision belief is more involved, but still computable. For simplicity we forgo further discussion of this case.

We next show why this process simplifies when the decision thresholds and evidence distributions are symmetric.

PROPOSITION 5.3. *When the distributions f_+ and f_- are symmetric and the agents have the same symmetric thresholds ($\pm\theta_{\pm} = \theta$), then $\text{Soc}_t = 0$ until private evidence leads one of the agents to make a decision. Thus there is no exchange of social information until one of the agents makes a decision.*

Proof. The argument is similar to that used to prove Proposition 4.3. We can proceed inductively again: If the two agents have not made a decision after the first observation, by symmetry this does not provide any evidence for either hypothesis $H = H^{\pm}$. Observing each other's decisions after this first observation hence results in no new information,

$$\text{Soc}_{1,1} = \log \frac{P(y_{1,0}^{(j)} \in \Theta|H^+)}{P(y_{1,0}^{(j)} \in \Theta|H^-)} = 0.$$

Therefore, the equilibration process terminates immediately, and both agents proceed to the next observation. As shown in Proposition 4.3, further observations provide no new social information, unless a decision is made. The two agents therefore continue to accumulate private information, until one of them makes a decision. $\square$

Fig. 5.2 provides examples of evidence accumulation in a two-agent recurrent network. In Fig. 5.2a,b, we illustrate the process with relatively small thresholds and show how the intervals Θ_n shrink at the end of each social evidence exchange (equilibration) following a private observation. Note that the sequence of intervals is the same in both examples because the social information exchange process is deterministic. In this example, equilibration occurs after two steps. In Fig. 5.2c,d, we provide an example with strongly asymmetric decision thresholds. The equilibration of social information between private observations requires several more steps, as shown in the inset of Fig. 5.2c. These examples also illustrate that the beliefs of the two agents do not have to converge, and they do not need to agree on a choice, in contrast to classic studies of *common knowledge* [2, 22, 33].

6. Accumulation of Evidence on General Networks. In networks that are not fully connected, rational agents need to take into account the impact of the decisions of agents that they do not observe directly. To do so they marginalize over all unobserved decision states. This computation can be complex, even when thresholds and evidence distributions are symmetric.

To illustrate we begin by describing the example of agents with symmetric decision thresholds and measurement distributions on a directed chain. Symmetry makes the computations more transparent, as the absence of a decision is not informative about the belief of any agent, hidden or visible. Social information is therefore only communicated when an agent directly observes a choice. Such an observation leads to a jump in the agent's belief, and can initiate a cascade of decisions down the chain [10].

We note that once an agent in the network makes a decision, symmetry can be broken: Agents know that all others who have observed the decision have evidence favoring the observed choice. As we have seen, once symmetry is broken even the absence of a choice can provide social information. In this case rational agents must equilibrate their beliefs using known survival probabilities, as discussed in the previous section. We do not analyze this case.

FIG. 5.2. The belief (LLR) of agents 1 and 2 in a recurrent network with asymmetric thresholds when (a) the agents agree, and (b) agents disagree on a final choice. Also shown are the intervals Θ_1, Θ_2 , resulting from the equilibration following the first two observations. Each agent knows that the LLR of its counterpart lies in this interval after equilibration ends. Although the beliefs evolve stochastically, the sequence Θ_1, Θ_2 is fixed across realizations. Here, $p = e/7$, $q = 1/7$, and s is determined from $s = 1 - p - q$. Also, $\theta_+ = 2$ for (a)-(b) and $\theta_+ = 20$ for (c) - (d). $\theta_- = -1$ for (a)-(b) and $\theta_- = -8$ for (c)-(d). Inset: Social information obtained from equilibration converges in seven steps at the indicated time.

6.1. Terminology and Notation. In a social network of N agents, we again assume that each agent makes a private observation at every time step. After incorporating the evidence from this observation the agent then updates its decision state and shares it with its neighbors. A directed edge from agent j to i , denoted by $j \rightarrow i$, means that agent j communicates its decision state to agent i . The set of neighbors that agent i observes is denoted by $N^{(i)} = \{j : j \rightarrow i\}$.

Agent i thus receives social information from all agents in $N^{(i)}$. However, agent i also needs to take into account decisions of *unobserved* agents. We define the set of all agents not visible to agent i as

$$U^{(i)} = \{j : j \notin N^{(i)}, \text{ and } j \neq i\}.$$

Thus all agents in the network, besides agent i , belong to $N^{(i)}$ or $U^{(i)}$. Let $W = \{1, 2, 3, \dots, N\}$. Then $W = N^{(i)} \cup U^{(i)} \cup \{i\}$ and $N^{(i)} \cap U^{(i)}, N^{(i)} \cap \{i\}, U^{(i)} \cap \{i\}$ all equal $\emptyset$.

We denote the set of decisions by the neighbors of agent i following the observation at time t by $d_t^{N^{(i)}} = \{d_t^{(k)} : k \in N^{(i)}\}$. Similarly, the set of the decisions by unobserved agents is denoted by $d_t^{U^{(i)}}$. More generally, $d_{1:t}^{N^{(i)}}$ denotes the sequence of decision states of the neighbors of agent i up to and including the decision following the observation at time t : $d_{1:t}^{N^{(i)}} = \{d_s^{N^{(i)}} : 1 \leq s \leq t\}$. We will see that in the case of symmetric thresholds and observations no equilibration occurs, so these decision states describe information available to each agent in the network completely until one of the agents makes a choice. At time t , the private and social observations obtained by agent i are therefore $I_t^{(i)} = \{\xi_{1:t}^{(i)}, d_{1:t-1}^{N^{(i)}}\}$. As before, we denote the private information by $\text{Priv}_{1:t}^{(i)} = \log \left(P(\xi_{1:t}^{(i)} | H^+) / P(\xi_{1:t}^{(i)} | H^-) \right)$.

6.2. Non-decisions. We next show that in social networks of agents with symmetric thresholds and measurement distributions, the agent's beliefs have two properties that simplify computations: An absence

of a decision is uninformative, and once a decision is observed, the resulting social information is additive.

PROPOSITION 6.1. *Assume all agents have symmetric thresholds and evidence distributions. If neither agent i , nor any of its neighbors in $N^{(i)}$ have made a decision by time t , then the belief of agent i is*

$$y_t^{(i)} = \log \frac{P(I_t^{(i)}|H^+)}{P(I_t^{(i)}|H^-)} = \log \frac{P(\xi_{1:t}^{(i)}, d_{1:t}^{N(i)} = 0|H^+)}{P(\xi_{1:t}^{(i)}, d_{1:t}^{N(i)} = 0|H^-)} = \text{Priv}_{1:t}^{(i)}.$$

Thus social information is only obtained from observing a choice.

Proof. The claim follows from an argument similar to that given in Sections 4 and 5. The main difference is that agent i here marginalizes over unobserved decision states.

By definition, the only direct social information agent i has is the set of decision states of its neighbors $d_{1:t}^{N(i)} = 0$ to time t . First we can split the private evidence from the social evidence conditioned on the agent's observations,

$$\begin{aligned} y_t^{(i)} &= \log \frac{P(I_t^{(i)}|H^+)}{P(I_t^{(i)}|H^-)} = \log \frac{P(\xi_{1:t}^{(i)}, d_{1:t}^{N(i)} = 0|H^+)}{P(\xi_{1:t}^{(i)}, d_{1:t}^{N(i)} = 0|H^-)} \\ &= \log \frac{P(\xi_{1:t}^{(i)}|H^+)}{P(\xi_{1:t}^{(i)}|H^-)} + \log \frac{P(d_{1:t}^{N(i)} = 0|\xi_{1:t}^{(i)}, H^+)}{P(d_{1:t}^{N(i)} = 0|\xi_{1:t}^{(i)}, H^-)}. \end{aligned}$$

We therefore need to show that

$$\log \frac{P(d_{1:t}^{N(i)} = 0|\xi_{1:t}^{(i)}, H^+)}{P(d_{1:t}^{N(i)} = 0|\xi_{1:t}^{(i)}, H^-)} = 0. \quad (6.1)$$

Marginalizing over all possible combinations of decision states for unseen neighbors, $d_{1:t}^{U(i)}$, and observations of neighbors, $\xi_{1:t}^{N(i)}$, yields

$$\log \frac{P(d_{1:t}^{N(i)} = 0|\xi_{1:t}^{(i)}, H^+)}{P(d_{1:t}^{N(i)} = 0|\xi_{1:t}^{(i)}, H^-)} = \log \frac{\sum_{\mathcal{C}_t^{N(i)}} P(d_{1:t}^{N(i)} = 0, \xi_{1:t}^{N(i)}, d_{1:t}^{U(i)}|\xi_{1:t}^{(i)}, H^+)}{\sum_{\mathcal{C}_t^{N(i)}} P(d_{1:t}^{N(i)} = 0, \xi_{1:t}^{N(i)}, d_{1:t}^{U(i)}|\xi_{1:t}^{(i)}, H^-)}, \quad (6.2)$$

where $\mathcal{C}_t^{N(i)}$ denotes all possible chains $(\xi_{1:t}^{N(i)}, d_{1:t}^{U(i)})$ of social and private observations that neighbors of agent i could have made without reaching a decision by time t . For a decision history of the unobserved agents $d_{1:t}^{U(i)}$, let $-d_{1:t}^{U(i)}$ be the opposite decision history, flipping the sign of each decision in the vector $d_{1:t}^{U(i)}$, and leaving non-decisions unaffected. The vector of zeroes equals its negation. Likewise, for a private observation history, $\xi_{1:t}^{N(i)}$, of the neighbors, let $-\xi_{1:t}^{N(i)}$ be the opposite sequence of observations. Symmetry in thresholds and measurement distributions implies $P(y_t^{(j)} = z|H^+) = P(y_t^{(j)} = -z|H^-)$ for all $z \in \Xi$. By symmetry,

$$P(d_{1:t}^{N(i)} = 0, \xi_{1:t}^{N(i)}, d_{1:t}^{U(i)}|\xi_{1:t}^{(i)}, H^+) = P(d_{1:t}^{N(i)} = 0, -\xi_{1:t}^{N(i)}, -d_{1:t}^{U(i)}|\xi_{1:t}^{(i)}, H^-). \quad (6.3)$$

Therefore in Eq. (6.2), for every chain $(\xi_{1:t}^{N(i)}, d_{1:t}^{U(i)})$ appearing in the numerator, its negation $(-\xi_{1:t}^{N(i)}, -d_{1:t}^{U(i)})$ also appears in the denominator. The sums in the numerator and denominator on the right hand side of Eq. (6.2) are equal, and hence Eq. (6.1) holds.

It remains to show that none of the agents engage in equilibration. Note that after its first private measurement, no agent i obtains further information from a non-decision of its neighbors. The equilibration process therefore does not start. The same holds after further private measurements by induction.

□

Thus, before decisions are made, agents independently accumulate evidence. Decision states are only informative when an agent makes a choice. As in the case of two agents, a choice by any agent will lead to

a jump in the belief of all observing neighbors.

Remark: Suppose that two neighboring agents, j and k , both observe an agent making a choice. Agent k now knows that the belief of agent j must have increased by a known amount due to the mutually observed decision. Agent k thus gains knowledge about the belief of agent j , even if j does not immediately decide. The symmetry is therefore broken, and henceforth the situation is similar to the asymmetric case discussed in Section 5. This introduces additional drift, and potentially the need for equilibration, until agent k or j make decisions. We thus concentrate on the dynamics up to the first decision, but discuss the general case in more detail in some of our examples.

6.3. Example of Marginalization: 3-Agents on a Line. We demonstrate the computations needed to account for unobserved agents using the example shown in Fig. 6.1. In this case a decision by agent 3 is not observed by any other agent. If agent 1 makes a decision first, then agent 2 continues to accumulate information as in the case of two unidirectionally coupled agents discussed in Section 3. It is therefore sufficient to only consider the case when agent 2 makes a decision first at time $T^{(2)}$.

As shown in Proposition 6.1, agent 3 has no social information before observing a decision by agent 2. After observing a decision by agent 2, agent 3 updates its belief by marginalizing over possible decision histories of agent 1:

$$P(d_{T^{(2)}}^{(2)} = \pm 1 | H^\pm) = P(d_{T^{(2)}}^{(2)} = \pm 1, d_{T^{(2)}}^{(1)} = 0 | H^\pm) + \sum_{s=1}^{T^{(2)}} \sum_{d=\pm 1} P(d_{T^{(2)}}^{(2)} = \pm 1, d_{s-1}^{(1)} = 0, d_s^{(1)} = d | H^\pm). \quad (6.4)$$

Intuitively, a choice of agent 2 can be triggered by either: (a) A private observation leading to the belief $y_t^{(2)}$ reaching one of the thresholds, $\theta_\pm$. This possibility is captured by all terms in Eq. (6.4) except

$$P(d_{T^{(2)}}^{(2)} = \pm 1, d_{T^{(2)}-1}^{(1)} = 0, d_{T^{(2)}}^{(1)} = \pm 1 | H^\pm); \quad (6.5)$$

or (b) The decision of agent 1 causing a jump in the belief $y_t^{(2)}$ above threshold θ_+ or below threshold θ_- . This second possibility is captured by the term in Eq. (6.5). An argument equivalent to that in Proposition 4.1 shows that the social information communicated in the first case is close to $\pm\theta$ for $d_{T^{(2)}}^{(2)} = \pm 1$. However, in the second case the belief of agent 2 at the time of decision is in the range $y_{T^{(2)}}^{(2)} \in [\theta, 2\theta]$ for $d_{T^{(2)}}^{(2)} = +1$ or $y_{T^{(2)}}^{(2)} \in [-2\theta, -\theta]$ for $d_{T^{(2)}}^{(2)} = -1$, modulo a small correction. Agent 3 updates its belief by weighting both possibilities, and hence increases its belief by an amount greater than θ . We provide details of this argument in Appendix E, and show that the excess increase in belief of agent 2 can be large, even when the increments in private information, $\text{Priv}_t^{(2)}$, are small.

In Fig. 6.2, we illustrate how including social information affects key metrics in decision-making: decision-time and accuracy. When p/q is not too large, and hence the signal-to-noise ratio (SNR) of observations is low, the exchange of social information impacts both the accuracy and response time of agents significantly. Agents in a social network choose the correct decision more often compared to isolated agents, as the impact of social information is significant. When p/q , and hence SNR, is large, including social information does

FIG. 6.1. Agents on a line. Agent 3 observes the decisions of agent 2, but not of agent 1. However, agent 3 still needs to take into account the possibility that a choice by agent 1 has caused a decision by agent 2.

FIG. 6.2. The performance of the three agents on a directed line is better than that of independent agents. (a) The fraction of trials for which all three agents make the correct choice. This quantity is larger when agents are allowed to exchange information (coupled), than when agents make decisions independently (uncoupled). (b) Time required for all three agents to make a decision. Average decision times are also smaller when agents exchange information. Here, $|\theta_{\pm}| = 30$, and ratio p/q determines the noisiness of the measurements, as described in Appendix C.

not benefit the agents strongly: accuracy and decision time of agents in a network are similar to isolated agents. Private information dominates the decision making process in these cases.

Marginalization becomes more complex when the unseen component, $U^{(i)}$, of the network is larger. For instance, if we consider a long, directed line of n agents, when the k -th agent makes a decision, the $k+1$ -th agent must marginalize over the preceding $k-1$ agents. If the resulting jump in the belief, $y_t^{(k+1)}$, exceeds 2θ , this triggers an immediate decision in agent $k+1$, and all successive agents. This is equivalent to herding behavior described in the economics literature [1, 5].

7. Three-Agent Cliques. In cliques, or all-to-all coupled networks, all agents can observe each others' decisions, and $U^{(i)} = \emptyset$ for all i . This simplifies the analysis, as no agent needs to marginalize over the decision states of unobserved agents. We start by discussing the case of three-agent cliques in some detail, and proceed to cliques of arbitrary size in the next section.

As we are assuming symmetry, social evidence is shared only after the first agent makes a choice. Without loss of generality, we assume that this was agent 3, and that $d_{T^{(3)}}^{(3)} = 1, d_{T^{(3)}}^{(1)} = d_{T^{(3)}}^{(2)} = 0$. There is a small probability that private information leads to a concurrent decision by multiple agents, and we describe this case at the end of the section. The social information obtained from a decision by agent 3 may or may not drive agent 1 or 2 to a decision. Both the presence and absence of a decision by either agent reveals further social information to its undecided counterpart, and we assume that all undecided agents wait until no further social information can be obtained before gathering further private information. We will see that, akin to the equilibration process, there can be a number of decision-making rounds before evidence accumulation continues.

There are three possible outcomes following the decision of agent 3 (See Fig. 7.1): If $d_{T^{(3)}}^{(3)} = 1$ and $y_{T^{(3)}}^{(i)} \geq 0$, for $i = 1, 2$, then both remaining agents decide immediately, $d_{T^{(3)},1}^{(i)} = 1$, and the evidence accumulation process stops. We therefore only examine cases where (i) $y_{T^{(3)}}^{(i)} < 0$ for $i = 1$ or $i = 2$ (but not both), and (ii) $y_{T^{(3)}}^{(i)} < 0$ for both $i = 1, 2$.

As in the equilibration process described previously, agents update their beliefs based on the observed decision state of their counterpart until no further information can be gained from the process. Before observing the decision of agent 3, and after the private observation at time $T^{(3)}$, the beliefs of agents $i = 1, 2$ are $y_{T^{(3)},0}^{(i)} = \text{Priv}_{1:T^{(3)}}^{(i)}$ by Proposition 6.1. They next account for the decision $d_{T^{(3)}}^{(3)} = 1$ that reaches both simultaneously:

$$y_{T^{(3)},1}^{(i)} = \text{Priv}_{1:T^{(3)}}^{(i)} + \log \frac{P(d_{T^{(3)}}^{(3)} = 1, d_{T^{(3)}-1}^{(3)} = 0, d_{T^{(3)}}^{(1,2)} = 0 | H^+)}{P(d_{T^{(3)}}^{(3)} = 1, d_{T^{(3)}-1}^{(3)} = 0, d_{T^{(3)}}^{(1,2)} = 0 | H^-)} \approx \text{Priv}_{1:T^{(3)}}^{(i)} + \theta. \quad (7.1)$$

The social information obtained from agent 3's decision should incorporate the increment $\epsilon_{T^{(3)}}^+$ discussed in

FIG. 7.1. (a) In a three-agent clique all three agents make independent observations, and observe each other's decisions. (b) Three main possibilities follow the decision of agent 3: (b1) If the beliefs of the undecided agents are both positive, both follow the decision of agent 3; (b2) If the decision causes only one of the remaining agents to decide (agent 2 here). This secondary decision leads to a further update in the belief of the remaining agent equal to $R_-(T^{(3)})$ (agent 1 here); (b3) If neither of the remaining agents decides, they observe each other's indecision, and update their belief by $R_-(T^{(1)}) < 0$. This update cannot lead to a decision, and both agents continue to accumulate private evidence. The dashed portion of the belief trajectory shows the intermediate steps in social information exchange: Each agent's belief consists of information due to its own private observation, a jump $\theta > 0$, followed by a jump $R_-(T^{(1)}) < 0$. No decision can be reached, and the two agents continue to accumulate information, but now have further knowledge about each other's belief. Cases in which the private evidence leads to a simultaneous decision by two or three agents are not shown.

Proposition 4.1. For a clearer exposition, we assume private evidence is weak, and $0 < |\epsilon_{T^{(3)}}^\pm| \ll 1$, and we thus use the approximation $\epsilon_{T^{(3)}}^\pm \approx 0$ from here on.

Any agent who remains undecided after the update given by Eq. (7.1) will update their belief iteratively in response to the decision information of their neighbor. We describe this process in the cases in which at least one agent remains undecided after observing the decision of agent 3.

Case (i) - One Agent Undecided. Without loss of generality, we assume $y_{T,0}^{(2)} \geq 0$ so that $d_{T,1}^{(2)} = 1$. After observing the decision of agent 3, the belief of agent 1 is $y_{T,1}^{(1)} = \text{Priv}_{1:T}^{(1)} + \theta$. After observing the decision of agent 2, a straightforward computation gives:

$$y_{T^{(3)},2}^{(1)} \approx \text{Priv}_{1:T}^{(1)} + \theta + \log \frac{P(d_{T^{(3)},1}^{(2)} = 1 | d_{T^{(3)},0}^{(3)} = 1, H^+)}{P(d_{T^{(3)},1}^{(2)} = 1 | d_{T^{(3)},0}^{(3)} = 1, H^-)} = \text{Priv}_{1:T^{(3)}}^{(1)} + \theta + R_+(T^{(3)}), \quad (7.2)$$

where

$$R_+(t) := \log \frac{P(y_t^{(2)} \in [0, \theta] | H^+)}{P(y_t^{(2)} \in [0, \theta] | H^-)}$$

since agent 2's belief must have been non-negative before observing the decision of agent 3. We prove in Proposition F.1 that $R_+(t) < \theta$. However, this increment in belief may be sufficient to lead to a decision by agent 1. If not, agent 1 continues to accumulate private evidence on the next time step.

We will estimate $R_+(t)$ in arbitrarily large cliques in section 8. In particular, $R_+(t)$ may be computed explicitly in the same way as $S_\pm(t)$ (See Appendix C).

Case (ii): Two Agents Undecided. If $y_{T^{(3)},0}^{(i)} < 0$ for $i = 1, 2$, then both agents remain undecided upon observing a decision by agent 3. After each observes this absence of a decision in its counterpart, it follows from the computations that lead to Eq. (7.2), that each updates its belief as

$$y_{T^{(3)},2}^{(i)} \approx \text{Priv}_{1:T^{(3)}}^{(i)} + \theta + R_-(T^{(3)}) \quad \text{with} \quad R_-(t) := \log \frac{P(y_t^{(\neg i)} \in (-\theta, 0] | H^+)}{P(y_t^{(\neg i)} \in (-\theta, 0] | H^-)}. \quad (7.3)$$

Due to symmetry, the new social information is equal for both agents.

Note that $R_-(t) \leq 0$ and also $|R_-(t)| < \theta$, as shown in Proposition F.1. Therefore Eq. (7.3) shows that this second belief update cannot lead to a decision, and $d_{T^{(3)},2}^{(i)} = 0$ for $i = 1, 2$. After this second step, both agents, $i = 1, 2$, know that

$$-\theta < y_{T^{(3)},0}^{(\neg i)} + \theta + R_-(T^{(3)}) < \theta \Rightarrow -2\theta - R_-(T^{(3)}) < y_{T^{(3)},0}^{(\neg i)} < -R_-(T^{(3)}).$$

Since $R_-(T^{(3)}) \in (-\theta, 0]$, this does not reveal any additional information, and the exchange of social information stops.

At the end of this exchange, both remaining observers know that the belief of the other is in the non-symmetric interval, $y_{T^{(3)},0}^{(\neg i)} \in (R_-(T^{(3)}), \theta + R_-(T^{(3)}))$. Therefore, future non-decisions become informative, and equilibration follows each private observation as in the case of asymmetric decision thresholds discussed in Section 5.

Concurrent Decisions. If the first decision is made simultaneously by two agents, the remaining agent receives independent social information from both. When the two deciding agents disagree, the social information they provide cancels. If the two agents agree on $H^\pm$, the undecided agent increases its belief by $\pm 2\theta$ and follows the decision of the other two.

The exchange of social information increases the probability that all three agents reach a correct decision (Fig. 7.2a), and decreases the time to a decision, both of a single agent, and all agents in the clique (See Fig. 7.2b). The addition of strong pulsatile increments of social information leads to many trials in which agents' beliefs attain values well beyond the decision thresholds $\pm\theta$, mapping to accuracy that is higher than would be obtained by beliefs at or slightly beyond threshold. This is particularly pronounced when the SNR of private measurements is low, and uncoupled agents' final beliefs are not pushed much beyond their decision thresholds. When private observations are more reliable, social information becomes less important.

In the next section, we will describe similar cascades of decisions in larger cliques. Note that as we increase the size of the network, the increase in accuracy and decrease in decision time due to the inclusion of social information both grow (See Fig. 7.2c,d). We next discuss the computations following a choice in such larger cliques.

8. Larger Cliques. We consider a clique of N agents who all have identical symmetric decision thresholds, and measurement distributions. No social information is exchanged until one or more agents makes a choice at time T . In what follows, we focus on the case of a single agent making the first decision. The case in which two or more agents simultaneously make the first decision can be easily analyzed as in the three agent clique, and it often leads to all agents deciding subsequently, so we do not discuss it further.

Without loss of generality, we assume agent 1 makes the first decision, and chooses H^+ . This means $y_{T^{(1)},0}^{(1)} \geq \theta$, and thus every other agent, i , updates their belief to

$$y_{T^{(1)},1}^{(i)} = y_{T^{(1)},0}^{(i)} + \theta, \quad i \neq 1.$$

As earlier, we assume that the excess information, $\epsilon_{T^{(1)}}^+$ (See Proposition 4.1 for definition), provided by a decision is negligible.

Upon observing a decision, the remaining agents stop making private observations and exchange social information until no further social information is available. Observing $d_{T^{(1)},0}^{(1)} = 1$ leads any agent i with belief $y_{T^{(1)},0}^{(i)} \geq 0$ to the same decision. We denote the set of these agents by A_1 , and call these the *agreeing agents*. We will see that there can be multiple waves of agreeing agents. The agents whose beliefs satisfy $y_{T^{(1)},0}^{(i)} < 0$ update their belief to $0 < y_{T^{(1)},1}^{(i)} < \theta$, but do not make a decision. We denote the set of these agents by U_1 , and call these *undecided agents*.

FIG. 7.2. (a) In a three-agent clique the probability that all three agents make the correct choice is larger when social information is exchanged. (b) The average time it takes all three agents decide is smaller when the agents communicate, since social information pushes agents closer to threshold. Both effects are more pronounced when the SNR of private evidence is low, i.e. p/q is small. (c) As the number of agents in a clique, N , increases, the probability that all agents in the clique make the correct choice grows. The difference is more pronounced when SNR is low. (d) Larger cliques provide more social information reducing the average decision times. Here, $|\theta_{\pm}| = 30$, and p, q are defined as in Section 4 and Appendix C. These simulations do not account for the situations where equilibration after the first decision is necessary.

Once agent 1 chooses H^+ at time $T^{(1)}$, the set of decisions from the first wave of agreeing agents follows at the social information exchange step, $(T^{(1)}, 1)$. These decisions are conditionally independent, given the observed decision of agent 1. Thus, each agreeing agent independently provides additional evidence for H^+ , while each undecided agent provides evidence for H^- . As in the case of three agents, the social information provided by an agreeing agent is $R_+(T^{(1)})$, and for a disagreeing agent is $R_-(T^{(1)}) = -R_+(T^{(1)})$, where $R_+(T^{(1)})$ is given in Eq. (E.3). The equality follows from our assumption of symmetry: See Appendix G for a proof. Note that $N = a_1 + u_1 + 1$, where a_1 and u_1 are the number of agreeing and undecided agents in the first wave following the decision of agent 1 at time $T^{(1)}$. All undecided agents thus update their belief as

$$y_{T,2}^{(j)} = y_{T,0}^{(j)} + \theta + (a_1 - (u_1 - 1))R_+(T^{(1)}), \quad j \in U_1. \quad (8.1)$$

Note, each undecided agent observes all agreeing agents, a_1 , and all undecided agents but itself.

We will see that rounds of social evidence exchange can follow. Let

$$R(t, (a, b)) := \log \left(\frac{P(y_t^{(j)} \in (a, b) | H^+)}{P(y_t^{(j)} \in (a, b) | H^-)} \right).$$

Note that $|R(t, (a, b))| \leq \theta$ by Proposition F.1, and that $R(t, [0, \theta)) = R_+(t)$, $R(t, (-\theta, 0]) = R_-(t)$.

Each remaining undecided agent has now observed the decision of the a_1 agreeing agents, and the decision of the other $u_1 - 1$ undecided agents besides itself. Eq. (8.1) implies several possibilities for what these agents do next:

1. If the number of agreeing agents, a_1 , exceeds the number of undecided agents, u_1 , by a sufficient amount, $(a_1 - (u_1 - 1))R_+(T^{(1)}) \geq 2\theta$, then all the remaining undecided agents, $j \in U_1$, go along with the decision of the first agent, and choose H^+ . Thus the second wave of agreeing agents encompasses the remainder of the network.

2. If the number of undecided agents, u_1 , exceeds the number of agreeing agents, a_1 , by a sufficient amount, $(a_1 - (u_1 - 1))R_+(T^{(1)}) \leq -2\theta$, then all undecided agents update their belief to $y_{T^{(1)},2}^{(j)} \leq -\theta$, $j \in U_1$. This is a somewhat counterintuitive situation: If sufficiently many agents remain undecided after observing the first agents' choice, then, after observing each other's absence of a decision, they all agree on the opposite. Thus a wave of agreement is followed by a larger wave of disagreement.
3. If $-2\theta < (a_1 - (u_1 - 1))R_+(T^{(1)}) < -\theta$, then some of the remaining agents may disagree with the original decision and choose H^- , while others may remain undecided. We call the set of disagreeing (contrary) agents C_2 , and the set of still undecided agents U_2 , and denote the sizes of these sets by c_2 , and u_2 respectively. Note that $U_1 = C_2 \cup U_2$. The agents in U_2 know that $y_{T^{(1)},0}^{(j)} \in (-\theta, ((u_1 - 1) - a_1)R_+(T^{(1)}) - 2\theta]$ for all $j \in C_2$, and $y_{T^{(1)},0}^{(j)} \in (((u_1 - 1) - a_1)R_+(T^{(1)}) - 2\theta, 0)$ for all $j \in U_2$. All agents in U_2 thus update their belief again to

$$y_{T^{(1)},2}^{(j)} = y_{T^{(1)},0}^{(j)} + \theta + a_1 R_+(T^{(1)}) + (u_2 - 1)R(T^{(1)}, (((u_1 - 1) - a_1)R_+(T^{(1)}) - 2\theta, 0)) + c_2 R(T^{(1)}, (-\theta, ((u_1 - 1) - a_1)R_+(T^{(1)}) - 2\theta)). \quad (8.2)$$

This update includes the social information obtained from the initial observation, θ , and from the agreeing agents in the first round, $a_1 R_+(T^{(1)})$. The last two summands in Eq. (8.2) give a refinement of the social information from the originally undecided agents in U_1 .

As a result of this second update some undecided agents in U_2 can make a choice. If so, the process repeats, until no undecided agents are left, or no decisions occur after an update. This process thus must terminate after a finite number of steps. This is akin to the equilibration of social information described earlier, but involves observed choices and occurs across the network. Importantly, the process is realization dependent, as it depends on the private evidence accumulated by the undecided agents. After the process is complete, symmetry is broken, and social information equilibration occurs after each further private measurement.

4. If $-\theta \leq (a_1 - (u_1 - 1))R_+(T^{(1)}) \leq 0$, then no agents in U_1 make a decision, and no undecided agent obtains any further social information. They thus continue to gather private information. Symmetry is broken as they know that $y_{T^{(1)},0}^{(j)} \in (-\theta, 0]$ for all remaining agents, $j \in U_1$.
5. If $0 < (a_1 - (u_1 - 1))R_+(T^{(1)}) < \theta$, some undecided agents may choose H^+ , and some may remain undecided. We call the first set A_2 and the second U_2 , and denote the sets' cardinality by a_2 and u_2 , respectively. In this case, $U_1 = A_2 \cup U_2$. All undecided agents $j \in U_2$ know that $y_{T^{(1)},0}^{(j)} \in [-(a_1 - (u_1 - 1))R_+(T^{(1)}), 0)$ for all $j \in A_2$, and $y_{T^{(1)},0}^{(j)} \in (-\theta, -(a_1 - (u_1 - 1))R_+(T^{(1)}))$ for all $j \in U_2$. They thus update their belief to

$$y_{T^{(1)},2}^{(j)} = y_{T^{(1)},0}^{(j)} + \theta + a_1 R_+(T^{(1)}) + (u_2 - 1)R(T^{(1)}, (-\theta, -(a_1 - (u_1 - 1))R_+(T^{(1)}))) + a_2 R(T^{(1)}, [-(a_1 - (u_1 - 1))R_+(T^{(1)}), 0)). \quad (8.3)$$

As a result of this update, some agents may make a choice. If so, the process continues as discussed in 3.

To summarize, before any agent in a clique makes a decision, each agent accumulates private evidence independently. A given agent will make the first decision, say H^+ . All other agents with positive evidence will follow this choice in a first wave of agreement. Undecided agents modify their evidence according to Eq. (8.1). How many agents agree a_n , disagree c_n , or remain undecided u_n in exchange n depends on the numbers of agreeing, disagreeing, and undecided agents from previous steps: $a_{1:n-1}$, $c_{1:n-1}$, and $u_{1:n-1}$.

Large System Size Limit. Let α_N be the probability that the first agent to decide in an N -agent clique makes the correct decision. We conjecture that as the clique size $N \rightarrow \infty$, the probability that the majority of the agents make the correct decision goes to 1; furthermore, all agents will make the correct decision with probability α_N . An asymptotic argument for this result in the special case of a biased random walk on an integer lattice is straightforward. However, a detailed proof of the general claim would be lengthy, so we save a thorough treatment of this problem for future work.

FIG. 8.1. (a) In a clique of size N , we assume agent 1 makes a decision which is observed by all other agents. Agents whose beliefs are positive (agreeing agents), follow the decision of this initial agent. The remaining undecided agents continue to exchange social information. (b) Average social information, $E[(a_1 - (u_1 - 1))R_+]$, available to undecided agents after observing the first wave of agreeing agents as a function of the first decision time for different clique sizes. (c) The probability that the majority of the group has positive belief as a function of the first decision time for different clique sizes. This is equivalent to the probability that the majority of the clique consists of agreeing agents when the first decision occurs. (d) The probability that no agent has made a decision as a function of decision time for different clique sizes. Here, $|\theta_{\pm}| = 20$, $p = e/10$, and $q = 1/10$. (e) The fraction of agents choosing correctly after the first wave following a decision (solid) and after equilibration (dashed) as a function of clique size for various values of p/q and θ .

In Fig. 8.1 we plot information about the behavior of cliques of various sizes, hinting at trends that emerge as $N \rightarrow \infty$. We assume the correct state is H^+ . As the size of the clique grows, so does the amount of information available to undecided agents after the first wave of decisions, *i.e.* the members of the set U_1 (See Fig. 8.1b). As clique size grows, the first decision occurs earlier (Figs. 8.1c-d). In particular, as N grows, the first decision time approaches $\theta/\log(p/q)$ —the minimum number of steps required to make a decision. Moreover, most agents accrue evidence in favor of the correct decision at an earlier time. By the time the first decision occurs, the majority of the clique will be inclined toward the correct decision. The ensuing first decision, if it is correct, will immediately cause the majority of the clique to choose correctly. However, note that as clique size grows, the probability that the initial decision is incorrect also grows. It is therefore not obvious that the asymptotic fraction of the clique that chooses correctly approaches 1: If the initial decision is incorrect, the majority of the clique will be undecided after incorporating social information from the initial decision (See Fig. 8.1c). However, the social information exchange described above can still lead the remainder of the clique to overcome the impact of the wrong initial decision. Numerical simulations suggest that the fraction of the clique that makes the correct choice does limit to 1, at least in the cases we have examined.

9. Discussion. There has been extensive work on the mathematics of evidence accumulation by individual agents [7, 9, 46, 56, 59]. The resulting models have been very successful in explaining experimental data and the dynamics of decision making [24, 55]. In parallel, economists have developed mathematical models of networks of rational decision makers [25, 38]. However, in the work of economists static models are the norm, and the temporal dynamics of decisions are not studied explicitly. In contrast, temporal dynamics have been extensively studied in the psychology literature, but predominantly in the case of subjects making decisions based on private evidence [32] (although see [3, 26, 29]). Here, we provide a bridge between these research areas by developing a normative models of evidence accumulation in social networks.

The beliefs of evidence-accumulators in a network have been modeled by diffusively coupled drift-diffusion equations [44, 52]. Such models approximate continuous information-sharing between agents [52], and have lead to a number of interesting insights. However, it is not always clear when linear coupling between neighbors approximates the normative process of evidence accumulation and exchange between rational

agents. In a related model, an agent observing a choice was assumed to experience a jump in its belief [10]. This study examined how the size of this jump affects the decisions of the collective. While not normative, this model displays interesting dynamics, and provides insight into decision making when agents over- or under-weight available evidence [6].

The belief dynamics of agents in our normative model differs from past accounts in several respects. First, we assume agents only communicate decision information, and do so normatively. In this case, coupling between agents is nonlinear and depends on a threshold-crossing process. We have shown that agents can exchange social evidence even in the absence of a decision when their measurement distributions or decision thresholds are not symmetric. Observing the decision state of a neighbor may result in further exchanges of social information with other neighbors. Furthermore, agents that do not have a direct view of all other agents in the network must marginalize over all possible decision states of the unobserved nodes. This can lead to complex computations, in contrast to the model discussed in [28], where a central agent sees all actions and optimizes a final decision assuming all other agents accumulate evidence independently.

The absence of choice is informative only in the case of asymmetries. This could be due to one choice being more rewarding than another [27, 39, 57]. For instance, the random dot motion task, which requires subjects to determine the dominant direction of randomly moving dots, has been extensively used in psychophysical and neurobiological studies. In this task subjects set their decision thresholds symmetrically when the reward and frequency of each choice is symmetric [24, 50]. However, when one drift direction is more likely or rewarded more, subjects exhibit accuracy and response time biases consistent with an asymmetric prior or thresholds [27, 39]. Recordings suggest neural activity representing decisions is dynamically modulated during the evidence accumulation process to reflect this bias [27, 43]. This is notably different from the static offsets apparent in the Bayesian model, but further investigation is needed to determine how priors are represented across the entire decision-making system [53].

In our model, unless threshold asymmetries are large, the social information communicated by the indecision of a single agent is weak. However, in large social networks the totality of such evidence provided by all other agents can be large compared to the private information obtained by an agent. Therefore, in large networks even small asymmetries can produce social information from indecisions that strongly influences decisions: Consider a new, desirable product that hits the market. If a large fraction of a social network is indecisive about purchasing the product, this can communicate potential issues with the product. This signal could be particularly strong if the product is particularly desirable upon its release. Another example is the case of decentralized detection: Consider a distributed network of independent environmental hazard sensors which only communicate their decision (e.g., whether the environment is hazardous or safe) to a central agent (fusion center), which combines their independent decisions [11]. The central agent only uses social information, but has a much more accurate estimate of the environmental state than any single sensor. Moreover, if decision thresholds are set asymmetrically, say so that the hazard threshold is much lower, the presence of indecisions in all sensors would suggest they are collecting evidence in favor of a safe environment. By computing the LLR of the survival probability from each independent sensor, the central agent could obtain strong evidence suggesting a safe environment even before receiving any decisions.

Our normative model is unlikely to accurately describe the decision making process of biological agents. The rationality of humans and other animals is bounded [51], while some of the computations that we have described are quite complex, and provide only marginal increases in the probability of making a correct choice. Thus biological agents are likely to perform only those computations – perhaps approximately – that provide the largest impact [13]. A normative model allows us to identify which computations are important, and which offer only fractional benefits.

A strong simplifying assumption in this work is that observers make uncorrelated measurements. This assumption will not hold in most situations, as members of a social group are likely to obtain dependent information. However, experimental evidence suggests that humans do not take into account such correlations when making decisions [17, 31] and behave as if the evidence gathered by different members of their social group is independent of their own. The types of models we describe may approximate the decision-making process of agents that make such assumptions and the impact of ignoring such correlations when they are present is an interesting avenue for future work. Our assumption that all agents in the network are identical can be relaxed. As long as all agents know each other’s measurement distributions and decision thresholds, social information can be computed as we have done in the present work, although notation becomes more cumbersome. The case in which agents do not know their neighbors’ measurement distributions

and thresholds is considerably more complex. In this case agents either marginalize over the unknown parameters of their neighbors, or try to infer them based on the temporal dynamics of their decision states.

Our findings are distinct, but related to previous work on herding and common knowledge. In this modeling work agents typically make a single private measurement, followed by an exchange, or propagation of social information [5, 33]. In recurrent networks, agents can announce their preference for a choice, until they reach agreement with their neighbors [2, 22, 23, 36]. This framework can be simplified so that agents make binary decisions, based solely on their neighbors' opinion, admitting asymptotic analyses of the cascade of decisions through networks with complex structure [60]. In contrast, we assume that agents accumulate private and social evidence, and make irrevocable decisions. The resulting decision-making processes are considerably different: For instance, in our model agents, and there is no guarantee that social learning occurs. Combining private and social evidence also makes it difficult to derive exact results, but we expect asymptotic formulas are possible for large cliques, and simpler assumptions on the decision process.

Key to understanding the collective behavior of social organisms is uncovering how they gather and exchange information to make decisions [13, 20]. Mathematical models allow us to quantify how different evidence-accumulation strategies impact experimentally observable characteristics of the decision-making process such as the speed and accuracy of choices [8], and the level of agreement among members of a collective [12]. Such models can thus lead the way towards understanding the decisions and disagreements that emerge in social groups.

Acknowledgements. This work was supported by an NSF/NIH CRCNS grant (R01MH115557) and an NSF grant (DMS-1517629). BK and KJ were supported by NSF grant (DMS-1662305). KJ was also supported by NSF NeuroNex grant (DBI-1707400). ZPK was also supported by an NSF grant (DMS-1615737).

Appendix A. Decomposition of an Agent's Belief into Private and Social Information. The following proposition shows that the belief of each agent is a sum of information from private measurements, and observed decisions, as claimed in Eq. (3.2).

PROPOSITION A.1. *Assume that in the network depicted in Fig. 3.1 agent 1 makes a choice at time $T^{(1)}$. If agent 2 has not yet made a decision, its belief is*

$$y_s^{(2)} = \begin{cases} \text{Priv}_{1:s}^{(2)} + \text{Soc}(d_s^{(1)} = 0), & s < T^{(1)} \\ \text{Priv}_{1:s}^{(2)} + \text{Soc}(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1), & s \geq T^{(1)}. \end{cases} \quad (\text{A.1})$$

Proof. The decisions of agent 1 are independent from the observations of agent 2, when conditioned on the state, H . Thus

$$P(\xi_{1:s}^{(2)}, d_{1:s}^{(1)} | H) = P(\xi_{1:s}^{(2)} | H) P(d_{1:s}^{(1)} | H).$$

Hence, taking the logarithm of the ratio of $P(\xi_{1:s}^{(2)}, d_{1:s}^{(1)} | H)$ gives

$$y_s^{(2)} = \text{Priv}_{1:s}^{(2)} + \text{Soc}(d_{1:s}^{(1)}).$$

It remains to show that $\text{Soc}(d_{1:s}^{(1)})$ simplifies to the form given in Eq. (A.1).

Consider $P(d_1^{(1)} = i_1, \dots, d_s^{(1)} = i_s | H)$ where $i_1, \dots, i_s \in \{-1, 0, 1\}$. Before agent 1 makes a decision, it communicates a decision state of 0, so if $s < T^{(1)}$:

$$P(d_1^{(1)} = 0, \dots, d_s^{(1)} = 0 | H) = P(d_s^{(1)} = 0 | H).$$

This equality holds because $d_s^{(1)} = 0$ implies that all previous decision states must also be 0.

Similarly, when agent 1 makes a decision at time $T^{(1)}$, we can write

$$P(d_1^{(1)} = 0, \dots, d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1, \dots, d_s^{(1)} = 1 | H) = P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H).$$

Again, this is because $d_1^{(1)} = 0, \dots, d_{T^{(1)}-2}^{(1)} = 0$ is implied by $d_{T^{(1)}-1}^{(1)} = 0$ and the values of the decision states after time $T^{(1)}$, $d_{T^{(1)}+1}^{(1)} = 1, \dots, d_s^{(1)} = 1$ are implied by $d_{T^{(1)}}^{(1)} = 1$. Thus

$$\text{Soc}(d_1^{(1)} = 0, \dots, d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1, \dots, d_t^{(1)} = 1) = \text{Soc}(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1),$$

and we note that the evidence from the decision state depends on the value of the decision state and the time when it was first non-zero, i.e., when agent 1 made a choice. $\square$

Appendix B. Proof of Proposition 4.1: Bounding Social Information from a Decision. We prove the case $d_{T^{(1)}} = +1$ in detail. The $d_{T^{(1)}} = -1$ is analogous. With $\Theta = (\theta_-, \theta_+)$, we have

$$P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^+) = P(\text{Priv}_1^{(1)}, \dots, \text{Priv}_{1:T^{(1)}-1}^{(1)} \in \Theta, \text{Priv}_{1:T^{(1)}}^{(1)} \geq \theta_+ | H^+).$$

We observe that $d_{T^{(1)}-1}^{(1)} = 0$ and $d_{T^{(1)}}^{(1)} = 1$ imply $y_{T^{(1)}-1}^{(1)} \in \Theta$ and $y_{T^{(1)}}^{(1)} \geq \theta_+$. Hence

$$\theta_- < \sum_{t=1}^{T^{(1)}-1} \log \frac{P(\xi_t^{(1)} | H^+)}{P(\xi_t^{(1)} | H^-)} < \theta_+ \leq \sum_{t=1}^{T^{(1)}} \log \frac{P(\xi_t^{(1)} | H^+)}{P(\xi_t^{(1)} | H^-)}.$$

Following [7], the last of these inequalities implies

$$e^{\theta_+} \prod_{t=1}^{T^{(1)}} P(\xi_t^{(1)} | H^-) \leq \prod_{t=1}^{T^{(1)}} P(\xi_t^{(1)} | H^+). \quad (\text{B.1})$$

We define the set of *legal* chains of beliefs as those that lead to a decision at time $T^{(1)}$, but not earlier,

$$\mathcal{C}_+(T^{(1)}) = \{\xi_{1:T^{(1)}}^{(1)} \mid \text{Priv}_{1:t}^{(1)} \in \Theta, 1 \leq t < T^{(1)}; \text{Priv}_{1:T^{(1)}}^{(1)} \geq \theta_+\}.$$

Note that $\mathcal{C}_+(T^{(1)}) \subseteq \Xi^{T^{(1)}}$. Hence,

$$P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^+) = \sum_{\xi_{1:T^{(1)}}^{(1)} \in \mathcal{C}_+(T^{(1)})} f_+(\xi_{1:T^{(1)}}^{(1)}),$$

where we are using the notation

$$f_+(\xi_{1:T^{(1)}}^{(1)}) := \prod_{t=1}^{T^{(1)}} f_+(\xi_t^{(1)}) = \prod_{t=1}^{T^{(1)}} P(\xi_t^{(1)} | H_+).$$

Every legal chain satisfies the inequality (B.1), so summing yields

$$\sum_{\xi_{1:T^{(1)}}^{(1)} \in \mathcal{C}_+(T^{(1)})} f_+(\xi_{1:T^{(1)}}^{(1)}) \geq e^{\theta_+} \sum_{\xi_{1:T^{(1)}}^{(1)} \in \mathcal{C}_+(T^{(1)})} f_-(\xi_{1:T^{(1)}}^{(1)}) = e^{\theta_+} P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^-).$$

Thus

$$P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^+) \geq e^{\theta_+} P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^-),$$

and so

$$\log \frac{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^+)}{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^-)} \geq \theta_+.$$

By noting $\xi_{1:T^{(1)}}^{(1)} \in \mathcal{C}_+(T^{(1)})$, we can directly compute

$$\log \frac{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^+)}{P(d_{T^{(1)}-1}^{(1)} = 0, d_{T^{(1)}}^{(1)} = 1 | H^-)} = \log \frac{\sum_{\xi_{1:T^{(1)}}^{(1)} \in \mathcal{C}_+(T^{(1)})} f_+(\xi_{1:T^{(1)}}^{(1)})}{\sum_{\xi_{1:T^{(1)}}^{(1)} \in \mathcal{C}_+(T^{(1)})} f_-(\xi_{1:T^{(1)}}^{(1)})} = \theta_+ + \epsilon_{T^{(1)}}^+,$$

where we have defined

$$\epsilon_{T^{(1)}}^+ := \log \frac{\sum_{\xi_{1:T^{(1)}} \in \mathcal{C}_+(T^{(1)})} f_+(\xi_{1:T^{(1)}})}{\sum_{\xi_{1:T^{(1)}} \in \mathcal{C}_+(T^{(1)})} f_-(\xi_{1:T^{(1)}})} - \theta_+.$$

We can obtain a simpler upper bound by noting that for any $\xi_{1:T^{(1)}}^{(1)} \in \mathcal{C}_+(T^{(1)})$, we have

$$\text{Priv}_{1:T^{(1)}}^{(1)} = \text{Priv}_{1:T^{(1)}-1}^{(1)} + \log \frac{f_+(\xi_{T^{(1)}})}{f_-(\xi_{T^{(1)}})} < \theta_+ + \log \frac{f_+(\xi_{T^{(1)}}^{(1)})}{f_-(\xi_{T^{(1)}}^{(1)})}.$$

which implies

$$\epsilon_{T^{(1)}}^+ \leq \sup_{\xi \in \Xi_{T^{(1)}}^+} \left(\log \frac{f_+(\xi)}{f_-(\xi)} \right),$$

where we define

$$\Xi_{T^{(1)}}^+ = \left\{ \xi_{T^{(1)}} \mid \xi_{1:T^{(1)}} \in \mathcal{C}_+(T^{(1)}) \right\}$$

by positivity of $\log(f_+(\xi_{T^{(1)}})/f_-(\xi_{T^{(1)}}))$, since $\xi_{T^{(1)}}$ triggers a decision, and $\log \frac{f_+(\xi)}{f_-(\xi)} \leq \sup_{\xi \in \Xi_{T^{(1)}}^+} \left(\log \frac{f_+(\xi)}{f_-(\xi)} \right)$ for all $\xi \in \Xi_{T^{(1)}}^+$. A similar argument can be used to derive bounds on $\text{Soc}_{T^{(1)}}^{(1)}$ for $d_{T^{(1)}}^{(1)} = -1$.

Appendix C. Explicit Formula for Social Information for Four-State Biased Random Walk.

We use standard first-step analysis [54] to calculate non-decision social evidence for $\theta_+ = 2$ and $\theta_- = -1$. In this case we have four belief states, $y_t^{(1)} \in \{-1, 0, 1, 2\}$, where $y_t^{(1)} = +2 = \theta_+$ and $y_t^{(1)} = -1 = \theta_-$ are absorbing boundaries. Let $P_+^l(t) := P(y_t^{(1)} = l | H^+)$. Then, for $\mathbf{P}_+(t) := (P_+^{-1}(t), P_+^0(t), P_+^1(t), P_+^2(t))^T$,

$$\mathbf{P}_+(t+1) = M(p, q)\mathbf{P}_+(t), \quad M(p, q) := \begin{pmatrix} 1 & q & 0 & 0 \\ 0 & s & q & 0 \\ 0 & p & s & 0 \\ 0 & 0 & p & 1 \end{pmatrix}$$

with initial condition $P_+^0(0) = 1$ and $s = 1 - p - q$.

Let $\mathbf{v}(t)$ be a vector of probabilities that $y_t^{(1)}$ is at a non-absorbing state, then we can solve for

$$\mathbf{v}(t) = \begin{pmatrix} P_+^0(t) \\ P_+^1(t) \end{pmatrix} = \tilde{M}(p, q)^t \mathbf{v}(0) = \tilde{M}(p, q)^t \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \tilde{M}(p, q) := \begin{pmatrix} s & q \\ p & s \end{pmatrix}.$$

$\tilde{M}(p, q)$ has eigenvalues $\lambda_1 := s - \sqrt{pq}$, $\lambda_2 := s + \sqrt{pq}$ with eigenvectors

$$\mathbf{e}_1 = \begin{pmatrix} -\sqrt{q/p} \\ 1 \end{pmatrix}, \quad \mathbf{e}_2 = \begin{pmatrix} \sqrt{q/p} \\ 1 \end{pmatrix},$$

so that

$$\mathbf{v}(0) = \frac{1}{2}\sqrt{\frac{p}{q}}[-\mathbf{e}_1 + \mathbf{e}_2] = -\frac{1}{2} \begin{pmatrix} -1 \\ \sqrt{p/q} \end{pmatrix} + \frac{1}{2} \begin{pmatrix} 1 \\ \sqrt{p/q} \end{pmatrix}.$$

Thus,

$$\mathbf{v}(t) = -\frac{\lambda_1^t}{2} \begin{pmatrix} -1 \\ \sqrt{p/q} \end{pmatrix} + \frac{\lambda_2^t}{2} \begin{pmatrix} 1 \\ \sqrt{p/q} \end{pmatrix}.$$

We calculate the eigenpairs for $\tilde{M}(p, q)^t$ by swapping p and q , and can thus obtain an expression for $\mathbf{v}(t)$ in that case. Then,

$$\begin{aligned}\text{Soc}_t^{(2)} &= \log \frac{(1, 1) \cdot (\tilde{M}(p, q)^t \mathbf{v}(0))}{(1, 1) \cdot (\tilde{M}(p, q)^t \mathbf{v}(0))} = \log \frac{(\lambda_1^t + \lambda_2^t) + \sqrt{p/q}(\lambda_2^t - \lambda_1^t)}{(\lambda_1^t + \lambda_2^t) + \sqrt{q/p}(\lambda_2^t - \lambda_1^t)} \\ &= \log \frac{(1 - \sqrt{p/q}) + (1 + \sqrt{p/q})(\lambda_2/\lambda_1)^t}{(1 - \sqrt{q/p}) + (1 + \sqrt{q/p})(\lambda_2/\lambda_1)^t}.\end{aligned}$$

For $d_t^{(1)} = 0$ we have

$$\lim_{t \rightarrow \infty} \text{Soc}_t^{(2)} = \log \frac{1 + \sqrt{p/q}}{1 + \sqrt{q/p}} = \log \sqrt{p/q} = \frac{1}{2} \log \frac{p}{q},$$

since $|\lambda_2| > \lambda_1$. For our choice of parameters, $\sqrt{p/q} = \sqrt{e}$ so

$$\lim_{t \rightarrow \infty} \text{Soc}_t^{(2)} = \frac{1}{2} \log e = \frac{1}{2}.$$

Note the recurring appearance of the quantity p/q in this calculation. This is a measure of the noisiness of the independent observations of the environment and is a proxy for the signal-to-noise ratio (SNR). In particular, note that if $p \gg q$, then $\lim_{t \rightarrow \infty} \text{Soc}_t^{(2)}$ is large. On the other hand, if observations are extremely noisy, i.e. $p \approx q$, then $\lim_{t \rightarrow \infty} \text{Soc}_t^{(2)} \approx 0$.

Interestingly, in the limit where all observations are informative ($s \rightarrow 0$), the social information will alternate indefinitely between two values. In this case, $\lambda_2/\lambda_1 \rightarrow -1$, and we have

$$\text{Soc}_t^{(2)} = \log \frac{(1 - \sqrt{p/q}) + (1 + \sqrt{p/q})(-1)^t}{(1 - \sqrt{q/p}) + (1 + \sqrt{q/p})(-1)^t} = \begin{cases} 0, & t \text{ even}, \\ \log(p/q), & t \text{ odd}. \end{cases}$$

This follows from the fact that the belief of agent 1 must be at lattice site 0 (1) after an even (odd) number of observations.

Appendix D. Proof of Proposition 5.2: Bounding Social Information from a Decision in a Recurrent Two-Agent Network. We use an argument similar to that in Proposition 4.1. First, consider the case in which $n = 0$, so private observation $\xi_{T^{(i)}}^{(i)}$ triggers the decision. If $d_{T^{(i)},0}^{(i)} = +1$, we know

$$y_{T^{(i)},0}^{(i)} = \text{Priv}_{1:T^{(i)}}^{(i)} + \text{Soc}_{T^{(i)}-1}^{(i)} \geq \theta_+,$$

and since $\text{Priv}_{1:T^{(i)}-1}^{(i)} + \text{Soc}_{T^{(i)}-1}^{(i)} < \theta_+$, then

$$y_{T^{(i)},0}^{(i)} = \text{Priv}_{1:T^{(i)}}^{(i)} + \text{Soc}_{T^{(i)}-1}^{(i)} < \theta_+ + \log \frac{f_+(\xi_{T^{(i)}}^{(i)})}{f_-(\xi_{T^{(i)}}^{(i)})}.$$

Marginalizing over all chains $\mathcal{C}_+(T^{(i)}, 0)$ such that $y_{s,m}^{(i)} \in \Theta$ for all $0 < s < T^{(i)}$, and corresponding $0 \leq m \leq N_s$, preceding the decision $d_{T^{(i)},0}^{(i)} = +1$ at $(T^{(i)}, 0)$, we thus find

$$\theta_+ \leq \log \frac{P(d_{T^{(i)},0}^{(j)} = 0, d_{T^{(i)}-1, N_{t-1}}^{(i)} = 0, d_{T^{(i)},0}^{(i)} = 1 | H^+)}{P(d_{T^{(i)},0}^{(j)} = 0, d_{T^{(i)}-1, N_{T^{(i)}-1}}^{(i)} = 0, d_{T^{(i)},0}^{(i)} = 1 | H^-)} < \theta_+ + \epsilon_{T^{(i)},0}^+,$$

where $N_{T^{(i)}-1}$ is the maximal substep in the social information exchange following the private observation at timestep $T^{(i)} - 1$. A similar argument shows that

$$\theta_- + \epsilon_{T^{(i)},0}^- < \log \frac{P(d_{T^{(i)},0}^{(j)} = 0, d_{T^{(i)}-1, N_{t-1}}^{(i)} = 0, d_{T^{(i)},0}^{(i)} = -1 | H^+)}{P(d_{T^{(i)},0}^{(j)} = 0, d_{T^{(i)}-1, N_{T^{(i)}-1}}^{(i)} = 0, d_{T^{(i)},0}^{(i)} = -1 | H^-)} \leq \theta_-,$$

where

$$\pm \epsilon_{T^{(i)},0}^{\pm} \leq \sup_{\xi \in \Xi_{T^{(i)}}^{\pm}} \left(\log \frac{f_+(\xi)}{f_-(\xi)} \right),$$

where we define

$$\Xi_{T^{(1)}}^{\pm} = \left\{ \xi_{T^{(1)}} \mid \xi_{1:T^{(1)}} \in \mathcal{C}_{\pm}(T^{(1)}, 0) \right\}.$$

On the other hand, suppose $d_{T^{(i)},n}^{(i)} = +1$, and $0 < n \leq N_{T^{(i)}}$ so that the decision is reached during the social information exchange following an observation. Then,

$$\text{Priv}_{1:T^{(i)}}^{(i)} + \text{Soc}_{T^{(i)}-1,n-1}^{(i)} < \theta_+,$$

which implies

$$\theta_+ \leq y_{T^{(i)},n}^{(i)} = \text{Priv}_{1:T^{(i)}}^{(i)} + \text{Soc}_{T^{(i)}-1,n-1}^{(i)} + (\text{Soc}_{T^{(i)},n}^{(i)} - \text{Soc}_{T^{(i)},n-1}^{(i)}) < \theta_+ + (\text{Soc}_{T^{(i)},n}^{(i)} - \text{Soc}_{T^{(i)},n-1}^{(i)}).$$

Marginalizing over all chains $\mathcal{C}_+(T^{(i)}, n)$ such that $y_{s,m}^{(i)} \in \Theta$ for all $0 < s \leq T^{(i)}$, and corresponding $0 \leq m \leq N_s$, preceding the decision at $(T^{(i)}, n)$,

$$\theta_+ \leq \log \frac{P(d_{T^{(i)},n}^{(j)} = 0, d_{T^{(i)},n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = 1 | H^+)}{P(d_{T^{(i)},n}^{(j)} = 0, d_{T^{(i)},n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = 1 | H^-)} < \theta_+ + \epsilon_{T^{(i)},n}^+,$$

and similarly for $\mathcal{C}_-(T^{(i)}, n)$, we have

$$\theta_- + \epsilon_{T^{(i)},n}^- < \log \frac{P(d_{T^{(i)},n}^{(j)} = 0, d_{T^{(i)},n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = -1 | H^+)}{P(d_{T^{(i)},n}^{(j)} = 0, d_{T^{(i)},n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = -1 | H^-)} \leq \theta_-,$$

where

$$\epsilon_{T^{(i)},n}^{\pm} \lesseqgtr (\text{Soc}_{T^{(i)},n}^{(i)} - \text{Soc}_{T^{(i)},n-1}^{(i)}).$$

Following the arguments in Theorem 5.1, we note then that

$$\begin{aligned} y_{T^{(i)},n+1}^{(j)} &= \log \frac{P(\xi_{1:T^{(i)}}^{(j)}, d_{1:T^{(i)},0:n}^{(j)} = 0, d_{1:T^{(i)},0:n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = \pm 1 | H^+)}{P(\xi_{1:T^{(i)}}^{(j)}, d_{1:T^{(i)},0:n}^{(j)} = 0, d_{1:T^{(i)},0:n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = \pm 1 | H^-)} \\ &= \log \frac{P(\xi_{1:T^{(i)}}^{(j)} | H^+)}{P(\xi_{1:T^{(i)}}^{(j)} | H^-)} + \log \frac{P(d_{T^{(i)},n}^{(j)} = 0, d_{T^{(i)},n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = -1 | H^+)}{P(d_{T^{(i)},n}^{(j)} = 0, d_{T^{(i)},n-1}^{(i)} = 0, d_{T^{(i)},n}^{(i)} = -1 | H^-)} \\ &= \text{Priv}_{1:T^{(i)}}^{(j)} + \text{Soc}_{T^{(i)},n+1}^{(j)}. \end{aligned}$$

Appendix E. Three agents on a directed line. Here we show that when agent 3 observes a decision by agent 2, at time $T^{(2)}$ the increment in $y_{T^{(2)}}^{(3)}$ exceeds θ , depending mostly on the thresholds θ , but less on the informativeness of individual observations.

First, we note that for any $s = 1, \dots, T^{(2)} - 1$,

$$\frac{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = \pm 1 | H^+)}{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = \pm 1 | H^-)}, \frac{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = 0 | H^+)}{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = 0 | H^-)} \in [e^{\theta}, e^{\theta + \epsilon_{T^{(2)}}^+}],$$

using arguments similar to those presented in Prop. 4.1 and Prop. 5.2. Therefore, the best estimate of these likelihood ratios can only bound them within finite intervals. However, such considerations muddle

the presentation, and in fact in the case for which the measurement distributions $f_+(\xi)$ and $f_-(\xi)$ are close, $\max_\xi |f_+(\xi) - f_-(\xi)| < \varepsilon$, then $\text{Priv}_t^{(2)}$ are small and so $\epsilon_{T^{(2)}}^\pm$ are also small. With this in mind, we write

$$\frac{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = \pm 1 | H^+)}{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = \pm 1 | H^-)} \approx \frac{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = 0 | H^+)}{P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = 0 | H^-)} \approx e^\theta. \quad (\text{E.1})$$

This is a common approximation in the literature on sequential probability ratio tests [58]. The situations where a decision from agent 1 does *not* immediately cause a decision from agent 2, conditioned on H^+ , are described by

$$\begin{aligned} \zeta_t^+ := & P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}}^{(1)} = 0 | H^+) + \sum_{s=1}^{T^{(2)}} P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = -1, d_{s-1}^{(1)} = 0 | H^+) \\ & + \sum_{s=1}^{T^{(2)}-1} P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = 1, d_{s-1}^{(1)} = 0 | H^+). \end{aligned}$$

Note, when agent 2 makes an H^+ decision immediately after agent 1 makes an H^- decision, ($d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}-1}^{(1)} = -1, d_{T^{(2)}-2}^{(1)} = 0$), it must have done so due to private information $\text{Priv}_t^{(2)}$ arriving immediately before the social information $\text{Soc}_{T^{(2)}}^{(2)}$ due to agent 1's opposite decision, which would otherwise drive $y_{T^{(2)}}^{(2)}$ far from the $+\theta$ threshold. Since we have assumed $\text{Priv}_t^{(2)}$ are small, this situation is extremely unlikely, so $P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}-1}^{(1)} = -1, d_{T^{(2)}-2}^{(1)} = 0 | H^+) \approx 0$. Using Eq. (E.1), we have

$$\begin{aligned} \zeta_{T^{(2)}}^+ \approx & e^\theta \cdot \left[P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^-) + \sum_{s=1}^{T^{(2)}} P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = -1, d_{s-1}^{(1)} = 0 | H^-) \right. \\ & \left. + \sum_{s=1}^{T^{(2)}-1} P(d_{T^{(2)}}^{(2)} = 1, d_s^{(1)} = 1, d_{s-1}^{(1)} = 0 | H^-) \right]. \end{aligned}$$

By Eq. (6.4), we have

$$\begin{aligned} P(d_{T^{(2)}}^{(2)} = 1 | H^+) &= \zeta_{T^{(2)}}^+ + P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^+) \\ P(d_{T^{(2)}}^{(2)} = 1 | H^-) &\approx e^{-\theta} \zeta_{T^{(2)}}^+ + P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^-), \end{aligned}$$

so that

$$\text{Soc}_{T^{(2)}}^{(3)} \approx \log \left(\frac{P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^+) + \zeta}{P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^-) + \zeta e^{-\theta}} \right). \quad (\text{E.2})$$

Note that

$$P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^\pm) = P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^\pm) P(d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^\pm),$$

and

$$P(d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^+) \approx e^\theta P(d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^-) \approx \frac{1}{1 + e^{-\theta}}.$$

Let

$$R_+(t) := \log \left(\frac{P(y_t^{(2)} \in [0, \theta] | H^+)}{P(y_t^{(2)} \in [0, \theta] | H^-)} \right) \geq 0. \quad (\text{E.3})$$

Note that we can increase $R_+(t)$ by increasing θ , without changing the measurements, or their distributions.

It follows that,

$$\begin{aligned}\text{Soc}_{T^{(2)}}^{(3)} &\approx \log \left(\frac{P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^+) P(d_{T^{(2)}-1}^{(1)} = 1 | H^+) + \zeta_{T^{(2)}}^+ e^\theta}{P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^+) P(d_{T^{(2)}-1}^{(1)} = 1 | H^+) e^{-R_+(T^{(2)})} + \zeta_{T^{(2)}}^+} \right) \\ &\approx \log \left(\frac{P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^+) + \zeta_{T^{(2)}}^+ (1 + e^{-\theta})}{P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^+) e^{-R_+(T^{(2)})} + \zeta_{T^{(2)}}^+ (1 + e^{-\theta})} \right) + \theta \geq \theta.\end{aligned}$$

We therefore see that increasing $R_+(T^{(2)})$ increases the magnitude of social information received from observing a choice by agent 2.

The impact of a decision $d_{T^{(2)}}^{(2)} = -1$ can be computed in a similarly, yielding

$$\text{Soc}_{T^{(2)}}^{(3)} \approx \log \left(\frac{P(y_{T^{(2)}}^{(2)} \in (-\theta, 0] | H^+) e^{-|R_-(T^{(2)})|} + \zeta_{T^{(2)}}^- (1 + e^{-\theta})}{P(y_{T^{(2)}}^{(2)} \in (-\theta, 0] | H^+) + \zeta_{T^{(2)}}^- (1 + e^{-\theta})} \right) - \theta \leq \theta,$$

where

$$\begin{aligned}\zeta_{T^{(2)}}^- &:= P(d_{T^{(2)}}^{(2)} = -1, d_{T^{(2)}}^{(1)} = 0 | H^-) + \sum_{s=1}^{T^{(2)}-1} P(d_s^{(2)} = -1, d_s^{(1)} = -1, d_{s-1}^{(1)} = 0 | H^-) \\ &\quad + \sum_{s=1}^{T^{(2)}} P(d_{T^{(2)}}^{(2)} = -1, d_s^{(1)} = 1, d_{s-1}^{(1)} = 0 | H^-),\end{aligned}$$

and

$$R_-(t) := \log \left(\frac{P(y_t^{(2)} \in (-\theta, 0] | H^+)}{P(y_t^{(2)} \in (-\theta, 0] | H^-)} \right) \leq 0. \quad (\text{E.4})$$

Appendix F. Social Information from Bounds on Neighbor's Beliefs. We next show that knowing that an agent's belief lies within a subset of the interval $\Theta = (-\theta, \theta)$ provides at most an amount of social evidence equal to θ .

In the proof we assume that a sequence of private observations results in a belief (LLR) that can lie exactly on the decision threshold θ , as in the examples in the text. The proof can be extended to cases when information gained from individual measurements is small, so that $0 < |\epsilon_t^\pm| \ll 1$. The proof in this case is equivalent, but somewhat more involved.

PROPOSITION F.1. *Let $-\theta < a \leq b < \theta$. If agent j has not made a decision and has accumulated only private information by time T , then*

$$\left| \log \frac{P(a \leq y_T^{(j)} \leq b | H^+)}{P(a \leq y_T^{(j)} \leq b | H^-)} \right| < \theta.$$

Proof. Define the subset $V_T(a, b) \subset \Xi^T$ of the product space of observations Ξ^T consisting of observation sequences that result in a belief contained in $[a, b]$ at time T :

$$V_T(a, b) = \{\xi_{1:T}^{(j)} \in \Xi^T : \text{Priv}_{1:t}^{(j)} \in (-\theta, \theta), \text{ for } t \leq T, \text{ Priv}_{1:T}^{(j)} \in [a, b]\}.$$

By definition, we can write

$$e^a \leq \prod_{t=1}^T \frac{f_+(\xi_t^{(j)})}{f_-(\xi_t^{(j)})} \leq e^b, \quad \forall \xi_{1:T} \in V_T(a, b),$$

which can be rearranged as

$$e^a \prod_{t=1}^T f_{-}(\xi_t^{(j)}) \leq \prod_{t=1}^T f_{+}(\xi_t^{(j)}) \leq e^b \prod_{t=1}^T f_{-}(\xi_t^{(j)}), \quad \forall \xi_{1:T}^{(j)} \in V_T(a, b).$$

Summing over all $\xi_{1:T}^{(j)} \in V_T(a, b)$ then yields

$$e^a \sum_{V_T(a,b)} \prod_{t=1}^T f_{-}(\xi_t^{(j)}) \leq \sum_{V_T(a,b)} \prod_{t=1}^T P_{+}(\xi_t^{(j)}) \leq e^b \sum_{V_T(a,b)} \prod_{t=1}^T P_{-}(\xi_t^{(j)}),$$

so that by noting

$$P(y_T^{(j)} \in [a, b] | H^{\pm}) = \sum_{V_T(a,b)} \prod_{t=1}^T P_{\pm}(\xi_t^{(j)}),$$

and rearranging we find

$$e^{-\theta} < e^a \leq \frac{P(y_T^{(j)} \in [a, b] | H^{+})}{P(y_T^{(j)} \in [a, b] | H^{-})} \leq e^b < e^{\theta},$$

which implies the desired result. $\square$

Appendix G. Proving Reflection Symmetry of Social Information from LLR Bounds.

PROPOSITION G.1. Assume agent i has not received any social information at time t , so their belief, $y_t^{(i)}$, is based on private information. Also assume that the decision thresholds, $\theta_{\pm} = \pm\theta$, and measurement distributions, $f_{+}(\xi) = f_{-}(-\xi)$, are symmetric. Let

$$R_{+}(t) := \log \frac{P(y_t^{(i)} \in [0, \theta] | H^{+})}{P(y_t^{(i)} \in [0, \theta] | H^{-})}, \quad \text{and} \quad R_{-}(t) := \log \frac{P(y_t^{(i)} \in (-\theta, 0] | H^{+})}{P(y_t^{(i)} \in (-\theta, 0] | H^{-})}.$$

Then $R_{-}(t) = -R_{+}(t)$.

Proof. Following from the argument in Proposition F.1, we can compute

$$P(y_t^{(i)} \in [0, \infty) | H^{\pm}) = \sum_{V_T^{+}} \prod_{t=1}^T f_{\pm}(\xi_t^{(i)}),$$

where we have defined

$$V_T^{+} = \{\xi_{1:T}^{(j)} \in \Xi^T : \text{Priv}_{1:s}^{(j)} \in (-\theta, \theta), \text{ for } s \leq t, \text{ Priv}_{1:t}^{(j)} \in [0, \infty)\}.$$

By symmetry, we know that for any $\xi_{1:t}^{(i)} \in V_T(a, b)^{+}$, there exists $-\xi_{1:t}^{(i)} \in V_T(a, b)^{-}$ where

$$V_T^{-} = \{\xi_{1:T}^{(j)} \in \Xi^T : \text{Priv}_{1:s}^{(j)} \in (-\theta, \theta), \text{ for } s \leq t, \text{ Priv}_{1:t}^{(j)} \in (-\infty, 0]\},$$

and vice versa. Therefore we can write

$$P(y_t^{(i)} \in (-\infty, 0] | H^{\pm}) = \sum_{V_T^{-}} \prod_{t=1}^T f_{\pm}(\xi_t^{(i)}) = \sum_{V_T^{+}} \prod_{t=1}^T f_{\pm}(-\xi_t^{(i)}) = \sum_{V_T^{+}} \prod_{t=1}^T f_{\mp}(\xi_t^{(i)}) = P(y_t^{(i)} \in [0, \infty) | H^{\mp}),$$

where the penultimate equality holds since $f_{+}(\xi) = f_{-}(-\xi)$. Therefore,

$$R_{-}(t) = \log \frac{P(y_t^{(i)} \in (-\theta, 0] | H^{+})}{P(y_t^{(i)} \in (-\theta, 0] | H^{-})} = \log \frac{P(y_t^{(i)} \in [0, \theta] | H^{-})}{P(y_t^{(i)} \in [0, \theta] | H^{+})} = -\log \frac{P(y_t^{(i)} \in [0, \theta] | H^{+})}{P(y_t^{(i)} \in [0, \theta] | H^{-})} = -R_{+}(t).$$

□

Appendix H. Agreement Information in Cliques of Size N . Let A_1 denote the set of agreeing agents and U_1 denote the set of undecided agents after the initial decision. We want to know what an undecided agent $j \in U_1$ (so $d_{T^{(1)},1}^{(j)} = 0$) learns by observing the distribution of A_1 and U_1 . To compute the evidence update for agent j first note that,

$$\begin{aligned} P(d_{T^{(1)},0}^{(1)} = 1, d_{T^{(1)},1}^{(i)}, d_{T^{(1)},1}^{(k)} | H) &= P(d_{T^{(1)},1}^{(i)}, d_{T^{(1)},1}^{(k)} | d_{T^{(1)},0}^{(1)} = 1, H) P(d_{T^{(1)},0}^{(1)} = 1 | H) \\ &= P(d_{T^{(1)},1}^{(i)} | d_{T^{(1)},0}^{(1)} = 1, H) P(d_{T^{(1)},1}^{(k)} | d_{T^{(1)},0}^{(1)} = 1, H) P(d_{T^{(1)},0}^{(1)} = 1 | H). \end{aligned}$$

for any pair of agents $i \neq j$ different from 1. Therefore,

$$y_{T^{(1)},1}^{(j)} = y_{T^{(1)},0}^{(j)} + \log \frac{P(d_{T^{(1)},0}^{(1)} = 1 | H^+)}{P(d_{T^{(1)},0}^{(1)} = 1 | H^-)}.$$

After observing this first wave of decision, all remaining undecided agents update their belief as,

$$\begin{aligned} y_{T^{(1)},2}^{(j)} &= y_{T^{(1)},1}^{(j)} + \log \frac{P(d_{T^{(1)},0}^{(1)} = 1 | H^+)}{P(d_{T^{(1)},0}^{(1)} = 1 | H^-)} + \sum_{i \in A_1} \log \frac{P(d_{T^{(1)},1}^{(i)} = 1 | d_{T^{(1)},0}^{(1)} = 1, H^+)}{P(d_{T^{(1)},1}^{(i)} = 1 | d_{T^{(1)},0}^{(1)} = 1, H^-)} \\ &\quad + \sum_{k \in U_1 - \{j\}} \log \frac{P(d_{T^{(1)},1}^{(k)} = 0 | d_{T^{(1)},0}^{(1)} = 1, H^+)}{P(d_{T^{(1)},1}^{(k)} = 0 | d_{T^{(1)},0}^{(1)} = 1, H^-)}. \end{aligned}$$

By conditional independence, this simplifies to:

$$\begin{aligned} y_{T^{(1)},2}^{(j)} &= y_{T^{(1)},0}^{(j)} + \theta + a_1 \log \frac{P(0 \leq y_{T^{(1)},0}^{(i)} < \theta | H^+)}{P(0 \leq y_{T^{(1)},0}^{(i)} < \theta | H^-)} + (u_1 - 1) \log \frac{P(-\theta < y_{T^{(1)},0}^{(k)} < 0 | H^+)}{P(-\theta < y_{T^{(1)},0}^{(k)} < 0 | H^-)} \\ &= y_{T^{(1)},0}^{(j)} + \theta + a_1 R_+ + (u_1 - 1) R_-, \end{aligned}$$

where $R_{\pm}$ are defined as in Eqs. 7.3 and E.3. By symmetry, $R_-(T^{(1)}) = -R_+(T^{(1)})$. Thus

$$y_{T^{(1)},2}^{(j)} = y_{T^{(1)},0}^{(j)} + \theta + (a_1 - (u_1 - 1)) R_+(T^{(1)}).$$

REFERENCES

- [1] D. ACEMOGLU, M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR, *Bayesian learning in social networks*, The Review of Economic Studies, 78 (2011), pp. 1201–1236.
- [2] R. J. AUMANN, *Agreeing to disagree*, The annals of statistics, (1976), pp. 1236–1239.
- [3] B. BAHAMI, K. OLSEN, P. E. LATHAM, A. ROEPSTORFF, G. REES, AND C. D. FRITH, *Optimally interacting minds*, Science, 329 (2010), pp. 1081–1085.
- [4] F. BALCI, P. SIMEN, R. NIYOGI, A. SAXE, J. A. HUGHES, P. HOLMES, AND J. D. COHEN, *Acquisition of decision making criteria: reward rate ultimately beats accuracy*, Atten. Percept. Psychophys., 73 (2011), pp. 640–657.
- [5] A. V. BANERJEE, *A simple model of herd behavior*, Q. J. Econ., (1992), pp. 797–817.
- [6] J. M. BECK, W. J. MA, X. PITKOW, P. E. LATHAM, AND A. POUGET, *Not noisy, just wrong: the role of suboptimal inference in behavioral variability*, Neuron, 74 (2012), pp. 30–39.
- [7] R. BOGACZ, E. BROWN, J. MOEHLIS, P. HOLMES, AND J. D. COHEN, *The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks*, Psychological Review, 113 (2006), pp. 700–765.
- [8] R. BOGACZ, E.-J. WAGENMAKERS, B. U. FORSTMANN, AND S. NIEUWENHUIS, *The neural basis of the speed–accuracy tradeoff*, Trends in neurosciences, 33 (2010), pp. 10–16.
- [9] J. R. BUSEMEYER AND J. T. TOWNSEND, *Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment*, Psychological review, 100 (1993), p. 432.
- [10] R. J. CAGINALP AND B. DOIRON, *Decision dynamics in groups with interacting members*, SIAM Journal on Applied Dynamical Systems, 16 (2017), pp. 1543–1562.
- [11] J.-F. CHAMBERLAND AND V. V. VEERAVALLI, *Decentralized detection in sensor networks*, IEEE Transactions on Signal Processing, 51 (2003), pp. 407–416.

- [12] L. CONRADT AND T. J. ROPER, *Consensus decision making in animals*, Trends in ecology & evolution, 20 (2005), pp. 449–456.
- [13] I. D. COUZIN, *Collective cognition in animal groups*, Trends in cognitive sciences, 13 (2009), pp. 36–43.
- [14] I. D. COUZIN, J. KRAUSE, N. R. FRANKS, AND S. A. LEVIN, *Effective leadership and decision-making in animal groups on the move*, Nature, 433 (2005), pp. 513–516.
- [15] M. H. DEGROOT, *Optimal statistical decisions*, vol. 82, John Wiley & Sons, 2005.
- [16] W. EDWARDS, *The theory of decision making*, (1954), pp. 380–417.
- [17] B. ENKE AND F. ZIMMERMANN, *Correlation neglect in belief formation*, CESifo Working Paper Series No. 4483, (2013).
- [18] R. FAGIN, J. Y. HALPERN, Y. MOSES, AND M. VARDI, *Reasoning about knowledge*, MIT press, 2004.
- [19] N. R. FRANKS AND T. RICHARDSON, *Teaching in tandem-running ants*, Nature, 439 (2006), pp. 153–153.
- [20] C. D. FRITH AND U. FRITH, *Mechanisms of social cognition*, Annual review of psychology, 63 (2012), pp. 287–313.
- [21] D. GALE AND S. KARIV, *Bayesian learning in social networks*, Games and Economic Behavior, 45 (2003), pp. 329–346.
- [22] J. GEANAKOPOLOS, *Common knowledge*, Handbook of game theory with economic applications, 2 (1994), pp. 1437–1496.
- [23] J. D. GEANAKOPOLOS AND H. M. POLEMARCHAKIS, *We can't disagree forever*, Journal of Economic Theory, 28 (1982), pp. 192–200.
- [24] J. I. GOLD AND M. N. SHADLEN, *The neural basis of decision making*, Annu. Rev. Neurosci., 30 (2007), pp. 535–574.
- [25] S. B. GOYAL, *Connections: An introduction to the economics of networks*, 2012.
- [26] S. P. HALLER, D. BANG, B. BAHRAMI, AND J. Y. LAU, *Group decision-making is optimal in adolescence*. bioRxiv 412726, 2018.
- [27] T. D. HANKS, M. E. MAZUREK, R. KIANI, E. HOPP, AND M. N. SHADLEN, *Elapsed decision time affects the weighting of prior probability in a perceptual decision task*, Journal of Neuroscience, 31 (2011), pp. 6339–6352.
- [28] M. KIMURA AND J. MOEHLIS, *Group decision-making models for sequential tasks*, SIAM Review, 54 (2012), pp. 121–138.
- [29] I. KRAJBICH, T. HARE, B. BARTLING, Y. MORISHIMA, AND E. FEHR, *A Common Mechanism Underlying Food Choice and Social Decisions*, PLoS Computational Biology, 11 (2015), p. e1004371.
- [30] D. LAZER, A. S. PENTLAND, L. ADAMIC, S. ARAL, A. L. BARABASI, D. BREWER, N. CHRISTAKIS, N. CONTRACTOR, J. FOWLER, M. GUTMANN, ET AL., *Life in the network: the coming age of computational social science*, Science (New York, NY), 323 (2009), p. 721.
- [31] G. LEVY AND R. RAZIN, *Correlation neglect, voting behavior, and information aggregation*, The American Economic Review, 105 (2015), pp. 1634–1645.
- [32] R. D. LUCE, *Response times: Their role in inferring elementary mental organization*, no. 8, Oxford University Press on Demand, 1986.
- [33] P. MILGROM AND N. STOKEY, *Information, trade and common knowledge*, Journal of economic theory, 26 (1982), pp. 17–27.
- [34] N. MILLER, S. GARNIER, A. T. HARTNETT, AND I. D. COUZIN, *Both information and social cohesion determine collective decisions in animal groups*, Proceedings of the National Academy of Sciences, 110 (2013), pp. 5263–5268.
- [35] E. MOSSEL, A. SLY, AND O. TAMUZ, *Asymptotic learning on bayesian social networks*, Probab. Theory Related Fields, 158 (2014), pp. 127–157.
- [36] E. MOSSEL AND O. TAMUZ, *Opinion exchange dynamics*, arXiv preprint arXiv:1401.4770, (2014).
- [37] M. MOUSSAÏD, N. PEROZO, S. GARNIER, D. HELBING, AND G. THERAULAZ, *The walking behaviour of pedestrian social groups and its impact on crowd dynamics*, PloS one, 5 (2010), p. e10047.
- [38] M. MUELLER-FRANK, *A general framework for rational learning in social networks*, Theoretical Economics, 8 (2013), pp. 1–40.
- [39] M. J. MULDER, E.-J. WAGENMAKERS, R. RATCLIFF, W. BOEKEL, AND B. U. FORSTMANN, *Bias in the brain: a diffusion model analysis of prior probability and potential payoff*, Journal of Neuroscience, 32 (2012), pp. 2335–2343.
- [40] M. NAGY, Z. AKOS, D. BIRO, AND T. VICSEK, *Hierarchical group dynamics in pigeon flocks*, Nature, 464 (2010), pp. 890–893.
- [41] M. E. NEWMAN, *The structure and function of complex networks*, SIAM review, 45 (2003), pp. 167–256.
- [42] S. A. PARK, S. GOÏAME, D. A. O'CONNOR, AND J. DREHER, *Integration of individual and social information for decision-making in groups of different sizes*, PLoS Biology, 15 (2017), p. e2001958.
- [43] M. L. PLATT AND P. W. GLIMCHER, *Neural correlates of decision variables in parietal cortex*, Nature, 400 (1999), p. 233.
- [44] I. POULAKAKIS, L. SCARDOVI, AND N. E. LEONARD, *Node classification in networks of stochastic evidence accumulators*, arXiv preprint arXiv:1210.4235, (2012).
- [45] R. M. RAAFAT, N. CHATER, AND C. FRITH, *Herding in humans*, Trends in cognitive sciences, 13 (2009), pp. 420–428.
- [46] A. E. RADILLO, A. VELIZ-CUBA, K. JOSIĆ, AND Z. P. KILPATRICK, *Evidence accumulation and change rate inference in dynamic environments*, Neural Computation, 29 (2017), pp. 1561–1610.
- [47] R. RATCLIFF AND G. MCKOON, *The diffusion decision model: theory and data for two-choice decision tasks*, Neural computation, 20 (2008), pp. 873–922.
- [48] T. SAATY, *How to make a decision: The analytic hierarchy process*, European J. Oper. Res., (1990), pp. 9–26.
- [49] T. D. SEELEY, S. CAMAZINE, AND J. SNEYD, *Collective decision-making in honey bees: how colonies choose among nectar sources*, Behavioral Ecology and Sociobiology, 28 (1991), pp. 277–290.
- [50] M. N. SHADLEN AND W. T. NEWSOME, *Motion perception: seeing and deciding*, Proceedings of the national academy of sciences, 93 (1996), pp. 628–633.
- [51] H. A. SIMON, *Bounded rationality*, in Utility and probability, Springer, 1990, pp. 15–18.
- [52] V. SRIVASTAVA AND N. E. LEONARD, *Collective decision-making in ideal networks: The speed-accuracy tradeoff*, IEEE Transactions on Control of Network Systems, 1 (2014), pp. 121–132.
- [53] C. SUMMERFIELD AND F. P. DE LANGE, *Expectation in perceptual decision making: neural and computational mechanisms*, Nature Reviews Neuroscience, 15 (2014), p. 745.
- [54] H. M. TAYLOR AND S. KARLIN, *An introduction to stochastic modeling*, Academic press, 2014.

- [55] M. USHER AND J. L. MCCLELLAND, *The time course of perceptual choice: the leaky, competing accumulator model.*, Psychological review, 108 (2001), p. 550.
- [56] A. VELIZ-CUBA, Z. P. KILPATRICK, AND K. JOSIĆ, *Stochastic models of evidence accumulation in changing environments*, SIAM Review, 58 (2016), pp. 264–289.
- [57] A. VOSS, K. ROTHERMUND, AND J. VOSS, *Interpreting the parameters of the diffusion model: An empirical validation*, Memory & Cognition, 32 (2004), pp. 1206–1220.
- [58] A. WALD, *Sequential tests of statistical hypotheses*, The annals of mathematical statistics, 16 (1945), pp. 117–186.
- [59] A. WALD AND J. WOLFOWITZ, *Optimum character of the sequential probability ratio test*, The Annals of Mathematical Statistics, (1948), pp. 326–339.
- [60] D. J. WATTS, *A simple model of global cascades on random networks*, Proceedings of the National Academy of Sciences, 99 (2002), pp. 5766–5771.
- [61] G. YAN, T. ZHOU, B. HU, Z.-Q. FU, AND B.-H. WANG, *Efficient routing on complex networks*, Physical Review E, 73 (2006), p. 046108.