

Heterogeneity Improves Speed and Accuracy in Social Networks

Bhargav Karamched^{1,2,*} MeStickler¹ ,



make private observations and gather social evidence by observing the choices of all other agents. They do not share private information but know the statistics of the observations each agent makes. A decision cannot be undone.

An example provides intuition: consider a group of people deciding between two products to buy. They study the products' specifications and read reviews, making a sequence of private observations. They also observe which product their friends choose. Each person combines private observations (product reviews) with social information (decisions of friends). They do not exchange information directly but know the type of information their friends gather, and thus how beliefs evolve [6,16]. Once purchased, the product cannot be returned.

Evolution of beliefs: We assume N agents accumulate noisy private observations and optimally combine them with information obtained from observing the decisions of their neighbors to choose between two hypotheses, H^+ or H^- . Either hypothesis is a priori equally likely to be correct. Each agent, i , makes decisions based on their belief, $y_i(t)$, which equals the log-likelihood ratio (LLR) between the hypotheses given all available evidence [21]. After a sequence of private observations, $\xi_i^{0:t}$, the belief is $y_i(t) = \frac{1}{\alpha} \log \frac{\mathbb{P}(\delta | H^+, \xi_i^{0:t})}{\mathbb{P}(\delta | H^-, \xi_i^{0:t})}$. If private observations are rapid and uncorrelated in time and between agents, beliefs evolve as

$$dy_i = \frac{1}{\alpha} dt \left[\alpha - \frac{\alpha}{2} dW_i \right] \quad (1)$$

where the sign of the drift equals that of the correct hypotheses, and $W_i(t)$ are independent, standard Wiener processes [22,23]. Each observer starts with no evidence, so $y_i(0) = 0$. We assume henceforth that H^+ is correct, and that $\alpha > 0$. When H^- is correct or $\alpha < 0$ the analysis is similar.

Each agent, i , sets a threshold, θ_i , and chooses H^+ (H^-) at time T_i if $y_i(T_i) \geq \theta_i$ ($y_i(T_i) \leq -\theta_i$), and $y_i(t) \in (-\theta_i, \theta_i)$ for $0 \leq t < T_i$. All other agents observe a decider's choice, but may not know their threshold. We consider omniscient agents who know each other's thresholds and the case of consensus bias where each agent assumes all others have the same threshold they do.

Belief updates from decision: Without loss of generality, we assume the belief of agent i is the first to reach threshold at time $t = T$ [Fig. 1(a)].

Until this decision, beliefs of all agents, $y_i(t)$ with $i = 1, \dots, N$, evolve independently according to Eq. (1). Upon observing the first decision, omniscient agents update

can leverage quick, unreliable decisions to improve the response of the population.

Dichotomous threshold distribution: The case of agents with either a high or a low threshold is tractable and sheds light on more general examples. Before a decision the belief of each agent evolves according to Eq. (1) with absorbing boundaries at $-\theta_i < \dots < \theta_i$. We assume that γN agents share threshold θ_{\min} and $\delta - \gamma N$ share threshold θ_{\max} for $-\theta_{\min} < \theta_{\max}$ and $\gamma \in \delta ; \rho$. The first decision is then likely made by an agent with a low threshold, and is thus fast but unreliable ([24], Sec. XII). We use the approximation $\mathbb{E}[\tau] \approx \theta_{\min} = \ln \delta \gamma N \rho$ which breaks down when $\gamma \ll \dots$, but works well otherwise [Fig. 2(c)].

A clique under consensus bias is homogeneous from an observer's perspective and thus behaves like a homogeneous population. Indeed, the expected size of the first

