Decision-making strategies of a coral reef fish given changing visual signal reliability

Cait Newport[1*], Lucas Wilkins[1], Guy Wallis[2], Joseph Bartram[1], Theresa Burt de Perera[1]

[1] Department of Zoology, University of Oxford, Oxford, United Kingdom

[2] Centre for Sensorimotor Performance, School of Human Movement and Nutrition Sciences, University of Queensland, Brisbane, Australia

Corresponding author: *caitlin.newport@biology.ox.ac.uk

Keywords

Adaptive decision making, learning, fish, cognition, vision

ABSTRACT

Making informed decisions is fundamental to intelligent life and is crucial to survival. The problem many organisms face, is how to make optimal decisions in the face of incomplete or conflicting information. There is mounting evidence that humans employ a Bayesian framework which optimally combines the statistical properties of current evidence with those of past experience. This, it is assumed, requires the operation of a sophisticated cortical system in which the statistics of past experience are stored. Here we demonstrate Bayesian updating behaviour in a fish performing a non-adaptive, visual discrimination task. We find that fish can learn probability distributions and apply this knowledge as uncertainty in the correct choice increases, ultimately increasing the overall frequency of correct choices. Our study underlines the ubiquity of Bayes-like behaviour in animals and, since fish have no cortex, implies that Bayesian integration can be performed by relatively simple neural circuitry. Furthermore, our results reveal that probabilistic decision rules can be used by fish when visual information is unreliable, indicating a possible mechanism for decision making given the inherent noise in incoming sensory information.

INTRODUCTION

Bayes theory has emerged as a framework to understand decision making in uncertain conditions (Chater et al. 2010; Perfors et al. 2011) and has become a cornerstone of theories of computational and cognitive neuroscience (Gershman 2015; Knill and Pouget 2004; Körding 2007; Körding and Wolpert 2006; Penny 2012; Tenenbaum and Griffiths 2002; Xu and Tenenbaum 2007). In humans, decisions based on the conscious or implicit likelihood of an outcome have been found to underpin a wide range of behaviours including sensory (e.g. Battaglia et al. 2003; Deneve and Pouget 2004; Knill and Pouget 2004) and spatial processing (e.g. Cheng et al. 2007), sensorimotor integration (Kording and Wolpert 2004; Körding and Wolpert 2006), and learning (e.g. Courville et al. 2006; Kruschke 2006; Kruschke 2008; Moens and Zénon 2019). This framework has proven tremendously powerful, broadly applicable, and has rapidly advanced our understanding of the rules that govern complex behaviour. In recent years, the application of Bayes theorem has been extensive in academic disciplines as diverse as economics, psychology, neuroscience, computer vision, and animal behaviour. To date, its application has been predominantly (although not exclusively) studied in humans and interpretations have typically been developed in the context of human neural processing with current models assuming a pivotal role for cortical systems (Funamizu et al. 2016; Lee and Mumford 2003; Rushworth and Behrens 2008; Yu and Dayan 2005). However, animals across all taxa, including those without cortices, face similar problems as humans (e.g. sensorimotor integration, learning, decision making) and evolution may favour similar underlying mechanisms despite differences in neural anatomy.

Bayesian inference is an application of Bayes' theorem which describes the integration of prior probability knowledge of a specific outcome with current incoming information, allowing the assessment of probability to be updated as more information becomes available. It provides a framework for optimal decision making in the face of unreliable information but requires that the decision-maker can learn and remember probabilistic information and alter their expectations based on context. To understand the generalizability of Bayesian inference,

we must determine whether some animals are constrained by their capacity for computational processing and/or they have developed rules that are more suitable to the types of decisions they face.

In non-human animals, empirical evidence of Bayesian inference has come from studies exploring whether animals change their behaviour given different prior sample information. In some cases, theoretical predictions of Bayesian inference have been compared to observed data of different behaviours including foraging (Alonso et al. 1995; Foley and Marjoram 2017; Olsson et al. 1999; Valone 1991; Valone and Brown 1989; van Gils et al. 2003), mate searching (Luttbeg and Warner 1999), collective behaviour (Gunji et al. 2018), and natal dispersal (Selonen and Hanski 2010). The behaviour of all animals tested (four birds, four mammals, one fish, one insect, one crab) is consistent with Bayesian inference, excluding the Gambel's quail (*Callipepla gambelii*) (Valone and Brown 1989). An alternative approach is to manipulate the priors and compare observed behaviours to specific hypotheses. Similarly to the previous type of studies, this has been mainly tested in the context of foraging (Biernaskie and Gegear 2007; Biernaskie et al. 2009; Lima 1984; Lima 1985; Marshall et al. 2013; Valone 1992; Valone and Giraldeau 1993), and one study using mating systems (Hunte et al. 1985). These studies broadly demonstrate that bees (Biernaskie and Gegear 2007; Biernaskie et al. 2009; Foley and Marjoram 2017), four species of bird, baboons (Marshall et al. 2013), and an amphipod can learn information about prior distributions, although in some cases experimental context played a role in learning (Marshall et al. 2013; Valone 1992). One criticism of the vast majority of previous studies is that the Bayes-like behaviour demonstrated may only appear to be Bayesian (Biernaskie et al. 2009; McNamara et al. 2006) and decision rules may actually be hard-wired through natural selection and shaped by context-specific factors (Ramírez and Marshall 2017). For example, animals adapted to stable foraging distributions may behave as if they know the prior distribution, but this behaviour would likely breakdown when animals are exposed to unusual situations (Biernaskie et al. 2009; McNamara et al. 2006). In most experiments, the distribution of food

items was not altered from their natural state, or distributions tested were similar to what might be naturally experienced. Biernaskie et al. (2009) addressed this methodological limitation by increasing the variation in resource distribution and reducing the learning time of priors, and found that the bees rapidly learned patch quality information and altered their behaviour in response. To test for Bayesian inference in animals robustly, a crucial next step is to test whether animals can integrate knowledge of prior probabilities in non-adaptive contexts (McNamara et al. 2006).

To understand the evolution of decision-making and the specific behaviours of different species, we need to know to what extent animals across taxa rely on probability-based decision making. Fish are an important group when considering the evolution of behaviour as they are highly varied taxonomically. In addition, they demonstrate complex behaviours such as cooperation (Bshary et al. 2006; Vail et al. 2014) and tool-use (Brown 2012; Dunn 2016; Pryor and Milton 2019), yet their brains are proportionally small and they lack a cortex (Northcutt 2002), the area of the mammalian brain generally associated with behavioural complexity. Relatively little is known about the underlying decision strategies which guide fish behaviour. There is ample evidence from psychophysical tests that a range of fish species can learn association-based if-then rules (e.g. Cheney et al. 2013; Newport et al. 2013; Schluessel et al. 2014a; Schluessel et al. 2014b; Siebeck et al. 2009; Sutherland 1969), can categorize stimuli based on some shared characteristic (Schluessel and Duengen 2014; Schluessel et al. 2012), and perhaps even learn abstract relationships (Newport 2020). Given the highly variable conditions of many aquatic ecosystems and the range of complex behaviours observed in fish, it seems reasonable to assume that their behaviours rely on equally complex decision rules. However, a combination of a relatively small number of simple decision rules which are applied given particular sensory input, may provide a solution that requires less cognitive power, whilst providing enough flexibility to respond to variable or changing conditions, within adaptive limits. If this were the case, then it would be expected

that fish decisions would not be optimized in conditions that are unusual or outside of their typical experience.

Using positive reinforcement, here we test if a species of fish can learn a simple associative rule (select the darkest stimulus) as well as an underlying probability distribution, and then combine the two sources of information whilst balancing uncertainty in the perceptual input. By doing so, the fish would maximize their probability of selecting the correct stimulus and would also maximise their overall number of food rewards. Our study differs from previous research as the task is abstract, and therefore non-adaptive. This enables us to test for Bayesian inference in a task that we would not expect the fish to have any innate priors for. By doing so, we are able to test, non-ambiguously, whether fish are able to make informed decisions using a Bayesian framework to deal with the inherent noise in the information that is received through sensory systems.

METHODS

Experimental overview

In two experiments, individual Picasso triggerfish (*Rhinecanthus aculteatus*) were presented with a four-alternative forced-choice task in which there was one dark grey circular target (S+) and three light grey circular distractors (S-) (Fig. 1). Grey stimuli were chosen because a straightforward change in brightness could be used to alter the discrimination difficulty, and because *R. aculeatus* can learn to discriminate them quickly and to a high degree of accuracy. Stimuli were presented in four possible positions: top right (TR), top left (TL), bottom right (BR) and bottom left (BL). The two experiments differed in terms of the target position distribution. In Experiment 1, there was a bias in the target position distribution allowing subjects to learn probability information. In Experiment 2, the target was equally likely to appear in any of the four positions. In Experiment 1, our hypothesis was that fish would fall back on their expectations linked to target location when exposed to conditions of increasing

stimulus uncertainty, as Bayesian inference would predict. However, in Experiment 2, in which there was no positional biased prior probability information, we did not expect the fish to have a preference for any particular position. Although this task shares some similarities with foraging as correct choices results in a food reward, the experiment presents a non-adaptive scenario to the fish. The food reward is dissociated from the stimuli as it is given to fish by a human experimenter after a choice has been made, and the stimuli are not naturally occurring or particularly similar to anything that is. In addition, the choice between four closely placed options would not be expected in a typical foraging scenario for this species.

Both experiments were divided into two training steps followed by a testing stage. The first training step taught fish to discriminate stimuli based on colour and in the second step fish were introduced to the target position distribution. In the second training step, the distribution of target positions was altered such that the likelihood of the target appearing at one location increased relative to the other three. All fish were given 12 sessions (10 trials each) in training step two, regardless of target selection accuracy. Of the 120 trials, in 72 trials (60%) the target was presented in the 'biased' position. In the remaining trials, the target appeared 16 times in each of the three other positions. Two fish were trained to each of the four possible positions (TR, TL, BR, BL). A two-tiered reward system was used in the hope that it would direct the attention of the fish to the secondary rule. When fish correctly selected S+ in the non-biased position, they were given a standard food pellet, but when the fish correctly selected the S+ and it was in the biased position, they were given a larger piece of cut-up bait fish. This reward procedure was used throughout Step 2 and for reinforcement trials in the testing stage. During this stage, S+ selection accuracy remained high (combined mean frequency for all sessions CI: 95 % ± 1.8 CI) and all fish progressed to the testing stage.

In the testing stage, S- was replaced with distractors that became increasingly difficult to distinguish from the target stimulus. Testing session included five reinforcement trials (R) and five probe trials (T1-T5). The reinforcement trials followed the same procedures as those of training Step 2 (S+ and S-) and were used as a reminder of the original task. Correct

selections were rewarded and the two-tiered reward system was maintained in these trials, therefore three of the five reinforcement trials were in the predetermined position. The remaining five trials were probe trials in which the difference in stimuli brightness became increasingly similar and eventually identical. In each testing session, one of each of the five treatments was presented. Probe trials were unrewarded and the biased position was not maintained. The order of reinforcement and testing trials was pseudorandom so that a particular trial type (R or T) would not appear in more than two consecutive trails. A total of 24 testing sessions were run.

In Experiment 2, the methods used in training, and testing were similar to that of Experiment 1, except that fish were not biased to a particular stimulus position in Step 2. In addition, there was a difference in reward value and frequency. For Fish 18K and 17R, the two-tiered reward system was maintained; in 60% of trials (selected pseudo-randomly), one of four positions was preassigned a higher reward value if chosen, with an equal frequency for all four positions. For the other two fish, the reward schedule was simplified so as all correct responses resulted in a single pellet. As in Experiment 1, the probe trials were not rewarded during testing for Fish 18K and 17R. However, Fish 24O and 25P, were rewarded. Fish maintained an S+ selection accuracy of 97% ± 1.5 CI in Step 2.

Under the training scenario, selecting stimuli based on its visual characteristics (i.e. the darkest stimulus) would lead to the highest reward rate (100%). Under the testing condition, as the stimuli become harder to discriminate, the likely reward rate decreases with the level of discrimination difficulty to chance (25%). However, by changing decision strategy to include position in their selection criteria, subjects in Experiment 1 could increase their chances of receiving a reward (60%). In this paradigm, two questions emerge. 1) Will the fish use position as a cue, and 2) if they do, will they change their decision strategy in such a way as to maximize their food rewards across all probe types, or will they only rely on position when visual cues are completely undetectable?

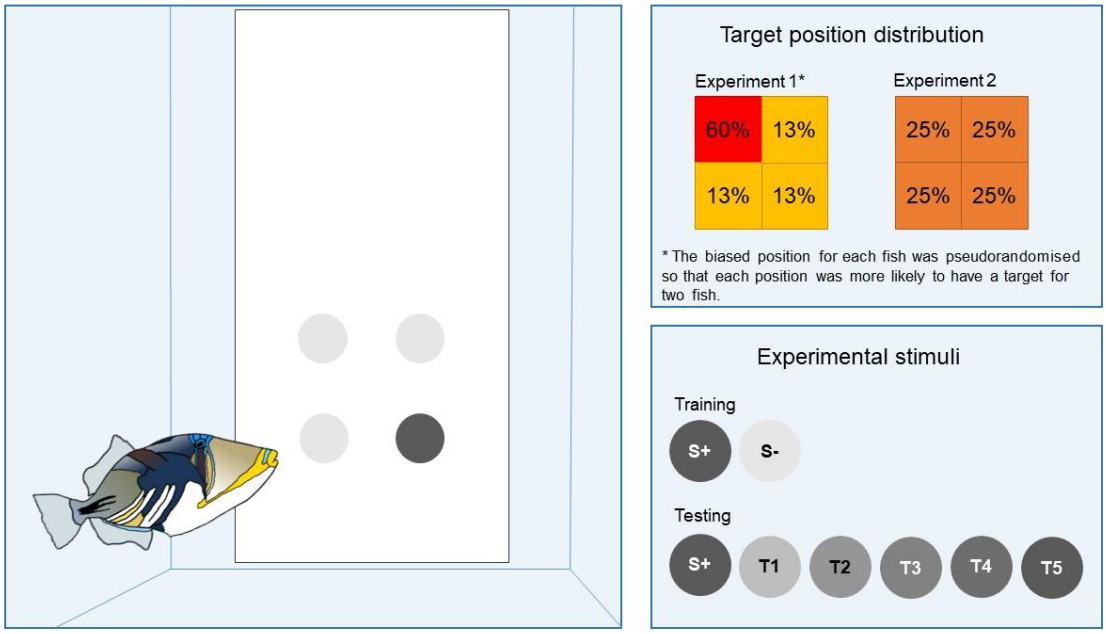

Figure 1: Summary of the experimental protocol. Selection of the target stimulus (S+) was rewarded with food, while the distractor stimuli (S-) was not. The figure is illustrative only and not rendered to scale.

Experimental procedure

At the beginning of an experimental session, a white opaque Perspex trapdoor was placed in the front third of the home tank (18cm from the front of the tank) and was used to separate the fish from the remainder of the tank. Once the door was in place, a white stimulus presentation board (40 x 20 cm) was inserted into the final quarter of the tank. There were four possible stimulus positions on the presentation board: top right (TR, 17cm from bottom, 5.5cm from right edge), top left (TL, 17cm from bottom, 5.5cm from left edge), bottom right (BR, 7cm from bottom, 5.5cm from right edge) and bottom left (BL, 7cm from bottom, 5.5cm from left edge). The shelter tube was then removed from the tank and the rock was placed behind the presentation board to ensure the fish had an unobstructed view of the stimuli once they went through the trapdoor.

A trial began with the fish behind the trap door. Once the door was raised, the fish swam into the open area and up to the presentation board and had to select one of four stimuli. If the fish choose the correct stimulus by biting it, they received either a food pellet (Formula One Marine Pellet, Ocean Nutrition, United States) or a piece of fish (Gamma Slice, Tropical Marine Centre, United Kingdom), depending on the experimental trial type (see below for details). If the fish choose incorrectly, they were not rewarded, the presentation board was removed from the aquarium, and the trial was terminated. After each trial, the fish swam back through the trap door and the door was closed in preparation for the next trial. One session consisting of 10 trials was run each weekday during the experimental period. Prior to experimentation, all fish were pre-trained. See Supplementary Materials Fig. 1 for a video example of the testing procedure. In training stage 1, fish were deemed to have learned the task once they achieved >75% accuracy in three consecutive sessions (10 trials/session). Each fish was given as many sessions as required to learn this task. In Experiment 1, all fish (n = 8) completed this stage within 4-11 sessions. In Experiment 2, all fish (n = 4) completed training Step 1 within 3-11 sessions.

Subjects

A total of 12 fish were used in this experiment; eight individuals were tested in Experiment 1 and another four were tested in Experiment 2. Individuals were purchased from local suppliers and all fish were cared for according to the code of practice for the care and use of animals for scientific purposes (Local Ethical Review Committee of Oxford University's Department of Zoology). The fish were housed in individual aquaria (35.5 x 60 x 31.5 cm) in a flow through marine aquarium system and fed a daily mixed diet of pellets and fish pieces as part of experiments. All fish had live rock and a PVC pipe (length: 15cm; diameter: 10.5cm) to serve as shelter. Three fish had previously participated in similar experiments using different training techniques but the remaining fish were naïve (Appendix 1). Prior to this experiment, all fish were at least pre-trained to swim through a trapdoor and to select white discs from a presentation board.

Stimuli

Stimuli were grey circular discs (2.5 cm diameter) that were printed on photographic paper (Canon matte photo paper), cut out using a hole-punch and laminated. Six stimuli types were made (Fig. 1). The darkest (S+) and brightest (S-) were used in training and reinforcement trials (RGB values: **S+**: 90 90 90, **S-**: 230 230 230). The remaining four stimuli were used as testing stimuli RGB values: **T1**: 190 190 190, **T2**: 150 150 150, **T3**: 130 130 130, **T4**: 110 110 110). A fifth testing condition (T5) used S+ stimuli for all four options. The reflectance of stimuli (candela/square meter) was measured using a Salford Electrical Instruments photometer (S+: 26.6, T4: 35.1, T3: 46.2, T2: 59.5, T1:92.2, S-: 127.3).

Statistical analysis

For experiment 1 probe T5, all stimuli presented were identical therefore it is impossible to state whether the fish had made a correct or incorrect decision based on the visual cue. For the sake of analysis, before testing began, one stimulus was randomly assigned as S+ for each trial. As these trials are unrewarded, the designation had no effect on the fish, however selection of this S+ was used to calculate the accuracy of those trials.

The selected stimulus colour and position was recorded for all trials. The results were tallied for each probe type and analysed using a generalised linear mixed model (GLMM) with a binomial distribution (log-link function) to test whether the selection frequencies of each stimulus type were significantly different from that expected due to chance (using the glmer function in the lme4 package; (Bates et al. 2015)). Selection frequencies of the target and biased position were included as fixed effects and fish ID was included as a random effect.

The relationship between the selection frequency of S+ and the selection of stimuli in either the biased position (experiment 1) or each of the four possible positions (experiment 2) was analysed using a linear regression (using the lm function in R; (R Core Team 2013)).

Bayesian probability matching model

To model the observed behaviour of the fish, a Bayesian probability matching model that accounts for prior probabilities, was created. Probability matching strategies do not result in the highest chance of being correct. To achieve this, the fish would have to choose the most likely outcome 100% of the time in T5. Instead, when probability matching is employed, a category is selected at a frequency that matches the probability assigned by the fish (e.g. 60% in T5 for Experiment 1, or 25% in Experiment 2).

We begin with the following Bayesian probability updating model:

$$\mathrm{p(r|s,c)} \propto \mathrm{p(s|r,c)\ p(r|c)}$$

Here, the variable r represents the position of the reward, s represents the position of the brighter disk (S+). The stimulus, reward and bias position can take one of four values: $\{\mathrm{top\ left, bottom\ left, bottom\ right, top\ right}\}$, and the contrast level, c can take one of six values: $\{R, T1, T2, T3, T4, T5\}$. We adopt the convention that $p(\cdot)$ represents a Bayesian probability, i.e. a belief, and other letters, e.g. $f(\cdot)$, to represent experimentally observed frequencies. In the equation above, then, the distribution $p(r)$ represents the prior probability of reward assigned by the fish to each of the positions. $\mathrm{p(r|s,c)}$ represents the posterior probability assigned by the fish, after being presented with S+ in position s, at contrast level c. The term $p(s|r,c)$ is the corresponding likelihood function.

The probability matching model implies that the observed choice frequencies in an experiment, $f(r|s,c)$, are equal to the posterior probability assigned by the fish, $p(r|s,c)$ and, that the fish's prior distribution of the reward positions, $p(r|c)$, matches the real-world frequencies of reward position during the training phase, $g(r)$, giving:

$$f(r|s,c) = p(r|s,c) \propto p(s|r,c)\ g(r)$$

In Experiment 1, $g(r)$ is equal to 3/5 for the biased position and 2/15 for each unbiased position and in Experiment 2 it is 1/4 for all positions. It should be noted that although we have used the letter r to represent reward location, the various functions of r mean different things. $g(r)$ is a frequency of reward location determined by the experiment, $p(r)$ and $p(r|s,c)$ represent the fish's beliefs about reward position, and $f(r|s,c)$ represents the choice made by the fish, which may be interpreted as reflecting where it considers the reward to be.

To assess the likelihood $p(s|r,c)$, a further experiment estimated the S+ selection accuracy for each probe type when fish motivation was high (see below for methods). In this experiment, $f(r)$ is uniform, so the previous equation can be simplified to $f(r|s,c) = p(s|r,c)$, meaning that the choice frequency directly reflects the fish's likelihood function.

During the testing phase, the S+ stimuli (and reward if given) were presented at one of the four positions, with a distribution of which we will call $h(s)$, which in our experiments was uniform. The frequency of correct choices for a given contrast level is then is given by the sum

$$f(\text{correct}, c) = \sum_{r=s} h(s) f(r|s,c)$$

Similarly, the probability of choosing the bias location, r_{bias}, is

$$f(\text{bias chosen}, c) = \sum_{s} h(s) f(r_{bias}|s,c)$$

S+ Selection Accuracy

Results from Experiment 2 as well as that of an additional four fish that previously participated in experiments not related to this study (n = 6 fish), were used to demonstrate that the selection accuracy of S+ decreased linearly as the difference in reflectance between S+ and all other stimuli decreased (see Supplementary Materials Fig. 2). All fish were trained and tested as described in Experiment 2, and the two additional fish followed the same reward protocol as Fish 24O and 25P. The reward allows for the fish to learn the stimuli, but it also serves to

ensure motivation stays high throughout a long experiment. As the aim of this test was to determine the maximum selection accuracy of the fish, if the fish had learned the stimuli within the 24 testing trials, it would only serve to strengthen the results. These results were used in our Bayesian probability matching model as an estimate of the discriminability of the stimuli in each probe condition.

RESULTS

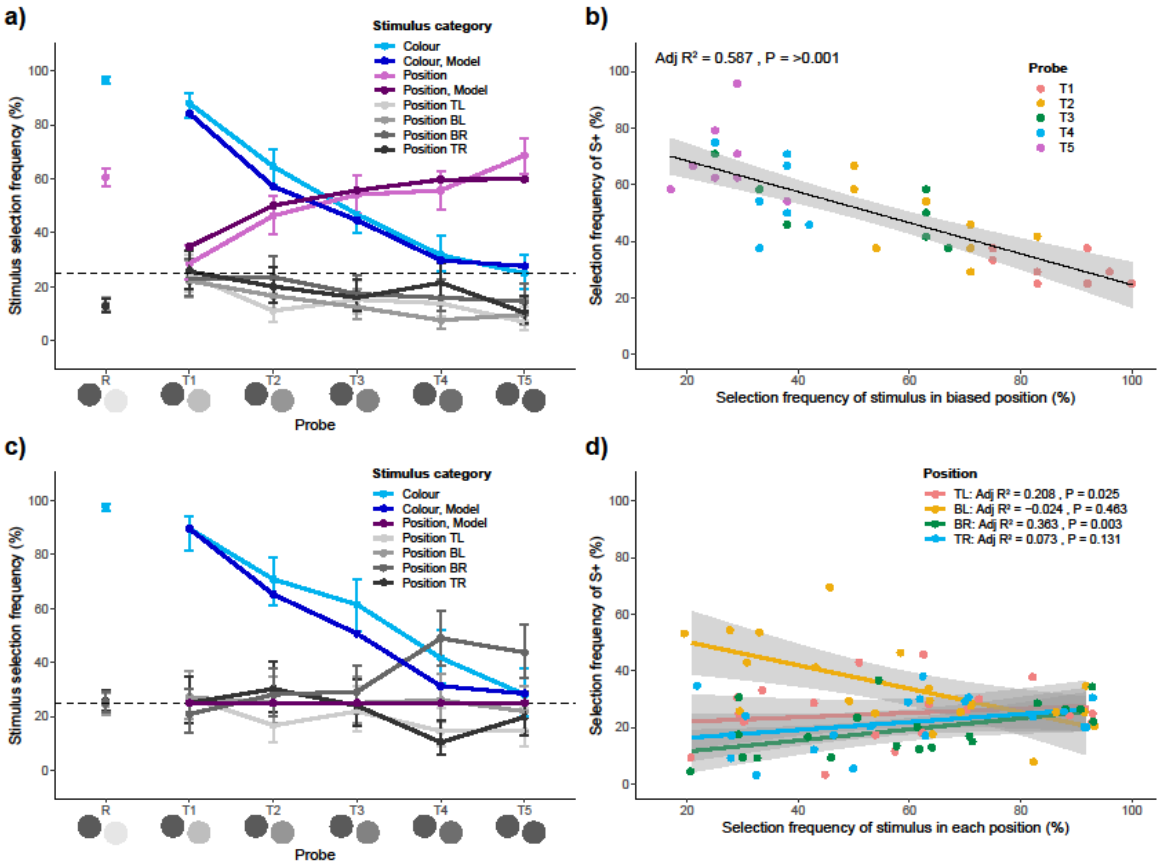

Figure 2: Stimulus selection frequency results of Experiments 1 and 2. In experiment 1, (**a, b**), the target stimulus was 4.5 times more likely to be in one of the four positions when the original conditioned stimuli were presented (S+ and S-). The four possible positions are Top Left (TL), Bottom Left (BL), Bottom Right (BR) and Bottom Left (BL). In experiment 2, (**c,d**), target stimulus was presented in all positions equally. **a)** The mean selection frequency of S+

and all stimuli positions for eight fish. The dashed line at 25% indicates a selection frequency consistent with chance. Error bars represent confidence intervals based on a generalized linear mixed effect model. Stimuli used in testing are shown below the figure where the top circles are S+ (not to scale). **b)** A linear regression analysis of S+ selection frequency versus selection of stimuli in the biased position. The shaded grey area indicates confidence intervals. **c)** The mean stimulus selection frequency for four fish, given no position bias. Symbols are the same as those in figure **a**. **d)** Linear regression analyses of S+ selection frequency versus each of four possible stimuli positions.

Experiment 1

Fish selected S+ at a frequency higher than chance for R (CI: 95.3 - 97.6%), T1 (CI: 82.6 - 91.9%), T2 (CI: 57.6 – 71.0%), T3 (CI: 39.9 – 53.9%) but not T4 (CI: 25.6 – 38.7%) or T5 (CI: 19.4 – 31.6%). See Supplementary Materials Table 1-2 for individual fish results for Experiments 1 and 2. While fish selected stimuli in the biased position at a frequency consistent with chance for R (CI: 57.5 – 63.7%) and T1 (CI: 22.7 – 35.4%), as stimuli became more similar, selection of stimuli in the biased position increased for T2 (CI: 39.4 – 53.4%), T3 (CI: 47.1 – 61.1%), T4 (CI: 48.6 – 62.6%) and T5 (CI: 61.9 – 74.9%). Note that S+ appears in the biased position in 60% of R trials. Selection of stimuli in all other positions is significantly different from the biased position (see Fig. 2a for further confidence intervals). In addition, there was a significant relationship between selection frequency of S+ and selection of stimuli in the biased position (Fig. 2b); as S+ selection accuracy decreased, selection of stimuli in the biased position increased ($R^2 = 0.587$, $p < 0.001$). Although the fish were unrewarded in the testing conditions, had the training reward schedule held true as the fish likely expected, they would have received a food reward at a rate of at least 50% across probe types, which is much higher than what is predicted by chance (25%).

Experiment 2

When the target was equally likely to appear in any of the four positions throughout training, selection accuracy of the target decreased similarly to what was observed in Experiment 1 (CI: R 96.2 – 98.9%, T1 81.7 – 94.3%, T2 61.0 – 79.1%, T3 51.4 – 70. 6%, T4 32.3 – 51.7%, T5 20.0 - 37.9%); however, fish did not begin selecting stimuli in a particular position until after T3, when the mean S+ selection accuracy began to decrease near to chance. In addition, the selection frequency of stimuli in either of the four positions was not generally above chance levels across treatments, except in the case of the bottom right position (Position BR) which was significantly higher than chance for T4 (CI: 39.1 – 58.9%) and T5 (CI: 34.2 – 53.8%). See Fig. 2c for further confidence intervals. There was a significant relationship between selection frequency of S+ and stimuli in Positions BR ($R^2 = 0.363$, $p = 0.003$) and TL ($R^2 = 0.208$, $p = 0.025$), but not Positions BL ($R^2 = -0.024$, $p = 0.463$) or TR ($R^2 = 0.073$, $p = 0.131$) (Fig. 2d).

DISCUSSION

The experiments in this study demonstrate two important findings: 1) Fish can learn probability distribution information, and 2) they can apply this knowledge in a novel, non-adaptive decision making task, even when the optimal response is uncertain. When the fish were presented with stimuli that they had learned to a high degree of accuracy, they chose the correct stimulus in almost all trials. As their positive stimulus selection error rate increased, the fish relied more on their previous knowledge of the correct response distribution. These results show that this highly adaptive, Bayesian behaviour is not dependent upon possession of a cortex, nor is its use likely to be limited to specific contexts or scenarios (e.g. foraging), as the fish used it in a task they are unlikely to have adapted to solve. These experiments have revealed a powerful mechanism for fish decision making when faced with unreliable incoming sensory information, which has a clear benefit to individuals.

By relying on knowledge of prior distributions, the fish increased the potential for receiving food rewards (Experiment 1). However, as has been observed in the foraging strategies of other species, the behaviour of the fish was not optimized to receive the most rewards possible (Bayes optimal). To achieve this, subjects should have selected stimuli in

the most likely position in all trials. Instead, fish appeared to use ‘probability matching’ and the stimulus selection rate matched the learned probability distribution. Probability matching has been previously observed in some, but not all (e.g. Herbranson and Schroeder 2010; Parducci and Polt 1958; Wilson and Rollin 1959) animals including fish (Behrend and Bitterman 1961; Behrend and Bitterman 1966; Bullock and Bitterman 1962; Perez-Escudero and de Polavieja 2011; Woodard and Bitterman 1973) and humans (e.g. Gaissmaier and Schooler 2008; Herbranson and Schroeder 2010; Rubinstein 2002). Why it occurs and its evolutionary significance is still debated. One possibility is that it is the result of a cognitive shortcut associated with cognitive limitations (Vulkan 2000), while another possibility is that the randomness of a sequence is not internalized leading to the use of alternative optimizing strategies (e.g. “win-stay, lose-shift”) (Gaissmaier and Schooler 2008). However, these explanations are typically considered in the context of human decision-making and may not be relevant to all animals if such divergent species do not share the same cognitive biases. These explanations both imply that probability matching is an artefact of a cognitive weakness but it may also provide some benefit to an individual as it allows for sampling of the environment and therefore an opportunity to update priors where conditions may be fluctuating. Whatever the cause, there is experimental evidence that the application of probability matching may be context dependent, and contingent on incentives, motivation, reinforcement, or understanding of the problem (Bitterman 1971; Börgers and Sarin 2000; Erev and Barron 2005; Rivas 2013; Shanks et al. 2002; Vulkan 2000; Wolford et al. 2004). In our Experiment 1, three of the eight fish choose the biased position in T5 at a higher rate than probability matching predicts (fish: 8M = 96%, 9N = 71%, 21P = 79%). Individual differences such as these have been observed in other studies even with very different experimental protocols (e.g. Bitterman 1971), suggesting that individual variation is a common occurrence. Although the function of probability matching is outside the scope of this study, for the purposes of our results, this behaviour provides clear evidence that the fish in Experiment 1 learned the probability distributions of the positive stimulus.

The individual variation observed in this, and many studies of fish behaviour, may in fact be a product of a Bayesian-like decision strategy. Probability-based choices can lead to more of a gradient of choices than other context specific decision rules such as “if-then” or “win-stay, lose-shift”, although this depends on the specifics of the rules learned. A probability-based decision rule may explain not only the range of behavioural outcomes observed in animals, but also the variability in their decisions even when faced with what appears to be identical choices (Beck et al. 2012), such as in a controlled laboratory environment. This is because an individual’s priors may differ, as well as the level of an individual’s uncertainty in a given trial, and whether they will make a choice that favours their expectation of a positive outcome or not. Our results also shows that individual past experience, and by extension, memory, plays an important role in fish decision-making. Hard-wired rules presumably have less demand on long-term memory, while the application of probability distributions would require some, and possibly quite a significant capacity for individual memory. How long these types of memories last, and whether or how quickly they can be overwritten by a period of acclimation to new circumstances, are interesting questions for future studies and would be valuable to our understanding of how fish adapt to persistent and stochastic changes in their environment. By showing that Bayesian mechanisms underpin decision making by fish, we can improve interpretation, modelling, and predictions of fish behaviour by including individual experience (priors) into models.

The fact that fish behaviour in our study was consistent with Bayesian inference when completing a non-adaptive task, suggests that this mechanism may underlie a wide range of decision tasks and not just specific ones such as foraging. In addition, our results may be particularly important to our understanding of sensory signal processing. In our experiments, uncertainty was introduced as the sensory input (visual brightness) approaches the limit of discrimination. Therefore, the results are relevant to our understanding of intermediate signal processing and may extend to the integration of multiple signal types or sources of information. For example, it suggests that fish can make optimized (but perhaps not optimal) decisions

about sensory input, even when the signal is weak or ambiguous. In a coral reef fish species that uses visual signals for a range of important behaviours under rapid changes in visual conditions, this capacity may be particularly important for survival. Therefore, evidence of Bayesian inference in fish decision-making is not only important to our understanding of the evolution of decision behaviour, but provides specific insights into the behaviour of the species tested.

In this study we have found that fish likely use a sophisticated decision strategy that relies on acquired knowledge of probability distributions, and updating stimulus choice based on current incoming sensory information. Using this strategy affords several benefits; fish can use previous experiences to inform current decisions, they can respond flexibly and dynamically to changing situations, and they can make optimized decisions in uncertain conditions. This one mechanism may underlie the response of fish to many seemingly complex decision scenarios, including unconscious sensory information processing, integration of spatial information, and sensorimotor integration, as well as more conscious choices between multiple alternatives or different behaviours. Our results, combined with that of previous studies, provide increasing evidence that Bayesian inference may be generally applicable across taxa and tasks, and is not limited to species with specialized neuroanatomy or adaptations to specific tasks.

ACKNOWLEDGEMENTS

We thank Hannah Smithson and Brian Rogers for advice on relative reflectance measurements and Alex Kacelnik for comments on the manuscript. This work was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement [659684], and a Leverhulme Trust Early Career Fellowship. The contents of the article reflect on the authors' views and not the views of the European Commission.

REFERENCES

Alonso, J. C., J. A. Alonso, L. M. Bautista, and R. Muñoz-Pulido. 1995. Patch use in cranes: a field test of optimal foraging predictions. Animal Behaviour 49:1367-1379.

Bates, D., M. Mächler, B. Bolker, and S. Walker. 2015. Fitting linear mixed-effects models using lme4. Journal of Statistical Software 67:1-48.

Battaglia, P. W., R. A. Jacobs, and R. N. Aslin. 2003. Bayesian integration of visual and auditory signals for spatial localization. Journal of the Optical Society of America A 20:1391-1397.

Beck, Jeffrey M., Wei J. Ma, X. Pitkow, Peter E. Latham, and A. Pouget. 2012. Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. Neuron 74:30-39.

Behrend, E. R., and M. E. Bitterman. 1961. Probability-matching in the fish. The American Journal of Psychology 74:542-551.

Behrend, E. R., and M. E. Bitterman. 1966. Probability-matching in the goldfish. Psychonomic Science 6:327-328.

Biernaskie, J. M., and R. J. Gegear. 2007. Habitat assessment ability of bumble-bees implies frequency-dependent selection on floral rewards and display size. Proceedings of the Royal Society B: Biological Sciences 274:2595-2601.

Biernaskie, J. M., S. C. Walker, and R. J. Gegear. 2009. Bumblebees Learn to Forage like Bayesians. American Naturalist 174:413-423.

Bitterman, M. E. 1971. Visual probability learning in the rat. Psychonomic Science 22:191-192.

Börgers, T., and R. Sarin. 2000. Naive Reinforcement Learning With Endogenous Aspirations. International Economic Review 41:921-950.

Brown, C. 2012. Tool use in fishes. Fish and Fisheries 13:105-115.

Bshary, R., A. Hohner, K. Ait-el-Djoudi, and H. Fricke. 2006. Interspecific Communicative and Coordinated Hunting between Groupers and Giant Moray Eels in the Red Sea. PLoS Biol 4:e431.

Bullock, D. H., and M. E. Bitterman. 1962. Probability-matching in the pigeon. The American Journal of Psychology 75:634-639.

Chater, N., M. Oaksford, U. Hahn, and E. Heit. 2010. Bayesian models of cognition. WIREs Cognitive Science 1:811-823.

Cheney, K. L., C. Newport, E. C. McClure, and N. J. Marshall. 2013. Colour vision and response bias in a coral reef fish. Journal of Experimental Biology 216:2967-2973.

Cheng, K., S. J. Shettleworth, J. Huttenlocher, and J. J. Rieser. 2007. Bayesian integration of spatial information. Psychological Bulletin 133:625-637.

Courville, A. C., N. D. Daw, and D. S. Touretzky. 2006. Bayesian theories of conditioning in a changing world. Trends Cogn Sci 10:294-300.

Deneve, S., and A. Pouget. 2004. Bayesian multisensory integration and cross-modal spatial links. J Physiol Paris 98:249-258.

Dunn, R. P. 2016. Tool use by a temperate wrasse, California sheephead *Semicossyphus pulcher*. Journal of Fish Biology 88:805-810.

Erev, I., and G. Barron. 2005. On adaptation, maximization, and reinforcement learning among cognitive strategies. Psychological Review 112:912-931.

Foley, B. R., and P. Marjoram. 2017. Sure enough: efficient Bayesian learning and choice. Animal Cognition 20:867-880.

Funamizu, A., B. Kuhn, and K. Doya. 2016. Neural substrate of dynamic Bayesian inference in the cerebral cortex. Nature Neuroscience 19:1682-1689.

Gaissmaier, W., and L. J. Schooler. 2008. The smart potential behind probability matching. Cognition 109:416-422.

Gershman, S. J. 2015. A unifying probalistic view of associative learning. PLOS Computational Biology 11:e1004567.

Gunji, Y.-P., H. Murakami, T. Tomaru, and V. Basios. 2018. Inverse Bayesian inference in swarming behaviour of soldier crabs. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 376:20170370.

Herbranson, W. T., and J. Schroeder. 2010. Are birds smarter than mathematicians? Pigeons (*Columba livia*) perform optimally on a version of the Monty Hall Dilemma. Journal of Comparative Psychology 124:1-13.

Hunte, W., R. A. Myers, and R. W. Doyle. 1985. Bayesian mating decisions in an amphipod, Gammarus lawrencianus Bousfield. Animal Behaviour 33:366-372.

Knill, D. C., and A. Pouget. 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. Trends in Neurosciences 27:712-719.

Körding, K. 2007. Decision theory: What "should" the nervous system do? Science 318:606-610.

Kording, K. P., and D. M. Wolpert. 2004. Bayesian integration in sensorimotor learning. Nature 427:244-247.

Körding, K. P., and D. M. Wolpert. 2006. Bayesian decision theory in sensorimotor control. Trends in Cognitive Sciences 10:319-326.

Kruschke, J. K. 2006. Locally Bayesian learning with applications to retrospective revaluation and highlighting. Psychol Rev 113:677-699.

Kruschke, J. K. 2008. Bayesian approaches to associative learning: From passive to active learning. Learning & Behavior 36:210-226.

Lee, T. S., and D. Mumford. 2003. Hierarchical Bayesian inference in the visual cortex. Journal of the Optical Society of America A 20:1434-1448.

Lima, S. L. 1984. Downy Woodpecker Foraging Behavior: Efficient Sampling in Simple Stochastic Environments. Ecology 65:166-174.

—. 1985. Sampling behavior of starlings foraging in simple patchy environments. Behavioral Ecology and Sociobiology 16:135-142.

Luttbeg, B., and R. R. Warner. 1999. Reproductive decision-making by female peacock wrasses: flexible versus fixed behavioral rules in variable environments. Behavioral Ecology 10:666-674.

Marshall, H. H., A. J. Carter, A. Ashford, J. M. Rowcliffe, G. Cowlishaw, and M. Boots. 2013. How do foragers decide when to leave a patch? A test of alternative models under natural and experimental conditions. The Journal of animal ecology 82:894-902.

McNamara, J. M., R. F. Green, and O. Olsson. 2006. Bayes' theorem and its applications in animal behaviour. Oikos 112:243-251.

Moens, V., and A. Zénon. 2019. Learning and forgetting using reinforced Bayesian change detection. PLoS Computational Biology 15:e1006713.

Newport, C. 2020. Concept learning in fish. Current Opinion in Behavioral Sciences.

Newport, C., G. Wallis, S. E. Temple, and U. E. Siebeck. 2013. Complex, context-dependent decision strategies of archerfish, *Toxotes chatareus*. Animal Behaviour 86:1265-1274.

Northcutt, R. G. 2002. Understanding vertebrate brain evolution. Integrative and Comparative Biology 42:743-756.

Olsson, O., U. Wiktander, N. M. Holmgren, and S. G. Nilsson. 1999. Gaining ecological information about Bayesian foragers through their behaviour. II. A field test with woodpeckers. Oikos:264-276.

Parducci, A., and J. Polt. 1958. Correction vs. noncorrection with changing reinforcement schedules. Journal of comparative and physiological psychology 51:492-495.

Penny, W. 2012. Bayesian Models of Brain and Behaviour. ISRN Biomathematics 2012:1-19.

Perez-Escudero, A., and G. de Polavieja. 2011. Collective Animal Behavior from Bayesian Estimation and Probability Matching. PLOS Computational Biology 7:e1002282.

Perfors, A., J. B. Tenenbaum, T. L. Griffiths, and F. Xu. 2011. A tutorial introduction to Bayesian models of cognitive development. Cognition 120:302-321.

Pryor, K. J., and A. M. Milton. 2019. Tool use by the graphic tuskfish *Choerodon graphicus*. Journal of Fish Biology 95:663-667.

R Core Team. 2013.R: A language and environment for statistical computing.R Foundation for Statistical Computing, Vienna, Austria.

Ramírez, J. C., and J. A. R. Marshall. 2017. Can natural selection encode Bayesian priors? Journal of Theoretical Biology 426:57-66.

Rivas, J. 2013. Probability matching and reinforcement learning. Journal of Mathematical Economics 49:17-21.

Rubinstein, A. 2002. Irrational diversification in multiple decision problems. European Economic Review 46:1369-1378.

Rushworth, M. F. S., and T. E. J. Behrens. 2008. Choice, uncertainty and value in prefrontal and cingulate cortex. Nature Neuroscience 11:389-397.

Schluessel, V., O. Beil, T. Weber, and H. Bleckmann. 2014a. Symmetry perception in bamboo sharks (*Chiloscyllium griseum*) and Malawi cichlids (*Pseudotropheus sp*.). Animal Cognition:1-19.

Schluessel, V., and D. Duengen. 2014. Irrespective of size, scales, color or body shape, all fish are just fish: object categorization in the gray bamboo shark *Chiloscyllium griseum*. Animal Cognition:1-11.

Schluessel, V., G. Fricke, and H. Bleckmann. 2012. Visual discrimination and object categorization in the cichlid *Pseudotropheus* sp. Animal Cognition:1-13.

Schluessel, V., H. Kraniotakes, and H. Bleckmann. 2014b. Visual discrimination of rotated 3D objects in Malawi cichlids (*Pseudotropheus sp.*): a first indication for form constancy in fishes. Animal Cognition 17:359-371.

Selonen, V., and I. K. Hanski. 2010. Decision making in dispersing Siberian flying squirrels. Behavioral Ecology 21:219-225.

Shanks, D. R., R. J. Tunney, and J. D. McCarthy. 2002. A re-examination of probability matching and rational choice. Journal of Behavioral Decision Making 15:233-250.

Siebeck, U. E., L. Litherland, and G. M. Wallis. 2009. Shape learning and discrimination in reef fish. Journal of Experimental Biology 212:2113-2119.

Sutherland, N. S. 1969. Shape discrimination in rat, octopus and goldfish: A comparative study. Journal of Comparative and Physiological Psychology 67:160-176.

Tenenbaum, J. B., and T. L. Griffiths. 2002. Generalization, similarity, and Bayesian inference. Behavioral and Brain Sciences 24:629-640.

Vail, A. L., A. Manica, and R. Bshary. 2014. Fish choose appropriately when and with whom to collaborate. Current Biology 24:R791-R793.

Valone, T. J. 1991. Bayesian and prescient assessment: foraging with pre-harvest information. Animal Behaviour 41:569-577.

—. 1992. Information for patch assessment: a field investigation with black-chinned hummingbirds. Behavioral Ecology 3:211-222.

Valone, T. J., and J. S. Brown. 1989. Measuring Patch Assessment Abilities of Desert Granivores. Ecology 70:1800-1810.

Valone, T. J., and L.-A. Giraldeau. 1993. Patch estimation by group foragers: what information is used? Animal Behaviour 45:721-728.

van Gils, Jan A., Ingrid W. Schenk, O. Bos, and T. Piersma. 2003. Incompletely Informed Shorebirds That Face a Digestive Constraint Maximize Net Energy Gain When Exploiting Patches. American Naturalist 161:777-793.

Vulkan, N. 2000. An Economist's Perspective on Probability Matching. Journal of Economic Surveys 14:101-118.

Wilson, W. A., Jr., and A. R. Rollin. 1959. Two-choice behavior of rhesus monkeys in a noncontingent situation. J Exp Psychol 58:174-180.

Wolford, G., S. E. Newman, M. B. Miller, and G. S. Wig. 2004. Searching for patterns in random sequences. Canadian Journal of Experimental Psychology 58:221-228.

Woodard, W. T., and M. E. Bitterman. 1973. Further experiments on probability learning in goldfish. Animal Learning & Behavior 1:25-28.

Xu, F., and J. B. Tenenbaum. 2007. Word learning as Bayesian inference. Psychological review 114:245-272.

Yu, A. J., and P. Dayan. 2005. Uncertainty, Neuromodulation, and Attention. Neuron 46:681-692.