

Bayesian Evidence Accumulation on Social Networks*

Bhargav Karamched[†], Simon Stolarczyk[†], Zachary P. Kilpatrick[‡], and Krešimir Josić[§]

Abstract. To make decisions we are guided by the evidence we collect and the opinions of friends and neighbors. How do we combine our private beliefs with information we obtain from our social network? To understand the strategies humans use to do so, it is useful to compare them to observers that optimally integrate all evidence. Here we derive network models of rational (Bayes optimal) agents who accumulate private measurements and observe the decisions of their neighbors to make an irreversible choice between two options. The resulting information exchange dynamics has interesting properties: When decision thresholds are asymmetric, the absence of a decision can be increasingly informative over time. In a recurrent network of two agents, the absence of a decision can lead to a sequence of belief updates akin to those in the literature on common knowledge. On the other hand, in larger networks a single decision can trigger a cascade of agreements and disagreements that depend on the private information agents have gathered. Our approach provides a bridge between social decision making models in the economics literature, which largely ignore the temporal dynamics of decisions, and the single-observer evidence accumulator models used widely in neuroscience and psychology.

Key words. decision making, probabilistic inference, social networks, stochastic dynamics, first passage time problems

AMS subject classifications. 62C10, 91D30

DOI. 10.1137/19M1283793

1. Introduction. Understanding how organisms use sensory and social information to make decisions is of fundamental interest in biology, sociology, and economics [12, 18, 22]. Psychologists and neuroscientists have developed a variety of experimental approaches to probe how humans and other animals make choices. Variants of the two-alternative forced choice (2AFC) task where an observer is asked to decide between two options based on noisy observations are particularly popular [7, 26, 52, 54]. The 2AFC task has motivated several mathematical models that successfully explain how humans use sequentially presented evidence to make decisions [9, 15, 62].

Most such evidence accumulation models take the form of drift-diffusion stochastic pro-

*Received by the editors August 26, 2019; accepted for publication (in revised form) May 24, 2020; published electronically August 18, 2020.

<https://doi.org/10.1137/19M1283793>

Funding: This work was supported by NSF/NIH CRCNS grant R01MH115557 and NSF grant DMS-1517629. The work of the first and fourth authors was supported by NSF grant DMS-1662305. The work of the fourth author was also supported by NSF NeuroNex grant DBI-1707400. The work of the third author was also supported by NSF grants DMS-1615737 and DMS-1853630.

[†]Department of Mathematics, University of Houston, Houston, TX 77204 (bkaramched@yahoo.com, simon.stolarczyk@gmail.com).

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

[§]Department of Mathematics, Department of Biology and Biochemistry, University of Houston, Houston, TX 77204, and Department of BioSciences, Rice University, Houston, TX 77005 (josic@math.uh.edu).

cesses that describe the information gathered by lone observers [7]. However, humans and many other animals do not live in isolation and may consider their neighbors' behavior as they make decisions. Animals watch each other as they forage [34, 39]. Stock traders, while not privy to all of their competitor's information, can observe each other's decisions. To make the best choices, many biological agents thus also take into account the observed choices of others [42, 47, 50].

Key to understanding the collective behavior of social organisms is uncovering how they gather and exchange information to make decisions [12, 22]. 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 [11]. Such models can thus lead the way towards understanding the decisions and disagreements that emerge in social groups.

Here we address the question of how idealized rational 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. Previous models of this type were not normative, as the belief update in response to an observed decision is a parameter [10]. Hence such models cannot describe the interplay between the decisions, nondecisions, and the belief updates we discuss here.

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 [13, 45]. Foraging animals may communicate their preferred feeding locations [21, 55] but not the evidence or process they used to decide. Traders can see their competitor's choice to buy or sell a stock but may not have access to the evidence that leads to these actions. Second, if agents communicate all the information they gathered to their neighbors, the problem is mathematically trivial, as every agent obtains all evidence provided to the entire network.

The behavior and performance of rational agents in a social network can depend sensitively on how information is shared [1, 40]. In some cases rational agents perform better in isolation than as part of a social network, even when they use all available information to decide on the most likely of several options. 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, 23]. In strongly connected networks, on the other hand, agents can repeatedly share their beliefs with one another and asymptotically come to agree on the most probable choice—a phenomenon known as *asymptotic learning*. [2, 24, 25, 40].

We show that such asymptotic learning typically does not occur in our model, as the decisions are immutable. We describe what information rational agents can obtain from

observing the decision states of their neighbors. Interestingly when decision thresholds are asymmetric, the absence of a decision in strongly connected networks can lead to a recursive belief update akin to that addressed in the literature on *common knowledge* [2, 24, 38]. We also show how a rational agent who only observes a portion of the network must marginalize over the decision states of all unobserved agents to correctly integrate an observed decision.

Social evidence exchange in larger networks becomes tractable when all agents can observe each other. We show that in such networks the first decision can lead to a wave of agreements. However, the agents who remain undecided can obtain further information by counting how many others have made a decision and how many remain on the fence. We show that this process can be self-correcting and lead to quicker and more accurate decisions compared to lone agents.

2. Definitions and setup.¹ We consider a model of a social network in which all agents are deciding between two options or hypotheses. To make a choice, agents gather both private (Priv) and social (Soc) information over time. We assume that private information comes from a sequence of noisy observations (measurements). In addition we assume that agents also gather social information by continuously observing each other's decisions. For instance, while foraging, animals make private observations but also observe each other's decisions [14, 39]. When deciding on whether or not to purchase an item or which of two candidates to vote for, people will rely on their own research but are also influenced by the decisions of their friends and acquaintances [1] as well as opinions on social networking sites [67]. In many of these situations agents do not share all information they gathered directly but only observe their neighbors' choices or the absence of such choices (e.g., not purchasing an item, or not going to the polls).

More precisely, we consider a set of agents who accumulate evidence to decide between two states, or hypotheses, H^+ or H^- . Each agent is rational (Bayesian): they compute the probability that either hypothesis holds based on all evidence they accrued. The agents make a decision once the conditional probability of one hypothesis, given all the accumulated observations, crosses a predetermined threshold [7, 66]. For simplicity, we assume that all agents in the network are identical, but we discuss how this condition can be relaxed.

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, 26, 51, 52, 62, 63, 66]. In the simplest setting 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 the probability of each measurement is given by the probability mass functions $f_\pm(\xi) := P(\xi|H^\pm)$. Observations, ξ_i , are drawn from the same set Ξ in either state H^\pm , and the two states are distinguished by the differences in probabilities of making certain measurements.

¹For the convenience of the reader a table of notation and symbols is available at <https://github.com/Bargo727/NetworkDecisions.git>, along with the code used to generate all figures.

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

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

where we define $\text{LLR}(\cdot) = \log \frac{P(\cdot|H^+)}{P(\cdot|H^-)}$; such log likelihood ratios will be used extensively in what follows. Also note that $y_0 = 0$, since both states are equally likely a priori. 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 $\theta_- < 0 < \theta_+$.

3. Multiple agents. In social networks not all agents will communicate with each other due to limitations in bandwidth or physical distances [35, 46, 68]. To model such exchanges, 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 [27]. Information flows along the direction of the arrow (Figure 1a), so agents observe neighbors in the direction opposite that of the arrows.

As in the case of a single observer, we assume that each agent, i , makes a sequence of noisy, identically distributed measurements, $\xi_{1:t}^{(i)}$, 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 incoming evidence simplifies calculations, but is unlikely to hold in practice. However, humans often treat redundant information as uncorrelated, thus making the same independence assumption as we do in our model [19, 36] (see section 9).

Each agent gathers social evidence by observing whether its neighbors have made a decision and 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 is still in the interval (θ_-, θ_+) .

For simplicity, we assume that the measurement distributions, f_\pm , and thresholds, θ_\pm , are identical across agents. The theory is similar if agents have different, but known, measurement distributions. The assumption that the distributions, $f_\pm(\xi)$, are discrete simplifies some convergence arguments. However, most evidence accumulation models take the form of continuous, drift-diffusion stochastic processes [7]. These models take the form of stochastic differential equations (SDEs) and approximate the discrete model well when many observations are required to reach a decision [63]. Such SDEs have been remarkably successful in describing the responses of humans and other animals under a variety of conditions [53].

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 Figure 1a,

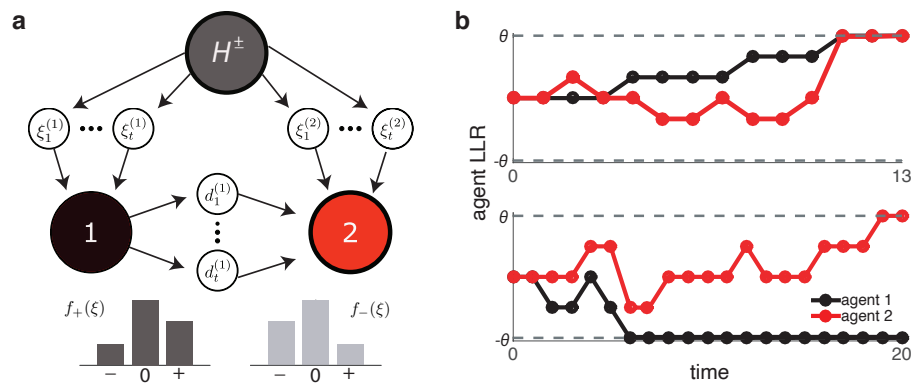


Figure 1. A pair of unidirectionally coupled agents deciding between the states H^+ and H^- on the basis of private and social evidence. (a) Schematic of the information flow in the network. Agent 1 accumulates its own observations, $\xi_{1:t}^{(1)}$, which result in a sequence of decision states, $d_{1:t}^{(1)}$, 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. Here, $H = H^+$.

where agent 2 observes agent 1 but not vice versa.

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, which jointly constitute $I_t^{(2)}$. Agents base their decisions on the computed LLR, or belief, $y_t^{(i)} = \log [P(H^+|I_t^{(i)})/P(H^-|I_t^{(i)})]$, and agent i makes a choice at time $T^{(i)}$ at which $y_t^{(i)}$ crosses threshold θ_+ or θ_- .

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

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

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

Assuming that the two agents make private observations synchronously, the information available to agent 2 at time t consists of its private observations and observed decision states of agent 1, $I_t^{(2)} = (\xi_{1:t}^{(2)}, d_{1:t}^{(1)})$. Importantly, we assume that an observer makes a private observation and then integrates all available social information *before* making its next private observation. Hence, evidence accumulation at time t consists of two steps: Agents first update their belief in response to a private observation. They then observe the decision states of other visible agents and update their belief again. We will see that in recurrent networks the exchange of social information may occur over several steps (see sections 5, 7, and 8).

We let $\text{Soc}_t^{(i)}$ be the social information that can be obtained by observing the decisions of

agent i up to time t so that

$$(3.2) \quad \text{Soc}_t^{(1)} = \text{LLR}(d_{1:t}^{(1)}).$$

The stochastic process $\text{Priv}_{1:t}$ is defined in the same way for both agents and is equivalent to that of a lone observer given in (2.1).

Belief updates rely on the assumption that the private observations of the two agents are conditionally independent and a straightforward application of Bayes' Theorem. In the present example each private observation leads agent 2 to update its belief to $y_{t,0}^{(2)}$ by adding $\text{Priv}_t^{(2)}$. After next observing the decision state of agent 1, the belief of agent 2 is a sum of the LLR corresponding to private observations ($\text{Priv}_{1:t}^{(2)}$) and the LLR corresponding to the social information ($\text{Soc}_t^{(1)}$):

$$(3.3) \quad y_{t,1}^{(2)} = \log \frac{P(H^+ | \xi_{1:t}^{(2)}, d_{1:t}^{(1)})}{P(H^- | \xi_{1:t}^{(2)}, d_{1:t}^{(1)})} = \underbrace{\sum_{l=1}^t \text{LLR}(\xi_l^{(2)})}_{\text{Priv}_{1:t}^{(2)}} + \underbrace{\text{LLR}(d_{1:t}^{(1)})}_{\text{Soc}_t^{(1)}}.$$

As decisions are immutable, the only information about agent 1's decision that agent 2 can use after a decision at time T is that agent 1's decision state switched from $d_{T-1}^{(1)} = 0$ to $d_T^{(1)} = \pm 1$ at $t = T$.

4. Social information. We next ask how much evidence is provided by observing the decision state of other agents in the directed network in Figure 1a and illustrate how non-decisions, $d_t^{(1)} = 0$, can be informative.

Decision evidence. Suppose agent 1 chooses H^+ at time T so that the belief of agent 2 at time $t \geq T$ is $y_{t,1}^{(2)} = \text{Priv}_{1:t}^{(2)} + \text{Soc}_t^{(1)} = \text{Priv}_{1:t}^{(2)} + \text{Soc}_T^{(1)}$. If $d_T^{(1)} = 1$, then agent 2 knows that $y_{T,0}^{(1)} \geq \theta_+$. Agent 1 has reached its decision based only on its private observations; hence, none of the social information obtained by agent 2 is redundant.

The belief of agent 1 at the time of the decision, T , could exceed the threshold. However, this excess is small if the evidence obtained from each measurement is small. The following proposition shows that, after observing a decision, agent 2 updates its belief by an amount close to θ_+ (Figure 1b).

Proposition 4.1. *If agent 1 decides at time T , and $d_T^{(1)} = 1$, then*

$$\theta_+ \leq \text{Soc}_T^{(1)} < \theta_+ + \epsilon_T^+, \quad \text{where} \quad 0 < \epsilon_T^+ \leq \sup_{\xi_{1:T} \in \mathcal{C}_+(T)} \text{LLR}(\xi_T),$$

and $\mathcal{C}_+(T)$ is the set of all chains of observations, $\xi_{1:T}$, that trigger a decision in favor of H^+ precisely at time T . An analogous result holds when $d_T^{(1)} = -1$.

Proof. Assume $d_T = 1$, and note that $d_{T-1}^{(1)} = 0$ and $d_T^{(1)} = 1$ imply $\sum_{t=1}^T \text{LLR}(\xi_t^{(1)}) \geq \theta_+$, so

$$(4.1) \quad e^{\theta_+} \prod_{t=1}^T f_-(\xi_t^{(1)}) \leq \prod_{t=1}^T f_+(\xi_t^{(1)}).$$

We define the set of all chains of observations, $\xi_{1:T}$, that lead to a decision precisely at T : $\mathcal{C}_+(T) = \{\xi_{1:T}^{(1)} \mid \text{Priv}_{1:t}^{(1)} \in (\theta_-, \theta_+), 1 \leq t < T; \text{Priv}_{1:T}^{(1)} \geq \theta_+\}$. Thus, $P(d_{T-1}^{(1)} = 0, d_T^{(1)} = 1 \mid H^+) = \sum_{\mathcal{C}_+(T)} f_+(\xi_{1:T}^{(1)})$, where $f_+(\xi_{1:T}^{(1)}) := \prod_{t=1}^T f_+(\xi_t^{(1)})$. Together with inequality (4.1) this yields

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

hence, $\text{LLR}(d_{T-1}^{(1)} = 0, d_T^{(1)} = 1) \geq \theta_+$. To obtain the upper bound on the social information inferred by agent 2 from the decision of agent 1 ($\text{Soc}_T^{(1)}$), we must bound the possible private information available to agent 1 ($\text{Priv}_{1:T}^{(1)}$). Note for any $\xi_{1:T}^{(1)} \in \mathcal{C}_+(T)$,

$$\text{Priv}_{1:T}^{(1)} = \text{Priv}_{1:T-1}^{(1)} + \text{LLR}(\xi_T^{(1)}) < \theta_+ + \text{LLR}(\xi_T^{(1)}),$$

which implies

$$\epsilon_T^+ \leq \sup_{\xi_{1:T} \in \mathcal{C}_+(T)} \text{LLR}(\xi_T),$$

by positivity of $\text{LLR}(\xi_T)$, since ξ_T triggers a decision, and $\text{LLR}(\xi_T) \leq \sup_{\xi_{1:T} \in \mathcal{C}_+(T)} \text{LLR}(\xi_T)$. A similar argument can be used to derive bounds on $\text{Soc}_T^{(1)}$ for $d_T^{(1)} = -1$. ■

Note that the evidence accumulation process will be analogous if the agents' thresholds differ. In particular, if agent 1's thresholds are θ_\pm^1 and agent 2's are θ_\pm^2 , then if agent 2 observes a decision by agent 1, it will update its LLR by approximately θ_\pm^1 , depending on the decision of agent 1.

In the following we will frequently assume that the two distributions $f_+(\xi)$ and $f_-(\xi)$ are close, and hence $\text{LLR}(\xi)$ is small for any ξ . Proposition 4.1 then implies that $\text{Soc}_T^{(1)} \approx \theta_+$. This is a common approximation in the literature on sequential probability ratio tests [65].

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 show this explicitly in the following.

Definition 4.2. *The measurement distributions $P(\xi \mid H^+) = f_+(\xi)$ and $P(\xi \mid H^-) = f_-(\xi)$ are symmetric if there exists an involution $\Phi : \Xi \rightarrow \Xi$, i.e., $\Phi = \Phi^{-1}$ such that $f_+(\xi) = f_-(\Phi(\xi))$ 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 [56]. In much of the decision making literature, for example, it is assumed that $\Xi = -\Xi$ and $\Phi(\xi) = -\xi$. In the following we let $\Phi(\xi) = -\xi$ when discussing distribution symmetry. 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 [29]. In examples we will assume that agents use asymmetric decisions thresholds ($\theta_+ \neq -\theta_-$) due to a known asymmetry in the 2AFC task. If the two thresholds differ, then an agent more readily adopts one of the two options.

When the measurement distributions are asymmetric then the agent can more easily obtain information about one of the two options. In such a situation, nondecision on the part of an agent will provide information to those agents observing it.

We call the social information agent 2 gathers before observing a choice by agent 1 the *nondecision evidence*. This social information is determined by the decomposition in (3.3) and (3.2). The *survival probability* of the stochastically evolving belief is given by

$$(4.2) \quad S_{\pm}(t) = P(d_t^{(1)} = 0 | H^{\pm}) = P(y_{s,0}^{(1)} \in (\theta_-, \theta_+), 0 \leq s \leq t | H^{\pm}).$$

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

$$(4.3) \quad \text{Soc}_t^{(1)} = \log \frac{S_+(t)}{S_-(t)} = \text{LLR}(T > t).$$

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.

When the threshold and measurement distributions are symmetric, then every sequence of observations $\xi_{1:t}^{(1)}$ favoring one decision has a matching sequence, $\Phi(\xi_1^{(1)}), \dots, \Phi(\xi_t^{(1)})$, providing equal evidence for the opposite decision since

$$\sum_{l=1}^t \text{LLR}(\xi_l^{(1)}) = - \sum_{l=1}^t \text{LLR}(\Phi(\xi_l^{(1)})).$$

By symmetry, both sequences of observations are equally likely, implying $S_+(t) = S_-(t)$. We have therefore shown that the following holds.

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

We showed above that when measurement distributions and thresholds are symmetric, nondecisions are uninformative in any network until the first decision is made. When thresholds and measurement distributions are not symmetric the absence of a decision of agent 1 can provide evidence for one of the choices. In particular, *social information from a nondecision is deterministic*, which we illustrate in the following example.

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 $\text{LLR}(\xi) \in \{+1, 0, -1\}$ with respective probabilities $\{p, s, q\}$ ($\{q, s, p\}$) when $H = H^+$ ($H = H^-$). In particular, we take $p/q = e$. This ensures that a decision based solely on private evidence results from a belief (LLR) exactly at integer thresholds.² Here we take $\theta_+ = 2$ and $\theta_- = -1$. Beliefs then evolve as biased random walks on the integer lattice, as demonstrated below.

²Unless we state otherwise, we choose p, q, s so that decisions based solely on private evidence result in an LLR that can reach the threshold exactly. In sections 6, 7, and 8 we modify p/q to study network dynamics for various levels of noise. Unless stated otherwise, we take $s = 0.5$.

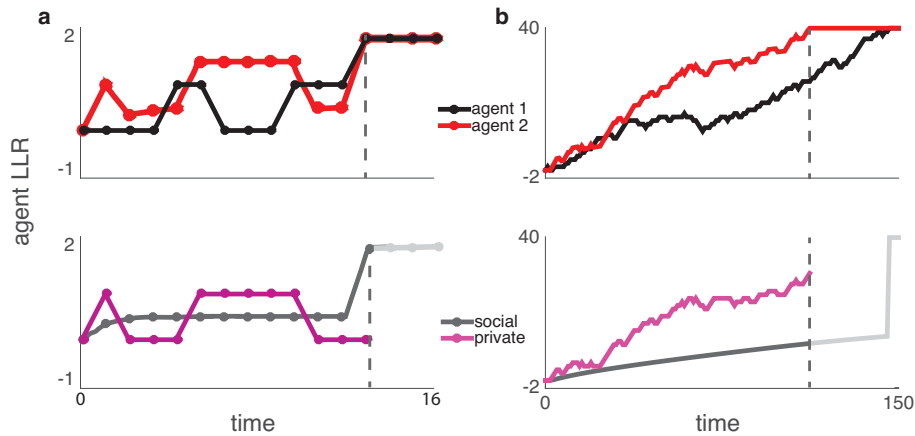


Figure 2. An example of the belief evolution in a two-agent, unidirectional network shown in Figure 1. (a) The belief of agent 1 is a random walk on the integers. For boundaries at $\theta_- = -1$ and $\theta_+ = 2$, initially two observations favoring H^+ are required for decision $d_2^{(1)} = 1$. The belief of agent 2 is the sum of stochastically evolving private evidence and deterministic social evidence until observing a choice. The belief of agent 2 is incremented by exactly θ_\pm when the agent observes a decision in this example. The social information equals the thresholds in this case, as by design agent 1 makes a decision when its belief equals θ_\pm . (b) The same processes with $\theta_- = -2$ and $\theta_+ = 40$. Here, $H = H^+$. Code to produce all the figures in the manuscript is available at <https://github.com/Bargo727/NetworkDecisions.git>.

Agent 1 makes only private observations, while agent 2 makes private observations and observes agent 1. Their beliefs are described by (2.1) and (3.3), respectively. Realizations of these processes are shown in Figure 2. The social evidence obtained by agent 2 is independent of the particular realization until agent 1 makes a decision, whereas private information is realization dependent. When thresholds are small, an expression for social evidence can be obtained explicitly.

Standard first-step analysis [61] yields nondecision social evidence, in our case $\theta_+ = 2$ and $\theta_- = -1$. Here 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 $\mathbf{P}_+(t) := (P_+^{-1}(t), P_+^0(t), P_+^1(t), P_+^2(t))^T$, the probabilities are updated according to

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

with initial condition $P_+^0(0) = 1$ and $s = 1 - p - q$. The rates p and q switch places in the $P_-^l := P(y_t^{(1)} = l | H^-)$ case.

We solve explicitly for the evolution of the probability of the unabsorbed states $\mathbf{v}_+(t) = (P_+^0(t), P_+^1(t))^T$ using the eigenpairs of the submatrix:

$$\mathbf{v}_+(t) = \frac{\lambda_+^t}{2} \begin{pmatrix} 1 \\ \sqrt{p/q} \end{pmatrix} - \frac{\lambda_-^t}{2} \begin{pmatrix} -1 \\ \sqrt{p/q} \end{pmatrix},$$

where $\lambda_{\pm} := s \pm \sqrt{pq}$, and for $\mathbf{v}_-(t) = (P_-^0(t), P_-^1(t))^T$ we have

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

Thus for $d_t^{(1)} = 0$,

$$\text{Soc}_t^{(1)} = \log \frac{(\lambda_+^t + \lambda_-^t) + \sqrt{p/q}(\lambda_+^t - \lambda_-^t)}{(\lambda_+^t + \lambda_-^t) + \sqrt{q/p}(\lambda_+^t - \lambda_-^t)} = \log \frac{(1 - \sqrt{p/q}) + (1 + \sqrt{p/q})(\lambda_+/\lambda_-)^t}{(1 - \sqrt{q/p}) + (1 + \sqrt{q/p})(\lambda_+/\lambda_-)^t},$$

and for $s > 0$,

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

since $|\lambda_+| > |\lambda_-|$. For our choice of parameters, $\sqrt{p/q} = \sqrt{e}$, so $\lim_{t \rightarrow \infty} \text{Soc}_t^{(2)} = \frac{1}{2}$.

When all observations are informative ($s = 0$), the social information will alternate indefinitely between two values. In this case, $\lambda_+/\lambda_- = -1$, and we have

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

since the belief of agent 1 must be at lattice site 0 (1) after an even (odd) number of observations.

These formulae reveal that social evidence is typically not accrued significantly beyond the first few observations when decision thresholds are small. This is because social information is bounded above by $\frac{1}{2} \log \frac{p}{q}$, which will only be large in cases that p/q is very large and decisions will almost always be rapid (occurring after two or three time steps). We show this for $\theta_- = -1$ and $\theta_+ = 2$ in Figure 2a, with $\text{Soc}_t^{(1)} \rightarrow 1/2$ as $t \rightarrow \infty$ for our choice of parameter values.

Note that p/q measures the noisiness of the independent observations of the environment and is a proxy for the signal-to-noise ratio (SNR). Generally speaking, if $p \gg q$, then $\lim_{t \rightarrow \infty} \text{Soc}_t^{(1)}$ is large. On the other hand, if observations are noisy, i.e., $p \approx q$, then $\lim_{t \rightarrow \infty} \text{Soc}_t^{(1)} \approx 0$. We investigate how varying p/q shapes network dynamics in more detail below and in sections 6, 7, and 8.

In general, social information, $\text{Soc}_t^{(1)}$, will converge to a value on the side of the threshold with larger absolute value 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 ratio $|\theta_+/\theta_-|$ with the increment in beliefs fixed at three different values of p/q . 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 Figure 3a). In particular, as p/q decreases so that private observations provide less information, the social information has an increasing impact on accuracy.

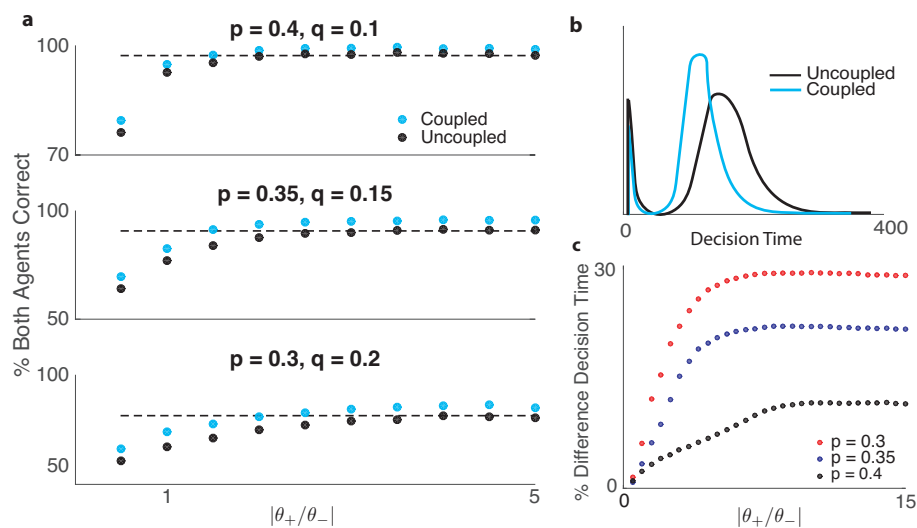


Figure 3. Statistics of decisions in a two-agent network with (coupled) and without (uncoupled) communication from agent 1 to agent 2. The lower threshold is fixed at $\theta_- = -2$ throughout. (a) The fraction of times both agents selected the correct choice as a function of asymmetry in the system 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 agent 2 to make a decision 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 .

Social information also affects decision times, particularly in the case of strongly asymmetric thresholds (Figure 3b). 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 $\theta_+ \gg -\theta_-$. As p/q increases, the difference in decision times between the agents decreases (see Figure 3c), as social information has a larger impact on decisions when private measurements are unreliable.

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 $NSoc_t^{(1)}$, where $Soc_t^{(1)}$ is as defined in (4.3). Thus the impact of nondecision 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 consider two agents that can observe and react to each other's choices. We again 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. Importantly, we assume agents

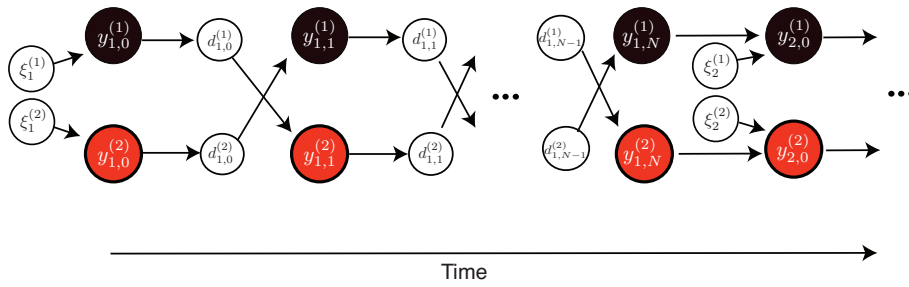


Figure 4. In a recurrent network of two agents the LLRs, $y_{t,n}^{(1)}$ and $y_{t,n}^{(2)}$, of the two observers are updated recursively. Agents update their belief after private observations, $\xi_t^{(i)}$, as well as observations of the subsequent series of decision states, $d_{t,n}^{(j)}$, of their 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.

add the information obtained from private observations first and then check whether or not they have sufficient information to make a decision. If not, agents gather social evidence from other agents in the network.

Unlike previous cases we considered, social information exchange can occur over several steps. We assume that agents exchange social information until they cannot learn anything new from each other and then make their next private observation (see Figure 4). Alternatively, we could assume that each social information exchange is followed by a private observation, but the analysis would proceed similarly.

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 is roughly bounded by the LLR increment needed to drive the previous estimate of social information across the threshold. We also show that the exchange of social information after an observation terminates in a finite number of steps either when indecisions provide 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 [38]. This body of work shows that rational agents that repeatedly announce their preferred choice must eventually reach agreement [2, 20, 24, 25, 43]. 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 describe the process of social exchange inductively, describing the basic case following an observation at $t = 1$ in some detail. The 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

$$(5.1) \quad \Pi = \{\text{LLR}(\xi) \mid \xi \in \Xi\}$$

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 one observes a decision or 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 subsequent private measurement.

We again associate a belief, $y_{1,n}^{(i)}$, with a corresponding decision state, $d_{1,n}^{(i)}$, as in (3.3). If neither of the first two private observations leads to a decision, then $d_{1,0}^{(i)} = 0$, and $y_{1,0}^{(i)} \in \Theta$ for $i = 1, 2$, where $\Theta \equiv (\theta_-, \theta_+)$. 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

$$\begin{aligned} y_{1,1}^{(i)} &= \text{LLR}(\xi_1^{(i)}, d_{1,0}^{(j)}) = \text{LLR}(\xi_1^{(i)}) + \text{LLR}(d_{1,0}^{(j)} = 0) \\ &= y_{1,0}^{(i)} + \text{LLR}(d_{1,0}^{(j)} = 0) = \text{Priv}_1^{(i)} + \underbrace{\text{LLR}(d_{1,0}^{(j)} = 0)}_{\text{Soc}_{1,1}} = \text{Priv}_1^{(i)} + \text{Soc}_{1,1}. \end{aligned}$$

We omit the superscripts on the social information, since $\text{Soc}_{1,1}^{(1)} = \text{Soc}_{1,1}^{(2)}$, as 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 section 4.

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 $j \neq i$, so $\theta_- < y_{1,0}^{(j)} + \text{Soc}_{1,1} < \theta_+$. As the initial private observation of agent j did not lead to a decision, agent i also knows $y_{1,0}^{(j)} \in \Theta$. 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,

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

Any initial measurement $\xi_1^{(j)}$ that led to a belief $y_{1,0}^{(j)} \notin \Theta_{1,1}$ would have also led 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. If an observation does not lead to a decision after the first exchange of social information, then $y_{1,0}^{(j)} = \text{Priv}_1^{(j)} \in \Theta_{1,1}$ implies that

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

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

To extend this argument we define $\Theta_{1,l-1} \equiv \bigcap_{n=0}^{l-1} (\Theta - \text{Soc}_{1,n})$ so that $\text{Priv}_1^{(i)} \in \Theta_{1,l-1}$ implies that agent i has not made a decision at step $l-1$ of the social exchange, $d_{1,l-1}^{(i)} = 0$. In this case we have

$$P(\xi_1^{(i)}, d_{1,0}^{(i)} = 0, \dots, d_{1,l}^{(i)} = 0, d_{1,0}^{(j)} = 0, \dots, d_{1,l}^{(j)} = 0 | H^+) = P(\xi_1^{(i)} | H^+) P(d_{1,l}^{(j)} = 0 | d_{1,l-1}^{(j)} = 0, H^+).$$

Thus if the l th exchange of social information results in no decision, each agent updates its belief as

$$(5.3) \quad y_{1,l}^{(i)} = \text{Priv}_1^{(i)} + \text{LLR}(y_{1,0}^{(j)} \in \Theta_{1,l-1}) \equiv \text{Priv}_1^{(i)} + \text{Soc}_{1,l},$$

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

This exchange of social information continues until one of the agents makes a choice or the absence of a decision does not lead to new social information [25]. The second case occurs at a step, N_1 , at which indecision provides no further information about the belief of either agent, so that

$$(5.4) \quad \Pi \cap \Theta_{1,N_1} = \Pi \cap \Theta_{1,N_1+1},$$

where Π is defined in (5.1). In this case, $\text{Soc}_{1,N_1} = \text{Soc}_{1,N_1+1}$, and if neither agent decides, 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 this exchange is complete by $\text{Soc}_1 = \text{Soc}_{1,N_1}$, and the resulting belief 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 $y_{1,0}^{(i)} = \text{Priv}_1^{(i)} \in (\bigcap_{n=0}^{N_1} (\Theta - \text{Soc}_{1,n})) \equiv \Theta_1$.

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 other's 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 $\text{Priv}_t^{(i)} = \text{LLR}(\xi_t^{(i)})$ be the private information obtained from an observation at time t . For the inductive step we assume that each equilibration process after the observation at time t either ends in a decision or allows each agent i to conclude that the accumulated social and private evidence of agent j was insufficient to cross a decision threshold. We will see that this means that $\text{Priv}_{1:t}^{(j)} \in \Theta_{1:t}$, where Θ_1 was defined above, and we define the other sets in the sequence recursively.

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 $j \neq i$. Let

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

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

$$(5.5) \quad \text{Soc}_{t,l} := \text{LLR} \left(\text{Priv}_{1:t-1}^{(i)} \in \Theta_{1:t-1} \text{ and } \sum_{k=1}^t \text{Priv}_k^{(j)} \in \Theta_{t,l-1} \right)$$

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*

$$(5.6) \quad y_{t,l}^{(i)} = \sum_{k=1}^t \text{Priv}_k^{(i)} + \text{Soc}_{t,l}$$

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 (5.6) is a random variable (depends on realization), while the social evidence is independent of realization and equal for both agents.

Proof. We prove the theorem using double induction. The outer induction is on the times of the private observations, t , while the inner induction is on the steps in the equilibration process. The basic case of the induction over t was proved in the previous section. Induction over t in this case follows similarly. To address the inner steps, 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 $d_{1:t,1:l}^{(i)} = 0$.

At $l = 0$, we assume that equilibration following private observation $\xi_t^{(i)}$ terminates at 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

$$y_{t,0}^{(i)} = \text{LLR}(\xi_{1:t+1}^{(i)}, d_{1:t,1:N_t}^{(i)} = 0, d_{1:t,1:N_t}^{(j)} = 0) = \sum_{k=1}^{t+1} \text{Priv}_k^{(i)} + \text{Soc}_{t+1,0},$$

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

Suppose no decision is 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

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

It follows that

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

where $\text{Soc}_{t+1,l+1}$ is defined in (5.5). This exchange of social information stops at $l = N_{t+1}$ when $\Pi \cap \Theta_{t+1,N_t} = \Pi \cap \Theta_{t+1,N_{t+1}}$, and neither agent learns anything further from the absence of a decision by its counterpart. ■

Belief update after 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 during the n th step of the subsequent social information exchange process, $d_{T,n}^{(i)} = +1$. Then agent $j \neq i$ updates its belief as*

$$y_{T,n+1}^{(j)} = \text{Priv}_{1:T}^{(j)} + \text{Soc}_{T,n+1}^+, \quad \text{with} \quad \theta_+ - \text{Soc}_{T,n} < \text{Soc}_{T,n+1}^+ < \theta_+ - \text{Soc}_{T,n} + \epsilon_{T,n}^+,$$

where we can bound

$$\epsilon_{T,0}^+ \leq \sup_{\xi_{1:T} \in \mathcal{C}_+(T,0)} \text{LLR}(\xi_T) \quad \text{and} \quad \epsilon_{T,n}^+ \leq (\text{Soc}_{T,n} - \text{Soc}_{T,n-1}), \quad n > 0,$$

and $\mathcal{C}_+(T,0)$ is the set of all chains of observations leading an agent to choose H^+ at time step $(T,0)$, given that the other agent did not decide before. An analogous result holds for $d_{T,n}^{(i)} = -1$.

A proof of Proposition 5.2 is given in Appendix A. This proposition shows that any social information obtained before observing a decision is subsumed in the information obtained from observing a decision and that the social information acquired from a neighbor's decision is roughly bounded by the increment needed to drive the social information across the threshold. After one of the agents makes a decision, the other agent continues gathering private observations until they make their own decision.

If the two agents have different decision thresholds, the expression for the postdecision belief is more involved but still computable. For simplicity we forgo further discussion of this case. On the other hand, when the thresholds and evidence distributions are symmetric, the evidence accumulation process is much simpler.

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.*

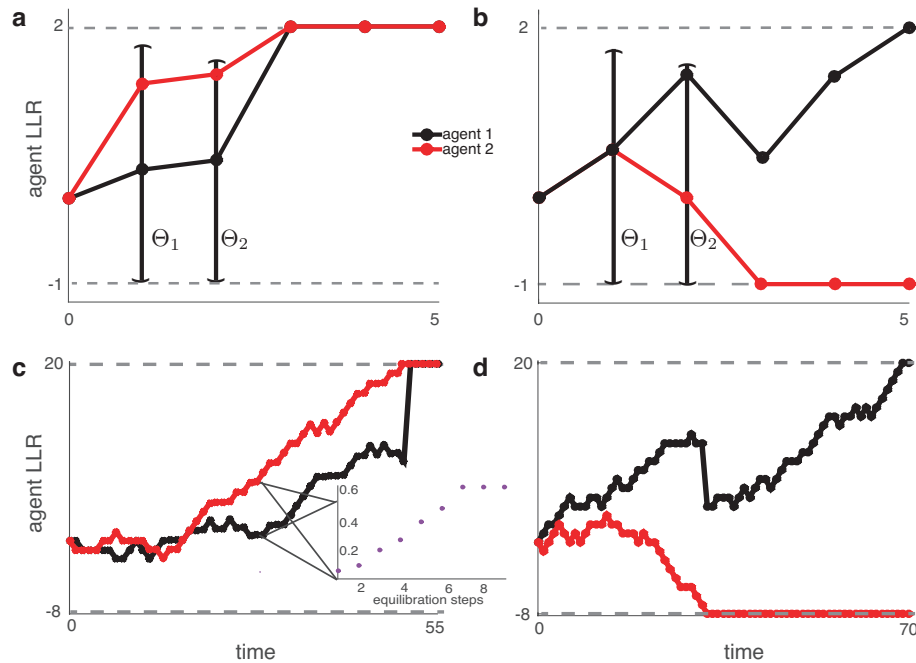


Figure 5. The belief (LLR) of agents 1 and 2 in a recurrent network with asymmetric thresholds when the agents (a) agree and (b) disagree on a final choice. Also shown are the intervals Θ_1, Θ_2 , resulting from the equilibration following the first two observations. Each agent knows that the LLR of 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. Here, $H = H^+$.

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,

$$\text{Soc}_{1,1} = \text{LLR}(y_{1,0}^{(j)} \in \Theta) = 0.$$

Therefore, the update 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. ■

Figure 5 provides examples of evidence accumulation in a two-agent recurrent network. In Figure 5a,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 ends after two steps. Figures 5c,d provide an example with strongly asymmetric decision thresholds. These examples also show that the beliefs of the two agents do not have to converge, and the

agents do not need to agree on a choice, in contrast to classic studies of *common knowledge* [2, 24, 38].

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. This results in a network evidence accumulation process akin to one when agents have asymmetric thresholds. As we have seen in sections 4 and 5, once symmetry is broken, even the absence of a choice provides social information, leading to belief equilibration. We therefore only describe evidence accumulation up to a first decision.

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 $\mathcal{N}^{(i)}$.

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

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

Thus all agents in the network, besides agent i , belong to $\mathcal{N}^{(i)}$ or $\mathcal{U}^{(i)}$. Therefore $\mathcal{N}^{(i)}$, $\mathcal{U}^{(i)}$, and $\{i\}$ form a partition of all nodes in the network.

We denote the set of decisions by the neighbors of agent i following the observation at time t by $d_t^{\mathcal{N}^{(i)}} = \{d_t^{(k)} : k \in \mathcal{N}^{(i)}\}$. Similarly, the set of the decisions by unobserved agents is denoted by $d_t^{\mathcal{U}^{(i)}}$. More generally, $d_{1:t}^{\mathcal{N}^{(i)}}$ denotes the sequence of decision states of the neighbors of agent i up to and including the decision following the observation at time t : $d_{1:t}^{\mathcal{N}^{(i)}} = \{d_s^{\mathcal{N}^{(i)}} : 1 \leq s \leq t\}$. 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 $I_t^{(i)} = \{\xi_{1:t}^{(i)}, d_{1:t}^{\mathcal{N}^{(i)}}\}$. As before, we denote the private information by $\text{Priv}_{1:t}^{(i)} = \text{LLR}(\xi_{1:t}^{(i)})$.

6.2. Nondecisions. In a social network of agents with symmetric thresholds and measurement distributions, two properties of decision information simplify computations: an absence of a decision is uninformative, and once a decision is observed, the resulting social information

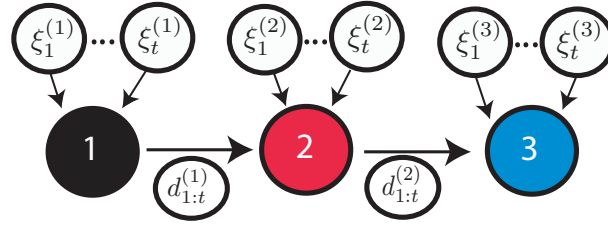


Figure 6. Agents on a line. Agent 3 observes the decisions of agent 2, but not of agent 1. However, agent 3 still needs to take into account the possibility that a choice by agent 1 has caused a decision by agent 2.

is additive. Therefore,

$$y_t^{(i)} = \text{LLR}(\xi_{1:t}^{(i)}, d_{1:t}^{\mathcal{N}^{(i)}} = 0) = \text{Priv}_{1:t}^{(i)}.$$

The proof of this equality follows an argument similar to that given in sections 4 and 5. The main difference is that in the present case each agent marginalizes over unobserved decision states. Symmetry implies that for each decision history of unobserved agents, $d_{1:t}^{\mathcal{U}^{(i)}}$, there is a corresponding opposite decision history, $-d_{1:t}^{\mathcal{U}^{(i)}}$, obtained by flipping the sign of each decision in the vector $d_{1:t}^{\mathcal{U}^{(i)}}$ and leaving nondecisions unaffected. By symmetry, both decision histories are equally probable, and their contributions cancel when agent i computes its belief. Hence no agent i obtains information from observing the absence of a decision by any of its neighbors, and equilibration never starts. As a result agents independently accumulate evidence, and their neighbors' decision states only become informative when one of the neighboring agents 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.

6.3. Example of marginalization: Three agents on a line. We demonstrate the computations needed to account for unobserved agents using the example shown in Figure 6. In this case a decision by agent 3 is not observed by any other agent. Agents 1 and 2 therefore accumulate information as described in section 3. To understand decisions in this network it is therefore sufficient to consider only the computations agent 3 must perform upon observing a decision by agent 2. We will see that observing a decision in a network that is not fully connected can result in a jump in the observer's LLR that is greater than the threshold of the observed decision. The additional information comes from the possibility that decisions of unobserved agents could have led to the decision of an observed agent.

Agent 3 has no social information before observing a decision by agent 2. After observing a decision by agent 2 at time $T^{(2)}$, agent 3 updates its belief by marginalizing over possible decision histories of agent 1:

(6.1)

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

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, θ_\pm , or (b) the decision of agent 1 causing a jump in the

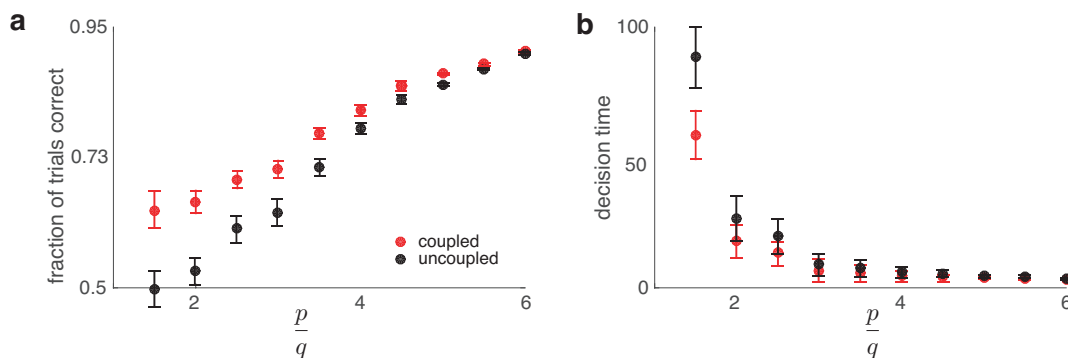


Figure 7. The performance of the three agents on a directed line is better than that of independent agents. (a) The fraction of trials for which all three agents make the correct choice. This quantity is larger when agents are allowed to exchange information (coupled) than when agents make decisions 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 section 4.

belief $y_t^{(2)}$ above threshold θ_+ or below threshold θ_- . Possibility (b) is captured by the term

$$(6.2) \quad P(d_{T(2)}^{(2)} = \pm 1, d_{T(2)-1}^{(1)} = 0, d_{T(2)}^{(1)} = \pm 1 | H^{\pm})$$

in (6.1), while possibility (a) corresponds to all the other summands.

An argument equivalent to that in Proposition 4.1 shows that the social information communicated in case (a) is close to $\pm\theta$ for $d_{T(2)}^{(2)} = \pm 1$. However, in the second case the belief of agent 2 satisfies $y_{T(2)}^{(2)} \in [\theta, 2\theta]$ 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 possibility (a) and (b), and hence increases its belief by an amount greater than θ . We show this explicitly in Appendix C.

In Figure 7, we illustrate how social information impacts decision time and accuracy. When p/q is not too large, and hence the SNR of observations is low, the exchange of social information impacts both the accuracy and response time of agents significantly, and agents make the correct decision more frequently than isolated agents. When the SNR is large, private information dominates decisions, and the impact of social information is small.

Marginalization becomes more complex when the unseen component, $\mathcal{U}^{(i)}$, 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 $(k+1)$ st agent must marginalize over the preceding $k-1$ agents. The computations in this case are similar to what we presented, but more complex. Assuming the likelihood ratio threshold for the k th agent choosing H^+ is e^{θ} (as in (C.1) for the calculation above), marginalizing over the decision states of the previous $k-1$ agents shows that the jump in the $(k+1)$ st agent's belief can far exceed θ . 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 $\mathcal{U}^{(i)} = \emptyset$ for all i . This simplifies the analysis, as no agent needs to

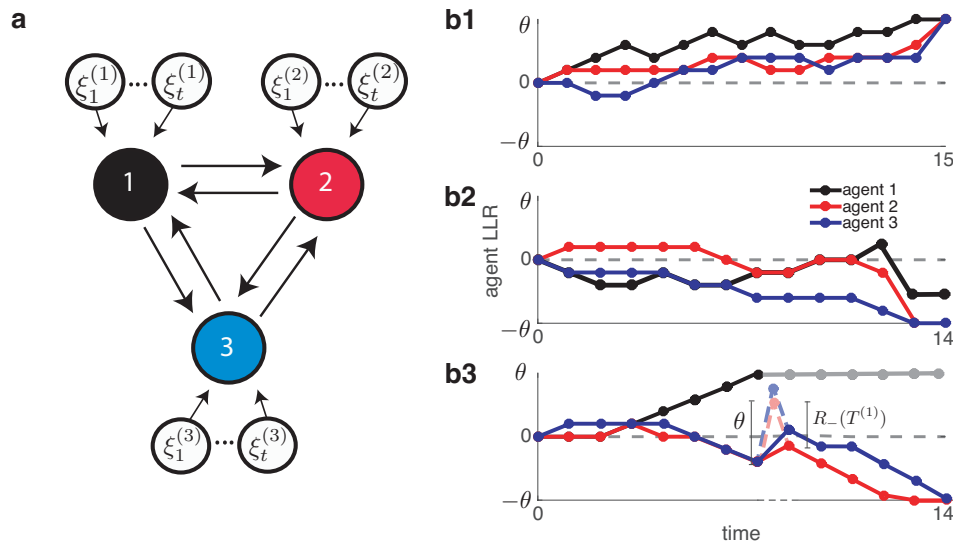


Figure 8. (a) In a three-agent clique all three agents make independent observations and observe each other's decisions. (b) Three main possibilities follow a decision: (b1) If the beliefs of the undecided agents are both positive, both follow the decision of the decider (agent 1 here). (b2) If the decision 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 (agent 1 here). (b3) If neither of the remaining agents decides, they observe each other's indecision and update their belief accordingly. 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. 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. Here, $H = H^+$.

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. There is a small probability that private information leads to a concurrent decision by multiple agents, and we consider this case at the end of the section. Observing a decision by an agent may or may not drive the undecided agents to a decision. Both the presence and absence of a decision by either remaining agent reveals further social information to its undecided counterpart. We will see that, akin to the equilibration process, there can be a number of steps of social evidence exchange. Again the remaining agents gather all social information before the next private measurements.

For concreteness we assume in the following without loss of generality that agent 1 makes a decision before agents 2 and 3. Agent 1's decision leads to a jump in the beliefs of agents 2 and 3 approximately equal to the threshold corresponding to the decision. An argument equivalent to the one presented in the previous section shows that $d_{T(1)}^1 = \pm 1$ leads to a jump approximately equal to $\pm\theta$ in the remaining agents' beliefs. As noted in our discussion following Proposition 4.1 we omit the correction $\epsilon_{T(1)}^+$ to this social information.

There are three possible outcomes: if $d_{T(1),0}^{(1)} = 1$ and $y_{T(1),0}^{(i)} \geq 0$ for $i = 2, 3$, then both remaining agents decide immediately, $d_{T(1),1}^{(i)} = 1$, and the evidence accumulation process stops.

We therefore only examine cases where (i) $y_{T^{(1)},0}^{(i)} < 0$ for $i = 2$ or $i = 3$ (but not both), and (ii) $y_{T^{(1)},0}^{(i)} < 0$ for both $i = 2, 3$.

Before observing agent 1's decision, and after the private observation at time $T^{(1)}$, the beliefs of agents $i = 2, 3$ are $y_{T^{(1)},0}^{(i)} = \text{Priv}_{1:T^{(1)}}^{(i)}$. After observing $d_{T^{(1)},0}^{(1)} = 1$ their beliefs are

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

Any agent who remains undecided after the update given by (7.1) updates their belief iteratively in response to the decision information of their neighbor.

Case (i): One agent undecided. Without loss of generality, we assume $y_{T^{(1)},0}^{(2)} \geq 0$ so that $d_{T^{(1)},1}^{(2)} = 1$. After observing agent 1's decision, the belief of agent 3 is $y_{T^{(1)},3}^{(3)} = \text{Priv}_{1:T^{(1)}}^{(3)} + \theta$. A straightforward computation following agent 2's decision gives

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

where

$$R_+(t) := \text{LLR}(y_t^{(2)} \in [0, \theta] | y_t^{(2)} \in (-\theta, \theta)) = \text{LLR}(y_t^{(2)} \in [0, \theta))$$

since agent 2's belief must have been nonnegative before observing agent 3's decision. Note that $R_+(t) < \theta$ by the following proposition.

Proposition 7.1. *Let $-\theta < a \leq b < \theta$. If agent j has not made a decision and has accumulated only private information by time T , then*

$$\left| \text{LLR}(a \leq y_T^{(j)} \leq b | y_T^{(j)} \in (-\theta, \theta)) \right| = \left| \text{LLR}(a \leq y_T^{(j)} \leq b) \right| < \theta.$$

To prove this, we assume an agent's belief lies in some subinterval (a, b) of $\Theta = (-\theta, \theta)$ and bound the social information obtained from this knowledge, concluding it provides at most an increment of social evidence of amplitude θ . The proof is provided in Appendix B.

Proposition 7.1 implies $R_+(t) < \theta$, but this increment in belief may be sufficient to lead to a decision by agent 3. 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 $S_{\pm}(t)$ given in (4.2).

Case (ii): Two agents undecided. If $y_{T^{(1)},0}^{(i)} < 0$ for $i = 2, 3$, then both agents remain undecided upon observing agent 3's decision. After each observes this absence of a decision in its counterpart, it follows from the computations that lead to (7.2) that each updates its belief as

$$(7.3) \quad y_{T^{(1)},2}^{(i)} \approx \text{Priv}_{1:T^{(1)}}^{(i)} + \theta + R_-(T^{(1)}),$$

where $R_-(t) \equiv \text{LLR}(y_t^{(-i)} \in (-\theta, 0])$. 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 7.1. Therefore (7.3) shows that this second belief update cannot lead to a decision, and $d_{T^{(1)},2}^{(i)} = 0$ for $i = 2, 3$.

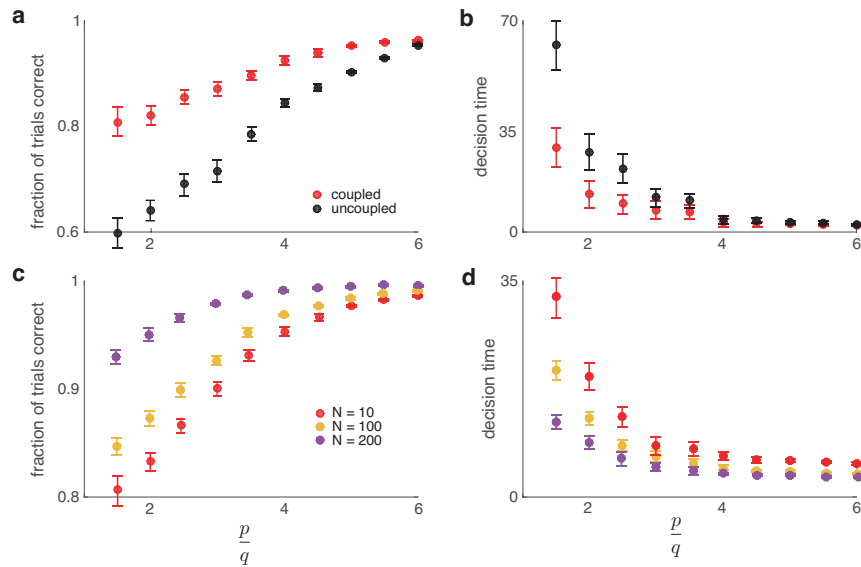


Figure 9. (a) In a three-agent clique the probability that all three agents make the correct choice is larger when social information is exchanged. (b) The average time it takes all three agents to decide is smaller when the agents communicate, since social information pushes agents closer to the threshold. 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.

After this second step, the agents cannot obtain further social information and proceed with their private measurements.

At the end of this exchange, both remaining observers know that the belief of the other is in the nonsymmetric interval, $y_{T(1),2}^{(-i)} \in (R_-(T^{(1)}), \theta + R_-(T^{(1)}))$. Therefore, future non-decisions become informative, and equilibration follows each private observation, as in the case of asymmetric decision thresholds discussed in section 5. One key difference is that in the present case the undecided agents have been influenced by the decider, and the asymmetry depends on the positions of the agents' LLR at the time of the decision.

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 (Figure 9a), and decreases the time to a decision, both of a single agent and all agents in the clique (see Figure 9b). This is particularly pronounced when the SNR of private measurements is low. With highly informative private observations, social information becomes less important.

8. Larger cliques. We next 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 . 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 analyzed as in the three-agent clique, and it often leads to all agents deciding subsequently. 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

$$y_{T^{(1)},1}^{(i)} = y_{T^{(1)},0}^{(j)} + \text{LLR}(d_{T^{(1)},0}^{(1)} = 1) \approx y_{T^{(1)},0}^{(i)} + \theta, \quad i \neq 1.$$

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*.

The decisions of agreeing agents 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 an undecided agent is $R_-(T^{(1)}) = -R_+(T^{(1)})$, where $R_+(T^{(1)})$ is given in (C.5). The equality follows from our assumption of symmetry, as we now show.

Proposition 8.1. *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 are symmetric. Let*

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

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

Proof. Following from the argument in Proposition 7.1, 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 $V_T^+ = \{\xi_{1:T}^{(i)} \in \Xi^T : \text{Priv}_{1:s}^{(i)} \in (-\theta, \theta) \text{ for } s \leq T, \text{ Priv}_{1:T}^{(i)} \in [0, \infty)\}$. By symmetry, we know that for any $\xi_{1:t}^{(i)} \in V_T^+$, there exists $-\xi_{1:t}^{(i)} \in V_T^-$, where $V_T^- = \{\xi_{1:T}^{(i)} \in \Xi^T : \text{Priv}_{1:s}^{(i)} \in (-\theta, \theta) \text{ for } s \leq T, \text{ Priv}_{1:T}^{(i)} \in (-\infty, 0]\}$, 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,

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

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 $T^{(1)}$. To compute the evidence update for undecided agent j after the first wave, note that

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

for any pair of agents $i \neq j$ different from 1. Therefore, after observing this first wave of decision, all remaining undecided agents update their belief as

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

By conditional independence, this simplifies to

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

where R_{\pm} are defined as in (7.3) and (C.5). By symmetry, $R_-(T^{(1)}) = -R_+(T^{(1)})$. Thus, all undecided agents update their belief as

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

Note that each undecided agent observes all a_1 agreeing agents and all other $u_1 - 1$ undecided agents.

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

$$R(t, (a, b)) := \text{LLR}(y_t^{(j)} \in (a, b) | y_t^{(j)} \in (-\theta, \theta)) = \text{LLR}(y_t^{(j)} \in (a, b)).$$

Note that $|R(t, (a, b))| \leq \theta$ by Proposition 7.1 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 indecision of the other $u_1 - 1$ undecided agents, excluding itself. Equation (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 , and satisfies $(a_1 - (u_1 - 1)) R_+(T^{(1)}) \geq \theta$, all the remaining undecided agents, $j \in U_1$, go along with the decision of the first agent and choose H^+ . 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 agent's choice, then, after observing each other's absence of a decision, they all agree on the opposite. 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 . We 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 beliefs again to

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

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 (8.2) give a refinement of the social information from the originally undecided agents in U_1 . As a result 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 it involves observed choices and occurs across the network. However, the process depends on the private evidence accumulated by the undecided agents and thus depends on realization.

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 we 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 beliefs to

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

If some agents decide as a result, the process continues as discussed in possibility 3.

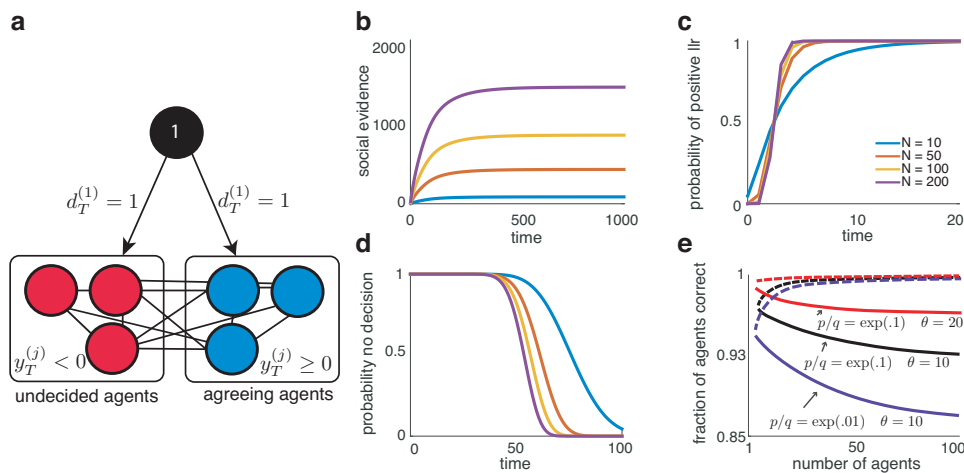


Figure 10. (a) In a clique of size N , agent 1 makes a decision which is observed by all other agents. Agents whose beliefs are positive (agreeing agents) follow the decision of this initial agent. The remaining undecided agents continue to exchange social information. (b) Average social information, $E[(a_1 - (u_1 - 1))R_+]$, available to undecided agents after observing the first wave of agreeing agents conditioned on the decision time. (c) The probability that the majority of the group has positive belief conditioned on a 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 conditioned on the 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 .

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 (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. In Figure 10 we plot information about the behavior of cliques of various sizes, hinting at trends that emerge as $N \rightarrow \infty$. 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 Figure 10b). As clique size grows, the first decision occurs earlier (Figures 10c–d). In particular, as N grows, the first decision time approaches $\gamma \equiv \theta / \log(p/q)$ —the minimum number of steps required to make a decision. Indeed, as N grows, there is a higher probability that multiple agents make an initial decision simultaneously. 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 an agent's 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 Figure 10c).

However, the social information exchange described above can still lead the remainder of the clique to overcome the impact of the wrong initial decision. To see this, note that as $N \rightarrow \infty$, a fraction p^γ of the agents will choose H^+ and another fraction q^γ of agents will choose H^- simultaneously at time γ . Specifically, on average $p^\gamma N$ agents choose H^+ at $t = \gamma$ and $q^\gamma N$ agents choose H^- at $t = \gamma$. As all belief trajectories are independent until this time, all undecided agents observing these decisions will increase their belief by $\Gamma \approx (p^\gamma - q^\gamma)N\theta$, with the approximation becoming exact in the limit of large N . Since $p^\gamma > q^\gamma$, for sufficiently large N , with high probability $\Gamma > 2\theta$, implying that all undecided agents at time γ choose H^+ . All in all, $q^\gamma N$ agents choose incorrectly while the rest of the clique chooses correctly.

9. Discussion. Our evidence accumulation model shows how to best combine streams of private and social evidence to make decisions. There has been extensive work on the mathematics of evidence accumulation by individual agents [7, 9, 51, 63, 66]. The resulting models have been very successful in explaining experimental data and the dynamics of decision making [26, 62], and there is evidence that they can explain the formation of social decisions [34]. In parallel, economists have developed mathematical models of networks of rational decision makers [27, 43]. However, economists predominantly consider static models as 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 [37] (although see [3, 28, 34]). Our model thus provides a bridge between approaches and ideas in several different fields.

The belief dynamics of agents we considered differs from past accounts in several respects. First, we assume agents only communicate decision information. 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. This can lead to complex computations, in contrast to the model discussed in [33], where a central agent sees all actions and optimizes a final decision assuming all other agents accumulate evidence independently.

The beliefs of evidence-accumulators in a network have been modeled by diffusively coupled drift-diffusion equations [49, 58]. Such models approximate continuous information-sharing between agents [58]. 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 provides insight into decision making when agents over- or underweight available evidence [6].

The absence of choice is informative only in the case of asymmetries. This could be due to one choice being more rewarding than another [29, 44, 64]. 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 [26, 56]. 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 [29, 44]. Recordings suggest neural activity representing decisions is dynamically modulated during the evidence accumulation process to reflect this bias [29, 48]. 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 [60].

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.

We assumed that all agents implement the same decision thresholds θ_{\pm} and that they all know this. This assumption can be relaxed in different ways: agents could use different decision thresholds $\theta_{\pm}^{1:N}$, which may or may not be known to the other agents in the network. When thresholds are known, much of our analysis carries over, with agents updating their belief in response to observing a decision state depending on who they observe. This increases the notational complexity, but the opinion exchange dynamics remain largely the same. When thresholds are not known, agents can also update their belief about, i.e., the posterior over, the possible decision thresholds of their neighbors over multiple decisions. Indeed, finding out which of your friends are rash to make decisions and which ones tend to be cautious is essential to properly weighing their opinions. However, this requires observations and interactions over multiple decisions—something we do not consider here. In addition, there is evidence that humans tend to treat others as if they make decisions in similar ways [3]. How humans learn each other's decision thresholds and how such learning impacts decision making has not been extensively explored and provides fruitful avenues for future work.

We have assumed that each observer maintains a constant, predetermined threshold while accumulating evidence. In many situations this is likely not the best policy. We could extend our modeling framework to thresholds that can change in response to accumulated private and social evidence as well as time pressure. Such approaches have been developed to describe single ideal observers accumulating evidence in uncertain environments and rely on dynamic programming to define the agent's decision policy [16]. This can result in decision thresholds that vary in time and that can be realization dependent [17]. In a social network, to optimize their own threshold, every agent needs to model the process by which other agents choose their threshold. This can lead to the need for recursive reasoning to set the thresholds appropriately. In practice one can use the setting of partially observed Markov processes to implement such computations [32].

Our Bayesian model is unlikely to accurately describe the decision making process of biological agents. The rationality of humans and other animals is bounded [57], 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 [12]. A

normative model allows us to identify which computations are important and which offer only fractional benefits.

Our analysis relied on the assumption that private observations received by the individual agents are independent. When this is not the case agents need to account for correlations between private observations in determining the value of social information. In general, we expect that an increased redundancy of neighboring agents' observations reduces the impact of social information. However, experimental evidence suggests that humans do not take into account such correlations when making decisions [19, 36] and behave as if the evidence gathered by different members of their social group is independent. Indeed, it can be shown that properly integrating social information from unseen agents can be NP-hard, although the computations can be tractable in some networks [30, 31]. Our model may approximate the decision making process of agents that make such assumptions. The impact of ignoring such correlations when they are present is an interesting avenue for future work [59].

Our findings are distinct from 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, 38]. In recurrent networks, agents can announce their preference for a choice, until they reach agreement with their neighbors [2, 24, 25, 41]. This framework can be simplified so that agents make binary decisions, based solely on their neighbors' opinions, admitting asymptotic analyses of the cascade of decisions through networks with complex structure [67]. 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 case 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 with simpler assumptions on the decision process.

Appendix A. Bounding social information from a decision in a recurrent two-agent network. Here we prove Proposition 5.2 using an argument similar to that in Proposition 4.1. First, consider the case in which $n = 0$, so the private observation $\xi_{T(i)}^{(i)}$ triggers the decision. If $d_{T,0}^{(i)} = +1$, we know

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

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

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

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

$$\theta_+ - \text{Soc}_{T-1,0} \leq \text{LLR}(d_{T,0}^{(j)} = 0, d_{T-1, N_{T-1}}^{(i)} = 0, d_{T,0}^{(i)} = 1) < \theta_+ - \text{Soc}_{T-1,0} + \epsilon_{T,0}^+,$$

where N_{T-1} is the maximal substep in the social information exchange following the private observation at time step $T - 1$.

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

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

which implies

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

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

$$\theta_+ - \text{Soc}_{T,n} \leq \text{LLR}(d_{T(i),n}^{(j)} = 0, d_{T(i),n-1}^{(i)} = 0, d_{T(i),n}^{(i)} = 1) < \theta_+ - \text{Soc}_{T,n} + \epsilon_{T,n}^+,$$

where

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

Following the arguments in Theorem 5.1, we note then that

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

The proof follows similarly for the case $d_{T,n}^{(i)} = -1$.

Appendix B. Proof of Proposition 7.1. To simplify 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. 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

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

which can be rearranged as

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

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

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

so that by noting $P(y_T^{(j)} \in [a, b] | H^{\pm}) = \sum_{V_T(a,b)} \prod_{t=1}^T f_{\pm}(\xi_t^{(j)})$ 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.

Appendix C. Social information from a decision by agent 2 exceeds θ for agent 3 in a three-agent unidirectional line. To show this, we note that for any $s = 1, \dots, T^{(2)} - 1$,

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

where the last approximation follows from again assuming $\text{Priv}_i^{(2)}$ is small. The situations where a decision from agent 1 does *not* immediately cause a decision from agent 2, conditioned on H^{+} , are described by

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

Using (C.1), we have

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

By (6.1), 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

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

Because the LLRs of distinct agents are independent before any agent makes a decision, we have

$$(C.3) \quad \begin{aligned} & P(d_{T^{(2)}}^{(2)} = 1, d_{T^{(2)}-1}^{(2)} = 0, d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^\pm) \\ &= P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^\pm) P(d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^\pm) \end{aligned}$$

and

$$(C.4) \quad P(d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^+) \approx e^\theta P(d_{T^{(2)}}^{(1)} = 1, d_{T^{(2)}-1}^{(1)} = 0 | H^-) \approx \frac{1}{1 + e^{-\theta}}.$$

Equation (C.3) omits the possibility that the beliefs of agents 1 and 2 cross their respective thresholds at the same time. However, this event has small probability when private observations are not very informative, and thresholds are reasonably large, so we do not consider this case in our calculations. Hence, in (C.3), $y_T^{(2)}$ represents the belief of agent 2 after observing a private signal at time T but before observing the decision of agent 1.

The first approximation in (C.4) follows from the fact that agent 1's decision H^+ triggers a jump in agent 2's belief approximately equal to θ . The second approximation in (C.4) is a standard result in the decision making literature and can be obtained by calculating the exit probability of a drift-diffusion process on a bounded interval through the boundary in the direction given by the drift [7]. Let

$$(C.5) \quad R_+(t) := \text{LLR}(y_t^{(2)} \in [0, \theta] | y_t^{(2)} \in (-\theta, \theta)) = \text{LLR}(y_t^{(2)} \in [0, \theta]) \geq 0.$$

Note that we can increase $R_+(t)$ by increasing θ without changing the measurement distributions: larger thresholds mean that more time is required to make a decision. In turn, longer decision times mean that the belief of an agent is more likely to lie in the half interval $[0, \theta]$ when $H = H^+$, resulting in a larger $\text{LLR}(y_t^{(2)} \in [0, \theta])$.

Applying (C.3) to (C.2) and suppressing the notation $d_{T^{(2)}-1}^{(1)} = 0$, we obtain

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

The two approximations in (C.4) and (C.5) then yield

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

It follows that

$$(C.6) \quad \text{Soc}_{T^{(2)}}^{(3)} \approx \log \left(\frac{P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^+) + \zeta_{T^{(2)}}^+(1 + e^{-\theta})}{P(y_{T^{(2)}}^{(2)} \in [0, \theta] | H^+)e^{-R_+(T^{(2)})} + \zeta_{T^{(2)}}^+(1 + e^{-\theta})} \right) + \theta \geq \theta.$$

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 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 $\zeta_{T^{(2)}}^-$ and $R_-(t)$ are defined equivalently to $\zeta_{T^{(2)}}^+$ and $R_+(t)$.

REFERENCES

- [1] D. ACEMOGLU, M. A. DAHLEH, I. LOBEL, AND A. OZDAGLAR, *Bayesian learning in social networks*, Rev. Econom. Stud., 78 (2011), pp. 1201–1236.
- [2] R. J. AUMANN, *Agreeing to disagree*, Ann. Statist., 4 (1976), pp. 1236–1239.
- [3] B. BAHRAMI, 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*, Quart. J. Econom., 107 (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*, Psychol. Rev., 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 Neurosci., 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*, Psychol. Rev., 100 (1993), pp. 432–459.
- [10] R. J. CAGINALP AND B. DOIRON, *Decision dynamics in groups with interacting members*, SIAM J. Appl. Dynam. Syst., 16 (2017), pp. 1543–1562, <https://doi.org/10.1137/16M1090922>.
- [11] L. CONRADT AND T. J. ROPER, *Consensus decision making in animals*, Trends Ecol. Evol., 20 (2005), pp. 449–456.
- [12] I. D. COUZIN, *Collective cognition in animal groups*, Trends Cognit. Sci., 13 (2009), pp. 36–43.
- [13] 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.
- [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*, John Wiley & Sons, 2004.
- [16] J. DRUGOWITSCH, R. MORENO-BOTE, A. K. CHURCHLAND, M. N. SHADLEN, AND A. POUGET, *The cost of accumulating evidence in perceptual decision making*, J. Neurosci., 32 (2012), pp. 3612–3628.
- [17] J. DRUGOWITSCH, R. MORENO-BOTE, AND A. POUGET, *Optimal decision-making with time-varying evidence reliability*, in Advances in Neural Information Processing Systems, 2014, pp. 748–756.
- [18] W. EDWARDS, *The theory of decision making*, Psychol. Bull., 51 (1954), pp. 380–417.

- [19] B. ENKE AND F. ZIMMERMANN, *Correlation neglect in belief formation*, Rev. Econom. Stud., 86 (2019), pp. 313–332.
- [20] R. FAGIN, J. Y. HALPERN, Y. MOSES, AND M. VARDI, *Reasoning about Knowledge*, MIT Press, 2004.
- [21] N. R. FRANKS AND T. RICHARDSON, *Teaching in tandem-running ants*, Nature, 439 (2006), pp. 153–153.
- [22] C. D. FRITH AND U. FRITH, *Mechanisms of social cognition*, Ann. Rev. Psychol., 63 (2012), pp. 287–313.
- [23] D. GALE AND S. KARIV, *Bayesian learning in social networks*, Games Econom. Behav., 45 (2003), pp. 329–346.
- [24] J. GEANAKOPOLOS, *Common Knowledge*, in Handbook of Game Theory with Economic Applications, Vol. 2, Elsevier, 1994, pp. 1437–1496.
- [25] J. D. GEANAKOPOLOS AND H. M. POLEMARCHAKIS, *We can't disagree forever*, J. Econom. Theory, 28 (1982), pp. 192–200.
- [26] J. I. GOLD AND M. N. SHADLEN, *The neural basis of decision making*, Annu. Rev. Neurosci., 30 (2007), pp. 535–574.
- [27] S. GOYAL, *Connections: An Introduction to the Economics of Networks*, Princeton University Press, 2012.
- [28] S. P. HALLER, D. BANG, B. BAHRAMI, AND J. Y. LAU, *Group Decision-Making Is Optimal in Adolescence*, bioRxiv preprint, <https://doi.org/10.1101/412726>, 2018.
- [29] 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*, J. Neurosci., 31 (2011), pp. 6339–6352.
- [30] J. HAZLA, A. JADBABAIE, E. MOSSEL, AND M. A. RAHIMIAN, *Bayesian Decision Making in Groups Is Hard*, preprint, <https://arxiv.org/abs/1705.04770>, 2017.
- [31] J. HAZLA, A. JADBABAIE, E. MOSSEL, AND M. A. RAHIMIAN, *Reasoning in Bayesian opinion exchange networks is PSPACE-hard*, Proc. Mach. Learn. Res., 99 (2019), pp. 1614–1648.
- [32] K. KHALVATI, S. A. PARK, S. MIRBAGHERI, R. PHILIPPE, M. SESTITO, J.-C. DREHER, AND R. P. RAO, *Modeling other minds: Bayesian inference explains human choices in group decision-making*, Sci. Adv., 5 (2019), eaax8783.
- [33] M. KIMURA AND J. MOEHLIS, *Group decision-making models for sequential tasks*, SIAM Rev., 54 (2012), pp. 121–138, <https://doi.org/10.1137/10079700X>.
- [34] I. KRAJBICH, T. HARE, B. BARTLING, Y. MORISHIMA, AND E. FEHR, *A common mechanism underlying food choice and social decisions*, PLoS Comput. Biol., 11 (2015), e1004371.
- [35] 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, 323 (2009), pp. 721–723.
- [36] G. LEVY AND R. RAZIN, *Correlation neglect, voting behavior, and information aggregation*, Amer. Econom. Rev., 105 (2015), pp. 1634–1645.
- [37] R. D. LUCE, *Response Times: Their Role in Inferring Elementary Mental Organization*, Oxford University Press, 1986.
- [38] P. MILGROM AND N. STOKEY, *Information, trade and common knowledge*, J. Econom. Theory, 26 (1982), pp. 17–27.
- [39] N. MILLER, S. GARNIER, A. T. HARTNETT, AND I. D. COUZIN, *Both information and social cohesion determine collective decisions in animal groups*, Proc. Natl. Acad. Sci. USA, 110 (2013), pp. 5263–5268.
- [40] E. MOSSEL, A. SLY, AND O. TAMUZ, *Asymptotic learning on Bayesian social networks*, Probab. Theory Related Fields, 158 (2014), pp. 127–157.
- [41] E. MOSSEL AND O. TAMUZ, *Opinion exchange dynamics*, Probab. Surv., 14 (2017), pp. 155–204.
- [42] 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), e10047.
- [43] M. MUELLER-FRANK, *A general framework for rational learning in social networks*, Theoret. Econom., 8 (2013), pp. 1–40.
- [44] 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*, J. Neurosci., 32 (2012), pp. 2335–2343.
- [45] M. NAGY, Z. AKOS, D. BIRO, AND T. VICSEK, *Hierarchical group dynamics in pigeon flocks*, Nature, 464 (2010), pp. 890–893.
- [46] M. E. J. NEWMAN, *The structure and function of complex networks*, SIAM Rev., 45 (2003), pp. 167–256.

- <https://doi.org/10.1137/S003614450342480>.
- [47] 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 Biol., 15 (2017), e2001958.
 - [48] M. L. PLATT AND P. W. GLIMCHER, *Neural correlates of decision variables in parietal cortex*, Nature, 400 (1999), pp. 233–238.
 - [49] I. POULAKAKIS, L. SCARDOVI, AND N. E. LEONARD, *Node Classification in Networks of Stochastic Evidence Accumulators*, preprint, <https://arxiv.org/abs/1210.4235>, 2012.
 - [50] R. M. RAAFAT, N. CHATER, AND C. FRITH, *Herding in humans*, Trends Cognit. Sci., 13 (2009), pp. 420–428.
 - [51] A. E. RADILLO, A. VELIZ-CUBA, K. JOSIĆ, AND Z. P. KILPATRICK, *Evidence accumulation and change rate inference in dynamic environments*, Neural Comput., 29 (2017), pp. 1561–1610.
 - [52] R. RATCLIFF AND G. MCKOON, *The diffusion decision model: Theory and data for two-choice decision tasks*, Neural Comput., 20 (2008), pp. 873–922.
 - [53] R. RATCLIFF AND J. N. ROUDER, *Modeling response times for two-choice decisions*, Psychol. Sci., 9 (1998), pp. 347–356.
 - [54] T. SAATY, *How to make a decision: The analytic hierarchy process*, European J. Oper. Res., 48 (1990), pp. 9–26.
 - [55] T. D. SEELEY, S. CAMAZINE, AND J. SNEYD, *Collective decision-making in honey bees: How colonies choose among nectar sources*, Behav. Ecol. Sociobiol., 28 (1991), pp. 277–290.
 - [56] M. N. SHADLEN AND W. T. NEWSOME, *Motion perception: Seeing and deciding*, Proc. Natl. Acad. Sci. USA, 93 (1996), pp. 628–633.
 - [57] H. A. SIMON, *Bounded rationality*, in *Utility and Probability*, Springer, 1990, pp. 15–18.
 - [58] V. SRIVASTAVA AND N. E. LEONARD, *Collective decision-making in ideal networks: The speed-accuracy tradeoff*, IEEE Trans. Control Network Syst., 1 (2014), pp. 121–132.
 - [59] S. STOLARCZYK, M. BHARDWAJ, K. E. BASSLER, W. J. MA, AND K. JOSIĆ, *Loss of information in feedforward social networks*, J. Complex Networks, 6 (2017), pp. 448–469.
 - [60] C. SUMMERFIELD AND F. P. DE LANGE, *Expectation in perceptual decision making: Neural and computational mechanisms*, Nature Rev. Neurosci., 15 (2014), pp. 745–756.
 - [61] H. M. TAYLOR AND S. KARLIN, *An Introduction to Stochastic Modeling*, Academic Press, 2014.
 - [62] M. USHER AND J. L. MCCLELLAND, *The time course of perceptual choice: The leaky, competing accumulator model*, Psychol. Rev., 108 (2001), pp. 550–592.
 - [63] A. VELIZ-CUBA, Z. P. KILPATRICK, AND K. JOSIĆ, *Stochastic models of evidence accumulation in changing environments*, SIAM Rev., 58 (2016), pp. 264–289, <https://doi.org/10.1137/15M1028443>.
 - [64] A. VOSS, K. ROTHERMUND, AND J. VOSS, *Interpreting the parameters of the diffusion model: An empirical validation*, Memory Cognit., 32 (2004), pp. 1206–1220.
 - [65] A. WALD, *Sequential tests of statistical hypotheses*, Ann. Math. Statist., 16 (1945), pp. 117–186.
 - [66] A. WALD AND J. WOLFOWITZ, *Optimum character of the sequential probability ratio test*, Ann. Math. Statist., 19 (1948), pp. 326–339.
 - [67] D. J. WATTS, *A simple model of global cascades on random networks*, Proc. Natl. Acad. Sci. USA, 99 (2002), pp. 5766–5771.
 - [68] G. YAN, T. ZHOU, B. HU, Z.-Q. FU, AND B.-H. WANG, *Efficient routing on complex networks*, Phys. Rev. E, 73 (2006), 046108.