

Follow the straggler: zebrafish use a simple heuristic for collective decision-making

Kevin Kadak, Noam Miller*.

Department of Psychology, Wilfrid Laurier University

* Corresponding author: Department of Psychology, Wilfrid Laurier University. 75 University Ave. West,
Waterloo, Ontario, N2L 3C5, Canada.

Email: nmiller@wlu.ca

ORCID: 0000-0001-9363-9919

This is a pre-publication version of a paper to be published in *Proceedings of the Royal Society B*. A link to the final version will be added here when it is available. The paper's DOI is: [10.1098/rspb.2020.2690](https://doi.org/10.1098/rspb.2020.2690)

Abstract

Animal groups often make decisions sequentially, from the front to the back of the group. In such cases, individuals can use the choices made by earlier ranks, a form of social information, to inform their own choice. The optimal strategy for such decisions has been explored in models which differ on, for example, whether or not agents take into account the sequence of observed choices. The models demonstrate that choices made later in a sequence are more informative but it is not clear if animals use this information or rely instead on simpler heuristics, such as quorum rules. We show that a simple rule “copy the last observed choice”, gives similar predictions to those of optimal models for most likely sequences. We trained groups of zebrafish to choose one arm of a Y-maze and used them to demonstrate various sequences to naïve fish. We show that the naïve fish appear to use a simple rule, most often copying the choice of the last demonstrator, which results in near-optimal choices at a fraction of the computational cost.

Keywords: Collective behavior; social information; optimal choice; zebrafish; quorum model

Background

Members of species that live in groups must often make collective decisions that may pit personal preferences against the need to maintain group cohesion. In some cases, individuals may have personal information – based on their own past experiences of the environment – that conflicts with their social information – the choices they observe other members of their group making. How best to balance these sources of information to arrive at a decision has been the subject of much theoretical research [1,2]. There has also been increasing interest in how animals actually combine or decide between information from these two sources [3, 4, 5].

In some species, groups are often spread out along the direction in which the group is traveling [6]. When this occurs, decisions on which direction to move in may be made more-or-less sequentially, from the front to the back of the group [7, 8]. Such cases have been extensively examined in the lab, often using groups of fish swimming in Y-mazes, as they afford an opportunity to accurately define what the social information available to each individual consists of: the choices made by the members of the group that are ahead of it and who therefore choose an arm of the maze before it does [8, 9, 10]. It is, in addition, possible that most decisions made within the context of a group are sequential to some degree. For example, individuals may sometimes delay making a decision in ambiguous cases to allow more social information to accumulate (e.g., [11]), or individuals may weight social information in relation to how recent it is, thus effectively sequentializing it [12]. In this case, even choices made by freely moving groups may depend on the use of sequential social information.

In addition to empirical studies, there have been several attempts to model the decision-making process in sequential-choice paradigms [8, 11-15]. Most of these models focus on the optimal choice that agents can make, given the social and personal information at their disposal (e.g., [14, 15]). These models fit experimental data well in most cases, but their specific predictions have rarely been tested. Here, we

empirically test a key distinction between two classes of widely-cited models of optimal choice: models that take choice sequence into account, and those that do not.

Pérez-Escudero and de Polavieja [13] proposed a model of optimal sequential choice based on a Bayesian method for combining personal and social information (see also [16]). This model has been tested empirically and provides a good fit for collective decision-making data from a wide range of species [8, 16-20]. However, the original model exists in two forms that differ in their assumptions about how agents encode social information. In the basic form of the model, agents' choices are based on the relative numbers of group members that they observe having made each possible choice before making their own choice. Assume, for example, an agent that observes four conspecifics choose between two options, A and B, before it makes its own choice. If an agent in this situation observes that 3 conspecifics chose option A and only 1 chose option B, the agent – in the absence of any personal information that differentiates the two options – should be biased towards choosing option A. We call this version of the model “sequence-unaware” (SU). This version of the model is the one most often cited and used to fit experimental data.

In a more complex version of the model, which we call the “sequence-aware” (SA) model, agents also take into account the order in which previous choices occurred. Assume the same situation and distribution of choices as above: 3 to A, 1 to B. One possible sequences of choices that could generate this outcome is for the first conspecific to choose option A, followed by one that chooses option B, and then two in a row choosing option A, a sequence we denote {ABAA}. Alternatively, an agent could observe the sequence {AAAB}, in which the final conspecific contradicts the choices of the earlier 3. According to the SA model, agents in these two situations will have different choice biases because, in addition to considering the number of individuals at each option (which is identical in both cases), they also consider the information each conspecific is likely to have had at the time of making their choice.

This information is used to estimate the confidence of each conspecific in their choice, a proxy for the strength of their own bias (see [13] for details). For example, in the first sequence, {ABAA}, the lone dissenting individual (the second one to choose) has little social information at the time it chooses, based on the choice of only one previous individual. Its choice therefore suggests that it has personal information that biases it in favor of option B, strongly enough to overcome one conspecific-worth of conflicting social information. In the second sequence, {AAAB}, the dissenting conspecific is assumed to have observed all three individuals ahead of it choose option A. This individual's decision to nonetheless choose option B, contradicting all that social information, suggests it has a very strong bias in favor of option B. Assuming that conspecifics' choices are honest signals, this implies that the dissenting individual in the second sequence may have more (or better) personal information than the dissenting individual in the first sequence. The focal agent should therefore weight the information it gleans from the dissenter more heavily in the second sequence and thus be less biased towards option A in the second case. Note that the simpler SU model makes no distinction between the two sequences {ABAA} and {AAAB} because the overall number of choices to each option is the same in both cases. We note that this sequence-aware model assume that agents observe the choices made by those ahead of them in the group, and that each of those choices is made independently (see [13]). The model is thus only relevant for predicting choices based on sequential observations.

Mann [14] has presented an alternative model of collective choice that also considers the sequence of choices a focal agent observes. This model, which we term the Rational Agent (RA) model, assumes that agents make rational choices, i.e., those most likely to lead to the best outcomes given the personal and (possibly sequential) social information at their disposal. In this model, as in the SA model, the sequence of previous choices that an agent observes has a large effect on its likelihood of choosing each possible option. Mann [21] recently extended this model to show that changes in environmental conditions, such as (perceptual) noise levels, can affect the types of rules that agents use. For example, agents

habituated to noisy environments, where perceptual information may be unreliable, should follow rules that more strongly favor the most recent decisions they observed. We note that all the fish in our experiments had spent several months in the lab; we speculate on the possible consequences of this for their collective decision-making, in line with the predictions of the RA model [14, 21], in the discussion.

Keeping track of a sequence of choices made by others may be cognitively demanding, especially as the sequence increases in length. In many cases where there is an optimal solution to a problem which is computationally difficult, animals (including humans) resort to simpler heuristics that will, under most circumstances, give a similar or identical answer to the full rule [22, 23]. Such “rules of thumb” also occur in social situations [24]. One such simple rule is a quorum rule, by which individuals will follow any sub-group that exceeds in number some threshold. Quorum-like decision-making rules have frequently been detected in the collective behaviors of animal groups (e.g., [25]), including fish [9, 10, 26, 27], and have previously been used to model data from sequential-choice Y-maze experiments [25]. We therefore also compared a quorum-based (QU) model to our data.

An equally simple alternative to quorum rules is to select a single individual to copy. In a sequence of choices, the final choice – the one made by the conspecific immediately in front of the focal individual – is the most informative because, as noted above, that individual itself had access to the most social information. Animals could, therefore, use the following simple rule in complex sequential-choice situations: copy, with some fidelity, the choice made by the most recent conspecific. We call this the last-choice (LC) model. There is some experimental evidence that fish choosing when to move between two shelters follow this sort of heuristic [12]. More complex versions of this type of rule, such as making a choice based on the last n conspecifics observed, have also been suggested [12].

Most other agent-based models of collective movement define a set of rules followed by all agents, often synchronously (i.e., all agents decide which direction to move in at the same moment, then all

locations are updated at once; [28]; but see [29]). These rules are often ‘zonal’, meaning that agents’ responses to conspecifics depend on their relative distances from each other. Some zonal models have been modified to include agents moving towards (or away from) an external target [30, 31]. Zonal models are explicitly non-sequential but, as noted above, it is nonetheless possible that in many cases animals in freely-interacting groups sequentialize their movement decisions [12, 32], potentially using one or another of the rules discussed above.

In their presentation of the SU and SA models, Pérez-Escudero and de Polavieja [13] noted that in most cases the SA model makes very similar predictions to the simpler SU model, and that both fit the data they compared them to equally well. This lack of difference occurs because sequences like {AAAB} seldom occur in empirical data from freely moving groups, since individuals rarely contradict each other’s choices to such a degree. Miller et al. [8] suggested that this is at least partly because individuals prefer to remain in a cohesive group, for reasons independent of the informational value of those choices about the environment (see also [5, 20]). This is yet another reason why researchers have mostly relied on simpler models that ignore choice sequences, such as the SU model or quorum-based rules. To the best of our knowledge, the distinction between the two classes of model has not previously been empirically tested. Here, we test these model using zebrafish (*Danio rerio*), a highly social species that is commonly used in experiments on social decision-making and collective choice [5, 16, 33]. We compared the choice behaviors of groups of zebrafish to four leading models of collective choice described above: The SA and SU models [13], the RA model [14], and a quorum model (e.g., [27]), as well as the simple heuristic LC model.

We trained two separate groups of zebrafish to consistently choose either the left or right arm of a Y-maze (Figure S1). We call these fish “demonstrators”. We then selected 4 demonstrators from the two groups and released them into the maze in a predetermined order, approximately 1 second apart,

allowing us to construct any desired sequence of social information. When all four demonstrators had chosen one of the maze arms, we released into the maze a naïve fish, which we call the “observer”, that had no previous experience of either the maze or the demonstrator fish, and recorded its choice of arm. By using trained demonstrators, we were able to expose observer fish to sequences that are unlikely to occur under natural conditions and measure patterns of observed choice for each sequence.

Methods

Subjects and housing

80 adult wild-type zebrafish, obtained from a local pet store (Big Al's, Kitchener, Ontario, Canada), served as observers (test subjects) in the experiment. A further 92 zebrafish served as demonstrators, half trained to choose the left arm of the maze, and half the right. Finally, a further 20 fish were used as a social reward in training the demonstrators (see below). All the fish were housed in a high-density rack (Pentair Aquatic Habitats), in 10 L tanks, in groups of no more than 10. Observers and demonstrators were never housed in the same tank. The water in the rack was maintained at 25 ± 2 °C, the salinity between 600 – 1100 ppm. The housing room was on a 12:12 h light cycle with lights on at 7 am each day. All testing was done between 10 am and 3 pm. Fish were fed ad lib. each day following the completion of trials, on defrosted brine shrimp or flake food. All procedures complied with Canada Council on Animal Care guidelines and were approved by the Wilfrid Laurier University Animal Care Committee.

Apparatus

Fish were tested in a Y-maze constructed of PVC boards. The maze was placed inside a 6' X 6' tank that was filled with water matching the characteristics of the housing tanks, to a depth of 10 cm. In the tank, outside the maze, were heaters, filters, and bubblers that maintained the water quality. The filters and

bubblers were turned off during trials. The walls of the maze were 22 cm high; other dimensions of the maze are given in Figure S1. At the end of each arm of the maze there was a wider area, blocked off with a transparent plastic barrier, to contain the reward fish (see below). Along the stem of the Y-maze, where fish began each trial, we attached a series of 4 remotely-controlled transparent doors, creating 4 separate start chambers along the base of the maze (Figure S1). Doors were made of transparent acrylic sheets that spanned the width of the maze. Each door was glued to the tray of a hacked CD-ROM drive that was mounted above the maze, such that closing the CD-ROM drive retracted (opened) the door. Drives were controlled from a set of switches placed outside the apparatus. A pair of fluorescent lights were mounted along the sides of the stem of the maze to increase visibility and make the stem brighter than the ends of the arms, which we expected would increase fish's motivation to move down the maze. For training trials, a blue plastic floating feeding ring was affixed to the end of each arm of the maze.

Above the maze, we mounted a video camera (Sony HDR-CX900) such that the entire maze was visible. A second camera (Logitech C920 Webcam) was mounted above the far end of the maze so that we could see into the start chambers. Videos from both cameras were recorded for all test trials. The tank that the maze was in was isolated from the rest of the room by a white shower curtain hanging from the ceiling.

Procedure

Tagging: To identify the fish, all demonstrators and observers were tagged with fluorescent visible implant elastomer (VIE; Northwest Marine Technology) injected under their skin at least one week before the start of the experiment. Demonstrators trained to go to opposite arms were tagged in different colors. Observers were given unique tag combinations for their tank, so they could be individually identified. Reward fish were not tagged.

Training trials: In the first phase of the experiment, demonstrators were trained to choose one arm of the maze. Groups of 6-10 demonstrators at a time were placed in the stem of the maze. Only the front door, closest to the choice arms, was closed. 6 floating food pellets (Hikari beta-bio gold) were placed into the feeding ring at the arm designated as correct for this group. At the end of that arm, behind the transparent barrier, 6 reward fish were placed, to serve as an additional social reward. Demonstrators were left in the start box for 1 minute and then released into the maze. Once all the fish had chosen an arm, they were gently removed from the maze and returned to their home tanks. Demonstrator fish received three training trials per day for 12 days before the start of testing.

Testing trials: for testing trials, no reward fish were present and the feeding rings and food were removed. All 4 doors in the stem of the maze were closed. We placed demonstrators in the front 3 compartments, chosen to create the desired sequence for each test. For example, to obtain the sequence {AABA}, we placed two fish trained to choose arm A in the first compartment, one trained to choose arm B in the second, and one trained to choose arm A in the third. In all trials, a single observer fish was placed in the back compartment, alone. Fish were left in the start compartments for 1 minute, after which the first door was opened. As soon as the fish in that compartment had moved towards the choice arms, the next door was opened, and so on. Door openings were separated by about a second, as all the fish moved quickly down the maze. Trials did not conclude until the observer fish entered an arm of the maze. At the end of the trial, all the fish were gently removed from the maze and returned to their respective home tanks.

Each observer was tested a maximum of 4 times (8 fish only completed one trial, 8 only completed 2 trials, 24 only completed 3 trials), with at least one week between trials. Each demonstrator was only used for one test trial per day. Demonstrators received additional rewarded training trials interspersed

with the test trials, to maintain their preferences. We completed a total of 257 test trials. The raw choice data are available on the Open Science Framework repository: <https://osf.io/59sd2/>.

In many cases, despite extensive training, demonstrators did not choose the arm they were trained to go to. We coded each trial by the sequence of choices that was actually observed. In some cases, if a demonstrator chose an arm and then left it (and sometimes chose a different arm) before the observer chose any arm, we considered only each demonstrator's last choice as part of the sequence. Any further movement between the arms that occurred after the observer made its first arm choice was ignored.

Analysis

Videos were manually coded by one of the authors (KK). For each video, we coded the sequence of choices made by the demonstrators, followed by the choice made by the observer. Data were entered into Microsoft Excel and were analyzed using *Mathematica* (Wolfram Research, version 10.0) and R [34]. For each possible sequence, we coded the proportion of trials on which the observer chose the majority arm, defined as the arm chosen by the majority of the demonstrators. In sequences without a defined majority (e.g., {AABB}), we designated the first choice as the majority arm.

We fit five models to our choice data. The SU and SA models were implemented in *Mathematica*, using equations and code given in [13]. We note that, in our experiments, observer fish have no differential personal information about the arms of the maze (i.e., no reason based on past personal experience of the maze to prefer one arm over the other), and are never rewarded in the maze, which considerably simplifies the models. The SU model, when there is no differential personal information (and assuming that the arms of the maze are inherently identical), has one free parameter, denoted s in [13]. The SA model, under the same conditions, also has one parameter, denoted α' in [13].

The RA model was implemented in R, using code provided by Richard Mann (personal communication). This model, when there is no differential personal information, depends only on the ratio of the noisiness of environmental information under the experimental conditions, denoted η in [14], and the equivalent noisiness under natural conditions, denoted v . For this model, we therefore fit the ratio η/v to our data.

The QU model was implemented in *Mathematica* as a simple rule in the form of a Hill equation: $P(x) = N_x^k / (N_x^k + N_y^k)$, where N_x and N_y represent the number of individuals already at each of the arms x and y , and k is a free parameter that we fit to our data, representing the steepness of the function.

Finally, in the LC model, an agent's probability of selecting the arm chosen by the last demonstrator is given by: $P(L) = 1 / (1 + e^{-a})$, where a is a parameter that determines how reliably agents copy choices. Note that this model – like the other models we considered – is probabilistic: fish do not always copy the choice of the last demonstrator.

We fit the single free parameter of each model to our experimental data and calculated the AIC for each model (Table 1).

We next compared the predictions of the different models for all possible sequences of choices up to a length of 8, independently of our data (in which we had 4 demonstrators on all trials). To derive the probability of a particular sequence of social information occurring, we iterated the models over each possible sequence. Since our observer fish have no differential personal information about the maze, we assume that the likelihood of the single-choice sequence {A} is 0.5. We can then derive the likelihood of the sequence {AB} by using each model's equation for the probability of choosing B given the sequence {A}, and so on for sequences of any length. We compared predictions of all possible sequences of lengths up to 8. For each sequence, we derived the probability that it would occur, by each of the four comparison models (SA, SU, RA and QU), and the choice probability of the agent observing that

sequence by both the comparison and LC models. The parameters used for the models were those that best fit our data (Table 1). We then summed the likelihoods of all sequences in which the LC and comparison models gave opposing predictions (one > 0.5 and one < 0.5).

Results

We first examined the choices of our observer fish for side bias. The proportion of each subject's choices to the left arm of the maze was calculated (subjects that only completed a single trial were removed from this analysis). A one-sample Wilcoxon signed-ranks test found no evidence of side bias in our fish ($W = 1175.5$, $p = 0.052$), so we combined sequences across the two sides of the maze. For example, we considered the sequence "right, left, left, left" to be the same as "left, right, right, right", and denote both {BAAA}.

Despite their training, demonstrators did not always choose the arm they were trained to. As a result, observers were exposed to some sequences more often than others (Figure 1, green numbers below the x-axis). We calculated the proportion of trials of each sequence on which observers chose the majority arm. For sequences where there was no majority, such as {AABB}, we designated the first demonstrator's choice the majority arm. We then regressed several models of sequential choice against our data. We chose five models: the original SU and SA models [13]; a recent model of social decision-making by fully rational agents [14], which we label RA; a quorum-rule model (QU); and our heuristic LC model, which states that observers simply match the choice of the last demonstrator in the sequence (see Methods for details). All models were fit to the raw choice data, weighted by the frequency with which each sequence occurred. Table 1 shows the best-fit parameters, model likelihoods, and AIC scores for each model. Figure 1 shows the data (black dots) and best-fit version of each model (lines).

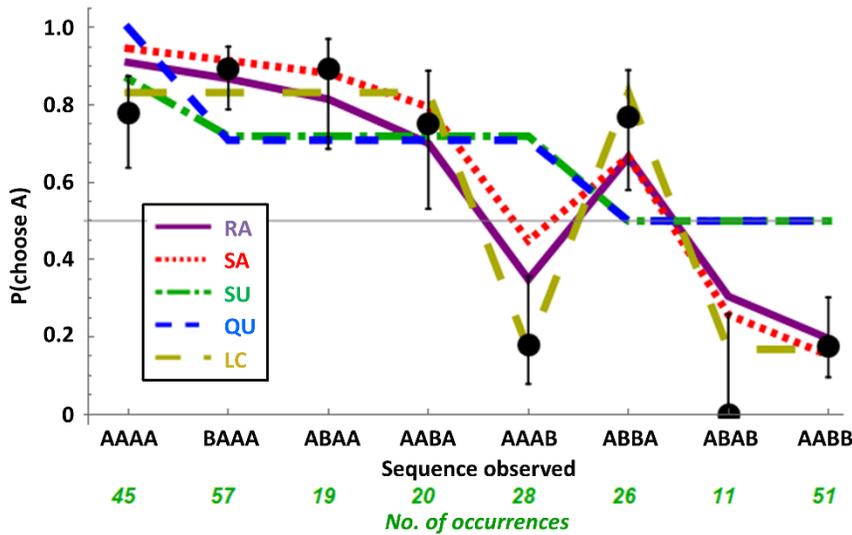


Figure 1. Experimental results and model fits. Black dots indicate the proportion of choices made by the observers to the majority arm for each demonstrated sequence. Lines show model fits for the Rational-Agent (RA; solid purple line), Sequence-Aware (SA; dotted red line), Sequence-Unaware (SU; dash-dotted green line), Quorum (QU; dashed blue line), and Last Choice (LC; long-dashed yellow line) models. Error bars show Wilson score intervals (binomial confidence intervals) on the empirical data. The thin horizontal grey line indicates chance levels. Green numbers below the x-axis show the number of times each sequence occurred in the data.

We found that models that depend on choice sequences (SA, RA, and LC) fit our data much better than the models that ignored sequences (SU and QU; AIC scores are given in Table 1). In particular, the SU and QU models failed to accurately predict choices on trials where there was no majority ({ABBA}, {ABAB}, and {AABB}; rightmost three points in Figure 1). In such cases, these models predict that the naïve observer fish should be equally likely to choose either arm. However, as the data show (black dots in Figure 1), fish in these situations displayed strong biases for one option or the other. Additionally, these models failed to predict the behavior of fish in the most extreme sequence, {AAAB}, where a single dissenting fish contradicts all previous choices. The simple models, in this case, predict a relatively strong preference for arm A, since there are three demonstrators there compared to only one in arm B. However, observers rarely chose arm A under these conditions (5 out of 28 trials).

Table 1. Model fits.

Model	Value	Log Likelihood	AIC
Rational-Agent (RA)	$\eta/v = 1.23$	-122.65	247.29
Sequence-Unaware (SU)	$s = 1.46$	-160.03	322.06
Sequence-Aware (SA)	$a' = 4.64$	-127.80	257.60
Quorum (QU)	$k = 0.81$	-204.81	411.62
Last-Choice (LC)	$a = 1.60$	-116.06	234.12

The table shows the best-fit parameter value (Value), Log Likelihood, and Akaike Information Criterion (AIC) score for each model's fit to the experimental data.

In all four of the sequences where the SU and QU models fail, the models that do account for sequence (SA, RA, and LC) predict the direction of the results correctly. All three models, for example, correctly predict that fish that observe the sequence {AAAB} are more likely than not to choose option B, since the dissenting demonstrator is signaling that it has a very strong bias in favor of B, and the models therefore weight that demonstrator's information more heavily. Similarly, the final demonstrator – who confirms or contradicts the most social information – has an outsized effect on the observer's responses in the three sequences that lack a majority choice in these models. We found that the LC model fit the data best of all five models that we tested (Table 1). This suggests that fish in our experiment might use a simple heuristic to decide which demonstrators to follow: copy the choice of the last demonstrator.

The parameter values that we found best fit the models to our data were generally similar to those previously reported for groups of fish. Miller et al. [8], using the SU model, found a value of $s = 1.5$ best fit their data on golden shiners (*Notemigonus crysoleucas*), close to the value we find for our data ($s = 1.46$). However, Pérez-Escudero and de Polavieja [13] found a slightly higher average value of $s = 2.5$ fit three datasets on three-spined sticklebacks (*Gasterosteus aculeatus*). Similarly, Pérez-Escudero and de Polavieja [13], using the SA model, reported a value of $\alpha' = 5$ for the same datasets, close to our value ($\alpha' = 4.64$).

To further explore the utility of the LC model, we compared the predictions of the model to those of the other four models. In other words, we asked: how often does following the final demonstrator lead to making the same choice as would be predicted by the other models? Figure 2 shows the predicted probability of choosing A by the SA (red circles), SU (green squares), RA (purple triangles), and QU (blue diamonds) models for all possible sequences of lengths up to 8 that end in a choice of option A. We weighted each sequence tested by its likelihood of occurring under that model (see Methods for details);

probability is indicated by size of symbol in Figure 2). The LC model returns a high chance of choosing option A for these sequences (about 80%, when fit to our data) and, as Figure 2 shows, the other models concur for the most common sequences (larger symbols). We counted the number of sequences in which each comparison model predicted a probability < 0.5 , i.e., in which it contradicted the prediction of the LC model. We then weighted this number by the probability that each sequence would occur (see Methods). The RA model always gave the same response as the LC model (i.e., the predictions made by the two models for any sequence were always either both > 0.5 or both < 0.5). In other words, at least for sequences of up to 8 individuals, following the last individual gives the same general answer (the choice bias is in the same direction) as computing the result of the RA optimal choice model. The SA and LC models gave conflicting predictions for 8 (of 255) sequences, representing 1.48% of likely events. The SU and QU models diverged from the LC model on 64 sequences, representing 11.25% and 0.005% of likely events, respectively ('unusual' sequences, which tend to be where the two models give divergent predictions, are far less likely to occur under the QU model). So, the LC model serves as a good heuristic for both sequence-aware models, but less so for the models that ignore sequence.

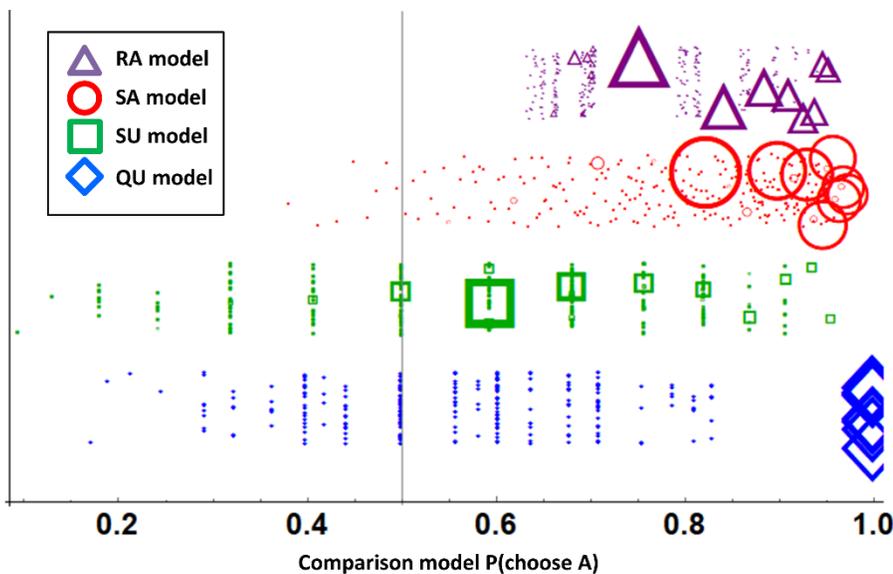


Figure 2. Comparing heuristic predictions. Comparison of the predictions of the LC model and those of the other four models for all possible sequences of length up to 8 that end in choice A (i.e., for which the LC model predicts choosing A; 255 sequences). RA model: purple triangles; SA model: red circles; SU model: green squares; QU model, blue diamonds. Positions of data along the y-axis are random, for ease of viewing. The position of each point along the x-axis represents the prediction of the corresponding comparison model. Symbol sizes reflect the likelihood of each sequence (the smallest symbols have been enlarged so that they are visible).

Discussion

By training zebrafish to choose one arm of a Y-maze and releasing them into the maze in a specific order, we were able to present naïve fish with controlled sequences of social information. We then fit five models to these data. We show, first, that sequence-aware models (SA, RA, and LC) fit our data much better than the simpler models (SU and QU). In other words, zebrafish, when making a choice between two options about which they possess no differential personal information, consider not only the number of individuals that have previously chosen each option, but also incorporate information about the sequence in which those choices were made. Their choices are affected by the information that each earlier-choosing conspecific, “demonstrator” in our experiment, was likely to have had. Demonstrators who contradict a larger amount of social information are assumed to have a stronger personal bias in favor of the option they selected, and their choice is therefore weighted more heavily in the observer’s own decision-making process.

The model that best fit our data was one according to which observer fish simply copied the choice of the last demonstrator they watched before making their own choice, a simple heuristic. Since the last demonstrator in any group has the most social information, its choice is the most informative according to all the sequence-based models. Thus, copying this choice – as our fish appeared to do – may be a computationally cheap way of achieving results similar to those given by sequence-dependent optimal choice models, such as the SA and RA models (as also suggested by [12]). Comparing the models to each other, we found that, at least for sequences of length up to 8 individuals, this simple heuristic suggested the same choice as the more complex models in almost every instance (always for the RA model, and 98.5% of the time for the SA model).

Mann [14, 21] noted that environmental information in the lab will tend to be less noisy (more reliable) than outside the lab, and we believe our lab to be no exception to this rule (e.g., subjects housing tanks

and the maze were bare of any plants or rocks; maze walls were uniformly white). Mann further suggested that in situations where there is less environmental noise (η) than in the “habitual” environment (v), social effects increase, including a preference for the most recent demonstrated choices, which is in keeping with our findings. We note that we found the best fit of the RA model to our data at a value of $\eta/v = 1.23$, suggesting that the noise in the maze was slightly higher than subjects’ habitual environment. It is possible that this is due to our fish having lived for a long time in the lab before testing, or reflects a generalized increased sociability in response to novelty (the maze; see e.g., [8]). Alternatively, if our fish are using a heuristic similar to the LC model, copying the last demonstrator they observe, this would also lead to a higher η/v value fitting the data best.

Our heuristic model only ever returns one of two probabilities, depending on the final choice the choosing agent observes, whereas the computationally sophisticated models return a range of probabilities. This suggests that using the LC or a similar heuristic entails a loss of flexibility in decision-making. For example, in most cases fish might possess some personal information about the choices before them based on past experience. Combining this personal information with social information – a process that all three computational models (SU, SA, and RA) address – may be more difficult when using a heuristic model, and might result in different choices. For example, upon observing the sequence {AAAB}, the SA model (using the best-fit parameter value for our data) predicts a 55% chance that the next individual will choose option B (this number is 68% for the RA model). These values are close enough to 50% that a small amount of personal information that favored choosing option A might be sufficient to alter the choice made by an agent. The LC model, for the same sequence, predicts a strong bias for option B (since that is the final observed choice in the sequence), which might be less likely to be overruled by personal information. We note, however, that this sequence is relatively unlikely to occur (using each model iteratively to predict the probability of the sequence, see Methods; $P = 0.026$ by the SA model, and $P = 0.025$ by the RA model). For most of the more likely sequences (larger symbols in

Figure 2), the computational models give much higher probabilities that are similar to the strength of bias in the LC model. There are also several different ways in which personal information could be combined with the predictions of the LC model, and we take no position on which of those is more likely. Our data, in which observers had no prior experience of the maze, afford no opportunity to test this aspect of the decision-making process.

Other computationally cheap heuristics exist that animals could use to make collective decisions. For example, animals could consider only the last two, or n , choices they observe (a family of heuristics of which our LC model is the simplest member; e.g., [12]). The SU model, which considers only the number of conspecifics at each choice, could also be considered a heuristic rule as it ignores some of the available information (the sequence of choices).

It is possible that some detail of our experimental paradigm, such as the timing between demonstrator choices, influenced the rule that our test fish used. Demonstrators in our experiment were spaced about 1 second apart, and the observer was released about 1 second after the last demonstrator had chosen an arm of the maze (which, in our maze, did not take very long). It is possible that increasing the gap in time between consecutive demonstrators or, more likely, between the demonstrators and the observer, would change the behavior of the observers. That is, following closely on the heels of the final demonstrator may encourage observers to simply copy that final choice, rather than integrating the information they gleaned from observing all the demonstrators. However, increasing the inter-choice interval might also challenge observers' memory, particularly of longer or more varied demonstrated sequences. We also note that the intervals we enforced between the fish's choices were already slightly longer than those commonly observed in freely-moving groups of fish (e.g., [8]).

Fish, and other animals that make collective decisions, may convey more information to other members of their group than simply which option they choose. For example, animals may communicate their

confidence in their choice [35, 36] by their body movements or the speed of their choice or, in humans, verbally [37]. Collective decision rules that take confidence information into account lead to more accurate choices, at least in humans [37, 38]. We did not measure the choice latencies of our demonstrators or observers, and cannot gauge the effect, if any, of such additional information on their choices. As far we are aware, no data exist on the kinds of behaviors fish might use to communicate confidence or whether they do so. Finally, we also note that our demonstrators often did not choose the arm they were trained to go to (Figure S3); it is possible that their swimming on such trials was somehow different and that this affected how strongly observers relied on those choices.

In summary, whether fish copy the last choice they observe or perform a more complex integration of social information sequences, they do not appear to be using sequence-independent decision rules, such as those implemented by the commonly-cited SU model [13] or by quorum rules [9, 25, 27]. Though the predictions of these models may be quite similar for most naturally occurring sequences of choices, since these are rarely strongly biased, we recommend further research to explore the consequences of sequence-dependent behavioral choice mechanisms, and of simple heuristics that may approximate the same behavior at a fraction of the computational cost.

Acknowledgments

This research was funded by the National Science and Engineering Research Council of Canada (NSERC) grant RGPIN-2016-06138 (to NM). The authors wish to thank Kelley Putzu for animal care; Alexandra Mahon, Rhyon Rodrigues, Chanpreet Grewal, and Hailey Katzman for assistance with running the experiments; and Alfonso Pérez-Escudero and members of the Collective Cognition Lab for helpful discussions of the work. We also wish to thank Richard Mann for extremely helpful comments on an earlier version of the manuscript.

Data Accessibility

All the data used in this article are available on our Open Science Framework repository:

<https://osf.io/59sd2/>.

Competing interests

The authors declare that they have no competing interests.

Author Contributions

Both authors conceptualized the study, analyzed the data, and wrote the paper. KK performed the experiments and coded the data from the videos.

References

1. Danchin É, Giraldeau L-A, Valone TJ, Wagner RH (2004). Public information: from nosy neighbors to cultural evolution. *Science*, 305:487-491.
2. Kendal RL, Boogert NJ, Rendell L, Laland KN, Webster M, Jones PL (2018). Social learning strategies: bridge-building between fields. *Trends in Cognitive Sciences*, 22: 651-665.
3. Pike TW, Laland KN (2010). Conformist learning in nine-spined sticklebacks' foraging decisions. *Biology Letters*, 6:466-468.
4. Kendal RL, Coolen I, van Bergen Y, Laland KN (2005). Trade-offs in the adaptive use of social and asocial learning. *Advances in the Study of Behavior*, 35:333-379.
5. Ayoub R, Armstrong E, Miller N (2019). Out of sight, out of mind: mechanisms of social choice in fish. *Animal Behavior*, 155:163-169.
6. Hemelrijk CK, Hildenbrandt H, Reinders J, Stamhuis EJ (2010). Emergence of Oblong School Shape: Models and Empirical Data of Fish. *Ethology*, 116: 1099-1112.
7. Bumann D, Krause J (1993). Front individuals lead in shoals of three-spined sticklebacks (*Gasterosteus aculeatus*) and juvenile roach (*Rutilus rutilus*). *Behavior*, 125:189–198.
8. Miller N, Garnier S, Hartnett A, Couzin ID (2013). Both information and social cohesion determine collective decisions in animal groups. *PNAS*, 110: 5263-5268.
9. Ward AJW, Sumpter DJT, Couzin ID, Hart PJB, Krause J (2008). Quorum decision making facilitates information transfer in fish shoals. *PNAS*, 105:6948–6953.
10. Sumpter DJT, Krause J, James R, Couzin ID, Ward AJW (2008). Consensus decision making by fish. *Current Biology*, 18:1773–1777.
11. Bikhchandani S, Hirshleifer D, Welch I (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100:992-1026.

12. Mann RP, Herbert-Read JE, Ma Q, Jordan LA, Sumpter DJT, Ward AJW (2014). A model comparison reveals dynamic social information drives the movements of humbug damselfish (*Dascyllus aruanus*). *J. R. Soc. Interface*, 11:20130794.
13. Pérez-Escudero A, de Polavieja GG (2011). Collective animal behavior from Bayesian estimation and probability matching. *PLoS Comp. Bio.*, 7(11):e1002282.
14. Mann RP (2018). Collective decision making by rational individuals. *PNAS*, 115:E10387-10386.
15. Kao AB, Miller N, Torney C, Hartnett A, Couzin ID (2014). Collective learning and optimal consensus decisions in social animal groups. *PLoS Computational Biology*, 10:e1003762.
16. Arganda S, Pérez-Escudero A, de Polavieja GG (2012). A common rule for decision making in animal collectives across species. *PNAS*, 109:20508-20513.
17. Aplin LM, Farine DR, Mann RP, Sheldon BC (2014). Individual-level personality influences social foraging and collective behaviour in wild birds. *Proceedings of the Royal Society B*, 281:20141016.
18. Eguíluz VM, Masuda N, Fernández-Gracia J (2015). Bayesian decision making in human collectives with binary choices. *PLoS One*, 10:e0121332.
19. Pérez-Escudero A, de Polavieja GG (2017). Adversity magnifies the importance of social information in decision-making. *J. R. Soc. Interface*, 14:20170748.
20. Pérez-Escudero A, Miller N, Hartnett AT, Garnier S, Couzin ID, de Polavieja GG (2013). Estimation models describe well collective decisions among three options. *PNAS*, 110:E3466-E3467.
21. Mann RP (2020). Collective decision-making by rational agents with differing preferences. *PNAS*, 117:10388-10396.
22. Laan A, de Sagredo RG, de Polavieja GG (2017). Signatures of optimal control in pairs of schooling zebrafish. *Proceedings of the Royal Society B*, 284: 20170224.
23. Gigerenzer G, Gaissmaier W (2011). Heuristic decision making. *Annual Review of Psychology*, 62:451-482.

24. Hutchinson JMC, Gigerenzer G (2005). Simple heuristics and rules of thumb: where psychologists and behavioural biologists might meet. *Behavioural Processes*, 69:97-124.
25. Sumpter DJT, Pratt SC (2009). Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society B*, 364:743,753.
26. Ward AJW, Herbert-Read JE, Sumpter DJT, Krause J (2011). Fast and accurate decisions through collective vigilance in fish shoals. *PNAS*, 108:2312-2315.
27. Ward AJW, Krause J, Sumpter DJT (2012). Quorum decision-making in foraging fish shoals. *PLoS ONE*, 7(3):e32411.
28. Parrish JK, Viscido SV, Grünbaum D (2002). Self-organized fish schools: an examination of emergent properties. *Biological Bulletin*, 202:296-305.
29. Bode NWF, Faria JJ, Franks DW, Krause J, Wood AJ (2010). How perceived threat increases synchronization in collectively moving animal groups. *Proceedings of the Royal Society B*, 277:3065-3070.
30. Couzin ID, Ioannou CC, Demirel G, Gross T, Torney CJ, Hartnett A, Conradt L, Levin SA, Leonard NE. Uninformed individuals promote democratic decisions in animal groups. *Science*, 334:1578-1580.
31. Ballerini M, Cabibbo N, Candelier E, Cavagna A, Cisbani E, Giardina I, Lecomte V, Orlandi A, Parisi G, Procaccini A, Viale M, Zdravkovic V (2008). Interaction ruling animal collective behavior depends on topological rather than metric distance: evidence from a field study. *PNAS*, 105:1232-1237.
32. Strandburg-Peshkin A, Twomey CR, Bode NWF, Kao AB, Katz Y, Ioannou CC, Rosenthal SB, Torney CJ, Wu HS, Levin SA, Couzin ID (2013). Visual sensory networks and effective information transfer in animal groups. *Current Biology*, 23:R709-R711.
33. Stevens JR, King AJ (2013). The lives of others: social rationality in animals. In: *Simple Heuristics in a Social World*, (ed. Ralph Hertwig, Ulrich Hoffrage, and the ABC Research Group). Oxford University Press.

34. R Core Team (2016). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria.
35. Korman A, Greenwald E, Feinerman O (2014). Confidence sharing: an economic strategy for efficient information flows in animal groups. *PLoS Computational Biology*, 10(10):e1003862.
36. Marshall JAR, Brown G, Radford AN (2017). Individual confidence-weighting and group decision-making. *Trends in Ecology and Evolution*, 32:636-645.
37. Kurvers RHJM, Wolf M, Naguib M, Krause J (2015). Self-organized flexible leadership promotes collective intelligence in human groups. *Royal Society Open Science*, 2:150222.
38. Bahrami B, Olsen K, Latham PE, Roepstorff A, Rees G, Frith CD (2010). Optimally interacting minds. *Science*, 329:1081-1085.

Electronic Supplementary Material

Proceedings of the Royal Society B

Follow the straggler: zebrafish use a simple heuristic for collective decision-making.

Kevin Kadak, Noam Miller.

Department of Psychology, Wilfrid Laurier University

DOI: 10.1098/rspb.2020.2690

Figure S1. Diagram of the Y-maze used in the experiment.

Figure S2. Photographs of the Y-maze used in the experiment.

Figure S3. Proportion of demonstrators that chose the rewarded arm on training trials, by training day.

Figure S1:

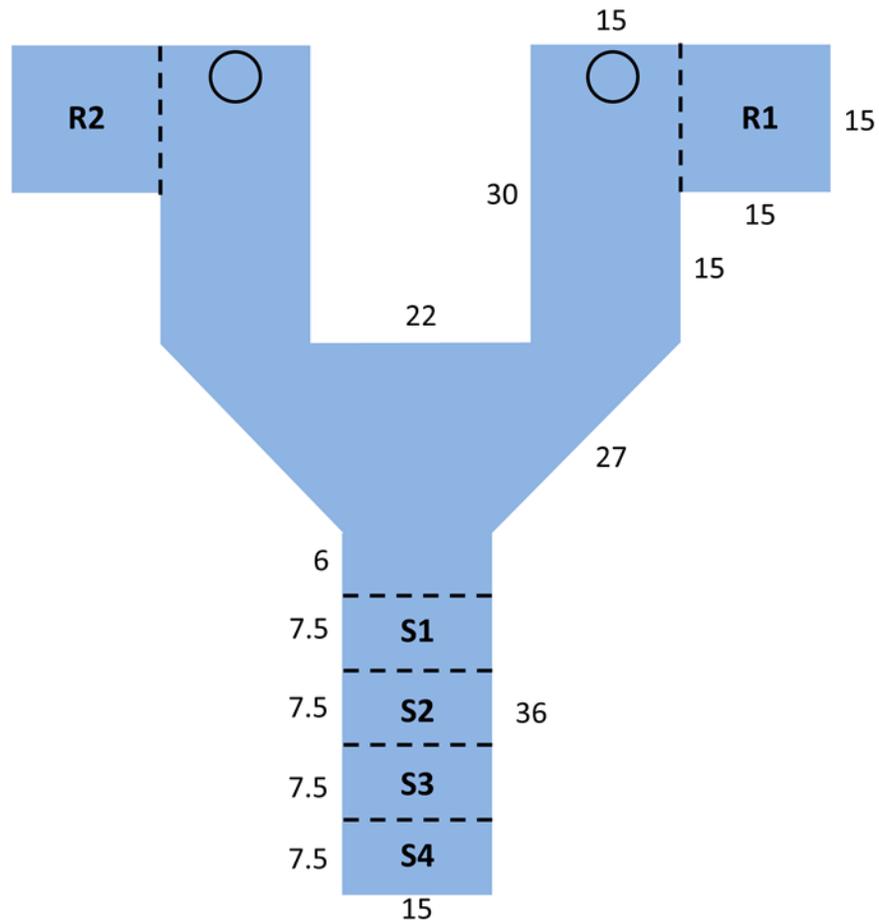
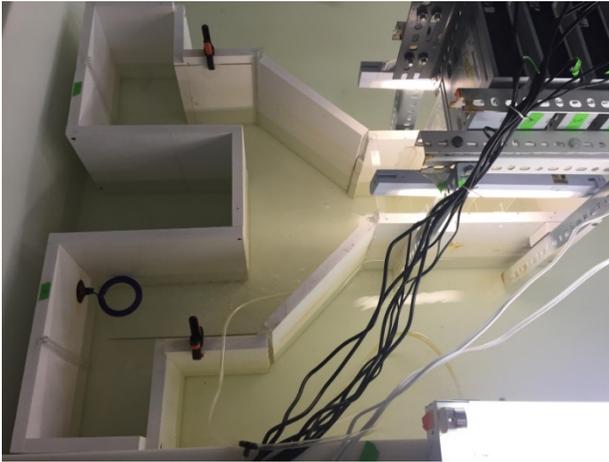


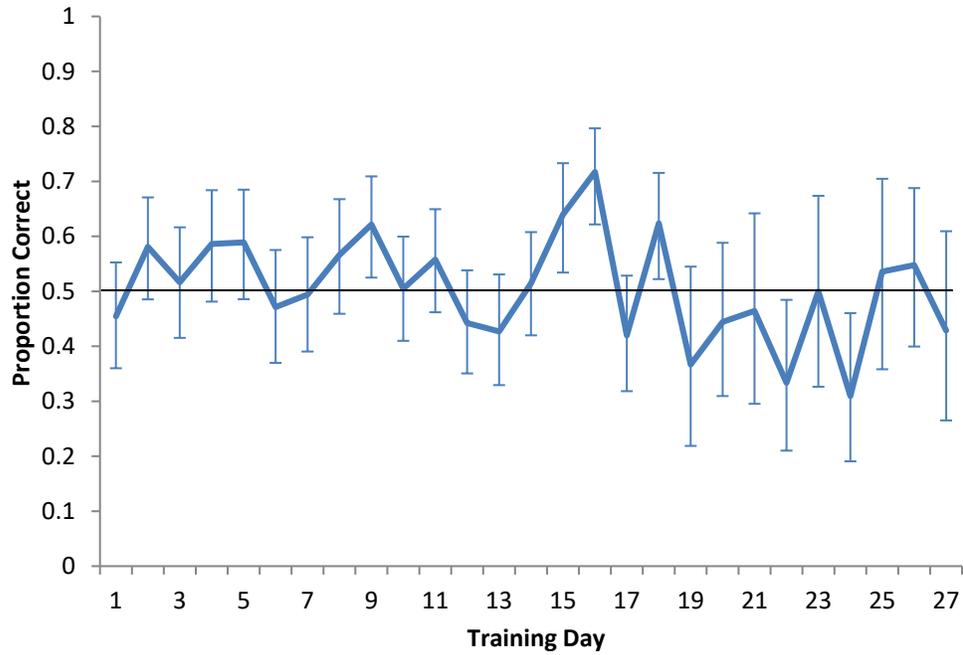
Diagram of the Y-maze used in the experiment. Numbers along the outside edges give the dimensions of the maze, in cm. The height of the maze walls was 22 cm. Dashed lines indicate transparent walls or doors. Circles at top indicate feeding rings (which were removed during test trials). R1, R2: reward zones; S1-S4: start boxes.

Figure S2:



Photographs of the Y-maze used in the experiment. Left panel: view of the maze from above showing the reward arms (at left), food reward ring, and CD-drives that operated the maze doors. Right panel: a view down the start arm of the maze showing the four transparent start doors (all open).

Figure S3:



Proportion of demonstrators that chose the rewarded arm on training trials, by training day. Starting from day 12, test trials which included observer fish were interspersed with the training trials. The horizontal black line shows chance levels; error bars show Wilson Score binomial confidence intervals.