



Collective decision making by rational individuals

Richard P. Mann^{a,1}

^aDepartment of Statistics, School of Mathematics, University of Leeds, Leeds LS2 9JT, United Kingdom

Edited by Simon A. Levin, Princeton University, Princeton, NJ, and approved September 25, 2018 (received for review July 11, 2018)

The patterns and mechanisms of collective decision making in humans and animals have attracted both empirical and theoretical attention. Of particular interest has been the variety of social feedback rules and the extent to which these behavioral rules can be explained and predicted from theories of rational estimation and decision making. However, models that aim to model the full range of social information use have incorporated ad hoc departures from rational decision-making theory to explain the apparent stochasticity and variability of behavior. In this paper I develop a model of social information use and collective decision making by fully rational agents that reveals how a wide range of apparently stochastic social decision rules emerge from fundamental information asymmetries both between individuals and between the decision makers and the observer of those decisions. As well as showing that rational decision making is consistent with empirical observations of collective behavior, this model makes several testable predictions about how individuals make decisions in groups and offers a valuable perspective on how we view sources of variability in animal, and human, behavior.

agent-based model | collective behavior | rational choice | social information | utility

Collective decision making is a ubiquitous task for social animal species, including humans (1). Whether deciding where to forage, which nest site to choose, or when to move, individual decisions are greatly informed by observing the choices that others make. As recently as 2008, Ward et al. (ref. 2, p. 1) were able to state that “little is known about the mechanisms underlying decision-making in vertebrate animal groups.” Since then, however, a large literature has explored the rules governing social information use in collective decisions across various taxa, for example in insects (3), fish (2, 4), birds (5, 6), and mammals (7, 8), including primates (9) and humans (10, 11). What links decision making in all of these groups is the presence of social reinforcement, with individuals demonstrating a strong preference for an option chosen by others, which increases with the number of others who have selected it. This reinforcement can be expressed as a social response function—the probability of selecting a given option conditioned on the number of other individuals that have previously chosen it. A large degree of variation has been observed in these social response functions, ranging from linear relationships (e.g., refs. 3 and 12) to strongly nonlinear “quorum” rules (13), where the apparent attractiveness of an option appears to increase exponentially with the number of individuals choosing it, before saturating as this number passes a quorum level. In addition to variation between taxa, studies have also highlighted how the same species can exhibit different patterns of collective behavior under different laboratory or field conditions (14–17), highlighting the potential importance of context-dependent social responses.

Complementing these empirical studies, mathematical theories have been developed to explain why these social decision rules take the form observed. For example, Easley and Kleinberg (18) proposed a toy model for understanding collective decision making in a group of rational agents. This model, illustrated by the example of an individual selecting a restaurant for a meal, demonstrated how easily an unbreakable consensus

decision could emerge, once the cumulative social information provided by past choices outweighs any new quality signal that an uncommitted individual might receive. More recent work has attempted to build a fully descriptive model of such collective decision making by considering the purportedly rational beliefs and decisions of agents exposed to the social information provided by choices of others (19, 20), and the studies have been successful in reproducing the observed response functions in a variety of taxa including insects (20), fish (19–21), and birds (5, 6). Recent extensions of these models have also considered how social responses might vary as a result of changes in environmental context (22).

However, while these models have had success in reproducing the observed features of collective decisions, this has been at the cost of internal consistency as theories of rational behavior. An agent’s decision involves two components: (i) an estimation stage, where the focal agent forms beliefs about the quality of its options, and (ii) a decision rule, which specifies how the agent acts based on those beliefs. In the first stage the models present a broadly coherent theory of estimation based on Bayesian updating. However, beliefs are restricted to statements about whether options are “good” or “bad” [or “best” (19)]. This binary categorization does not fully capture the range of possibilities that individuals face in expected-utility-maximizing, rational behavior (23). Choices made under uncertainty are characterized by risk–reward trade-offs, and it is not clear how these can be translated into a simple good or bad dichotomy or how decisions should be made on the basis of such a classification. The second stage of these models, the decision rule, introduces further departures from rationality. Here agents are assumed to select options probabilistically based on the results of the estimation stage. Such nondeterministic behavior is inconsistent with the idea of individuals as rational agents. This problem emerges as a result of a confusion regarding the sources of

Significance

Humans and animals use the decisions they observe others making as valuable information about the quality of potential choices. We anticipate that behavioral rules that better approximate the decisions of a rational, expected fitness-maximizing agent will be favored by natural selection. In this paper I show that many aspects of collective decision making can be predicted through considering what a rational agent would do. I also show that a theory of strictly rational agents explains much of the variability between and within species and makes sometimes counterintuitive predictions that can be experimentally tested. Most importantly, my model shows that how collective decision making is observed and measured is critically important in interpreting and understanding behavior.

Author contributions: R.P.M. designed research, performed research, and wrote the paper.

The author declares no conflict of interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹ Email: r.p.mann@leeds.ac.uk.

Published online October 15, 2018.

observational uncertainty in empirical studies. Because the decisions made are typically not predictable with certainty by an observer or experimenter, they are themselves deemed to be stochastic. Instead, this uncertainty can be understood by incorporating the viewpoint of the observer into the theory. The observer makes measurements of the physical and social environment that only imperfectly capture the information observed by the focal decision maker (or not observed, for example in the case of visual occlusion). From this one can recognize that the inability of the observer to predict individuals' decisions arises from the limited access the observer has to the information driving those actions, not from fundamentally stochastic behavior by the agents themselves.

Why should we be concerned about departures from rationality in these models? Given established critiques of rational-agent models (24, 25), should we not be more concerned about whether these models are consistent with empirical measurements? To this objection there are two responses. The first one is that rational-agent models provide a baseline from which to measure departures from rationality. These departures are interesting because they indicate either where an evolutionary process has been unable to produce an optimal solution or where other factors, such as cognitive cost, have produced a trade-off. Such departures can be detected only if we understand what genuinely rational behavior looks like. The second response concerns the purported goals of previous work. One can find a reasonable empirical match to any observed social response by choosing an appropriate parameterization of a sufficiently flexible mathematical function. It has been the explicit goal of theoretical work in this area (19, 20) to understand how these responses emerge from logical consideration of the individuals' own estimations and actions as rational decision makers. My goal here is to fully explore the consequences of the rationality assumption to show where data can (and cannot) be explained by this fundamental principle.

This paper develops a model of collective decision making by identical, rational agents, based on three fundamental principles: (i) Individuals behave as expected-utility-maximizing agents, based on their own beliefs about the world; (ii) each individual's beliefs are generated from the public and private information that it has access to, using Bayesian probability updating; and (iii) empirical observation of individuals' actions is undertaken by an observer that has its own private information as well as the public social information on which to base predictions and interpretations of individual behavior. The resulting model reproduces the key successful aspects of previous research, while making additional, testable predictions about social information use that are not accounted for in existing theory.

Theory

Consider the classic paradigm of a group faced with a sequential, binary decision. That is, a sequence of n identical individuals choose between option A and option B and can see the choices made by those ahead of them. Such a context is well approximated empirically by, for example, Y-maze experiments (e.g., refs. 2 and 26) where individuals are asked to choose between two competing arms of a maze. I develop a mathematical framework for calculating the optimal choice for each individual, based on private information that the individual alone observes and public information constituted by the observable choices made by others. I also derive mathematical expressions for the probability that an outside observer (such as an experimental scientist) will observe an individual making a particular choice, conditioned on what that observer can know about the system and the focal individual. As noted above, incorporating the observer explicitly is key to understanding the source of observational uncertainty in a fundamentally deterministic model.

Rational Choice. I start from the assumption that each of the two options has a true utility, U_A for option A and U_B for option B. These utilities may also be understood as fitness consequences of the decision in terms of evolutionary adaptation (27). Since the individuals are assumed to be identical, these utilities are the same for all. These true utilities are unknown, but each individual, $i \in 1, \dots, n$, can estimate the utility of each choice based on the specific information, I_i , that it possesses. Following the rules of Bayesian-rational decision making (23), I assert that individual i will choose option A if and only if the expected utility of A is greater than that of B, according to i 's estimation. Let $x = U_A - U_B$ be the true difference in utilities; then

$$P(i \rightarrow A|I_i) = \begin{cases} 1, & \text{if } \int_{-\infty}^{\infty} xp(x|I_i)dx > 0 \\ 0, & \text{otherwise} \end{cases}, \quad [1]$$

where $p(x|I_i)$ is a probability density representing individual i 's personal belief about x .

Private and Public Information. What information does individual i have? I assume that I_i is composed of two parts, private and public. First, there is direct sensory information that i can perceive from the two options. For example, a foraging individual may perceive differing food odors from A and B, or a prey animal may see differing patterns of shadows that suggest one choice is more likely to lead to a predator. This is the individual's private information. Second, if $i > 1$, then individual i can see the choices made by any other individual j , where $j < i$. This is public information—it is available to all individuals who still wait to make their choice. In common with previous work (19, 20) I make the important assumption that the choices of others provide information about the relative utilities of A and B, but do not influence the true values of these utilities. That is, an option does not become good simply because others have chosen it. In making this assumption I exclude phenomena such as predation–dilution effects (28), where the presence of conspecifics is itself desirable, or foraging competition (29), where the presence of other individuals lowers the utility of a given option.

Prior Belief. Before an individual receives any information regarding the utility difference, x , I assume that they have no reason to favor option A or B (any such reason should count as private information). I ascribe to them a prior belief regarding the values that x may take and by symmetry center this on zero. I further assume that by environmental habituation (either genetic or experience) they have an intrinsic idea of the scale of possible utility differences between competing choices. In this paper I assume that prior beliefs follow a normal distribution, and without loss of generality we can measure utilities in units that set this variance of this distribution to one:

$$p(x) = \frac{1}{\sqrt{2\pi}} \exp(-x^2) = \phi(x). \quad [2]$$

Hereafter I use $\phi(x)$ to refer to the standard normal distribution density function. Throughout this paper I assume that information and expectations about the environment are normally distributed. This assumption is likely to hold well for low-level sensory information such as detecting food or predators, but may be less appropriate for more cognitively advanced tasks. The model development detailed here can be followed for any alternative distribution of interest.

Private Information. Individual i has access to private information that gives a noisy estimate of x . This may be via visual, olfactory, or other sensory stimuli, but here I model this as an abstract quantity, Δ_i , that is generated stochastically by the environment

based on the real utility difference x , and a noise variance ν^2 due to both the latent sources of environmental noise (e.g., air currents disrupting olfactory gradients) and limitations of an individual's sensory apparatus. Mathematically, Δ_i is normally distributed with mean x and variance ν^2 :

$$p(\Delta_i|x) = \phi((\Delta_i - x)/\nu). \quad [3]$$

Individual i revises its belief about x in the light of its private information using Bayes' rule:

$$p(x|\Delta_i) \propto p(\Delta_i|x)p(x) \propto \phi((\Delta_i - x)/\nu)\phi(x). \quad [4]$$

Consider the case of the first decision maker, $i = 1$. This individual has no public, social information to draw on, and bases the estimate of x entirely on its private information, Δ_1 . For the purposes of making a rational decision, the important quantity for the individual to evaluate is the sign of the expected utility difference, $\mathbb{E}(x|\Delta_1)$:

$$\begin{aligned} \mathbb{E}(x|\Delta_1) &= \int_{-\infty}^{\infty} xp(x|\Delta_1)dx \\ &\propto \int_{-\infty}^{\infty} x\phi((\Delta_1 - x)/\nu)\phi(x)dx. \end{aligned} \quad [5]$$

It is clear that $\mathbb{E}(x|\Delta_1) > 0$ if and only if (iff) $\Delta_1 > 0$ and therefore that the sign of individual 1's private information dictates which option it will choose.

Social Information. Having determined that the first decision maker uses the sign of its private information to make a choice, I now consider the case of the second and subsequent decision makers. Individual 2 begins its estimation of x in the same manner as individual 1, by updating its original prior belief (which is identical for all agents), using its own private information, Δ_2 :

$$p(x|\Delta_2) \propto \phi((\Delta_2 - x)/\nu)\phi(x). \quad [6]$$

What information does the choice of individual 1 provide to individual 2? Since the choice of individual 1 does not change the true utilities of the options, its choice can influence the estimation of individual 2 only by giving information about the private information that individual 1 received. If individual 1 were to communicate its private information directly to individual 2, then the second individual could update its belief based on these new data. However, imagine that individual 1 has chosen option A. Individual 2 does not know what private information individual 1 has received, but can only infer from the observed resulting choice that $\Delta_1 > 0$. Therefore, individual 2 must consider all possible values of the Δ_1 that the first individual may have observed, weighted by probability, and adjust its belief accordingly. Let $C_1 = 1$ indicate that individual 1 chose option A (and conversely $C_1 = -1$ for option B); then

$$\begin{aligned} p(x|C_1 = 1, \Delta_2) &\propto p(x|\Delta_2)P(C_1 = 1|x) \\ &\propto p(x|\Delta_2)P(\Delta_1 > 0|x) \\ &\propto \phi((\Delta_2 - x)/\nu)\Phi(x/\nu)\phi(x), \end{aligned} \quad [7]$$

where $\Phi(z) = \int_{-\infty}^z \phi(t)dt$ is the cumulative distribution function of the standard normal distribution. Similarly, if individual 1 had chosen B, then individual 2 would make the estimation

$$p(x|C_1 = -1, \Delta_2) \propto \phi((\Delta_2 - x)/\nu)\Phi(-x/\nu)\phi(x). \quad [8]$$

The decision of individual 2 is now governed by its expected value of x . Define Δ_2^* such that

$$\begin{aligned} \mathbb{E}(x|C_1, \Delta_2^*) &= 0 \\ \Rightarrow \int_{-\infty}^{\infty} x\phi((\Delta_2^* - x)/\nu)\Phi(C_1x/\nu)\phi(x)dx &= 0. \end{aligned} \quad [9]$$

Individual 2 will now choose option A iff $\Delta_2 > \Delta_2^*$, implying that $\mathbb{E}(x|C_1, \Delta_2) > 0$.

Subsequent Decisions. To complete our view of how social information is used we need to consider the viewpoint of the third individual (assuming $n > 2$). This individual must consider not only the decisions made by individuals 1 and 2 in conjunction with its own private information, but also the order in which these decisions were made. As with all individuals, I begin by updating the universal prior, using individual 3's private information:

$$p(x|\Delta_3) \propto \phi((\Delta_3 - x)/\nu)\phi(x). \quad [10]$$

Individual 3 can also update its belief based on the choice of individual 1, in exactly the same manner as individual 2:

$$p(x|C_1, \Delta_3) \propto \phi((\Delta_3 - x)/\nu)\Phi(C_1x/\nu)\phi(x). \quad [11]$$

Now individual 3 needs to update its belief based on the decision made by individual 2, C_2 :

$$p(x|C_1, C_2, \Delta_3) \propto P(C_2|x, C_1)p(x|C_1, \Delta_3). \quad [12]$$

To evaluate the first term on the right-hand side, individual 3 needs to adopt the viewpoint of individual 2 and calculate the critical value Δ_2^* based on C_1 . Then, from individual 3's perspective, the probability of choice C_2 is

$$\begin{aligned} P(C_2 = 1|x, C_1) &= P(\Delta_2 > \Delta_2^*|x) \\ &= \Phi((x - \Delta_2^*)/\nu). \end{aligned} \quad [13]$$

Similarly, $P(C_2 = -1|x, C_1) = \Phi(-(x - \Delta_2^*)/\nu)$. Hence, individual 3 updates its belief based on C_2 to

$$\begin{aligned} p(x|C_1, C_2, \Delta_3) &\propto \phi((\Delta_3 - x)/\nu)\Phi(C_1x/\nu) \\ &\quad \times \Phi(C_2(x - \Delta_2^*)/\nu)\phi(x). \end{aligned} \quad [14]$$

As with individual 2, we can thus evaluate a critical value, Δ_3^* , defined by

$$\int_{-\infty}^{\infty} x\phi((\Delta_3^* - x)/\nu)\Phi(C_1x/\nu)\Phi(C_2(x - \Delta_2^*)/\nu)\phi(x)dx = 0. \quad [15]$$

Individual 3 will now choose option A iff $\Delta_3 > \Delta_3^*$. By iteratively proceeding in a similar fashion we can determine the belief of individual i , based on its private information and the observed choices of previous individuals, C_1, \dots, C_{i-1} ,

$$\begin{aligned} p(x|C_1, \dots, C_{i-1}, \Delta_i) &\propto \phi((\Delta_i - x)/\nu)\phi(x) \\ &\quad \times \prod_{j=1}^{i-1} \Phi(C_j(x - \Delta_j^*)/\nu), \end{aligned} \quad [16]$$

and

$$\int_{-\infty}^{\infty} xp(x|C_1, \dots, C_{i-1}, \Delta_i^*)dx = 0, \quad [17]$$

and I define $\Delta_1^* = 0$.

Observation. So far I have discussed how each individual uses private and public information to make a rational, expected-utility-maximizing decision. Now I consider the perspective as an observer of this process. As an observer, one is able to observe the same public, social information available to the individuals themselves—the sequence of decisions. However, the observer's viewpoint differs in two ways. First, the observer has no access to the private information of any individual. Second, the observer may have knowledge about the true environmental conditions. For example, the observer may have designed an experiment such that $x = U_A - U_B = 0$, e.g., a Y maze with symmetrical arms. Furthermore, especially in a laboratory setting, the observer may have altered the environment, such that noise levels differ from those that the individuals are habituated to.

Assume that true values of x and ν^2 and also the noise variance of the experimental environment, η^2 , are known. How can we calculate the probability that individual i will make a specific choice? First, we can follow the calculations above for that individual, to determine the critical value Δ_i^* , based on the observed previous decisions. We can then evaluate the probability, conditioned on the known x and η^2 , that individual i 's private information will exceed this value:

$$P(C_i = 1 | x, \eta, C_1, \dots, C_{i-1}) = P(\Delta_i > \Delta_i^* | x, \eta) = \Phi((x - \Delta_i^*)/\eta). \quad [18]$$

Although this equation provides the probability that the observed decision will be a particular option, it implies no non-deterministic behavior; the uncertainty encoded by the probability is purely a consequence of the observer not sharing the same information as the focal decision maker. Note that the information provided by C_1, \dots, C_{i-1} is encoded in the calculated value of Δ_i^* . Since each critical value depends iteratively on those before, this is determined by the order of decisions made, as well as the aggregate numbers choosing A and B.

Unordered Social Information. Observers not only have differing information from the individuals under observation, they also make choices about how to measure and record behavior. As an example of this, the majority of previous studies have largely ignored the precise order of previous decisions made when measuring social responses. For comparison with this previous work, we can consider what we, the observer, would predict about the decision of individual i , conditioned on knowing only the number of previous individuals choosing A (n_A) and B (n_B). This requires us to consider the set of all possible sequences, \mathcal{C} , that obey the result in n_A, n_B and to sum over the probability that each of these is the sequence that led to the current arrangement. This summation, combined with the decision rule for each specific sequence derived above, then gives the probability that the next choice will be either A or B, conditioned on this unordered observation of social information:

$$P(C_i | n_A, n_B, x, \eta) = \sum_{\mathcal{C}} P(C_i | x, \eta, C_1, \dots, C_{i-1}) \times P(C_1, \dots, C_{i-1} | x, \eta). \quad [19]$$

Note that this calculation assumes that the individuals themselves are aware of the order in which decisions were made, but that the observer has been unable to record these or has chosen not to do so.

Conflicting Information. Several experimental studies have investigated scenarios where a conflict is introduced between an indi-

vidual's private and social information, to identify the relative strengths of the two factors. For example, in such an experiment each individual may be trained in advance to associate food with one of two or more different colors or patterns. Individuals with different trained associations are then placed in a group and presented with a decision where each option has a color or pattern signal (e.g., ref. 30).

This scenario can be simulated by giving each agent private information drawn from a mixture of two conflicting distributions,

$$p(\Delta_i | x) = \frac{1}{2} [\phi((\Delta_i - z/2)/\eta) + \phi((\Delta_i + z/2)/\eta)], \quad [20]$$

where z is the magnitude of the conflict in information, and η is again the experimental noise level. These two parameters indicate respectively how strong the training has indicated the utility difference is (e.g., the amount or quality of food provided) and the reliability of the signal (e.g., whether the food was always provided in the same quantities). With each individual's private information drawn from this mixture distribution, a simulation of the group's aggregate behavior can follow as above. It should be noted here I am still assuming that the individual decision makers have identical utility functions; the conflict between them is solely on the level of the information they have received and not one of differing preferences.

Results

In this section I consider a variety of possible experimental and field study scenarios that illustrate the key predictions of the model.

Role of Environmental Signal-to-Noise Ratio. I begin by considering an experimental field study in the habitual environment of the decision makers. In this case the decision makers' previous experience gives them reliable prior information about typical signal and noise levels in their environment, while the observer can experimentally control the true utilities of possible choices. I analyzed the expected behavior when the decision makers are confronted with a symmetrical binary choice in which the true utility difference is zero: $U_A - U_B = 0$. I calculated the probability that the next decision maker would choose option A, conditioned on different previous decision sequences, and for a range of environmental noise/signal ratios. This analysis, illustrated in Fig. 1, shows that environmental noise levels have little effect on the predicted choice probabilities, displaying only a slight inflection at a ratio of one. This result can be understood intuitively by noting that higher ambient noise levels reduce the reliability of both private and public information—the decision maker should trust its own information less, since it may result from noise, but should also recognize that the decisions of others are also more likely to be incorrect. These two effects almost perfectly balance in this analysis for a wide range of possible noise/signal ratios. Since the environmental noise/signal ratio has little effect on behavior, I set it to a value of $\nu = 1$ henceforth for simplicity.

Observed Social Interaction Rules. My model gives the probability that a focal individual will make a given choice conditioned on any ordered sequence of previous decisions. However, in practice researchers either are often unable to observe this precise sequence or choose to ignore the details of the order in which decisions were made, focusing instead only on the number of individuals who have previously chosen A or B in aggregate. To make comparisons between the theory developed here and previous work, I therefore calculated the expected observations on this aggregated level by considering all possible sequences of previous choices that could have led to an aggregate state n_A, n_B




Fig. 1. Consistency of predicted decisions across environmental noise/signal ratios, for a range of possible observed past decision sequences. Each line is labeled with the corresponding sequence of past decisions, from left to right. The black line shows cases with one previous decision maker, blue lines cases with two, and red lines cases with three, respectively.

(Eq. 19). To illustrate the predicted observations a researcher would make in such an experiment, we consider two hypothetical experiments, each with 10 individuals and in which the ambient noise level matches that of the habitual environment. The first experiment uses symmetric options ($x = U_A - U_B = 0$). The predicted observations made in these experiments are shown in Fig. 2. Fig. 2*A* shows the probability that a focal decision maker will choose option A, conditioned on known n_A and n_B . Contour lines of equal probability show a radial pattern that is suggestive of Weber's law of relative differences (31). To illustrate this further, Fig. 2*B* shows how the probability of choosing option A varies with the relative proportion of previous decisions: $n_A/(n_A + n_B)$. We see that the results averaged over all possible sequences (red circles) show a very close linear trend. However, this apparently simple relationship results from the weighted average of sequence-specific probabilities, shown with black circles, where larger points indicate more probable sequences. Evaluating the probability of generating each possible sequence of decisions, we can determine the probability for the final value of n_A after all decisions have been made (Fig. 2*C*). This exhibits the classic U-shaped distribution that is characteristic of observations in many collective decision-making studies (e.g., ref. 13). At first glance this result conflicts with the pattern in Fig. 2*B*—a sequence of decisions made according to Weber's law would be expected to result in a uniform distribution of equal probabilities for final values of n_A (32). This apparent contradiction is resolved by noting that the linear relationship in Fig. 2*B* does not hold as a decision rule in its own right, but only as the average behavior aggregated over many possible sequences of decisions using the true behavioral rules shown in Eqs. 16 and 17. This highlights how apparently straightforward analysis of empirical data may lead to erroneous conclusions about underlying behavioral mechanisms.

In addition to a symmetric experimental setup, I also consider a hypothetical experiment in which one option is objectively better than the other, for example through the presence of food

(e.g., ref. 33) or the absence of a predator (e.g., ref. 26). In this example I assume that option A is better and set $x = U_A - U_B = \sqrt{2/\pi}$. This value is equivalent to the average absolute difference in utilities in the habitual environment and thus represents a “typical” decision for the agents to make. I made predictions of the observed decisions made by 10 agents as in the symmetric case, the results of which are shown in Fig. 2*D–F*. In this case we see that, as expected, decisions systematically favor the higher-utility option. As above, the decisions observed as a function of n_A and n_B hide a broader variety of social contexts defined by the ordering of decisions (Fig. 2*E*).

Context Specificity. The hypothetical experiment above was assumed to take place in an environment where noise levels were the same as the decision makers' habitual experience. I showed that this habitual noise level did not in itself have a strong influence on predicted decisions—individuals habituated to noisy environments should be no more or less likely follow one another in their own environment than those from less noisy habitats. But what if individuals are removed from their own habitat and placed in an unfamiliar environment? As an example, consider collecting fish which usually shoal in somewhat murky and strongly odored rivers or lakes and placing these in clear water in a uniform, plastic experimental arena (e.g., refs. 13, 26, 34–36, and many others). What impact might this change have on their behavior? One possibility is that a severe change of environment may lead to erratic or pathological behavior as a result of distress or disorientation, which I do not account for here. Another possibility is that the fish continue to follow social interaction rules that have evolved to be near optimal in their own habitat, without accounting for the changed context.

In the relatively sterile laboratory conditions, the ambient levels of noise such as stray odors may be far lower than in the wild. In combination with a symmetrical experimental setup as above ($U_A = U_B$), this means that an individual fish is less likely to observe private information of a sufficient magnitude to contradict the apparent social information provided by its conspecifics and will therefore be more inclined to aggregate with and follow these other individuals than it would in the wild. Conversely, a strongly lit laboratory environment may introduce a greater intensity of visual noise simply by virtue of the greater overall intensity of visual stimulation. I explicitly calculate the effect that experimental noise has using Eq. 18, varying the ratio between experimental and habitual noise levels. That is, I assume that the agents continue to act rationally in the belief that noise levels remain at those of their habitual environment, while in fact the noise levels depart from this baseline. As shown in Fig. 3, I find as expected that lower experimental noise levels increase the tendency to follow the majority (Fig. 3*A* and *B*), resulting in a greater aggregate consensus (Fig. 3*E* and *F*). Higher noise levels reduce the weight of social information (Fig. 3*C* and *D*) and thus prevent consensus from emerging (Fig. 3*G* and *H*). Conversely this means that for a given experimental noise level, individuals from noisier habitual environments are expected to aggregate more strongly and behave more socially than those from habitats with less noise.

Dynamic Social Information. As noted previously, the predictions my model makes about the decision a focal individual will make, conditioned on the available public information, depend strongly on the precise order in which previous decisions were made. To investigate this further, I now focus on the relative importance of the most recent decision in particular. Using Eqs. 16–18, I calculated the probability, in a symmetric experimental setup, that a focal individual will choose option A, conditioned on previous sequences of decisions of the form $C = \{-1, -1, \dots, 1\}$, that is, sequences in which the most recent previous decision

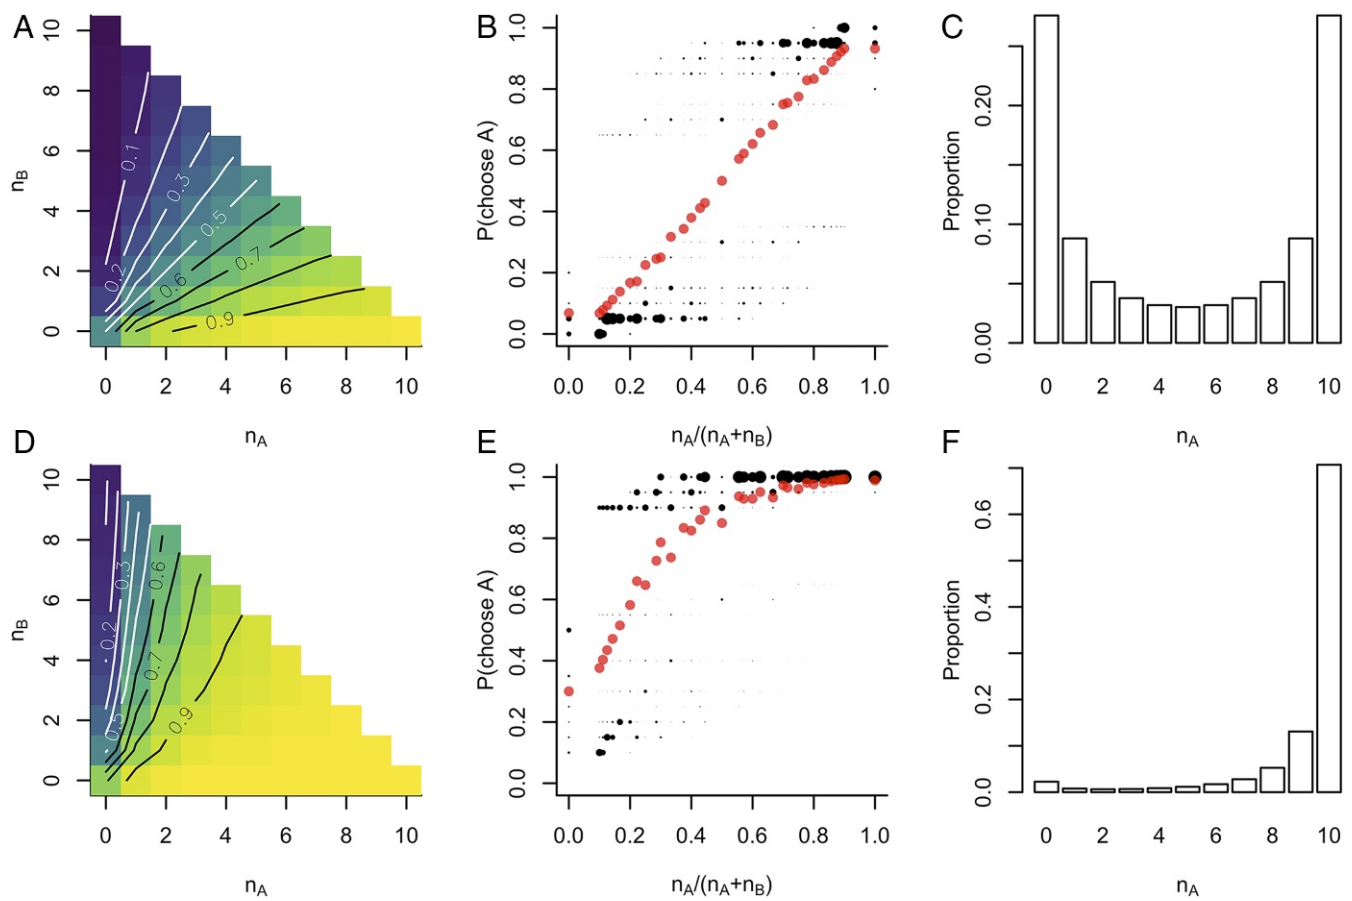


Fig. 2. Results of two hypothetical studies of collective decision making. (A–C) Predicted observations made from a symmetric experimental setup ($x = 0$). (D–F) Predicted observations from an experiment in which option A is superior by an amount typical in the habitual environment ($x = \sqrt{2/\pi}$). (A and D) The probability, from the perspective of an external observer, that a focal agent will choose option A, conditioned on the number of agents, n_A , n_B , previously choosing options A and B, respectively. (B and E) The probability of the focal agent selecting option A against the proportion of previous agents selecting A. Red circles indicate the average across all possible sequences of previous choices, while black circles indicate probabilities conditioned on specific sequences (discretized to intervals of 0.05), with larger points indicating more likely sequences. The average trend shows the relationship that would be observed in an experiment where sequence information was discarded. (C and F) The probability of possible aggregate outcomes, defined as the number of agents in total that will select option A, showing the high probability of consensus decisions and of collectively choosing the higher-utility option in the asymmetric scenario.

was option A, after a series of individuals choosing B. In most models such sequences—with many individuals choosing B and only one choosing A—would result in a high probability that the focal individual would choose B, following the majority. By contrast, as shown in Fig. 4, I predict that the focal individual is most likely to choose A, regardless of the weight of the majority for B. While longer series of earlier choices for B do make the probability to select A somewhat lower, this nonetheless always remains above 0.5. This result might initially appear counterintuitive: Why should the social information provided by a large majority of other individuals be outweighed by one recent decision? However, this neatly illustrates the consequences of taking seriously the idea of identical, rational agents. The focal individual, observing the most recent decision, must conclude that the individual making that decision has observed private information which is sufficient to outweigh all of the previous public social information. The focal individual cannot observe this private information directly, but since the agents are identical, it can infer that had it seen this information itself, it would have made the same decision. Therefore, it must conclude, before observing its own private information, that the available public information is now in favor of A. Since, in a symmetric implementation of the model with $U_A = U_B$, the focal individual's private information is

equally likely to favor either choice, the observational prediction is that the most recent decision will be followed on the majority of occasions, regardless of the overall number of previous choices made for either option.

Conflicting Information. I considered a scenario of 10 individuals that receive conflicting private information according to Eq. 20. I evaluated the probability of each possible aggregate outcome at each of 20 different magnitudes of conflict between $z=0$ and $z=10$ and five different experimental noise levels of $\eta=1/3, 2/3, 1, 3/2$, and 3. The results, shown in Fig. 5, show the degree of group consensus for each scenario, between zero (individuals split between two options equally) and one (all individuals choosing the same option). As already shown above, consensus is not guaranteed even when conflict is zero, and noisier experimental setups tend to reduce consensus. Increasing the magnitude of conflict decreases the expected degree of consensus. With sufficient conflict consensus breaks down entirely and each individual simply follows its own private information, leading to consensus values in line with those expected from a binomial distribution (dashed line in Fig. 5). For high noise conditions the decline of consensus is gradual from an initially low value, whereas in low noise conditions there is a clearer transition

from consensus to independent decision making. When training information is highly reliable compared with that found in the habitual environment ($\eta = 1/3$), this transition is very sharp, and the intermediate range between full consensus and completely independent decisions is very narrow.

Discussion

I have developed a model of collective animal decision making based on perfectly rational individual decisions by identical individuals in the context of private and public information. Using this model I have explored the consequences of rational decision making, from the perspective of an observer who also has only partial information about individuals under observation. The results shown here demonstrate fundamental similarities both to earlier models of collective decision making and to the key features of observed behavior across a variety of taxa (20).

Formulating this model focused attention on the underappreciated role of the observer and experimental context in understanding why animals under study make particular observed choices. Specifically, I have shown that even when agents themselves are purely rational (and therefore act deterministically on their own information), their actions appear random to the observer as a result of the agents' private information. Furthermore, the observer potentially influences both the behavior of the individuals under study (through control of experimental conditions) and the interpretation of behavior observed (through the choice of what to measure). Both of these aspects of observation are underappreciated in the collective behavior literature and are potentially responsible for a substantial proportion of the variance in empirical observations. It should be noted that this perspective does not imply that the actual process being observed is dependent on the observer. Two different observers making measurements of the same experiment will observe the same decisions being made, but they may come to different conclusions depending on what they know about the experimental setup and what they choose to measure.

This model predicts that the decision-making process is context specific. To the degree that laboratory conditions represent a lower noise environment than in the wild, I anticipate that observed social tendencies will be more pronounced than in the wild. In human behavior, this offers an explanation for why individuals exposed to social information in laboratory experiments exhibit stronger effect sizes than those exposed in more naturalistic environments (16); the laboratory environment is subject to less spurious information that can contradict the social information presented. A recent study of context-dependent collective behavior in sticklebacks also found that these fish were more cohesive in featureless environments than those with more distractions such as food or plant cover. The model also offers an ecological explanation for the differing social behaviors observed in different species in similar contexts. For example, Aron et al. (37) found that Argentine ants (*Linepithema humile*) showed a stronger preference for social information over private information compared with garden ants (*Lasius niger*) and reflected that this may be explained by the differing ecology of these two species: Argentine ants are restless migrators feeding on novel food sources (a high noise/signal environment), while garden ants are sedentary and feed on well-established food sources (a low noise/signal environment). In similar laboratory conditions, *L. humile* therefore arguably experiences a greater reduction in noise relative to its habitual environment, leading to a prediction of stronger social behavior. Similarly, Wright et al. (38) found that wild-strain zebrafish exhibited a stronger shoaling tendency than laboratory-strain specimens when both were tested in the same laboratory environment. This was attributed to differences in predation risk but could also reflect informational differences in each strain's habitual environment.

One should, therefore, be careful when interpreting differences in laboratory behavior between species from different environments, as this may betoken differing contrasts between the laboratory and the wild, rather than different habitual levels

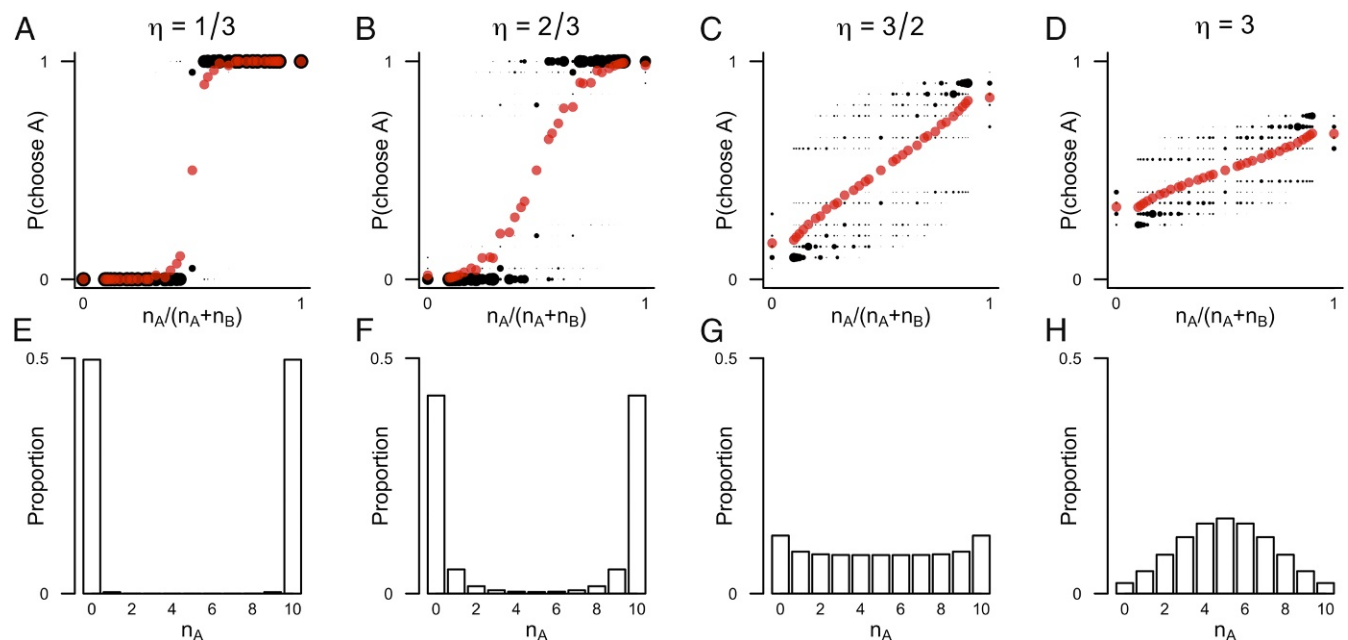


Fig. 3. The effect of varying the experimental noise level. (A–D) The social response function for hypothetical experiments with noise levels in the ratios $\eta = 1/3, 2/3, 3/2,$ and 3 relative to those of the habitual environment. As in Fig. 2, red circles indicate the average observed response, while black circles indicate specific sequences of past decisions, with size indicating the relative probability of each sequence. (E–H) The corresponding proportions of aggregate outcomes for each experimental noise level. Social response and aggregate cohesion are stronger in experiments with noise levels lower than those of the habitual environment ($\eta < 1$) and correspondingly weaker in experiments with greater noise levels ($\eta > 1$).

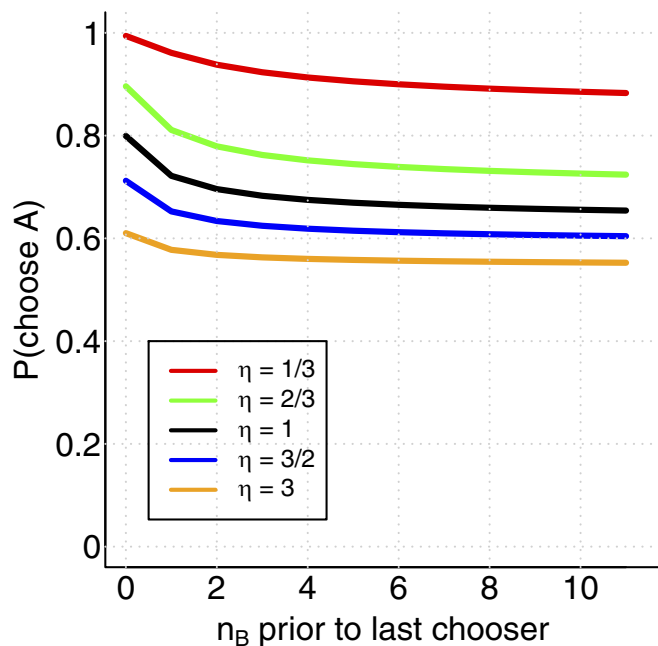


Fig. 4. Predicted probability for a focal individual to choose option A, conditioned a sequence of past decisions of the form B, . . . , B,A, evaluated for a range of experimental to habitual noise ratios (η). In all cases the probability to choose A remains above 0.5, regardless of the size of the majority choosing B.

of sociality. With the development of increasingly advanced tracking technology, researchers have recently reoriented toward studying animal behavior in the wild (39–41) and to studying human social behavior outside of the laboratory (42–44). This study supports this trend; collective behavior in the wild may vary significantly from that in the laboratory, and understanding natural behavior thus requires studying the animals in their habitual environment. An interesting corollary of this finding is that social behavior may be expected to change and become more apparently rational in the laboratory over time in cognitively plastic species such as humans, as they habituate to the new environment. Indicative results of such an effect in a related domain have been shown for example by Burton-Chellew, Nax, and West (45), who found that initially “irrational” prosocial behavior by players in a public goods game become more “rationally” self-serving as the game was repeated many times under the same laboratory conditions. Whether or not similar plasticity is seen in the laboratory use of social information is worthy of further study.

I also investigated how context specificity affects behavior when individuals are given conflicting information, for example through training before the experiment. I found that in laboratory conditions where complete consensus decision making is the norm (low noise relative to the habitual environment), the reaction of individuals to conflicting information is predicted to be strongly nonlinear with respect to the magnitude of the conflict, with a sharp transition between consensus and independent decision making. However, this transition, which might be observed as a critical threshold in the laboratory, would be less clearly observed in more natural conditions ($\eta = 1$), where the model predicts a more gradual decline in the degree of consensus achieved as the magnitude of conflict is increased. Again, this highlights how behaviors observed in laboratory experiments may not be directly translatable to wild behavior. Furthermore, I have shown that collective decision making with conflicting information depends on both the magnitude and the reliability of the information individuals receive, whereas previous studies have

often treated these distinct informational features ambiguously in force-based models of collective decision making (e.g., refs. 46 and 47).

Comparing the model predictions and experimental studies highlighted a further important data analysis consideration. I measured how individual decisions varied in the context of how many others had previously chosen different options, but without any information about the sequence of those choices. In the symmetric case, agents exhibited an apparent social interaction that depended linearly on the number of other individuals choosing either option—a Weber’s law response function. The linearity of this relationship is apparently at odds with the strong tendency to consensus at the aggregate level. Perna et al. (3) noted the same apparent conflict in their experimental study of trail formation in ants. They proposed that the conflict could be resolved through the introduction of stochastic noise in the decision-making process. In contrast, I have shown that this conflict can be resolved within a rational model by focusing on the importance of ordering in the sequence of previous decisions. The sequence of decisions has often been ignored in previous work or relegated to additional material not central to the study’s key insights; this model forces us to recognize the central place that ordering has in understanding rational social behavior. As shown by Perna, Grégoire, and Mann (48), inferred social responses depend intimately on how social context and behavior are measured. The results shown here should reiterate that so-called “model-free” data-driven analysis (49) is an illusion—even when no model is specified, it is implicit in the choices made by researchers regarding what to measure.

Looking more closely at the effect of ordering also revealed testable predictions about the use of social information for which few existing data are available. Specifically, the model predicts that social information associated with the most recent decision makers should have an overwhelming impact on the focal agent. Excluding their own private information, focal agents

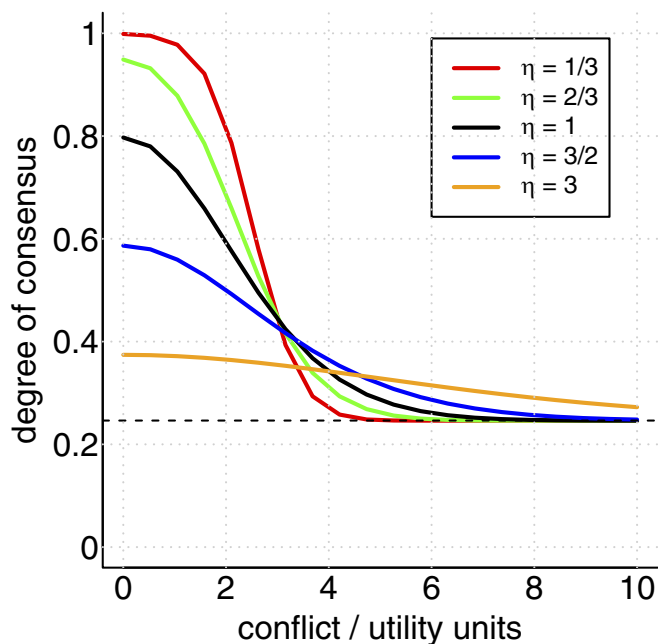


Fig. 5. The average degree of consensus ($(|n_A - n_B| / (n_A + n_B))$) in groups when individuals receive conflicting information, as a function of the magnitude of information conflict and the experimental noise level (η). The dashed line shows the expectation from a binomial distribution where each individual chooses independently.

should always conclude that the total social information favors the most recent decision made, precisely because they must believe that, as identical rational agents, they would have made the same decision given the same information. From an observational perspective this means that we should expect to see the most recent decision being more predictive of the next decision than the aggregate numbers of previous choices. To the best of my knowledge no experiment has previously tested this specific hypothesis. However, some suggestive evidence of such an effect has been seen in at least one previous study (50), where decisions by humbug damselfish to switch between two coral regions were best predicted by the most recent movement of a conspecific. Preferential following of recent decision makers may also help to explain the sensitivity of groups to changes of movement by relatively few initiators and the corresponding prevalence of “false alarms” in groups of prey animals (51). However, reliably inferring such a behavioral rule from observational data is difficult—potentially unknown information driving the most recent decision maker’s choice may also be influencing the focal individual. The prediction that the most recent decisions provide the most salient social information could be tested experimentally by inducing a conflict between the most-recently observable decision and the majority of previous decisions, for example

through the use of artificial substitute conspecifics (e.g., refs. 52 and 53).

A model of rational behavior should not be mistaken for a detailed understanding of biological cognition: Behavior results from biological processes that are subject to evolutionary pressure and physical constraints, and understanding these biological mechanisms will be important in gaining further understanding of how animals cognitively represent and process social information (54). Real-world animal species, including humans, only approximate perfect rationality (24, 55) and typically only in contexts in which adaptation has taken place. Nonetheless, the results of this study serve an important purpose. Assumptions of rationality and optimality are an important tool in understanding adaptive behavior. These assumptions, when posited (19, 20), should be followed to their logical conclusion. Otherwise, it is impossible to determine, by comparing predictions and empirical data, whether or not observed behavior supports them. Precisely because departures from rationality are of such profound interest to biologists, economists, and psychologists, it is important to be precise in identifying what those departures are.

ACKNOWLEDGMENTS. I am grateful to Graham Budd, Alfonso Pérez-Escudero, and Viktoria Spaiser for valuable feedback on the manuscript.

- Conradt L, List C (2009) Group decisions in humans and animals: A survey. *Philos Trans R Soc Lond B Biol Sci* 364:719–742.
- Ward AJW, Sumpter DJT, Couzin ID, Hart PJB, Krause J (2008) Quorum decision-making facilitates information transfer in fish shoals. *Proc Natl Acad Sci USA* 105:6948–6953.
- Perna A, et al. (2012) Individual rules for trail pattern formation in Argentine ants (*Linepithema humile*). *PLoS Comput Biol* 8:e1002592.
- Ward AJW, Krause A, Sumpter DJT (2012) Quorum decision-making in foraging fish shoals. *PLoS One* 7:e32411.
- Farine DR, Aplin LM, Garroway CJ, Mann RP, Sheldon BC (2014) Conspecifics and heterospecifics influence decision-making in multi-species songbird flocks. *Anim Behav* 95:173–182.
- Aplin LM, Farine DR, Mann RP, Sheldon BC (2014) Individual-level personality influences social foraging and collective behaviour in wild birds. *Proc R Soc B* 281:20141016.
- Gall GEC, Strandburg-Peshkin A, Clutton-Brock T, Manser MB (2017) As dusk falls: Collective decisions about the return to sleeping sites in meerkats. *Anim Behav* 132:91–99.
- Walker RH, King AJ, McNutt JW, Jordan NR (2017) Sneeze to leave: African wild dogs (*Lycaon pictus*) use variable quorum thresholds facilitated by sneezes in collective decisions. *Proc R Soc B* 284:20170347.
- Strandburg-Peshkin A, Farine DR, Couzin ID, Crofoot MC (2015) Shared decision-making drives collective movement in wild baboons. *Science* 348:1358–1361.
- Faria JJ, Krause S, Krause J (2010) Collective behavior in road crossing pedestrians: The role of social information. *Behav Ecol* 21:1236–1242.
- Gallup AC, et al. (2012) Visual attention and the acquisition of information in human crowds. *Proc Natl Acad Sci USA* 109:7245–7250.
- Mann RP, Faria J, Sumpter DJT, Krause J (2013) The dynamics of audience applause. *J R Soc Interf* 10:20130466.
- Sumpter DJT, Pratt SC (2009) Quorum responses and consensus decision making. *Philos Trans R Soc B Biol Sci* 364:743–753.
- Hoare DJ, Couzin ID, Godin JGJ, Krause J (2004) Context-dependent group size choice in fish. *Anim Behav* 67:155–164.
- Burns ALJ, Herbert-Read JE, Morrell LJ, Ward AJW (2012) Consistency of leadership in shoals of mosquitofish (*Gambusia holbrooki*) in novel and in familiar environments. *PLoS One* 7:e36567.
- Jerit J, Barabas J, Scott C (2013) Comparing contemporaneous laboratory and field experiments on media effects. *Public Opin Q* 77:256–282.
- Jolles JW, Laskowski KL, Boogert NJ, Manica A (2018) Repeatable group differences in the collective behaviour of stickleback shoals across ecological contexts. *Proc R Soc B* 285:20172629.
- Easley D, Kleinberg J (2010) *Networks, Crowds, and Markets: Reasoning About a Highly Connected World* (Cambridge Univ Press, New York).
- Pérez-Escudero A, De Polavieja GG (2011) Collective animal behavior from Bayesian estimation and probability matching. *PLoS Comput Biol* 7:e1002282.
- 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* 109:20508–20513.
- Pérez-Escudero A, et al. (2013) Estimation models describe well collective decisions among three options. *Proc Natl Acad Sci USA* 110:E3466–E3467.
- Pérez-Escudero A, de Polavieja GG (2017) Adversity magnifies the importance of social information in decision-making. *J R Soc Interf* 14:20170748.
- Neumann JV, Morgenstern O (2007) *Theory of Games and Economic Behavior* (Princeton Univ Press, Woodstock, Oxfordshire, United Kingdom).
- Kahneman D (2012) *Thinking, Fast and Slow* (Penguin Books, London).
- Fawcett TW, Hamblin S, Giraldeau L-A (2012) Exposing the behavioral gambit: The evolution of learning and decision rules. *Behav Ecol* 24:2–11.
- 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.
- Kacelnik A, Marsh B (2002) Cost can increase preference in starlings. *Anim Behav* 63:245–250.
- Hamilton WD (1971) Geometry for the selfish herd. *J Theor Biol* 31:295–311.
- Fretwell SD, Stevan Calver J (1969) On territorial behavior and other factors influencing habitat distribution in birds. *Acta Biotheoretica* 19:37–44.
- Miller N, Garnier S, Hartnett AT, Couzin ID (2013) Both information and social cohesion determine collective decisions in animal groups. *Proc Nat Acad Sci USA* 110:5263–5268.
- Weber EH (1834) *De Pulsu, Resorptione, Auditu et Tactu: Annotationes Anatomicae et Physiologicae* (CF Koehler, Leipzig, Germany).
- Mann RP, Garnett R (2015) The entropic basis of collective behaviour. *J R Soc Interf* 12:20150037.
- Clément RJG, et al. (2017) Collective decision making in guppies: A cross-population comparison study in the wild. *Behav Ecol* 28:919–924.
- Herbert-Read JE, et al. (2011) Inferring the rules of interaction of shoaling fish. *Proc Natl Acad Sci USA* 108:18726–18731.
- Katz Y, Tunstrom K, Ioannou CC, Huepe C, Couzin ID (2011) Inferring the structure and dynamics of interactions in schooling fish. *Proc Natl Acad Sci USA* 108:18720–18725.
- Gautrais J, et al. (2012) Deciphering interactions in moving animal groups. *PLoS Comput Biol* 8:e1002678.
- Aron S, Beckers R, Deneubourg J-L, Pasteels JM (1993) Memory and chemical communication in the orientation of two mass-recruiting ant species. *Insectes Sociaux* 40:369–380.
- Wright D, Nakamichi R, Krause J, Butlin RK (2006) QTL analysis of behavioral and morphological differentiation between wild and laboratory zebrafish (*Danio rerio*). *Behav Genet* 36:271–284.
- Krause J, et al. (2013) Reality mining of animal social systems. *Trends Ecol Evol* 28:541–551.
- Hughey LF, Hein AM, Strandburg-Peshkin A, Jensen FH (2018) Challenges and solutions for studying collective animal behaviour in the wild. *Philos Trans R Soc B* 373:20170005.
- King AJ, Fehlmann G, Biro D, Ward AJ, Fürtbauer I (2018) Re-wilding collective behaviour: An ecological perspective. *Trends Ecol Evol* 33:347–357.
- Nathan E, Pentland A (2006) Reality mining: Sensing complex social systems. *Pers Ubiquitous Comput* 10:255–268.
- Levitt SD, List JA (2007) What do laboratory experiments measuring social preferences reveal about the real world?. *J Econ Perspect* 21:153–174.
- Spaiser V, Luzatti D, Gregoriou A, Advancing sustainability: Using smartphones to study environmental behaviour in a field-experimental setup. *Data Sci*, 10.3233/DS-180014.
- Burton-Chellew MN, Nax HH, West SA (2015) Payoff-based learning explains the decline in cooperation in public goods games. *Proc R Soc Lond B Biol Sci* 282:20142678.
- Couzin ID, Krause J, Franks NR, Levin SA (2005) Effective leadership and decision-making in animal groups on the move. *Nature* 433:513–516.
- Couzin ID, et al. (2011) Uninformed individuals promote democratic consensus in animal groups. *Science* 334:1578–1580.

48. Perna A, Grégoire G, Mann RP (2014) On the duality between interaction responses and mutual positions in flocking and schooling. *Movement Ecol* 2:22.
49. Butail S, Mwaffo V, Porfiri M (2016) Model-free information-theoretic approach to infer leadership in pairs of zebrafish. *Phys Rev E* 93:042411.
50. Mann RP, et al. (2014) A model comparison reveals dynamic social information drives the movements of humbug damselfish (*Dascyllus aruanus*). *J R Soc Interf* 11:20130794.
51. Rosenthal SB, Twomey CR, Hartnett AT, Wu HS, Couzin ID (2015) Revealing the hidden networks of interaction in mobile animal groups allows prediction of complex behavioral contagion. *Proc Natl Acad Sci USA* 112:4690–4695.
52. Faria JJ, et al. (2010) A novel method for investigating the collective behaviour of fish: Introducing 'robotfish'. *Behav Ecol Sociobiol* 64:1211–1218.
53. Stowers JR, et al. (2017) Virtual reality for freely moving animals. *Nat Methods* 14:995–1002.
54. Jones M, Love BC (2011) Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behav Brain Sci* 34:169–188.
55. Pérez-Escudero A, Rivera-Alba M, de Polavieja GG (2009) Structure of deviations from optimality in biological systems. *Proc Natl Acad Sci USA* 106:20544–20549.