

SOCIAL SCIENCES

Wise or mad crowds? The cognitive mechanisms underlying information cascades

Alan N. Tump^{1*}, Timothy J. Pleskac^{1,2}, Ralf H. J. M. Kurvers¹

Whether getting vaccinated, buying stocks, or crossing streets, people rarely make decisions alone. Rather, multiple people decide sequentially, setting the stage for information cascades whereby early-deciding individuals can influence others' choices. To understand how information cascades through social systems, it is essential to capture the dynamics of the decision-making process. We introduce the social drift–diffusion model to capture these dynamics. We tested our model using a sequential choice task. The model was able to recover the dynamics of the social decision-making process, accurately capturing how individuals integrate personal and social information dynamically over time and when their decisions were timed. Our results show the importance of the interrelationships between accuracy, confidence, and response time in shaping the quality of information cascades. The model reveals the importance of capturing the dynamics of decision processes to understand how information cascades in social systems, paving the way for applications in other social systems.

INTRODUCTION

In many situations—be they financial investments, consumer choices, or simply crossing the street—one is generally not making a decision alone. Rather, there are many other people present each making their own decisions. In such situations, people can observe the choices of others and use that information to inform their own decisions. Early-deciding individuals can thereby trigger information cascades, in which later-deciding individuals adopt earlier choices, potentially creating a situation where, in the extreme case, everyone does what everyone else is doing, even at the expense of abandoning their private information (1–3).

Yet, for a myriad of reasons—from limited time and computational resources to biases in the decision process—people's choices do not always perfectly reflect the true state of the world. Information cascades can thus promote both positive and negative outcomes: In online environments, for example, both true and fake news can spread quickly (4); in offline environments, the behavior of initial pedestrians crossing a road can amplify both safe and risky behaviors in other pedestrians (5). Understanding the conditions leading to positive and negative information cascades is crucial across many domains, including financial markets (6), consumer preferences (7), political opinion formation (8), and opinion dynamics in social networks (9).

To understand the conditions underlying positive and negative information cascades, we need to model the timing of individual decisions and how individuals integrate personal and social information (i.e., other people's decisions) dynamically over time. We currently, however, lack a detailed understanding of the individual decision process in sequential choice paradigms. Many models of information cascades assume a random decision order and are thus ill-equipped to predict who will respond earlier and why [e.g., (1, 2, 10–12)]. When models of information cascades do refer to the timing of decisions, they do so often from an optimal Bayesian perspective based on the quality of each individual's private information

and typically model only the order in which decisions arrive [e.g., (13–15)]. Yet, we know that people's actual choice behavior often deviates systematically from optimal Bayesian models (16–18). Moreover, arguably, a fully descriptive model needs to account for the speed of people's choices.

To address these shortcomings, we developed a dynamic theory of social decision-making by focusing on each individual's decision process. As a basis, we took a well-established modeling framework of individual decision-making that models decisions as a dynamic process in which information is accumulated as evidence over time until a threshold is reached [e.g., (19, 20)]. This evidence accumulation process has been successful in accounting for a wide range of decisions in domains including perception (21), memory (19), categorization (22), preference (23), and inference (17) and has successfully been applied to analyze the influence of static social information (24, 25). We extended this evidence accumulation framework by showing how the choices of others are integrated with personal information and together accumulated as evidence. This approach provides a process-level account of the choices and response times (RTs) of individuals in dynamic social systems. We tested the model in an empirical study. Findings showed that participants self-organize based on the quality of their personal information so that late deciders benefit from observing the choices of early deciders. Fitting the model to the data allowed us to test several hypotheses about how individuals simultaneously combine personal and social information and how they time their decision in groups. In addition, we reveal mechanisms leading to the amplification of correct or incorrect cascading information.

Social drift–diffusion model

Models of the evidence accumulation process during decision-making include the drift-diffusion model (DDM) (19, 26), the linear ballistic accumulator model (27), and the leaky competing accumulator model (20). Most of these models can, in principle, be extended to model a social system. Here, we focus on the DDM, arguably the most successful framework for accounting for human choice behavior, including some of the most basic aspects of the decision process, such as the speed-accuracy trade-off (21, 28), the construction of preferences (23), the formation of confidence judgments (17), the emergence of

Copyright © 2020
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

¹Max Planck Institute for Human Development, Center for Adaptive Rationality, Lentzeallee 94, 14195 Berlin, Germany. ²Department of Psychology, University of Kansas, Jayhawk Blvd., Lawrence, KS 66045, USA.

*Corresponding author. Email: tump@mpib-berlin.mpg.de

response biases (29, 30), and how attention guides the evidence accumulation process (31).

According to the DDM, people faced with a choice between two options, A or B, base their choice on an internal level of evidence. Initially, people can have a prior tendency and lean toward either option. This is modeled as an initial level of evidence. Over time, people extract further information about the options and accumulate this information as evidence. This accumulation gives rise to an evolving (latent) level of evidence, as depicted by the jagged line in Fig. 1A. The jaggedness arises because each sample of evidence is noisy (i.e., the stimuli itself and the cognitive and neural processes introduce variability into the evidence accumulation). Once a choice threshold has been reached, a decision is made. If the accumulated evidence reaches the upper threshold, option A is selected; if it crosses the lower threshold, option B is selected. The time it takes for the evidence to reach either threshold is the predicted RT. In the social DDM, we modify this framework to cover multiple individuals accumulating evidence at the same time (Fig. 1A). In this case, the evidence comes from two sources: personal information, gathered

from sampling the physical environment (e.g., for visual or auditory cues), and social information, gathered by observing the behavior of others (32).

Formally, we denote the cumulative evidence at time point t as $L(t)$. At the start, individuals may favor one option over the other, described by their relative start point. Often, the start point represents initial bias toward either option [e.g., (28)]. Here, in part because of the experimental design, the relative start point β is determined by people's initial evidence and expressed by their choice and confidence from an initial stage of the decision task (see Table 1). Such an approach is consistent with current models of confidence (17, 33). At each time step Δt , the current state of evidence $L(t)$ is updated by sampling new evidence until a decision is made (i.e., until the level of evidence reaches a choice threshold θ)

$$L(t + \Delta t) = L(t) + [\delta_p + \delta_s(t)] \times \Delta t + \sqrt{\Delta t} \times \epsilon \quad (1)$$

where ϵ is Gaussian white noise (i.e., the diffusion process) with a mean of 0 and a variance of 1. The parameters δ_p and $\delta_s(t)$ correspond




Fig. 1. Illustration of the social DDM and the experimental paradigm. (A) A generic example of the social DDM with five individuals, each represented by a jagged line. The start point of each individual indicates the personal evidence accumulated up to that point. At the start, no individual exceeds the choice threshold, and social information is absent, implying no social drift (as indicated by the horizontal arrow). Individuals who begin close to either of the thresholds (red lines) are likely to choose early, providing social information for undecided individuals. This social information affects the rate of evidence accumulation, with the drift rate shifting toward the choice threshold favored by the majority (as indicated by the arrow pointing upward). (B) Stages of the predator detection task. During the personal phase, individuals briefly observe a grid of “sharks” and “tuna.” They then make a personal decision whether to “stay” or “escape” and report their confidence in that decision. In the subsequent social phase, they are asked to make a second decision on whether to “stay” or “escape,” but now, they can freely time their decisions and simultaneously observe the choices of others before doing so. Last, the correct answer is displayed, and the next trial begins (with 40 trials in total).

Table 1. Description of the parameters of the social DDM.

Model feature	Parameter	Description
Nondecision time	τ	Response latency (e.g., motor RT). The parameter τ describes the time relative to the individual's fastest response.
Relative start point	$\beta = \frac{1}{1 + e^{-a(C-b)}}$	The relative start point is a function of the confidence in the personal choice C , which ranges from highly confident but incorrect to highly confident and correct (Fig. 4B). Parameter a determines how sensitive the start point is to changes in confidence; b captures other factors besides confidence in the personal decision that affect the start point.
Personal drift rate	δ_p	The average rate of evidence accumulation supporting the personal choice (Fig. 4C).
Social drift rate	$\delta_s(t) = s \times M(t)^q$	The social drift rate describes the impact of social information, with s being a scaling parameter that influences the strength of the social drift rate and q being a parameter that shapes the power function describing the relationship of majority size $M(t)$ and social drift rate (Fig. 4D).
Choice threshold	θ	The amount of evidence an individual has to accumulate to make a decision; θ ($-\theta$) reflects the correct (incorrect) choice threshold (Fig. 4E).

to the strength of the personal and social information uptake, respectively. Personal information uptake describes the integration of information extracted directly from the physical environment, as well as the evaluation of information from memory. Social information is defined as the majority size $M(t)$ of the individuals who have decided at time point t [see also (2, 3)]

$$M(t) = N_A(t) - N_B(t) \quad (2)$$

where $N_A(t)$ and $N_B(t)$ are the number of individuals who have already decided for option A or B, respectively. As the impact of majority size on social information use has formerly been successfully described by a power law [see also (34)], we model its impact on the social drift rate with a power function

$$\delta_s(t) = s \times M(t)^q \quad (3)$$

The parameter s is a scaling factor that influences the strength of the social drift; q governs the shape of the power function. When $q = 1$, each additional choice for the majority option has the same influence on the social drift rate (i.e., a linear effect); when $q > 1$ ($q < 1$), each

additional choice for the majority option has an increasingly stronger (weaker) impact on the social drift rate. Note that, in contrast to the individual drift rate, the social drift rate can vary over time (indicated by the changing direction of the arrows in Fig. 1A). By incorporating a social drift into the classical DDM, the social DDM can account for individuals being emitters and receivers of social information and thereby capture the dynamic information exchange among group members.

In sum, the social DDM characterizes (i) how individuals incorporate personal information with the parameters β and δ_p , (ii) how individuals incorporate social information depending on the majority size via the parameters s and q , and (iii) individuals' willingness to wait for social information with the parameter θ (see Table 1 for all parameter descriptions).

Predator detection task

We tested the social DDM in an empirical study (see Fig. 1B; see Methods for full details). In brief, participants were divided into groups of varying sizes ("small": $n = 3$ individuals; "medium": $n = 7$ to 10 individuals; or "large": $n = 15$ to 17 individuals). Each group of participants was seated together in a single room, facing a large screen. Participants were asked to imagine being a fish in a school facing a choice between two alternatives—namely whether to escape or stay—depending on the presence of predators, in this case, sharks. They were instructed to escape when five or more sharks were present and to stay when four or fewer sharks were present. At each trial, participants were shown—for 2 s—a grid with a varying number (3, 4, 6, or 7) of sharks hidden among harmless fish. Participants first made a personal choice on whether to "stay" or "escape" and then reported their confidence in that choice on a scale from 50 to 100%. They then entered the social phase, in which they had a maximum of 20 s to make a second decision on whether to "stay" or "escape," but without seeing the grid again. Instead, the display showed a count of the number of choices for each option. Participants were free to enter their choice at any point in time; they could thus respond early (thereby providing social information) or wait to observe the decisions of others. However, they could only decide once. Last, we provided feedback on the correct choice.

RESULTS

Empirical results: Groups show beneficial self-organization according to information quality

Participants achieved an accuracy of 74% in their personal choice (Fig. 2A), and participants reporting higher confidence in their personal choice were also more accurate [$\beta = 3.82$, credible interval (CI) = 3.35 to 4.28; Fig. 2B]. Participants were thus—at least partly—aware of the quality of their personal information. We fitted a two-stage dynamic signal detection (2DSD) model (17) to the choice, RT, and confidence data from the personal phase (see Supplementary Results and Discussion). The close correspondence between the model and the data suggests that a drift-diffusion process is a good description of the decision process during this stage of the experiment.

With an average accuracy of 79%, participants' choices during the social phase, where they had the opportunity to wait for social information before choosing again, were more accurate ($\beta = 0.3$, CI = 0.20 to 0.39; Fig. 2A). The reported level of confidence in their personal choice predicted their likelihood to improve ($\beta = -4.27$, CI = -4.88 to -3.68 ; Fig. 3A): participants reporting the lowest




Fig. 2. Choice accuracy and the relationship between personal accuracy and confidence. (A) Accuracy of the personal and social choices. Individuals, on average, achieved a higher decision accuracy during the social choice as compared to the personal choice. Each line connects a participant's average accuracy during the personal and social choice ($n = 141$ participants). (B) Participants reporting a higher confidence in their personal choice were more likely to be correct in their personal choice. The points and error bars reflect the mean and the 95% CIs of the posterior distribution from the Bayesian logistic regression model.

confidence level improved in more than 15% of trials; whereas the most confident, in less than 1% of trials. Why do unconfident participants achieve such higher gains from the social process? There appear to be two mechanisms that underly this. First, participants reporting lower confidence waited longer before making a decision during the social phase ($\beta = -4.86$, CI = -5.22 to -4.5 ; Fig. 3B). Second, participants partly adopted the decisions of others ($\beta = 0.62$, CI = 0.57 to 0.67 ; Fig. 3C): the larger the majority for the opposing option, the more likely participants were to change their decision. Individuals rarely changed their minds if the majority agreed with their personal decision. As fig. S1 shows, participants followed both correct and incorrect majorities, highlighting the importance of the accuracy of early-deciding participants for triggering positive/negative information cascades. Figure 3D shows the consequences of these patterns: Participants whose personal choices were accurate (and confident) tended to respond early in the social phase, whereas those whose choices were inaccurate (and unconfident) tended to wait longer, as illustrated by the downward trend of the blue dots (slope: $\beta = -0.16$, CI = -0.18 to -0.14). The latter participants increased their accuracy during the social phase through social influence, as illustrated by the higher yellow dots compared to the corresponding blue dots at higher RTs (interaction: $\beta = 0.11$, CI = 0.09 to 0.13).

Participants in groups, thus, appeared to self-organize according to information quality, with confident and accurate participants deciding early, thereby providing high-quality information for the less confident and less accurate participants who decided later. This beneficial self-organization depended on two crucial aspects: (i) a positive relationship between confidence and accuracy of personal choice across group members and (ii) a negative relationship between confidence and RT during the social choice phase. As Fig. 3E illustrates, groups showed the highest improvement when both conditions were met, and this occurred in most of the trials. Improvement was credibly lower for all other conditions (table S1).

Model results: The cognitive mechanisms driving self-organization

To understand the processes leading to the self-organization of groups, we need to understand the cognitive mechanisms underlying in-

dividuals' dynamic integration of personal and social information over time. To this end, we developed the social DDM (Fig. 1A and Table 1), which allowed us to test competing hypotheses on how participants integrate personal and social information over time. We examined three model features: (i) Individuals base their start point on their personal decision and reported confidence. (ii) When participants start drifting, they drift toward the correct option, their initially chosen option, or neither of the two. (iii) When social information becomes available, participants drift toward the option favored by the majority. We tested several candidate models composed of various combinations of these three features and used the deviance information criterion [DIC; (35)] to compare their performance where lower values indicate a better performing model. Figure 4A shows the models' DIC values relative to that of the best model (see also table S2). In the following, we present the results of the model with the lowest DIC (see table S3 for parameter estimates). Last, to test how the cognitive mechanisms were affected by group size, we compared the different group sizes (see table S4).

Individuals incorporate personal information via start point and drift rate

Participants incorporated their personal information (i.e., personal choice and confidence) during the social decision process in two distinct ways. First, consistent with current models of choice and confidence judgments (17, 33), they shifted their start point toward their initially chosen option: Individuals who reported higher confidence in the (in)correct option started closer to the threshold of the (in)correct option (small: $a = 4.20$, CI = 3.11 to 5.35 ; medium: $a = 3.42$, CI = 2.81 to 4.07 ; large: $a = 3.90$, CI = 3.46 to 4.37 ; Fig. 4B). This implies that individuals with high confidence in their personal choice were more likely to decide in favor of this option and to do so fast. Second, participants drifted toward the threshold of their initially chosen option (small: $\delta_p = 0.65$, CI = 0.45 to 0.86 ; medium: $\delta_p = 0.62$, CI = 0.50 to 0.75 ; large: $\delta_p = 0.53$, CI = 0.47 to 0.59 ; Fig. 4C). Both processes were independent of group size (table S4). In sum, across all group sizes, highly confident participants started close to the choice threshold of their initially chosen option and, on top of that, drifted toward that option, whereas participants with low confidence started out unbiased (i.e., in the middle between the thresholds).

Individuals incorporate social information via drift rate

We found that the drift rates were credibly influenced by the majority (Fig. 4D). The larger the majority favoring an option, the more strongly participants drifted toward that option. The shape of the relationship between majority size and social drift rate (the q parameter) differed between group sizes (small versus medium: $q = 0.82$, CI = 0.22 to 1.44 ; medium versus large: $q = 0.27$, CI = 0.14 to 0.41). In small groups, the drift rate increased with increasing majority size following a convex shape. In larger groups, each additional individual voting for the majority had less impact than the preceding one, and this function followed a concave shape. Accordingly, the influence of a single individual was larger in small groups than in large groups. Comparing the strength of the personal drift rate (i.e., toward the choice threshold of the initially chosen option) to the social drift rate (i.e., toward the option favored by the majority) showed that a majority of approximately two is required to counteract an individual's tendency to drift toward the choice threshold reflecting their initial choice. This highlights participants' tendency to give personal information more weight than social information. Corroborating this finding, Fig. 3C shows that a majority of approximately four participants in favor of the opposing option is required to induce a 50%

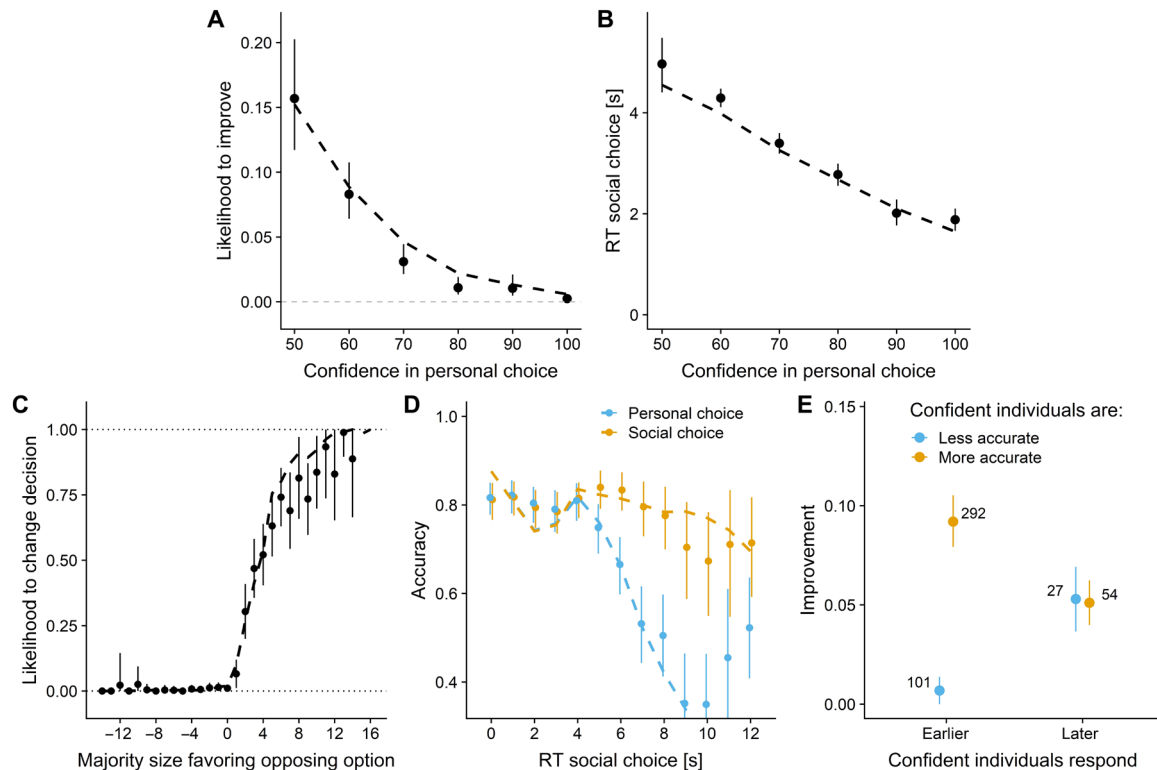


Fig. 3. Empirical results and predictions of the social DDM. Participants reporting higher confidence in their personal choice (A) improved less and (B) responded earlier during the social choice. (C) The larger the majority favoring the opposing option, the more likely participants were to change their decision. (D) The choices of participants who responded later in the social choice were less accurate in the personal choice (declining blue dots) but improved more in the social choice (indicated by the increasing difference between blue and yellow dots at later RTs). For visualization purposes, RTs are binned by rounding to the closest integer. RTs greater than 13 s (less than 1%) were assigned to the 12-s bin. (A to D) The dashed lines show the choices and RTs predicted by the social DDM, accurately capturing all relationships. For frequency distributions, see fig. S2. (E) Participants improved most when more confident individuals were more accurate (i.e., positive confidence-accuracy correlation; yellow dots) and responded earlier (i.e., negative accuracy-RT correlation; see Methods for details). Numbers indicate the number of trials. For all panels, the points and error bars depict the mean and the 95% CIs of the posterior distribution of the Bayesian regression model.

likelihood of changing a participant's decision. Last, we found that participants' willingness to wait for social information, captured by the threshold parameter θ , did not differ between group sizes (Fig. 4E).

Model predictions: The social DDM captures the self-organizing dynamics

The model described above was able to recover all the key features of the dynamics of the social decision-making process. The dashed lines in Fig. 3 show the model predictions of the social DDM. In line with the empirical data, the social DDM predicts that unconfident participants wait longer before making a decision (Fig. 3B), that individuals are increasingly likely to follow the majority as the size of that majority increases (Fig. 3C), and that participants whose personal choices were inaccurate wait longer and improve more during the social phase (Fig. 3D). As a result, participants with low confidence in their personal choice improved most (Fig. 3A). We investigated the validity of the model with a parameter recovery analysis (see Supplementary Information). For all parameters, the generating and recovered parameters were highly correlated, implying that each parameter describes a distinct mechanism. Furthermore, all recovered parameter estimates were close to the generating parameters, affirming the validity of the magnitude of the parameter estimates as captured by the social DDM (fig. S3).

Model simulations: The mechanisms driving positive/negative cascades

Next, we used the social DDM and its ability to capture the integration of personal and social information to investigate what psychological factors shape the quality of information cascades. To do so, we implemented the social DDM within an agent-based simulation and simulated a group of 10 agents who, at the start, hold personal information by assigning them correct or wrong choices and associated confidences. They then complete the social phase formalized by the social DDM (see Eqs. 1 to 3 and Methods for full details). To create agents whose behavior resembles the observed behavior of the participants, the agents' behavior was parameterized using the mean posterior estimates of medium-sized groups. We systematically varied three key features. First, as illustrated above, the beneficial self-organization is expected to be strongly influenced by the relationship between confidence and personal accuracy; hence, we systematically varied the confidence-accuracy relationship by assigning either higher, similar, or lower confidence to agents with a correct choice compared to agents with a wrong choice. Second, it is well known that the level of collective improvement depends on the average mean personal accuracy (2); therefore, we varied the agents' personal accuracy by assigning 30, 50, or 70% correct personal choices to agents. Third, a key feature of the social DDM is the choice threshold

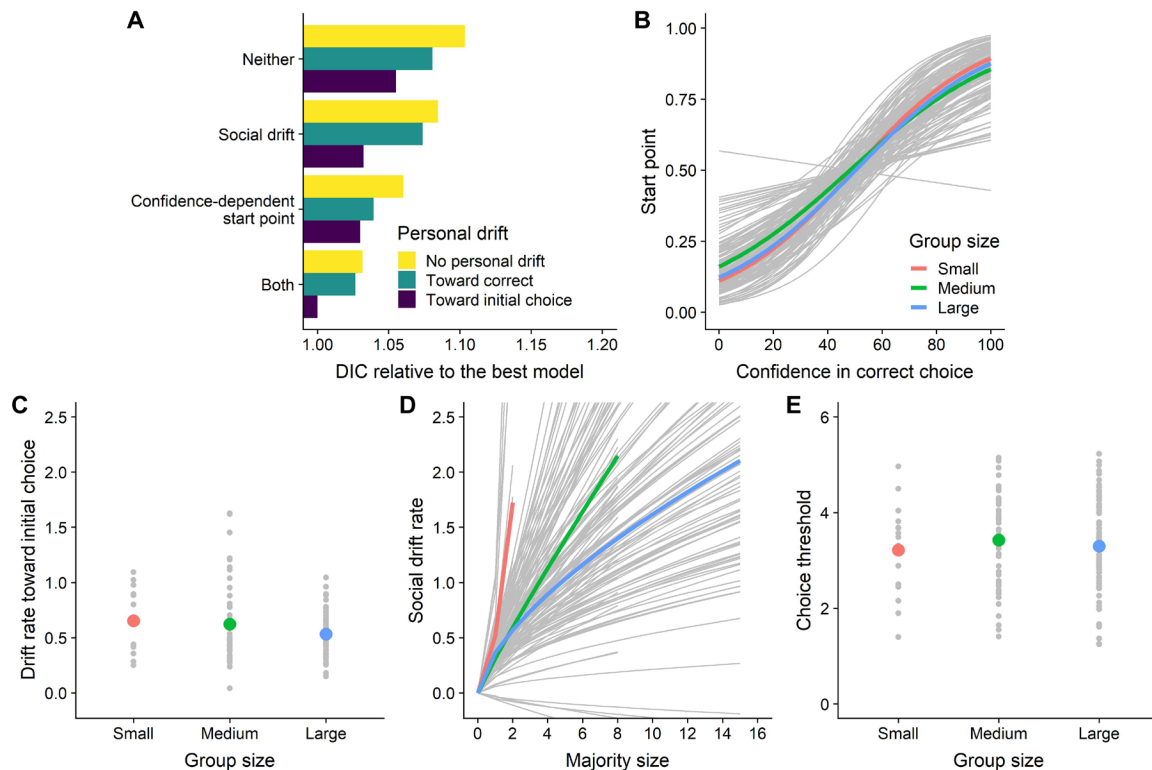


Fig. 4. Model comparison and individual- and group-level fittings of the social DDM for different group sizes. (A) The DIC values of all models relative to the model with the lowest DIC. The model with the lowest DIC (i.e., preferred model) features a (i) confidence-dependent start point, (ii) drift toward the initially chosen option, and (iii) social drift. (B) Participants reporting higher confidence in the correct/incorrect choice started closer to the correct/incorrect choice threshold at y value of 1/0. (C) Evidence tended to drift toward the choice threshold of the option chosen during the personal phase. (D) The larger the majority favoring an option, the more strongly participants drifted toward the choice threshold favored by the majority. Participants in smaller groups had a stronger drift given the same majority size. (E) The choice threshold θ , reflecting a participant's willingness to wait for social information, did not differ between group sizes. Gray lines/dots represent individual-level fittings; colored lines/dots represent the estimates on a group size level. Group size ranged from small ($n = 3$) to medium ($n = 7$ to 10), to large ($n = 15$ to 17).

(θ) as this governs the willingness to wait for social information. We varied the willingness to wait for social information by assigning a low, intermediate, or high choice threshold to all group members. Figure 5 shows the results of these simulations.

Across all levels of personal accuracy, we observed the key importance of the confidence-accuracy relationship: Whenever this relationship is positive (yellow dots), we observed an improvement. The strength of improvement is, in turn, mediated by the choice threshold: Higher choice thresholds allowed agents to improve more. This can be explained because high choice thresholds allow for a better match between confidence and decision time (i.e., beneficial self-organization). Mirroring these results, we found detrimental effects when confident individuals were less accurate (blue dots), and this effect was stronger the higher the choice threshold. That is, perhaps counterintuitively when agents were more cautious, they were also more susceptible to bad social information creating negative information cascades. When there was no relationship between confidence and accuracy, we observed beneficial (detrimental) effects when personal accuracy was above (below) chance, and again, these effects were strongest at high choice threshold. These simulations illustrate the importance of a positive confidence-accuracy relationship, an overall willingness to wait for social information and high personal accuracy, for the emergence of positive information cascades and, hence, wise crowds.

DISCUSSION

We have shown that the behavior of individuals in a social sequential decision-making task can be described by an evidence accumulation process whereby personal and social information is integrated until a decision is made, formalized by the social DDM. The model accurately predicts decision time and choice by taking personal information, social information, and the willingness to wait for social information into account. It successfully captured all the interrelationships of the key behavioral results of the social phase, thereby revealing the cognitive underpinnings of the group-level self-organization according to information quality. Measuring how individuals process personal and social information affords a deeper understanding of how individuals in a social environment cope with the complex problem of evaluating personal information and how they time their decision and incorporate social information.

During the social decision-making process, individuals incorporated personal information in two ways: At the start of the process, they adjusted their subjective level of evidence to their confidence (i.e., they adjusted their start point), and during the process, they reinforced their “belief” in their original choice over time (i.e., they drifted toward the choice threshold of their personal choice). We also found evidence for such “belief reinforcement” over time in the personal phase (see 2DSD model analysis in Supplementary Information). The reinforcement of initial beliefs can potentially have a

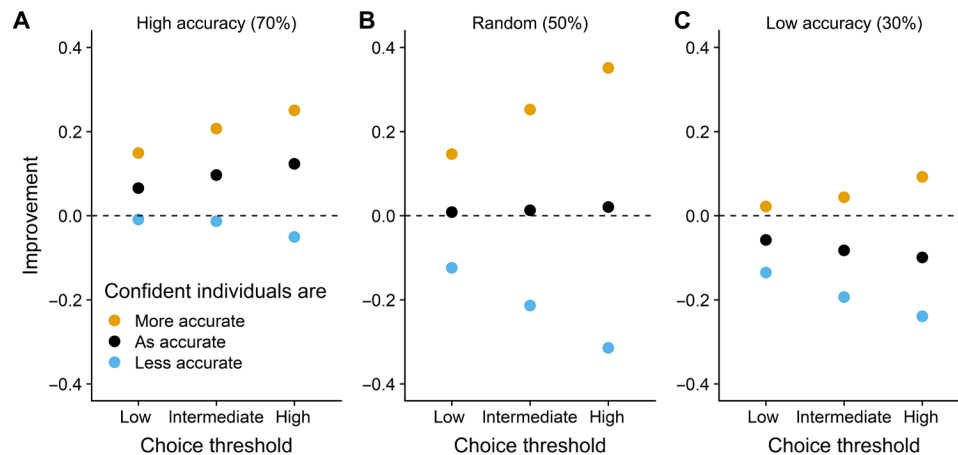


Fig. 5. Agent-based simulations of the social DDM. The predicted improvement for different choice thresholds, for situations in which confident agents are more accurate (yellow dots), as accurate (black dots), or less accurate (blue dots) than unconfident agents. **(A)** Groups with a high average personal accuracy improved, unless confident agents were less accurate than unconfident ones. **(B and C)** Groups with a personal accuracy of 50 and 30% only improved when confident agents were more accurate than unconfident ones. At all levels of personal accuracy, a higher choice threshold strengthened the positive (negative) impact of confident individuals being more (less) accurate.

large influence in real-world social choices. Because individuals generally gather personal information before receiving social information, reinforcement of initial beliefs can lead to situations where even strong counterfactual social information may no longer prove persuasive [i.e., confirmation bias (36, 37)]. Many studies have found that individuals indeed weight personal information more strongly than social information, a phenomenon called egocentric discounting [e.g., (38–41)]. In almost all of these studies, participants made a personal judgment before receiving social information. When the order was reversed, the influence of social information indeed increased (42). Our finding of belief reinforcement provides a compelling explanation for stronger egocentric discounting, simply by providing personal information first. Future studies could test whether increasing the length of the delay between personal choice and provision of social information reduces the influence of social information, as predicted by the social DDM.

When looking at how social information entered the evidence accumulation process, we found that individuals incorporated social information by drifting toward the choice threshold favored by the majority. The larger the majority size, the more strongly individuals drifted toward that majority choice. For medium- and large-sized groups, the relationship followed a concave power function, where each additional individual voting for the majority choice had less additional impact on the drift rate. Such saturating influence is consistent with the findings of earlier studies (34, 43, 44). In groups of three, the relationship followed a convex function. Weighting single choices less with increasing group size is probably an adaptive strategy: In larger groups, waiting for further decisions avoids confirming fast but wrong choices, as others can still correct initial mistakes. In small groups, fast but wrong choices will also occur, but because there are few others to correct those choices, there is little point in delaying a response via a reduced social drift rate. Future studies could use the social DDM to investigate more complex strategies of social information integration. We used a power function to model social information uptake. There is empirical evidence coming from psychophysics that the relationship between stimulus intensity and perceived intensity often follows a power law [Steven's power law

(45)], and this has also been reported for social information use (34, 43). However, individuals may also use the so-called quorum responses (12, 46), which down-weight small minorities, but ramp up social information use once a critical threshold of social information is reached. Other potential strategies include weighting choices according to others' RTs (3) or strong conformity in the presence of unanimity among others. Another crucial aspect shaping the characteristics of information cascades is the amount of information conveyed in the choices of predecessors (47). Only detailed understanding of the applied strategies allows us to identify choices that are expected to be mostly imitative and, hence, uninformative or which ones are mostly based on personal information and, hence, informative for later-deciding individuals. We believe that our framework provides a tool to test these competing strategies of social information use in social dynamic decision-making environments.

The social DDM can also characterize other features of the dynamics of the social decision-making process. Beyond capturing how social information affects the accumulation of evidence, it also captures an individual's willingness to wait for social information via the decision threshold parameter θ . Thus, the model is able to distinguish, for instance, between individuals who are sensitive to majorities but unwilling to wait for social information and individuals who may be interested in observing the decisions of others but put more weight on their own personal information. The capacity to unify these different facets of social decision-making within a single theoretical framework is a long-standing goal of social decision-making in the areas of collective animal behavior (10, 12) and social psychology (34). Future studies could investigate the interrelationships between the different parameters and potential links to established personality measures.

Previous studies have provided evidence for both positive information cascades, such as knowledgeable individuals leading others to resources or safety (48–50), and negative ones, such as the spread of fake news, mobbing, or stampedes (2, 51). Here, we have shown the importance of two key aspects promoting positive information cascades. First is a positive confidence-accuracy relationship across group members. In many contexts, confidence is a valid cue for accuracy

(52, 53). The strongest association of confidence and accuracy across group members arises when all individuals are more confident when they are more accurate and when their confidence scales are well aligned [i.e., a given level of confidence implies the same level of accuracy across individuals; see also (54)]. The second key aspect promoting positive information cascades is a negative relationship between confidence and RT, meaning that more confident individuals respond faster. Our simulations demonstrate the importance of the overall willingness to wait for social information for groups to coordinate their responses according to their confidences and, hence, improve. Several further mechanisms in the social DDM are expected to influence this relationship—for example, how individuals adjust their start point depending on their confidence. If confident individuals do not start closer to a choice threshold, they are not expected to respond earlier. Also, interindividual differences in model parameters such as choice thresholds or personal and social drift rate could negatively affect the confidence-RT relationship.

The quality of information cascades is shaped by the relationship between accuracy and RT, whereby it is crucial for positive information cascades that accurate individuals respond faster than inaccurate individuals. The social DDM framework allows us to predict the quality of information cascades based on individual or task characteristics. For example, if individuals differ in their ability to solve a task (e.g., individual differences in drift rates), those with higher ability are expected to make faster, more accurate decisions than the less competent ones, triggering positive information cascades. In contrast, when individuals differ systematically in their speed-accuracy trade-off [e.g., differences in threshold separation; (55)], and groups harbor both fast, but inaccurate, individuals and slow, but accurate, individuals, we expect relatively many fast errors, triggering negative information cascades.

The social DDM not only allows to identify cognitive underpinnings of the sequential decision process. By providing a descriptive model with well-informed model characteristics via the measured parameter estimates, it also allows to conduct agent-based simulations with processes closely resembling the participants' behavior. Here, we used simulations to investigate how personal accuracy, confidence-accuracy correlations, and the decision threshold are expected to influence the groups' ability to improve. This framework can be used in future studies to investigate other model features to generate further predictions that can, in turn, be tested in empirical studies.

Because the DDM has been successful in accounting for behavioral phenomena across a wide range of tasks, our extension to social environments opens up new possibilities for studying a range of social and collective phenomena. It makes it possible to measure how individuals combine personal and social information and time their decisions whenever decisions are made sequentially and the choices are—at least partially—observable by others. We hope that future work will apply and extend the social DDM to areas such as dynamics in consumer preferences (7), emergency evacuations (56), and social media (4) or to areas of animal social and collective behavior such as predator detection and mate choice (57).

METHODS

Experimental procedure

Participants were 141 students from the Wageningen University (The Netherlands) and the University of Bielefeld (Germany). Participants were divided into 16 groups, with group size ranging from

small ($n = 3$ individuals, 5 groups) to medium ($n = 7$ to 10 individuals, 6 groups), to large ($n = 15$ to 17 individuals, 5 groups). We aimed for a similar number of groups per group size (see also table S5). This study was approved by the Institutional Review Board of the Max Planck Institute for Human Development. Before participation, each participant signed an informed consent form. Each group of individuals was seated on chairs facing a large screen. They were confronted with the following binary decision task: Individuals briefly (for 2 s) observed an image of a shoal of 72 stylized fish (tuna and sharks aligned in an 8 by 9 grid; see Fig. 1B). Participants were instructed to choose “escape” if there were five or more sharks and “stay” if there were four or fewer. The number of sharks present was three, four, six, or seven, and each number was repeated 10 times, resulting in 40 trials. Treatment order was randomized. After observing a shoal of fish, individuals had 5 s to report their personal decision and an additional 5 s to report their confidence in their personal decision. Participants were instructed to use confidence as the subjective probability of being correct on a scale from 50 to 100%. In the subsequent social phase, participants made a second decision on the same image. During this phase, they received social information in the form of the number of group members who had already decided on a particular option, displayed on the screen. The social information was first updated after 3 s and then iteratively every 2 s (i.e., at 3, 5, 7, 9, ..., 19 s). The social phase lasted 20 s. A countdown timer on the screen indicated the remaining choice time. Participants made all decisions using a wireless keypad. Afterward, we provided feedback on the correct choice. Participants received 0 points for an incorrect decision and 100 points for a correct decision. To avoid a scenario in which all participants waited until the last second for social information, we introduced a small cost of one point per second for correct decisions during the social phase. The members of each group with the highest payoff got a small reward in kind. Before the 40 study trials, participants completed two test trials to familiarize themselves with the procedure. These results were excluded from the analyses.

Statistical analysis

We used Bayesian hierarchical generalized linear models with the “brms” package (58) to analyze the empirical data. The parameter estimates were generated by running five Markov chain Monte Carlo (MCMC) simulations in parallel with 5000 iterations, of which the first 2500 were discarded as burn-in to reduce autocorrelations. To analyze the difference in the accuracy of personal and social choices (Fig. 2A), we fitted choice correct (yes/no) as a binary response variable and type of choice (personal/social) as a population-level effect (i.e., fixed effect). In this model (and all following models, unless stated otherwise), we included individual and group identity as group-level effects (i.e., mixed effects). We ran separate models to investigate how confidence related to (i) personal accuracy (Fig. 2B), (ii) likelihood to improve (Fig. 3A), and (iii) RT during social choice (Fig. 3B). “Personal accuracy” (correct/incorrect) and “likelihood to improve” (yes/no) were fitted as binomial response variables and “RT during social choice” as an exponentially modified Gaussian (ex-Gaussian) distributed response variable. Confidence was included as a population-level effect in all three models. To investigate whether the majority size affected the likelihood of an individual changing its decision (Fig. 3E), we fitted the likelihood to change the decision as a binary response variable (yes/no) and majority size favoring the opposing option as a population-level effect. To analyze the relationship between RT in the social phase and accuracy of personal

and social choices (Fig. 3D), we used decision correct (yes/no) as a binary response variable and type of choice (personal/social) in interaction with RT as a population-level effect.

To investigate how the interrelationships between confidence, accuracy, and RT affected improvement (Fig. 3E), we first calculated—for each group and trial—the Spearman’s correlation coefficients of confidence and accuracy as well as of confidence and RT. We converted these coefficients into dichotomous variables, with the correlation coefficient being either 0 and above or below 0. We excluded trials in which all individuals reported identical choices or confidences, because it was impossible to calculate correlation coefficients for these. We treated all four possible combinations of correlations as different levels of a single factor. We included the factor as a population-level effect and improvement as response variable. In this model, group identity was the only group-level effect. As statistical summary, we report the mean of the posterior distributions and the 95% CIs (see table S1 for the results of the regression models). To visualize the results (Figs. 2 and 3), while accounting for the hierarchical structure of the data, we reran the regression models, treating the continuous variables as categorical data. Unless stated otherwise, the points and error bars reflect the mean and the 95% CI of the posterior distribution. Visual inspection of the Markov chains and the Gelman-Rubin statistic (\hat{R}) indicated that all Markov chains converged.

Social DDM: Model parameter estimation

To understand the dynamics of the social phase, we developed the social DDM (Fig. 1A and Table 1). The model features decisions with variable drift rates to obtain choice and RT predictions. We calculated the probability density function of RTs and associated choice probabilities of the drift-diffusion process by implementing an extended version of a Markov chain approach (59). A detailed description of how to implement the Markov chain approach can be found in (59).

The model assumes that the state space of the decision maker’s evidence L is ranging from the lower choice threshold $-\theta$ (reflecting the wrong decision) to the upper threshold θ (reflecting the correct decision), with a step size of Δ and k being the number of steps to reach the choice threshold from a neutral start point

$$L = [-k\Delta, -(k - 1)\Delta, \dots, -\Delta, 0, \Delta, \dots, (k - 1)\Delta, k\Delta] \quad (4)$$

where $\theta = k\Delta$.

Each time step h , the evidence states change with probabilities given by a $m \times m$ transition probability matrix P , with $m = \frac{2 \times \theta}{\Delta} + 1$. The elements $p_{1,1} = 1$ and $p_{m,m} = 1$ are the two absorbing states and reflect the choice thresholds. The other elements of P with $1 < i < m$ are

$$p_{i,j} = \begin{cases} \frac{1}{2\alpha} \left(1 - \frac{u}{\sigma^2} \sqrt{h} \right) & \text{if } j = i - 1 \\ \frac{1}{2\alpha} \left(1 + \frac{u}{\sigma^2} \sqrt{h} \right) & \text{if } j = i + 1 \\ 1 - \frac{1}{\alpha} & \text{if } j = i \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

with σ^2 being the diffusion coefficient and $u = \delta_p + \delta_s$ being the total drift rate, whereby δ_p and δ_s are drift rates reflecting the accumulation of personal and social information, respectively (see Table 1).

The parameter $\alpha > 1$ improves the approximation of the continuous time process. We set $\alpha = 1.3$, $\sigma^2 = 1$, and $h = 0.005$. The transition probability matrix in its canonical form

$$P = \begin{bmatrix} P_I & 0 \\ R & Q \end{bmatrix} = \begin{matrix} & \begin{matrix} 1 & m & & 2 & 3 & \dots & m-2 & m-1 \end{matrix} \\ \begin{matrix} 1 \\ m \end{matrix} & \left(\begin{array}{c|cccccc} 1 & 0 & & 0 & 0 & \dots & 0 & 0 \\ 0 & 1 & & 0 & 0 & \dots & 0 & 0 \end{array} \right) \\ \begin{matrix} 2 \\ 3 \\ 4 \\ \vdots \\ m-3 \\ m-2 \\ m-1 \end{matrix} & \left(\begin{array}{c|cccccc} p_{12} & 0 & & p_{22} & p_{23} & \dots & 0 & 0 \\ 0 & 0 & & p_{32} & p_{33} & \dots & 0 & 0 \\ 0 & 0 & & 0 & p_{43} & \dots & 0 & 0 \\ \vdots & \vdots & & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & & 0 & 0 & \dots & p_{m-3,m-2} & 0 \\ 0 & 0 & & 0 & 0 & \dots & p_{m-2,m-2} & p_{m-3,m-1} \\ 0 & p_{m-1,m} & & 0 & 0 & \dots & p_{m-1,m-2} & p_{m-1,m-1} \end{array} \right) \end{matrix} \quad (6)$$

with P_I being a 2×2 matrix with the two absorbing states and R a $(m - 2) \times 2$ matrix containing the transition probabilities that eventually lead to the absorbing states in a single transition. Q is a $(m - 2) \times (m - 2)$ matrix including the remaining transition probabilities. The initial evidence state of the process is represented by Z an $m - 2$ vector containing the initial probability distribution. The initial start point β is a function of confidence and choice and relative to the upper (correct) threshold

$$\beta = \frac{1}{1 + e^{-a(C-b)}} \quad (7)$$

with a and b being free parameters. C is the reported confidence in the correct choice and is scaled from 0 (i.e., highly confident and wrong) to 1 (i.e., highly confident and correct). We set the distribution of the initial evidence states by $Z_{\beta^*} = 1$, with $\beta^* = \beta(m - 3) + 1$. Because β^* is not always an integer, we avoid rounding errors by giving most probability mass $1 - (\beta^* - \text{round}(\beta^*))$ to $Z_{\text{round}(\beta^*)}$ and the rest $\beta^* - \text{round}(\beta^*)$ the closest integer of β^* . For example, if the process starts unbiased (i.e., $\beta = 0.5$) and $m = 7$, then $\beta^* = 3$ and $Z = [0,0,1,0,0]$. However, if $\beta = 0.55$, then $\beta^* = 3.2$ and $\beta^* - \text{round}(\beta^*) = 0.2$ and therefore $Z = [0,0,0.8,0.2,0]$. We account for variable drift rates by updating the transition probabilities of Q at $t = (3,5,7,9, \dots, 19)$ s, reflecting the iterative updated social information (see the “Experimental procedure” section). With Q_n containing the transition probabilities at time point $t = nh$, we can calculate the probability of choosing the correct or wrong option after n time steps

$$[P(\text{correct} | n), P(\text{wrong} | n)] = Z \times Q_1 \times Q_2 \times Q_3 \dots Q_n \times R - \tau \times t_{\min} \quad (8)$$

with τ being the nondesideration time relative to the fastest response of the individual t_{\min} . By varying the transition probabilities of Q_n with changing δ_s , we are able to account for varying social information over time.

Integrating the social DDM into a Bayesian estimation technique, namely a differential evolution–MCMC algorithm, enables us to sample posterior probability densities of the model parameters (see Table 1). The differential evolution–MCMC is an extension of the Metropolis-Hastings algorithm where proposals are generated by taking the Markov states of parallel computed chains into account

(60). To estimate the effect of group size while controlling for individual differences, we used a hierarchical framework. Each parameter was fitted on an individual level but was simultaneously informed by a higher-order group-level prior, a normal distribution described by two hyper parameters (i.e., mean and variance), which were informed by the individual fittings. To estimate the posterior probability densities, we ran 24 chains in parallel, each with a chain length of 20,000 including a burn-in period of 10,000 and a thinning factor of 10 to reduce autocorrelations. The tuning parameter (γ) was set to $= 2.38/\sqrt{2d}$, with d being the dimensionality of the posterior, which was $d = 2$ for the hyperparameters and $d = 7$ for the individual parameters [see (60)]. To further improve the mixing of the parallel chains, we included deterministic and probabilistic (i.e., relying on the Metropolis-Hastings probability) migration steps where chain states are swapped across parallel chains (60). We performed the deterministic migration step with a probability of 5% where we first determine a random number of $n = 2, 3, \dots, 24$ chains and then sample n chains without replacement. We then swap the parameter set in a cyclic fashion where the set of the first sampled chain moves to the second, the second to the third, and so on, until the last set moves to the first set. A deterministic migration step strongly improves the mixing behavior of chains but does not resolve the frequent problem of differential evolution–MCMC algorithms that outlier chains hardly converge. We, therefore, additionally implemented a probabilistic migration step, which was carried out with a probability of 10%. For the probabilistic version, we swapped proposal states instead of accepted states between chains, which therefore still relied on the Metropolis-Hastings probability to be accepted. Thereby, we sampled two parallel chains and interchanged a single random parameter state.

We used the social DDM to compare competing hypotheses on how individuals integrate personal and social information. More specifically, we examined three model features: (i) Individuals base their start point on their personal choice and reported confidence; (ii) individuals drift toward the correct option, their initially chosen option, or neither of the two; and (iii) individuals drift toward the option favored by the majority. We compared the performance of models composed of the various combinations of these three features using the DIC (35). To investigate the effect of group size on the collective dynamics, we categorized groups as small (3 individuals), medium (7 to 10 individuals), or large (15 to 17 individuals) and fitted the parameters separately for each group size. As a statistical summary, we report the mean of the posterior distributions and the 95% CI. We excluded all observations for which personal choice, social choice, confidence, or RT of social choice were missing (~8%) and if the RT was below 0.1 s (~6%). The code can be accessed at <https://osf.io/ejfm4/>.

Social DDM: Predictions

To analyze the predictions (i.e., choices and RTs) of the social DDM, we generated decisions by sampling from the probability density functions produced by the model using the mean of the individual-level posterior distribution as model estimates. The probability density function was computed for each individual and trial by taking into account the individual model estimates, the personal choice, the reported confidence, and the social information observed by the individual at a given trial. We then sampled 10 choices and RTs to account for stochasticity. The model predictions are shown as dashed lines in Fig. 3.

Social DDM: Simulations

To analyze how behavioral features affect the quality of information cascades, we conducted agent-based simulations of the social DDM using a group size of 10 agents (“medium” group size). First, agents start with a relative tendency toward an option β determined by the assigned personal choices and confidences. Agents then iteratively gather further information over time described by Eqs. 1 to 3 until they reach a choice threshold θ . If not stated otherwise, we used the mean of the posterior estimate of medium-sized groups to assign parameters to the agents (see table S3). We systematically varied three key parameters. (i) We varied the confidence-accuracy relationship by assigning different confidences to agents with correct or wrong personal choices. The confidences assigned to the correct/wrong agents were either 100%/50%, 75%/75%, or 50%/100%, reflecting correct agents being more, as, or less confident than wrong ones. (ii) We varied average personal accuracy by assigning different proportions of correct and wrong personal choices to the agents. The average accuracy ranged from 30% (minority correct) to 50% (random) to 70% (majority correct). (iii) Last, we varied the height of the choice threshold θ , ranging from 75% (low) to 100% (intermediate) to 125% (high) of the mean of the posterior estimate of medium-sized groups. The nonddecision time was set to 0 as it is relative to a participant’s fastest response and is irrelevant for simulated agents. For each combination of average personal accuracy, confidence-accuracy relationship, and threshold, we ran 10,000 repetitions and reported the average improvement of agents after the social phase.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/6/29/eabb0266/DC1>

[View/request a protocol for this paper from Bio-protocol.](#)

REFERENCES AND NOTES

1. L. R. Anderson, C. A. Holt, Information cascades in the laboratory. *Am. Econ. Rev.* **87**, 847–862 (1997).
2. S. Bikhchandani, D. Hirshleifer, I. Welch, Learning from the behavior of others: Conformity, fads, and informational cascades. *J. Eco. Perspect.* **12**, 151–170 (1998).
3. C. Frydman, I. Krajbich, Using Response Times to Infer Others’ Beliefs: An Application to Information Cascades (May 8, 2019). Available at SSRN: <https://ssrn.com/abstract=2817026>.
4. S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online. *Science* **359**, 1146–1151 (2018).
5. J. J. Faria, S. Krause, J. Krause, Collective behavior in road crossing pedestrians: The role of social information. *Behav. Ecol.* **21**, 1236–1242 (2010).
6. I. Welch, Herding among security analysts. *J. Financ. Econ.* **58**, 369–396 (2000).
7. Y.-F. Chen, Herd behavior in purchasing books online. *Comput. Hum. Behav.* **24**, 1977–1992 (2008).
8. M. Battaglini, Sequential voting with abstention. *Games Econ. Behav.* **51**, 445–463 (2005).
9. F. Xiong, Y. Liu, Opinion formation on social media: An empirical approach. *Chaos* **24**, 013130 (2014).
10. J.-L. Deneubourg, S. Aron, S. Goss, J. M. Pasteels, The self-organizing exploratory pattern of the Argentine ant. *J. Insect Behav.* **3**, 159–168 (1990).
11. J. K. Goeree, T. R. Palfrey, B. W. Rogers, R. D. McKelvey, Self-correcting information cascades. *Rev. Econ. Stud.* **74**, 733–762 (2007).
12. D. J. Sumpter, S. C. Pratt, Quorum responses and consensus decision making. *Phil. Trans. R. Soc. B* **364**, 743–753 (2008).
13. C. Chamley, D. Gale, Information revelation and strategic delay in a model of investment. *Econometrica* **62**, 1065–1085 (1994).
14. J. Zhang, Strategic delay and the onset of investment cascades. *Rand J. Econ.* **28**, 188–205 (1997).
15. A. Ziegelmeyer, K. B. My, J.-C. Vergnaud, M. Willinger, Strategic Delay and Rational Imitation in the Laboratory, in *Papers on Strategic Interaction* (Max Planck Institute of Economics, Strategic Interaction Group, 2005).

16. R. Hertwig, T. J. Pleskac, T. Pachur, The Center for Adaptive Rationality, *Taming uncertainty* (MIT Press, 2019).
17. T. J. Pleskac, J. R. Busemeyer, Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychol. Rev.* **117**, 864–901 (2010).
18. A. Tversky, D. Kahneman, Judgment under uncertainty: Heuristics and biases. *Science* **185**, 1124–1131 (1974).
19. R. Ratcliff, A theory of memory retrieval. *Psychol. Rev.* **85**, 59–108 (1978).
20. M. Usher, J. L. McClelland, The time course of perceptual choice: The leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592 (2001).
21. R. Ratcliff, P. L. Smith, A comparison of sequential sampling models for two-choice reaction time. *Psychol. Rev.* **111**, 333–367 (2004).
22. R. M. Nosofsky, T. J. Palmeri, An exemplar-based random walk model of speeded classification. *Psychol. Rev.* **104**, 266–300 (1997).
23. J. R. Busemeyer, J. T. Townsend, Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* **100**, 432–459 (1993).
24. M. Germar, A. Schlemmer, K. Krug, A. Voss, A. Mojzisch, Social influence and perceptual decision making: A diffusion model analysis. *Pers. Soc. Psychol. Bull.* **40**, 217–231 (2014).
25. U. Toelch, F. Panizza, H. R. Heekeren, Norm compliance affects perceptual decisions through modulation of a starting point bias. *R. Soc. Open Sci.* **5**, 171268 (2018).
26. R. Ratcliff, G. McKoon, The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922 (2008).
27. S. D. Brown, A. Heathcote, The simplest complete model of choice response time: Linear ballistic accumulation. *Cogn. Psychol.* **57**, 153–178 (2008).
28. A. Voss, K. Rothermund, J. Voss, Interpreting the parameters of the diffusion model: An empirical validation. *Mem. Cognit.* **32**, 1206–1220 (2004).
29. F. P. Leite, R. Ratcliff, What cognitive processes drive response biases? A diffusion model analysis. *Judgm. Decis. Mak.* **6**, 651–687 (2011).
30. T. J. Pleskac, J. Cesario, D. J. Johnson, How race affects evidence accumulation during the decision to shoot. *Psychon. Bull. Rev.* **25**, 1301–1330 (2018).
31. I. Krajbich, A. Rangel, Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 13852–13857 (2011).
32. B. G. Galef, K. N. Laland, Social learning in animals: Empirical studies and theoretical models. *BioScience* **55**, 489–499 (2005).
33. S. Yu, T. J. Pleskac, M. D. Zeigenfuse, Dynamics of postdecisional processing of confidence. *J. Exp. Psychol. Gen.* **144**, 489–510 (2015).
34. B. Latané, The psychology of social impact. *Am. Psychol.* **36**, 343–356 (1981).
35. D. J. Spiegelhalter, N. G. Best, B. P. Carlin, A. Van Der Linde, Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Series B Stat. Methodol.* **64**, 583–639 (2002).
36. J. Klayman, Varieties of confirmation bias. *Psychol. Learn. Motiv.* **32**, 385–418 (1995).
37. R. S. Nickerson, Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* **2**, 175–220 (1998).
38. B. Jayles, H.-R. Kim, R. Escobedo, S. Cezera, A. Blanchet, T. Kameda, C. Sire, G. Theraulaz, How social information can improve estimation accuracy in human groups. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 12620–12625 (2017).
39. R. P. Larrick, J. B. Soll, Intuitions about combining opinions: Misappreciation of the averaging principle. *Manage. Sci.* **52**, 111–127 (2006).
40. A. Novaes Tump, M. Wolf, J. Krause, R. H. J. M. Kurvers, Individuals fail to reap the collective benefits of diversity because of over-reliance on personal information. *J. R. Soc. Interface* **15**, 20180155 (2018).
41. I. Yaniv, E. Kleinberger, Advice taking in decision making: Egocentric discounting and reputation formation. *Organ. Behav. Hum. Decis. Process.* **83**, 260–281 (2000).
42. D. J. Koehler, T. A. Beaugregard, Illusion of confirmation from exposure to another's hypothesis. *J. Behav. Decis. Mak.* **19**, 61–78 (2006).
43. R. Bond, Group size and conformity. *Group Process. Intergroup Relat.* **8**, 331–354 (2005).
44. S. Milgram, L. Bickman, L. Berkowitz, Note on the drawing power of crowds of different size. *J. Pers. Soc. Psychol.* **13**, 79–82 (1969).
45. S. S. Stevens, On the psychophysical law. *Psychol. Rev.* **64**, 153–181 (1957).
46. R. H. J. M. Kurvers, M. Wolf, J. Krause, Humans use social information to adjust their quorum thresholds adaptively in a simulated predator detection experiment. *Behav. Ecol. Sociobiol.* **68**, 449–456 (2014).
47. B. Çelen, S. Kariv, Distinguishing informational cascades from herd behavior in the laboratory. *Am. Econ. Rev.* **94**, 484–498 (2004).
48. J. R. G. Dyer, A. Johansson, D. Helbing, I. D. Couzin, J. Krause, Leadership, consensus decision making and collective behaviour in humans. *Philos. Trans. R. Soc. B. Biol. Sci.* **364**, 781–789 (2009).
49. R. H. J. M. Kurvers, M. Wolf, M. Naguib, J. Krause, Self-organized flexible leadership promotes collective intelligence in human groups. *R. Soc. Open Sci.* **2**, 150222 (2015).
50. I. Watts, M. Nagy, T. Burt de Perera, D. Biro, Misinformed leaders lose influence over pigeon flocks. *Biol. Lett.* **12**, 20160544 (2016).
51. R. M. Raafat, N. Chater, C. Frith, Herding in humans. *Trends Cogn. Sci.* **13**, 420–428 (2009).
52. R. Hertwig, Tapping into the wisdom of the crowd—with confidence. *Science* **336**, 303–304 (2012).
53. B. Bahrami, K. Olsen, D. Bang, A. Roepstorff, G. Rees, C. Frith, What failure in collective decision-making tells us about metacognition. *Philos. Trans. R. Soc. Lond B Biol. Sci.* **367**, 1350–1365 (2012).
54. D. Bang, C. D. Frith, Making better decisions in groups. *R. Soc. Open Sci.* **4**, 170193 (2017).
55. L. Chittka, P. Skorupski, N. E. Raine, Speed–accuracy tradeoffs in animal decision making. *Trends Ecol. Evol.* **24**, 400–407 (2009).
56. M. Moussaïd, M. Kapadia, T. Thrash, R. W. Sumner, M. Gross, D. Helbing, C. Hölscher, Crowd behaviour during high-stress evacuations in an immersive virtual environment. *J. R. Soc. Interface* **13**, 20160414 (2016).
57. E. Danchin, L.-A. Giraldeau, T. J. Valone, R. H. Wagner, Public information: From nosy neighbors to cultural evolution. *Science* **305**, 487–491 (2004).
58. P.-C. Bürkner, brms: An R package for Bayesian multilevel models using Stan. *J. Stat. Softw.* **80**, 1–28 (2017).
59. A. Diederich, J. R. Busemeyer, Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *J. Math. Psychol.* **47**, 304–322 (2003).
60. B. M. Turner, P. B. Sederberg, S. D. Brown, M. Steyvers, A method for efficiently sampling from distributions with correlated dimensions. *Psychol. Methods* **18**, 368–384 (2013).
61. R Core Team, *R: A Language and Environment for Statistical Computing* (R Foundation for Statistical Computing, 2019).
62. Stan Development Team, *RStan: The R Interface to Stan* (2018). R package version 2.18.2.

Acknowledgments: We thank M. Wolf for helpful feedback, S. Goss for editing the manuscript, O. Krüger for hosting the experiments, and all tutors of the 2017 “Basismodul Biologie” at the University of Bielefeld, Germany. **Funding:** R.H.J.M.K. acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG; German Research Foundation, grant number: KU 3369/1-1) and Germany’s Excellence Strategy—EXC 2002/1 “Science of Intelligence”—project number 390523135. The funders had no role in the study design, data collection and analysis, and decision to publish or preparation of the manuscript. **Author contributions:** Conceptualization: A.N.T., R.H.J.M.K., and T.J.P.; methodology: A.N.T., R.H.J.M.K., and T.J.P.; software: A.N.T.; data collection and curation: A.N.T.; formal analysis: A.N.T., R.H.J.M.K., and T.J.P.; writing—original draft: A.N.T. and R.H.J.M.K.; writing—reviewing and editing: A.N.T., R.H.J.M.K., and T.J.P.; supervision: R.H.J.M.K. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** The data and code to implement all analysis can be accessed at <https://osf.io/ejfm4/>. All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. Additional data related to this paper may be requested from the authors.

Submitted 24 January 2020

Accepted 29 May 2020

Published 15 July 2020

10.1126/sciadv.abb0266

Citation: A. N. Tump, T. J. Pleskac, R. H. J. M. Kurvers, Wise or mad crowds? The cognitive mechanisms underlying information cascades. *Sci. Adv.* **6**, eabb0266 (2020).

Wise or mad crowds? The cognitive mechanisms underlying information cascades

Alan N. Tump, Timothy J. Pleskac and Ralf H. J. M. Kurvers

Sci Adv **6** (29), eabb0266.
DOI: 10.1126/sciadv.abb0266

ARTICLE TOOLS

<http://advances.sciencemag.org/content/6/29/eabb0266>

SUPPLEMENTARY MATERIALS

<http://advances.sciencemag.org/content/suppl/2020/07/13/6.29.eabb0266.DC1>

REFERENCES

This article cites 57 articles, 6 of which you can access for free
<http://advances.sciencemag.org/content/6/29/eabb0266#BIBL>

PERMISSIONS

<http://www.sciencemag.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2020 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).