



This is a repository copy of *Asynchrony rescues statistically optimal group decisions from information cascades through emergent leaders*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/197508/>

Version: Published Version

Article:

Reina, A. orcid.org/0000-0003-4745-992X, Bose, T. orcid.org/0000-0003-0014-4090, Srivastava, V. orcid.org/0000-0002-9786-0159 et al. (1 more author) (2023) Asynchrony rescues statistically optimal group decisions from information cascades through emergent leaders. Royal Society Open Science, 10 (3). 230175. ISSN 2054-5703

<https://doi.org/10.1098/rsos.230175>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Research



Cite this article: Reina A, Bose T, Srivastava V, Marshall JAR. 2023 Asynchrony rescues statistically optimal group decisions from information cascades through emergent leaders. *R. Soc. Open Sci.* **10**: 230175. <https://doi.org/10.1098/rsos.230175>

Received: 13 February 2023

Accepted: 21 February 2023

Subject Category:

Mathematics

Subject Areas:

behaviour/cognition/computational biology

Keywords:

collective decision-making, Bayesian brain, information cascades, emergent leaders

Author for correspondence:

Andreagiovanni Reina

e-mail: andreagiovanni.reina@ulb.be

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.6456121>.

Asynchrony rescues statistically optimal group decisions from information cascades through emergent leaders

Andreagiovanni Reina^{1,2}, Thomas Bose²,
Vaibhav Srivastava³ and James A. R. Marshall^{2,4}

¹Institute for Interdisciplinary Studies on Artificial Intelligence (IRIDIA), Université Libre de Bruxelles, Brussels 1050, Belgium

²Department of Computer Science, University of Sheffield, Sheffield, S1 4DP, UK

³Department of Electrical and Computer Engineering, Michigan State University, East Lansing, MI 48824-1226, USA

⁴Opteran Technologies Limited, Sheffield, UK

AR, 0000-0003-4745-992X; TB, 0000-0003-0014-4090; VS, 0000-0002-9786-0159; JARM, 0000-0002-1506-167X

It is usually assumed that information cascades are most likely to occur when an early but incorrect opinion spreads through the group. Here, we analyse models of confidence-sharing in groups and reveal the opposite result: simple but plausible models of naive-Bayesian decision-making exhibit information cascades when group decisions are synchronous; however, when group decisions are asynchronous, the early decisions reached by Bayesian decision-makers tend to be correct and dominate the group consensus dynamics. Thus early decisions actually rescue the group from making errors, rather than contribute to it. We explore the likely realism of our assumed decision-making rule with reference to the evolution of mechanisms for aggregating social information, and known psychological and neuroscientific mechanisms.

1. Introduction

Information cascades, where individuals follow others' decisions regardless of self-sourced evidence, are usually assumed to occur in asynchronous decision-making, in which early decisions tend to be incorrect and dominate the decision dynamics so that the group decision is incorrect. Previous work assumed cascades to happen only when the first responding individual exerts disproportionate influence on other group members [1–4]. The converse assumption

would be that synchronous group decision-making mechanisms offer the best protection from information cascades.

Here, we explore the optimal pooling of information in synchronous and asynchronous group decision-making mechanisms. A standard assumption in behavioural ecology, psychology and neuroscience is that individuals apply optimal probabilistic computational rules where possible (e.g. [5–10]). If optimal computation is infeasible, it is argued that rules that approximate optimal computations in typically encountered scenarios will be used. Similarly, in behavioural ecology and psychology, research focuses on the optimal pooling of information by groups (e.g. [11,12]).

In evolutionary terms, the neurological mechanisms to process asocial environmental information must have developed earlier than sociality appeared. Thus, we assume that when group living began, evolution led to the adaptation of pre-existing Bayesian heuristics [7] to also process social information. We study the implications of such assumptions on group decisions in two scenarios: collective detection of an instantaneous signal with synchronous interaction among individuals, and continuous environmental sampling with asynchronous interaction. Our analysis shows that for the synchronous case, in which there are no early decisions, decision-making is unstable and negative information cascades are observed. In the asynchronous case, however, early decisions tend to be correct and lead to positive information cascades. This observation is the opposite to the usual assumption that early decisions are erroneous and lead to negative information cascades, showing how group leaders can spontaneously emerge for the benefit of collective decisions.

1.1. Problem formulation

We study the problem of a single-shot collective decision in which N individuals pool information to make a decision on the correct state of the world S . We assume the choice is binary, i.e. there are two possible states of the world $S \in \{S^+, S^-\}$. We assume that each state of the world has a prior probability, $P(S^+)$ and $P(S^-)$. We assume that the cost matrix for classifications is symmetric, i.e. the cost of an error, as well as the reward for the correct classification, is the same for either state of the world. We consider two types of collective decisions: *signal detection* and *sequential sampling* (figure 1).

In signal detection, at time t_0 , individuals are exposed to a signal emitted by an instantaneous event, which they categorize as S^+ or S^- (figure 1a). Through optimal signal detection theory, individuals compare the estimated signal with a threshold as described in §1.2.1. Therefore, each individual i , at time t_0 , has an independent opinion $x_i^{t_0} \in \{S^+ = 1, S^- = -1\}$ on the true state of the world S and the relative confidence $c_i^{t_0}$ on the accuracy of its opinion (figure 1b). Every individual i repeatedly exchanges its opinion and confidence with its neighbours M_i defined by a communication network topology \mathcal{G} , which can be static or time-varying. Individuals, at each synchronous social interaction, update their opinion and confidence in order to determine the correct state S (figure 1c).

In sequential sampling, each agent integrates evidence from the environment over time in order to correctly classify the state of the world. As in neuroscientific studies [16], the statistically optimal process of evidence integration is represented as a *drift diffusion model* (DDM), [17,18], which describes the evolution over time of the individual i 's decision evidence $y_i(t)$ as a biased Brownian motion process that is governed by two terms: the drift A_i and the diffusion W (figure 1d). The former term models the evidence integration towards the correct decision, while the latter term models the noise in the integration process (implemented as a Wiener process with standard deviation σ equal for every agent). Individuals integrate evidence $y_i(t)$ until one of the two thresholds $z^+ > 0$ or $z^- < 0$ has been reached (i.e. we model a free-response scenario). The individual's decision corresponds to the sign of the crossed threshold, or equivalently of the integrated evidence, i.e. $x_i^t = \text{sign}(y_i(t))$. Individuals making a decision communicate it to their neighbours M_i on \mathcal{G} (figure 1e,f), who combine the received information with their accumulated evidence as described in §1.2.2.

1.2. Weighted Bayes Consensus

We formulate how naive Bayes-optimal individuals can employ the statistically optimal Bayes' rule [19] to update their opinion x_i^{t+1} and confidence c_i^{t+1} from their neighbours' opinions in the two scenarios considered, collective signal detection and collective sequential sampling. We describe updates as naive-Bayes because they neglect correlations in social information [20].

1.2.1. Collective signal detection

In signal detection, individuals form an opinion in favour of either $S^+ = 1$ or $S^- = -1$ by comparing the estimated signal with a threshold specific to each agent (figure 1a). We assume that each agent i has an

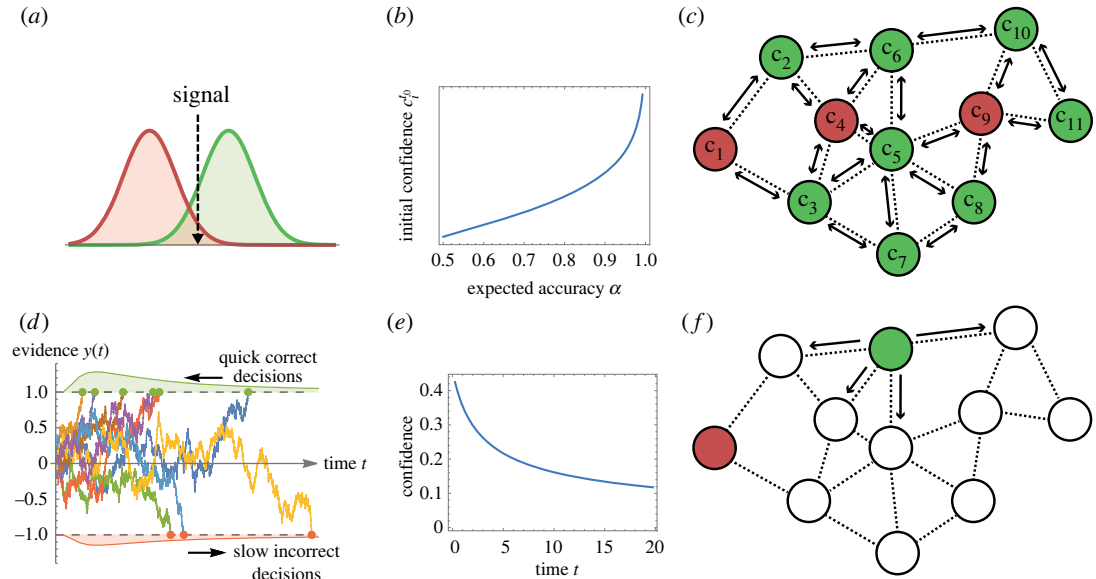


Figure 1. We consider two types of collective decisions—(a–c) signal detection and (d–f) sequential sampling—characterized, respectively, by synchronous and asynchronous social interactions. (a) An instantaneous event at time t_0 produces a signal that all individuals estimate and compare with a threshold to make a decision between the red and green alternatives (signal detection theory). (b) We assume that each agent has an estimate of its accuracy (e.g. through previous experience), with which it can estimate its confidence $c_i^{t_0}$ as the log-odds ratio, equation (1.1), Marshall *et al.* [11]. (c) Individuals synchronously exchange opinions and confidence c_i with their nearest neighbours (i.e. information spreads on a random geometric graph [13]), and, in order to reach a consensus decision, they update their opinions by locally optimal Bayesian integration of confidence-weighted votes (Weighted Bayes Consensus rule). The arrows indicate bidirectional synchronous interactions, the colours are the individuals' opinions. (d) In sequential sampling, each individual optimally integrates noisy evidence from the environment until it has enough information to make a decision. This process is modelled as a drift diffusion model (DDM). The graphics show examples of DDM trajectories for drifts sampled from a random distribution biased towards the correct decision as positive drift. The expected decision time is shorter for correct decisions (positive threshold) and longer for incorrect decisions (negative threshold), because, as indicated in [14], errors are in most cases caused by low drift-diffusion ratios which take longer, on average, to reach the decision threshold than DDMs with high drift-diffusion ratios which lead in most cases to correct decisions. (e) When the individual does not know its DDM's drift but can only estimate its expected sampling ability, its confidence (computed with equation (1.4)) is high when the accumulated evidence hits the decision threshold early (a quick decision is a proxy of higher DDM's drift-diffusion ratio, in agreement with neurological mechanisms [15]) and low when it hits the threshold late. (f) An individual (node) only communicates once it makes a decision, which it communicates to its neighbours (in the graphics, the green node with one-way communication arrows; the red node has reached its decision earlier and does not continue communicating).

estimate of its accuracy α_i in determining the true state of the world. This information can have been acquired by the individual, for example through previous experience and decisions in the same environment. In this experimental scenario, at time t_0 , each individual synchronously makes an independent estimate $x_i^{t_0}$ of the world's state. Following optimal signal detection theory [11], each agent i can also compute its confidence $c_i^{t_0}$ as the log-odds ratio

$$\text{confidence} = \ln\left(\frac{\text{accuracy}}{\text{error rate}}\right), \quad (1.1)$$

where the error rate is the complementary probability of being correct, i.e. $1 - \text{accuracy}$, thus $c_i^{t_0} = \ln(\alpha_i / (1 - \alpha_i))$, figure 1b.

After the individual decisions, every iteration $t > t_0$ the individuals share with each other their opinion x_i^t and confidence c_i^t , and use the received information to update their new opinion x_i^{t+1} and new confidence c_i^{t+1} (figure 1c). Statistically optimal individuals compute the new aggregate opinion x_i^{t+1} following optimal confidence weighting theory presented in [11,21,22] as

$$x_i^{t+1} = \text{sign}\left(x_i^t c_i^t + \sum_{j \in M_i} (x_j^t c_j^t)\right). \quad (1.2)$$

This leaves, however, the problem of how individuals update their confidence in their new opinion. We start by noting that the neighbours' confidences can be used to derive the neighbours' accuracies (or, equivalently, vice versa). Assuming that all agents update their confidence through the same computation, the inverse of the confidence computation of equation (1.1) gives the accuracy of each neighbour (see *Material and methods*). We label this update rule as Weighted Bayes Consensus, and in *Material and methods* we show it can be reduced to linear summation

$$c_i^{t+1} = \left| x_i^t c_i^t + \sum_{j \in M_i} (x_j^t c_j^t) + \pi(x_i^{t+1}) \right|, \quad (1.3)$$

where the operator $|\cdot|$ is the absolute value, and $\pi(x_i^{t+1})$ is the log of the prior ratio in favour of $S = x_i^{t+1}$, i.e. $\pi(x_i^{t+1}) = \ln(P(S^+)/P(S^-))$ for $x_i^{t+1} = +1$, and the reciprocal of the log argument for $x_i^{t+1} = -1$. Therefore, the Weighted Bayes Consensus rule is a simple linear update rule for both opinion (equation (1.2)) and confidence (equation (1.3)).

1.2.2. Collective sequential sampling

In the sequential sampling scenario, individual i makes a decision $x_i^t = \text{sign}(y_i(t))$ when the integrated evidence $y_i(t)$ reaches the threshold $z^+ > 0$ or $z^- < 0$ in favour of the positive or negative world state hypotheses, respectively, at time t (figure 1d). The thresholds are optimally set as a function of the priors $P(S^+)$ and $P(S^-)$, and the cost matrix, following [16] (for details, see text ST1 in the electronic supplementary material). Note that the thresholds are set to an equal and fixed value for all individuals, as every individual has the same knowledge at the beginning of the decision-making process. We assume that agents know the integration noise σ , the cost matrix, and the world state priors $P(S^+)$ and $P(S^-)$, which are the same for the entire population, in agreement with previous theory [7]. Individuals do not know their drift A_i —which represents the individual's accuracy in sampling the state of the world [16]—but they know the random distribution from which the drifts' magnitudes, \hat{A}_i , are sampled (assuming no systematically misinformed individuals, the sign of A_i is always equal to the correct state S). In other terms, individuals know the group accuracy distribution, but do not know the accuracy of any specific individual.

An individual j communicates its decision $x_j^t \in \{-1, +1\}$ to its neighbours M_j once—when its integrated evidence $y_j(t)$ reaches either threshold (z^+ , z^-), see figure 1f. The information about another individual reaching threshold is additional evidence that the neighbours can use during their continuous evidence integration. Therefore, an individual i that receives the neighbour's decision x_j^t at time t and has not yet made a decision (i.e. $z^- < y_i(t) < z^+$), integrates x_j^t as a 'kick' k into its evidence accumulator $y_i(t)$. The optimal size of this kick corresponds to the neighbour's confidence in its decision, which depends both on the quantity of integrated evidence and the integration time (figure 1e). In general, quick decisions are considered as an indication of high confidence (due to a high drift-diffusion ratio \hat{A}_i/σ), conversely slow decisions are likely to be influenced by high levels of noise (low \hat{A}_i/σ); see figure 1d and [14]. Note that there is no difference if the decision-maker computes its own confidence (k) and sends this information, or if every agent infers k once receives a neighbour's decision. Assuming identical thresholds in the population and simultaneous start of evidence integration, an agent receiving a neighbour's decision has information on both the integration time t (i.e. communication time) and on the integrated quantity $y_j(t) = z^{\text{sign}(x_j^t)}$ (i.e. z^+ for $x_j^t = +1$ and z^- for $x_j^t = -1$). Therefore, the optimal kick size, for example assuming $x_j^t = +1$, is

$$k = \ln \left(\frac{P(S^+ | \text{1st hit threshold is } z^+ \text{ at } t)}{P(S^- | \text{1st hit threshold is } z^+ \text{ at } t)} \right). \quad (1.4)$$

Again, applying Bayesian theory [19], we obtain

$$\begin{aligned} k &= \ln \left(\frac{P(S^+ | z^+, t)}{P(S^- | z^+, t)} \right) \\ &= \ln \left(\frac{P(t | z^+, S^+)}{P(t | z^+, S^-)} \right) + \ln \left(\frac{P(z^+ | S^+)}{P(z^+ | S^-)} \right) + \ln \left(\frac{P(S^+)}{P(S^-)} \right). \end{aligned} \quad (1.5)$$

The three terms of the r.h.s. of equation (1.5), respectively, are the log-odds of the first passage time of the DDM through z^+ at t , the log-odds of hitting z^+ before z^- , and the log-odds of the prior on the states of the world. The precise DDM parameters are unknown to the individual, thus, as proposed in [15], the

individual averages the probability for any possible DDM weighted by the prior probability of such DDM parameters to manifest. See *Material and methods* for the detailed derivation and figure 1e for a graphical illustration of equation (1.5).

More sophisticated agents could aggregate social information with more advanced computation that uses the absence of decision from neighbours as informative data [23]. Similarly, an individual could refine the computation of equation (1.5) by observing the social network on its neighbours and treating differently the case in which the neighbour makes a decision based solely on its personal information or after receiving social information [20]. Such nuanced calculations are likely to be unrealistic to be implemented in the brain, hence in our study, we assume naive individuals that neglect previous social interactions. Signals from neighbours making their decisions are treated independently, thus each neighbour is implicitly considered as the first decider, consistent with the naive-Bayes assumption. We base our assumptions on the argument that mechanisms for optimal evidence integration of asocial cues have been co-opted to the social case, without any refinement.

2. Results

We quantified the effect of the proposed rules on a group of N individuals that cooperate through social signalling with each other. In both tested scenarios—collective signal detection and sequential sampling—we assumed individuals communicate on a partially connected network, i.e. each individual i has a limited number of neighbours $M_i < N$. We conducted our tests on random geometric graphs (RGG) [13], which are constructed by locating the N nodes at uniform random locations in a unit square, and connecting two nodes when their Euclidean distance is smaller than δ . The value of δ determines the average degree connectivity κ —that is, the average number of neighbours each individual has. We chose to study interaction on an RGG topology as it closely relates to systems embedded in a physical environment, thus matching the characteristics of several biological systems. Results for other types of network topologies are reported in the electronic supplementary material.

2.1. Synchronous updates lead to negative information cascades

As noted above, the naive Bayes-optimal signal detection rule, Weighted Bayes Consensus, gives linear updating of both decisions (equation (1.2)) and confidence (equation (1.3)). In *Material and methods*, we show that such linear updating of confidence leads to an unstable process on the agent network; this means that decisions will be precipitated more rapidly than in stable processes, but at the expense of accuracy. In figure 2, we numerically compare the speed and accuracy of Weighted Bayes Consensus against the Belief Consensus algorithm [24], through which every individual, by iteratively averaging weighted opinions over its neighbourhood, computes the weighted mean of the entire population (find a detailed description of the algorithm in §5.3). On the one hand, Weighted Bayes Consensus is the locally optimal solution, as individuals apply the Bayes-optimal signal detection rule on information locally available at each moment; on the other hand, Bayes Consensus is the globally optimal solution, as after a number of iterations every individual computes the global weighted average (equation (1.2) computed on every member), which corresponds to the optimal solution to the collective signal detection problem [11]. In both cases, optimality is defined in terms of accuracy only, assuming naive individuals. In the *Discussion*, we consider the relevance of these algorithms for natural systems; for now, we note that, as group heterogeneity varies, a speed–accuracy trade-off is described (figure 2). Compared with the Belief Consensus algorithm, the Weighted Bayes Consensus is dominated on group accuracy but takes on average a shorter time to reach consensus. This comparison makes it possible to appreciate the effect of the unstable dynamics of the Weighted Bayes Consensus in contrast to the slower but stable dynamics of the Belief Consensus algorithm. Figure 2 also shows that the group improves in collective accuracy with increasing heterogeneity σ_w as a consequence of higher mean individual accuracy (see also figure SF1 in the electronic supplementary material).

2.2. Asynchronous updates prevent negative information cascades through the emergence of informed leaders

In collective sequential sampling, individuals can be assumed to incur a cost that is a linear function of ω_e for erroneous decisions (assuming that correct decisions incur no cost) and ω_t for the time taken to make their decision. This can be defined according to the Bayes risk [16,25], enabling agents to set optimally

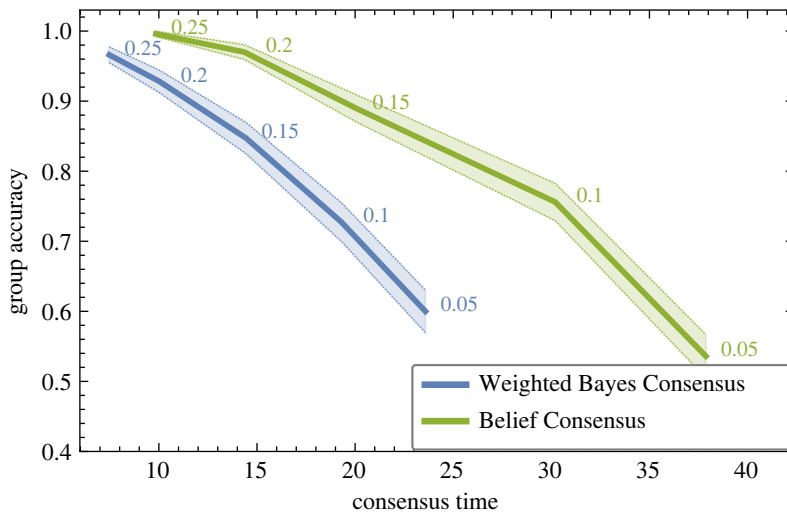


Figure 2. Via synchronous updates of individuals' confidence through the Weighted Bayes Consensus rule—which neglects correlation of social information—the group reaches a consensus in less time than the optimal strategy (Belief Consensus). However, quick runaways can lead to erroneous decisions, as shown by a lower group accuracy. Here, we show the results for 10^3 simulations of $N = 50$ individuals that have individual accuracy α drawn from a normal distribution $\mathcal{N}(\mu_\alpha = 0.5, \sigma_\alpha)$ (flipping to $1 - \alpha$ when $\alpha < 0.5$), and varying heterogeneity σ_α (shaded areas are 95% confidence intervals). The Weighted Bayes Consensus rule (WBC, blue lines) has a lower group accuracy than the Belief Consensus algorithm (BC, green lines); however, it is quicker (heterogeneity level σ_α is indicated next to the curve; group accuracy is computed as the proportion of runs with unanimous agreement for S^+).

their decision thresholds in order to minimize expected cost [16,26] (see text ST1 in the electronic supplementary material). For collective decisions in sequential collective decision-making we find that, contrary to the synchronous case, larger numbers of information cascades are triggered by the best decision-makers (figure 3). This is because, on average, the best individuals are expected to reach their decision threshold quicker than others (figure 1*d*, [14]). Such early signals cause a larger response than delayed decisions (figure 3*a*). The resulting effect is that the best individuals—those that are more accurate because they have a higher signal-to-noise ratio \tilde{A}/σ —more often trigger a cascade of decisions in the group (figure 3*b*), and the best decision-makers' cascades are typically larger than the ones triggered by the inferior individuals (figures 3*c*; electronic supplementary material, SF6). Therefore, we observe that on average the best decision-makers have the highest influence on the group, acting as emergent group leaders as a direct consequence of a combination of psychological and neuroscientific mechanisms [14–16].

2.3. Model comparison

The goal of this study is to show that, contrary to common intuition [1–4], early decisions can have a beneficial impact on the collective dynamics by triggering positive information cascades, even in populations of naive-Bayesian agents, whereas in the absence of temporal ordering among decisions (synchronous scenario), naive-Bayesian agents can frequently suffer negative information cascades. Despite being biologically unrealistic (as further discussed in §4), the collective dynamics in the synchronous scenario can be rescued by a simple change in the individual behaviour, by averaging neighbours' opinions rather than summing them (Belief Consensus algorithm). Our analysis also explains the causes of our results. In particular, we compute the mathematical stability and instability of the synchronous scenario systems when there is perpetual integration of social information and we indicate how confidence can be inferred from the decision speed based on known neuroscientific mechanisms [15,16].

Here, we explicate similarities and differences in the two models and in the assumptions on which the two scenarios are based. Both scenarios describe how individuals integrate social information in order to improve their own world's estimate. Both scenarios are also based on the same assumptions that individuals are naive because they neglect correlations in social information and locally integrate social evidence through Bayes-optimal rules according to the information they have access to. Hence,

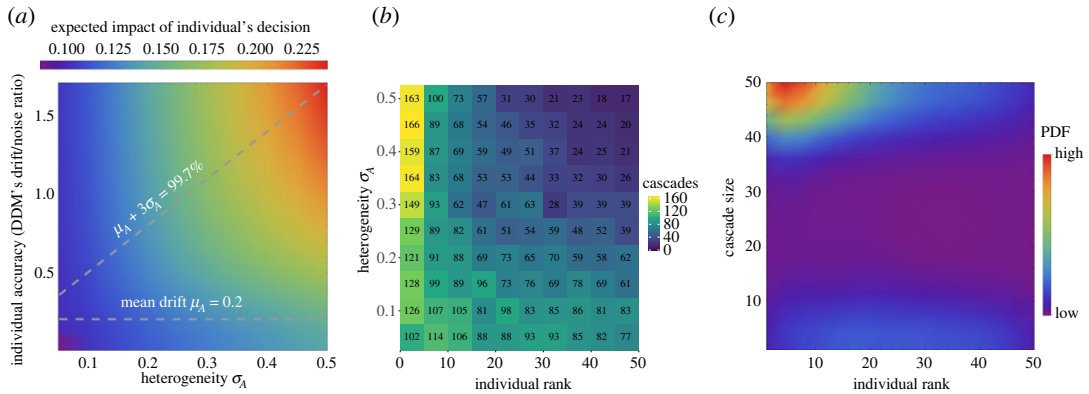


Figure 3. Emergent leaders from psychological and neuroscientific mechanisms. (a) We report the expected impact of an individual decision on its neighbours. In diverse groups, the most accurate individuals are expected to have a large impact on others. We computed the expected decision time for each DDM with noise $\sigma=1$ and drift sampled from the normal distribution $\mathcal{N}(\mu_A, \sigma_A)$ (with $\mu_A=0.2$, and σ_A varied on the x -axis, with $3\sigma_A$ indicated with dashed lines). The threshold z is set to optimize the Bayes risk with costs $\omega_t=1$ and $\omega_e=100$. Higher drifts are expected to reach the threshold earlier [16], and earlier reactions are considered a sign of higher confidence [15]. For each case, we visualize the kick size k (equation (1.5)) at the expected decision time, normalized by the threshold z , i.e. the colour indicates k/z . Therefore, larger values bring the individual closer to its decision threshold. (b) In diverse groups, the best individuals—that is, with higher drift/noise ratio, $\tilde{\mu}/\sigma$ —more often trigger a cascade. We sort (on the x -axis) the individuals in decreasing order of drift/noise ratio and report the number of cascades each individual triggers. We count as a cascade the triggering of a sequence of at least $N/10$ decisions. The results are from 500 simulation runs of a group of $N=50$ individuals communicating on a sparse network (connected random geometric graph with average degree $\kappa=10$), and with drift sampled from $\mathcal{N}(\mu_A=0.2, \sigma_A)$. In more homogeneous groups (low σ_A), cascades are almost equally likely to be triggered by any individual. Instead, in highly heterogeneous groups (high σ_A), cascades are predominantly caused by the best individuals. (c) The most accurate individuals trigger the largest cascades. We show the probability density function (PDF) for each individual, sorted in decreasing drift/noise ratio order, to trigger a cascade of different sizes (on the y -axis). The PDF is computed from 500 runs for the case of $\sigma_A=0.5$ and the same parameters of panel (b). Thus, in summary, leaders emerge in heterogeneous groups as their decisions are followed (a) strongly, (b) more frequently and (c) by a larger portion of the population.

the considered strategies are optimal in terms of accuracy on the presumption of naive individuals; we further discuss the biological relevance of our assumptions in §4. Notwithstanding the strong similarities, the two scenarios differ in terms of the environmental information the individuals integrate, how and when they communicate with one another, and, consequently, in the rules to combine their opinion and confidence with the ones of their neighbours (figure 1). As a consequence of such differences, the performance of the two scenarios cannot be compared directly, rather we show how (mis-)information cascades have a different impact on the two scenarios. Comparing quantitatively the speed-accuracy results of both scenarios is impractical. In fact, in figure 2, we only analyse the runs of the signal detection scenario that reached unanimous agreement; however, a condition of unanimous consensus is rare in the sequential sampling scenario, because individuals do not change their decision once they reached a threshold. Although there is no consensus, electronic supplementary material, figure SF2a,b shows that, in sequential sampling, a large majority of the group makes correct decisions, more frequently than in an asocial condition. The objective of our analysis is to show the negative impact of quick runaways in the synchronous signal detection scenario (figure 2), and explain that the situation is the opposite in the asynchronous sequential sampling scenario where the large majority of information cascades are triggered by individuals making correct early decisions (figure 3). These results generalize to different network topologies and all tested parameters, as shown in electronic supplementary material, figures SF3, SF4, SF5 and SF6.

3. Previous work

Collective decision-making in groups of individuals that update their opinion beliefs has been widely investigated, commencing with the seminal model of DeGroot [27]. Collective decision-making models have been investigated in the social sciences in the form of social learning and in engineering as consensus-averaging algorithms. We briefly review previous relevant approaches.

3.1. Non-Bayesian social learning

A large amount of work has investigated social learning [1–4,28–41] in which individuals update their beliefs with a Bayes-optimal rule that assumes correlation neglect, also referred to as non-Bayesian social learning. The correlation neglect assumption is that individual agents do not take account of the fact that incorporating neighbours' social information with their own iteratively leads to correlated information. Instead, when individuals know the full network topology, they can apply the actual Bayes-optimal update rule, as in [42–44], or approximations of it [45,46], although even with full information doing so may be computationally prohibitive. In studies of non-Bayesian social learning, various aspects have been analysed, such as the conditions for polarization of the population [35,41], or how information cascades can be the result of non-Bayesian update of local beliefs [1–4]. Studies showed how correlation neglect can improve the performance of voting systems [47] or lead to the formation of extremists [37,38,41,48]. In these studies, individuals sequentially make their rational decision based either on all previous individual decisions [1–3] or only the previous [4]. As more individuals make the same decision, the probability the next individuals will ignore their personal opinion and follow the social information becomes higher [49]. Individuals neglect correlation of information and the ordering of previous decisions, which can have determining effects on the collective dynamics, as shown in [50]. In our work, we do not externally impose the ordering of votes, rather we test both synchronous simultaneous voting and asynchronous signalling with the ordering determined by the environmental sampling dynamics.

3.2. Consensus averaging algorithms

As a form of social learning, consensus averaging algorithms allow the nodes of a network, each having a numeric value, to compute in a decentralized way the average of all these values. Therefore, through these decentralized algorithms, each agent on a sparse graph can converge on the same average confidence value. The Belief Consensus algorithm [24] uses a linear function, while other averaging algorithms employ nonlinear [51–54] or heterogenous functions [55]. The advantage of consensus averaging algorithms is a guarantee of convergence in a relatively small number of time steps. Consensus-averaging opinion dynamics models have also shown unbounded increases in individual agent confidence, leading to the formation of extremists in populations [56–62].

3.3. Optimal evidence accumulation

The dynamics of a network of optimal evidence accumulators has been investigated in the form of coupled DDMs [63] in which each accumulator can access the state of its neighbours prior to reaching its own decision. Accessing the internal state of other agents is biologically implausible, and accordingly, in our work, neighbours only share their decision when the decision threshold is reached. A similar recent study, [20], has derived the theory to allow optimal decision-makers, modelled as DDMs, to update their evidence based on neighbours' decisions (once the neighbour's evidence reaches the decision threshold). However, this work makes the biologically unrealistic assumptions that agents are truly Bayes-optimal and do not use correlation neglect as a computational short cut. These assumptions require the agents to know the complete communication topology \mathcal{G} in order to compute 'second-order' evidence integration over the behaviour of the neighbours of neighbours. The calculations rapidly become very intricate. Additionally, in integrating only neighbours' decisions, but not the time take to reach those decisions, the agents modelled by Karamched *et al.* [20] neglect an important information source, which we incorporate into our model. In agreement with previous analysis [64], our model predicts that the mean collective cost (computed from decision time and errors) decreases by increasing group heterogeneity and group connectivity (see figure SF2 in the electronic supplementary material).

4. Discussion

We have shown analytically that for synchronous decisions locally optimal Bayesian integration of weighted votes, in order to reach a group decision, is described by an unstable linear dynamical system in which erroneous decisions dominate. As shown numerically in comparison with an existing linear consensus algorithm with guaranteed convergence, this results in faster decisions but at the expense of group decision accuracy. By contrast, when decisions are asynchronous early decisions tend to be correct, and hence, through confidence-signalling, leaders can spontaneously emerge from the best informed members of a group and

precipitate fast and accurate group decisions. That animal groups exploit the skills of the best individuals has already been observed [65,66]; however, in our analysis, group leaders emerge from social interactions as the consequence of applying confidence mechanisms from neuroscience [15,16] to social dynamics.

Our results can be interpreted through the lens of ‘information cascades’ in decision-making groups of humans and other animals, in which early erroneous information is assumed to dominate (e.g. [67–70]). In contrast to this accepted view, however, negative information cascades occur when decisions are synchronous, so there are no ‘early’ decisions, but the move to asynchronous decisions actually results in early decisions being correct more often than incorrect and, correspondingly, leads to positive rather than negative information cascades on average. Our predictions are consistent with the empirical observations of collective decision-making in fish [71], in which the first fish making a decision is generally no less accurate than later fish. Despite standard theory on sequential choices suggesting the first decision-maker should perform worse, empirical results [71] and our analysis indicate the opposite: early responses are the consequence of having access to better information, and thus acting on that information sooner. Correct and early responders can be individuals with better abilities to discriminate between environmental stimulus and noise, either due to systematic higher capabilities [65,66] or to occasional access to a better information source (e.g. due to a better position) [71]. While our model is based on confidence mechanisms from neuroscience [15,16], we do not exclude the possibility that in some species decision order may also be determined by individual traits, such as boldness or impulsivity [72].

Our analysis assumes that optimal rules for asocial information integration may have been co-opted to social scenarios where they are non-optimal, since they neglect correlated information. In the literature, correlation neglect has been studied under different names, such as ‘bounded rationality’ [33], ‘imperfect recall’ [40,73], ‘persuasion bias’ [30,35] or ‘naive inference’ [38]. Such correlation neglect has been observed in experiments with humans, which are cognitively advanced organisms that could, in principle, solve the correlation problem but still neglect to do so [74–77]. Thus, since natural selection acts at the level of the individual rather than the group [78], our results may help provide a normative explanation for such apparently non-adaptive behavioural outcomes. Indeed, evidence of maladaptive social information leading to suboptimal group decision-making has been reported in several species via empirical observations [69,70,79–83] and theoretical models [84,85].

As noted, a superior solution to decision-making under correlation neglect exists for the synchronous decision case in the form of the Belief Consensus algorithm which averages rather than sums information from neighbours. Changing to use this method of evidence integration would be straightforward even for selection acting on individuals within groups, since the behavioural selection is at the level of the individual, and membership of a group in which decisions are reached more effectively is individually advantageous. If evolutionary stable, this change of strategy would globally improve collective decision-making, but would not contradict our results, as interactions are synchronous and there are no early decisions. It is important noting that, regardless of which strategy has a higher selective advantage, in any case, the synchronous decision model is a very unrealistic abstraction of biological reality. By contrast, for the more realistic scenario of asynchronous decisions, avoiding correlation neglect is informationally and computationally very demanding [74–77], hence the heuristic of applying naive-Bayesian evidence integration to social information is highly plausible, and under this reasonable assumption early decisions tend to precipitate positive rather than negative information cascades, in contradiction to previous assumptions.

5. Material and methods

Our method applies Bayes’ rule [19] to specify how the individual i should compute a Bayes-optimal integration of its M_i neighbours’ opinions to update its opinion x_i^{t+1} and its confidence c_i^{t+1} .

5.1. Integrating neighbours’ confidence into collective signal detection

Each individual i communicates to its neighbours M_i its opinion x_i^t and its confidence c_i^t . Assuming all individuals use the same computation of equation (1.1) to derive their confidence, its inverse gives the accuracy of each neighbour

$$\alpha_i^t = \frac{e^{c_i^t}}{1 + e^{c_i^t}}. \quad (5.1)$$

Given the set of received votes $V_i^t = \{X_i^t, K_i^t\}$ as the combination of received opinions $X_i^t = \{x_i^t, x_{j \in M_i}^t\}$ at time t and the set of accuracies $K_i^t = \{\alpha_i^t, \alpha_{j \in M_i}^t\}$ from equation (5.1), the agent i can compute its confidence from the probability that the aggregated opinion x_i^{t+1} is correct (i.e. the true state of the world $S \in \{S^+, S^-\}$ is equal to the individual's opinion $x_i^{t+1} \in \{S^+ = +1, S^- = -1\}$). The new confidence c_i^{t+1} corresponds to the log-odds of being correct rather than incorrect,

$$\begin{aligned} c_i^{t+1} &= \ln \left(\frac{P(\text{my opinion is correct} | \text{received votes})}{P(\text{my opinion is incorrect} | \text{received votes})} \right) \\ &= \ln \left(\frac{P(S = x_i^{t+1} | V_i^t)}{P(S = -x_i^{t+1} | V_i^t)} \right). \end{aligned} \quad (5.2)$$

Neglecting information correlations, a statistically optimal individual can compute the probability $P(S = x_i^{t+1} | V_i^t)$ that the aggregated opinion x_i^{t+1} is correct given the received votes $V_i^t = \{X_i^t, K_i^t\}$ using Bayes' rule as

$$P(S = x_i^{t+1} | V_i^t) = \frac{P(V_i^t | S = x_i^{t+1})P(S = x_i^{t+1})}{P(V_i^t)}. \quad (5.3)$$

where the probability of observing the votes V_i^t assuming $S = x_i^{t+1}$ corresponds to a simple multiplication of probabilities as

$$P(V_i^t | S = x_i^{t+1}) = \prod_{j \in \{M_i, i\}} k_j, \quad k_j = \begin{cases} \alpha_j^t & \text{if } x_j^t = x_i^{t+1} \\ 1 - \alpha_j^t & \text{if } x_j^t = -x_i^{t+1} \end{cases}. \quad (5.4)$$

From equations (5.1) and (5.4), we have that if $x_j^t = x_i^{t+1}$ then

$$\alpha_j^t = \frac{e^{c_j^t x_j^t}}{1 + e^{c_j^t x_j^t}},$$

and if $x_j^t = -x_i^{t+1}$, then

$$1 - \alpha_j^t = 1 - \frac{e^{-c_j^t x_j^t}}{1 + e^{-c_j^t x_j^t}} = \frac{e^{c_j^t x_j^t}}{1 + e^{c_j^t x_j^t}}.$$

Therefore, for equation (5.4), irrespective of the sign of x_j^t we have that

$$k_j^+ = \frac{e^{c_j^t x_j^t}}{1 + e^{c_j^t x_j^t}} \quad \text{and} \quad k_j^- = \frac{1}{1 + e^{c_j^t x_j^t}}. \quad (5.5)$$

Using the above simplification, the update of equation (5.2) becomes

$$\begin{aligned} c_i^{t+1} &= \left| \ln \left(\prod_{j \in M_i} e^{c_j^t x_j^t} \right) + \pi(x_i^{t+1}) \right| \\ &= \left| x_i^t c_i^t + \sum_{j \in M_i} (x_j^t c_j^t) + \pi(x_i^{t+1}) \right|, \end{aligned} \quad (5.6)$$

where equation (5.6) corresponds to equation (1.3) in the main text.

5.2. Sequential sampling scenario

In the sequential sampling scenario, an individual i that is integrating evidence and receives at time t a decision x_j^t from its neighbour j , updates its evidence variable $y_i(t)$ by k , which it computes with equation (1.5). The first two terms of this equation are the log-odds of the first passage time of the DDM through the threshold for x_j^t at t and the log-odds of hitting the threshold for x_j^t before the one for $-x_j^t$. If, without loss of generality we assume $x_j^t = 1$, the first-passage time through z^+ is computed, following the results of [86], as

$$P(t | z^+, A, \sigma) = e^{-(A^2 t / 2\sigma^2) - (Az^+ / \sigma^2)} (1 + e^{2Az^+ / \sigma^2}) \theta \left(t, \frac{z^+}{\sigma}, \frac{2z^+}{\sigma} \right), \quad (5.7)$$

where the function $\theta(t, u, v)$ is defined as

$$\theta(t, u, v) = \sum_{k=-\infty}^{+\infty} \frac{v - u + 2kv}{\sqrt{2\pi t^3}} e^{-((v-u+2kv)^2/2t)}.$$

Instead, the probability of hitting z^+ before z^- is

$$P(z^+|A, \sigma) = 1 - \frac{1}{e^{(2Az^+/\sigma^2)} + 1}, \quad (5.8)$$

as from [16].

The individual does not know the drift rate but only knows the random distribution from which the drift is sampled. Therefore, the individual integrates all possible drifts over the given random distribution and equation (1.4) can be rewritten as

$$\begin{aligned} \text{kick size} = k &= \ln \left(\frac{P(S^+ | \text{1st hit threshold is } z^+ \text{ at } t)}{P(S^- | \text{1st hit threshold is } z^+ \text{ at } t)} \right) \\ &= \ln \left(\frac{\int_0^\infty P(t|z^+, A, \sigma, S^+) P(z^+|A, \sigma, S^+) P(A|\sigma, S^+) dA}{\int_0^\infty P(t|z^+, A, \sigma, S^-) P(z^+|A, \sigma, S^-) P(A|\sigma, S^-) dA} \right). \end{aligned} \quad (5.9)$$

Recall that S^+ determines the sign of A , and therefore

$$P(t|z^+, A, \sigma, S^+) = P(t|z^+, +\tilde{A}, \sigma) \quad \text{and}$$

$$P(t|z^+, A, \sigma, S^-) = P(t|z^+, -\tilde{A}, \sigma),$$

and equivalently applies for equation (5.8).

5.3. Analytical comparison

We compare the dynamics of the proposed Weighted Bayes Consensus rule and the linear consensus averaging algorithm from the literature, Belief Consensus [24]. Belief Consensus is a decentralized algorithm which allows each agent on a sparse graph to converge on the same average value [24]. Each agent i runs the algorithm by repeatedly integrating information received from its neighbours M_i . The algorithm implements linear updates that provably converge on global consensus in a finite number of time steps. The algorithm is defined as

$$\begin{aligned} y_i^{t+1} &= x_i^t c_i^t + \epsilon \sum_{j \in M_i} (x_j^t c_j^t - x_i^t c_i^t), \\ x_i^{t+1} &= \text{sign}(y_i^{t+1}), \quad c_i^{t+1} = |y_i^{t+1}|, \end{aligned} \quad (5.10)$$

where $x_i^t \in \{-1, +1\}$ is the option selected by agent i at time t , $c_i^t \geq 0$ is its confidence at that time defined according to equation (1.1), M_i are its neighbours, and ϵ is a parameter. Given the Laplacian matrix L of the connectivity graph \mathcal{G} , in order to guarantee convergence the parameter ϵ must be chosen so that $(I - \epsilon L)$ is a doubly stochastic matrix (where I is the identity matrix of appropriate dimensions). Metropolis–Hastings matrices are among the state-of-the-art techniques to compute ϵ in a decentralized fashion using the local neighbourhood only [87].

We focus on the dynamics of $y_i^t = x_i^t c_i^t$, where $x_i^t = \text{sign}(y_i^t)$ and $c_i^t = |y_i^t|$, as from equations (5.10) and (5.6). Let \mathbf{y}^t be the vector of y_i^t s. Given a graph \mathcal{G} without self-loops, we denote its adjacency matrix by A . Using this notation, we can rewrite equation (5.6) as

$$\mathbf{y}^{t+1} = (\mathbf{I} + A)\mathbf{y}^t, \quad (5.11)$$

where \mathbf{I} is the identity matrix of appropriate dimensions. Similarly, we can rewrite the Belief Consensus as

$$\mathbf{y}^{t+1} = F\mathbf{y}^t, \quad (5.12)$$

where F is a row stochastic matrix.

Both the Belief Consensus (equation (5.10)) and the Weighted Bayes Consensus (equation (5.11)) are linear dynamical systems. It is known that if the underlying graph is connected, the dynamics of equation (5.12) converge to the average of the initial values of \mathbf{y} , i.e. to $(1/N) \sum_{i=1}^N y_i^0$, where N is the number of agents [24]. This convergence is a consequence of the fact that for a connected graph, the matrix R has one

eigenvalue at 1 with associated eigenvector $\mathbf{1}_N$, and all remaining eigenvalues are inside a unit disc centred at the origin. In the context of hypothesis testing, the aggregate log-odds (log-odds of all agents pooled together) is compared against a single threshold. In this sense, the dynamics of equation (5.12) yields the correct statistic at each node which can be compared against the correct threshold, which in our case is zero, (i.e. we need to simply determine the sign of y_i^t). Note that for the consensus \mathbf{y}^t is always bounded.

The dynamics of equation (5.11) replace the action of averaging with the neighbours with the action of simply adding the value of the neighbours to the current agent's value. Note that the dynamics of equation (5.11) are unstable for most graphs, i.e. the value of \mathbf{y}^t grows unboundedly. The agents ignore this instability as the opinion x_i^t is determined only by the sign of y_i^t . The underlying idea is that the projection of the initial condition onto the eigenvector associated with the largest eigenvalue will dominate after a small initial transient, and will be indicative of the sign of the average pooled statistic. However, the eigenvector associated with the largest eigenvalue of $I + A$ is not the ones vector $\mathbf{1}_N$ except for regular graphs. Except for regular graphs, the dominant mode of \mathbf{y}^t will not be associated with the average statistic and will not yield the desired accuracy, but since \mathbf{y}^t will grow exponentially, it will be very quick in reaching a region in which the sign of y_i^t will be stable.

Note that because we are only interested in the sign of the average of the initial conditions, we could also leverage instability to reach quicker decisions in the case of the dynamics of equation (5.12). In equation (5.10), we could destabilize equation (5.12) by introducing the tuneable parameter $\epsilon > 0$ as follows:

$$\mathbf{y}^{t+1} = (\epsilon I + F)\mathbf{y}^t, \quad (5.13)$$

where I is the identity matrix of appropriate dimensions. This dynamics will have the dominant eigenvalue of $1 + \epsilon$, and associated eigenvector $\mathbf{1}_N$. Hence the dominant (unstable) mode will correspond to the average of initial conditions.

Data accessibility. Data and relevant code for this research work are stored in GitHub: <https://github.com/DiODEProject/DecisionsOnNetworks> and have been archived within the Zenodo repository: <https://doi.org/10.5281/zenodo.7032373> [88].

The data are provided in electronic supplementary material [89].

Authors' contributions. A.R.: conceptualization, data curation, formal analysis, funding acquisition, investigation, methodology, software, visualization, writing—original draft, writing—review and editing; T.B.: investigation; V.S.: formal analysis, writing—review and editing; J.A.R.M.: conceptualization, formal analysis, funding acquisition, investigation, methodology, project administration, supervision, writing—original draft, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. We thank Naomi Leonard for helpful discussions. This study was partially funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 647704). A.R. also acknowledges support from the Belgian F.R.S.-FNRS, of which he is a Chargé de Recherches.

References

- Bikhchandani S, Hirshleifer D, Welch I. 1992 A theory of fads, fashion, custom, and cultural change as informational cascades. *J. Polit. Econ.* **100**, 992–1026. (doi:10.1086/261849)
- Banerjee AV. 1992 A simple model of herd behavior. *Q. J. Econ.* **107**, 797–817. (doi:10.2307/2118364)
- Smith L, Sorensen P. 2000 Pathological outcomes of observational learning. *Econometrica* **68**, 371–398. (doi:10.1111/1468-0262.00113)
- Çelen B, Kariv S. 2004 Observational learning under imperfect information. *Games Econ. Behav.* **47**, 72–86. (doi:10.1016/S0899-8256(03)00179-9)
- Laplace PS. 1812 *Theorie Analytique des Probabilités*. Paris, France: Ve Courcier.
- Knill DC, Pouget A. 2004 The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* **27**, 712–719. (doi:10.1016/j.tins.2004.10.007)
- McNamara JM, Green RF, Olsson O. 2006 Bayes' theorem and its applications in animal behaviour. *Oikos* **112**, 243–251. (doi:10.1111/j.0030-1299.2006.14228.x)
- Friston K. 2010 The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138. (doi:10.1038/nrn2787)
- Trimmer PC, Houston AI, Marshall JA, Mendl MT, Paul ES, McNamara JM. 2011 Decision-making under uncertainty: biases and Bayesians. *Anim. Cogn.* **14**, 465–476. (doi:10.1007/s10071-011-0387-4)
- Pouget A, Beck JM, Ma WJ, Latham PE. 2013 Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* **16**, 1170. (doi:10.1038/nn.3495)
- Marshall JA, Brown G, Radford AN. 2017 Individual confidence-weighting and group decision-making. *Trends Ecol. Evol.* **32**, 636–645. (doi:10.1016/j.tree.2017.06.004)
- Bahrami B, Olsen K, Latham PE, Roepstorff A, Rees G, Frith CD. 2010 Optimally interacting minds. *Science* **329**, 1081–1085. (doi:10.1126/science.1185718)
- Penrose M. 2003 *Random geometric graphs*. Oxford Studies in Probability, no. 5. Oxford, UK: Oxford University Press.
- Ratcliff R, Smith PL, Brown SD, McKoon G. 2016 Diffusion decision model: current issues and history. *Trends Cogn. Sci.* **20**, 260–281. (doi:10.1016/j.tics.2016.01.007)
- Kiani R, Shadlen MN. 2009 Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* **324**, 759–64. (doi:10.1126/science.1169405)
- Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. 2006 The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice

- tasks. *Psychol. Rev.* **4**, 700–765. (doi:10.1037/0033-295X.113.4.700)
17. Ratcliff R. 1978 A theory of memory retrieval. *Psychol. Rev.* **85**, 59. (doi:10.1037/0033-295X.85.2.59)
 18. Ratcliff R, McKoon G. 2008 The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922. (doi:10.1162/neco.2008.12.06-420)
 19. Bayes T. 1763 An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, F. R. S. communicated by Mr. Price, in a letter to John Canton, A. M. F. R. S. *Phil. Trans. R. Soc. Lond.* **53**, 370–418. (doi:10.1098/rstl.1763.0053)
 20. Karamched B, Stolarczyk S, Kilpatrick ZP, Josić K. 2020 Bayesian evidence accumulation on social networks. *SIAM J. Appl. Dyn. Syst.* **19**, 1884–1919. (doi:10.1137/19M1283793)
 21. Nitzan S, Paroush J. 1982 Optimal decision rules in uncertain dichotomous choice situations. *Int. Econ. Rev.* **23**, 289–297. (doi:10.2307/2526438)
 22. Boland PJ. 1989 Majority systems and the Condorcet jury theorem. *J. R. Stat. Soc. D (The Statistician)* **38**, 181–189. (doi:10.2307/2348873)
 23. Trimmer PC, Houston AI, Marshall JA, Bogacz R, Paul ES, Mendl MT, McNamara JM. 2008 Mammalian choices: combining fast-but-inaccurate and slow-but-accurate decision-making systems. *Proc. R. Soc. B* **275**, 2353–2361. (doi:10.1098/rspb.2008.0417)
 24. Olfati-Saber R, Franco E, Frazzoli E, Shamma JS. 2006 Belief consensus and distributed hypothesis testing in sensor networks. In *Networked embedded sensing and control* (eds PJ Antsaklis, P Tabuada), pp. 169–182. Berlin, Germany: Springer. (doi:10.1007/11533382_11)
 25. Wald A, Wolfowitz J. 1948 Optimum character of the sequential probability ratio test. *Ann. Math. Stat.* **19**, 326–339. (doi:10.1214/aoms/1177730197)
 26. Edwards W. 1965 Optimal strategies for seeking information: models for statistics, choice reaction times, and human information processing. *J. Math. Psychol.* **2**, 312–329. (doi:10.1016/0022-2496(65)90007-6)
 27. DeGroot MH. 1974 Reaching a consensus. *J. Am. Stat. Assoc.* **69**, 118. (doi:10.1080/01621459.1974.10480137)
 28. Bala V, Goyal S. 1998 Learning from neighbours. *Rev. Econ. Stud.* **65**, 595–621. (doi:10.1111/1467-937X.00059)
 29. Bala V, Goyal S. 2001 Conformism and diversity under social learning. *Econ. Theory* **17**, 101–120. (doi:10.1007/PL00004094)
 30. DeMarzo PM, Vayanos D, Zwiebel J. 2003 Persuasion bias, social influence, and unidimensional opinions. *Q. J. Econ.* **118**, 909–968. (doi:10.1162/00335530360698469)
 31. Chamley CP. 2004 *Rational herds: economic models of social learning*. Cambridge, UK: Cambridge University Press.
 32. Banerjee A, Fudenberg D. 2004 Word-of-mouth learning. *Games Econ. Behav.* **46**, 1–22. (doi:10.1016/S0899-8256(03)00048-4)
 33. Golub B, Jackson MO. 2010 Naive learning in social networks and the wisdom of crowds. *Am. Econ. J.: Microecon.* **2**, 112–149. (doi:10.1257/mic.2.1.112)
 34. Jackson MO. 2011 An overview of social networks and economic applications. In *Handbook of Social Economics*, vol. 1 (eds J Benhabib, A Bisin, MO Jackson), pp. 511–585. San Diego, CA: Elsevier B.V.
 35. Corazzini L, Pavesi F, Petrovich B, Stanca L. 2012 Influential listeners: an experiment on persuasion bias in social networks. *Eur. Econ. Rev.* **56**, 1276–1288. (doi:10.1016/j.eurocorev.2012.05.005)
 36. Jadbabaie A, Molavi P, Sandroni A, Tahbaz-Salehi A. 2012 Non-Bayesian social learning. *Games Econ. Behav.* **76**, 210–225. (doi:10.1016/j.geb.2012.06.001)
 37. Ortoleva P, Snowberg E. 2015 Overconfidence in political behavior. *Am. Econ. Rev.* **105**, 504–535. (doi:10.1257/aer.20130921)
 38. Gagnon-Bartsch T, Rabin M. 2016 Naive social learning, mislearning, and unlearning. *Mimeo*. See <https://scholar.harvard.edu/files/rabin/files/gagnon-bartschrabin2016.pdf>.
 39. Mossel E, Tamuz O. 2017 Opinion exchange dynamics. *Probab. Surveys* **14**, 155–204. (doi:10.1214/14-PS230)
 40. Molavi P, Tahbaz-Salehi A, Jadbabaie A. 2018 A theory of non-Bayesian social learning. *Econometrica* **86**, 445–490. (doi:10.3982/ECTA14613)
 41. Levy G, Razin R. 2018 Information diffusion in networks with the Bayesian peer influence heuristic. *Games Econ. Behav.* **109**, 262–270. (doi:10.1016/j.geb.2017.12.020)
 42. Acemoglu D, Dahleh MA, Lobel I, Ozdaglar A. 2011 Bayesian learning in social networks. *Rev. Econ. Stud.* **78**, 1201–1236. (doi:10.1093/restud/rdr004)
 43. Acemoglu D, Bimpikis K, Ozdaglar A. 2014 Dynamics of information exchange in endogenous social networks. *Theor. Econ.* **9**, 41–97. (doi:10.3982/TE1204)
 44. Mossel E, Olsman N, Tamuz O. 2016 Efficient Bayesian learning in social networks with Gaussian estimators. In *2016 54th Annual Allerton Conf. on Communication, Control, and Computing (Allerton)*, pp. 425–432. IEEE.
 45. Gale D, Kariv S. 2003 Bayesian learning in social networks. *Games Econ. Behav.* **45**, 329–346. (doi:10.1016/S0899-8256(03)00144-1)
 46. Mossel E, Tamuz O. 2010 Iterative maximum likelihood on networks. *Adv. Appl. Math.* **45**, 36–49. (doi:10.1016/j.aam.2009.11.004)
 47. Levy G, Razin R. 2015 Correlation neglect, voting behavior, and information aggregation. *Am. Econ. Rev.* **105**, 1634–1645. (doi:10.1257/aer.20140134)
 48. Levy G, Razin R. 2015 Does polarisation of opinions lead to polarisation of platforms? The case of correlation neglect. *Q. J. Polit. Sci.* **10**, 321–355. (doi:10.1561/100.00015010)
 49. Arganda S, Pérez-Escudero A, de Polavieja GG. 2012 A common rule for decision making in animal collectives across species. *Proc. Natl Acad. Sci. USA* **109**, 20 508–20 513. (doi:10.1073/pnas.1210664109)
 50. Mann RP. 2018 Collective decision making by rational individuals. *Proc. Natl Acad. Sci. USA* **115**, E10 387–E10 396. (doi:10.1073/pnas.1811964115)
 51. Wang L, Xiao F. 2010 Finite-time consensus problems for networks of dynamic agents. *IEEE Trans. Autom. Control* **55**, 950–955. (doi:10.1109/TAC.2010.2041610)
 52. Jia P, MirTabatabaei A, Friedkin NE, Bullo F. 2015 Opinion dynamics and the evolution of social power in influence networks. *SIAM Rev.* **57**, 367–397. (doi:10.1137/130913250)
 53. Amelkin V, Bullo F, Singh AK. 2017 Polar opinion dynamics in social networks. *IEEE Trans. Autom. Control* **62**, 5650–5665. (doi:10.1109/TAC.2017.2694341)
 54. Ye M. 2019 *Opinion dynamics and the evolution of social power in social networks*. Cham, Switzerland: Springer International Publishing.
 55. Liu J, Ye M, Anderson BD, Basar T, Nedic A. 2018 Discrete-time polar opinion dynamics with heterogeneous individuals. In *2018 IEEE Conf. on Decision and Control (CDC)*, pp. 1694–1699, IEEE.
 56. Marvel SA, Kleinberg J, Kleinberg RD, Strogatz SH. 2011 Continuous-time model of structural balance. *Proc. Natl Acad. Sci. USA* **108**, 1771–1776. (doi:10.1073/pnas.1013213108)
 57. Dandekar P, Goel A, Lee DT. 2013 Biased assimilation, homophily, and the dynamics of polarization. *Proc. Natl Acad. Sci. USA* **110**, 5791–5796. (doi:10.1073/pnas.1217220110)
 58. Martins ACR, Galam S. 2013 Building up of individual inflexibility in opinion dynamics. *Phys. Rev. E* **87**, 042807. (doi:10.1103/PhysRevE.87.042807)
 59. La Rocca CE, Braunstein LA, Vazquez F. 2014 The influence of persuasion in opinion formation and polarization. *EPL (Europhysics Letters)* **106**, 40004. (doi:10.1209/0295-5075/106/40004)
 60. Balenzuela P, Pinasco JP, Semeshenko V. 2015 The undecided have the key: interaction-driven opinion dynamics in a three state model. *PLoS ONE* **10**, e0139572. (doi:10.1371/journal.pone.0139572)
 61. Pinasco JP, Semeshenko V, Balenzuela P. 2017 Modeling opinion dynamics: theoretical analysis and continuous approximation. *Chaos, Solitons Fractals* **98**, 210–215. (doi:10.1016/j.chaos.2017.03.033)
 62. Woolcock A, Connaughton C, Merali Y, Vazquez F. 2017 Fitness voter model: damped oscillations and anomalous consensus. *Phys. Rev. E* **96**, 032313. (doi:10.1103/PhysRevE.96.032313)
 63. Srivastava V, Leonard NE. 2014 Collective decision-making in ideal networks: the speed-accuracy tradeoff. *IEEE Trans. Control Netw. Syst.* **1**, 121–132. (doi:10.1109/TCNS.2014.2310271)
 64. Karamched B, Stickle M, Ott W, Lindner B, Kilpatrick ZP, Josić K. 2020 Heterogeneity improves speed and accuracy in social networks. *Phys. Rev. Lett.* **125**, 218302. (doi:10.1103/PhysRevLett.125.218302)
 65. Morand-Ferron J, Quinn JL. 2011 Larger groups of passerines are more efficient problem solvers in the wild. *Proc. Natl Acad.*

- Sci. USA* **108**, 15898. (doi:10.1073/pnas.1111560108)
66. Ioannou CC. 2017 Swarm intelligence in fish? The difficulty in demonstrating distributed and self-organised collective intelligence in (some) animal groups. *Behav. Process.* **141**, 141. (doi:10.1016/j.beproc.2016.10.005)
 67. Couzin ID. 2009 Collective cognition in animal groups. *Trends Cogn. Sci.* **13**, 36–43. (doi:10.1016/j.tics.2008.10.002)
 68. Anderson LR, Holt CA. 1997 Information cascades in the laboratory. *Am. Econ. Rev.* **87**, 847–862.
 69. Giraldeau L, Valone TJ, Templeton JJ. 2002 Potential disadvantages of using socially acquired information. *Phil. Trans. R. Soc. B* **357**, 1559–1566. (doi:10.1098/rstb.2002.1065)
 70. Rieucou G, Giraldeau LA. 2011 Exploring the costs and benefits of social information use: an appraisal of current experimental evidence. *Phil. Trans. R. Soc. B* **366**, 949–957. (doi:10.1098/rstb.2010.0325)
 71. Ward AJW, Herbert-Read JE, Sumpter DJT, Krause J. 2011 Fast and accurate decisions through collective vigilance in fish shoals. *Proc. Natl Acad. Sci. USA* **108**, 2312–2315. (doi:10.1073/pnas.1007102108)
 72. Jolles JW, Boogert NJ, Sridhar VH, Couzin ID, Manica A. 2017 Consistent individual differences drive collective behavior and group functioning of schooling fish. *Curr. Biol.* **27**, 2862–2868.e7. (doi:10.1016/j.cub.2017.08.004)
 73. Piccione M, Rubinstein A. 1997 On the interpretation of decision problems with imperfect recall. *Games Econ. Behav.* **20**, 3–24. (doi:10.1006/game.1997.0536)
 74. Kallir I, Sonsino D. 2009 The neglect of correlation in allocation decisions. *Southern Econ. J.* **75**, 1045–1066. (doi:10.1002/j.2325-8012.2009.tb00946.x)
 75. Eyster E, Wezsacker G. 2010 Correlation neglect in financial decision-making. *SSRN Electronic Journal.* (doi:10.2139/ssrn.1735339)
 76. Eyster E, Rabin M, Wezsacker G. 2015 An experiment on social mislearning. *SSRN Electronic Journal.* (doi:10.2139/ssrn.2704746)
 77. Enke B, Zimmermann F. 2017 Correlation neglect in belief formation. *Rev. Econ. Stud.* **86**, 313–332. (doi:10.1093/restud/rdx081)
 78. Williams GC. 1966 *Adaptation and natural selection: a critique of some current evolutionary thought*. Princeton, NJ: Princeton University Press.
 79. Laland KN, Williams K. 1998 Social transmission of maladaptive information in the guppy. *Behav. Ecol.* **9**, 493–499. (doi:10.1093/beheco/9.5.493)
 80. Pongrácz P, Miklósi Á, Kubinyi E, Topál J, Csányi V. 2003 Interaction between individual experience and social learning in dogs. *Anim. Behav.* **65**, 595–603. (doi:10.1006/anbe.2003.2079)
 81. Nocera JJ, Forbes GJ, Giraldeau L-A. 2006 Inadvertent social information in breeding site selection of natal dispersing birds. *Proc. R. Soc. B* **273**, 349–355. (doi:10.1098/rspb.2005.3318)
 82. Rieucou G, Giraldeau L-A. 2009 Persuasive companions can be wrong: the use of misleading social information in nutmeg mannikins. *Behav. Ecol.* **20**, 1217–1222. (doi:10.1093/beheco/arp121)
 83. Avarguès-Weber A, Lachlan R, Chittka L. 2018 Bumblebee social learning can lead to suboptimal foraging choices. *Anim. Behav.* **135**, 209–214. (doi:10.1016/j.anbehav.2017.11.022)
 84. Dechaume-Moncharmont F-X, Dornhaus A, Houston AI, McNamara JM, Collins EJ, Franks NR. 2005 The hidden cost of information in collective foraging. *Proc. R. Soc. B* **272**, 1689–1695. (doi:10.1098/rspb.2005.3137)
 85. Grüter C, Leadbeater E. 2014 Insights from insects about adaptive social information use. *Trends Ecol. Evol.* **29**, 177–184. (doi:10.1016/j.tree.2014.01.004)
 86. Srivastava V, Feng SF, Cohen JD, Leonard NE, Shenav A. 2017 A martingale analysis of first passage times of time-dependent Wiener diffusion models. *J. Math. Psychol.* **77**, 94–110. (doi:10.1016/j.jmp.2016.10.001)
 87. Bullo F. 2019 *Lectures on network systems*, 1.3 edn. With contributions by J. Cortes, F. Dorfler, and S. Martinez. Seattle, WA: Kindle Direct Publishing.
 88. Reina A, Bose T, Srivastava V, Marshallb JAR. 2022 DiODEProject/DecisionsOnNetworks: source code of Reina et al. The Royal Society Open Science (2022). Zenodo Repository. (doi:10.5281/zenodo.7032373)
 89. Reina A, Bose T, Srivastava V, Marshallb JAR. 2023 Asynchrony rescues statistically optimal group decisions from information cascades through emergent leaders. Figshare. (doi:10.6084/m9.figshare.c.6456121)