# Bridging the Gap

## From Uncertain Information to Adaptive Behaviour in Fish




Joseph George Bartram

St. John's College and the Department of Zoology

University of Oxford

Supervised by

Dr. Theresa Burt de Perera

Thesis submitted for the degree of

*Master of Science by Research*

Trinity Term 2020

# Declaration

All work presented in this thesis is primarily my own. Four co-authors made contributions to one or more chapters, as follows:

Theresa Burt de Perera supervised the entire project, and specifically contributed to experiment design in Chapter 2

Cait Newport designed the pilot experiment in Chapter 2 and collected a subset of the data in the pilot experiment.

I collaborated with Ollie Padget to write the code and create the charts which were used to assess bias position preference in Chapter 2. Tom Hart wrote the code used to simulate null expectations of position repeats in Chapter 3.

# Acknowledgements

I would like to set aside this space to thank all those who contributed to the writing of this thesis in all its many iterations, and who have supported me in ways both large and small throughout this long and difficult journey.

Firstly, this thesis was made possible by funding from the Natural Environment Research Council, and by conference travel funds provided by St. John's College.

None of this would have been possible without my supervisor Theresa Burt de Perera, who has been tireless advocate for this thesis, and endlessly resourceful in finding silver linings in the murky clouds of its many pitfalls.

I would like to extend my sincere gratitude to everyone in the OxNav research group for their advice and feedback. A shoutout especially to both Vicki Davis and Izzy Watts, who welcomed me into the FishNav lab group, and taught me which end of a siphon was which. I cannot give enough thanks to either Cecilia Karlsson or Sarah Skeels, who shared in my victories and commiserated my failures, and without whose community the lab would have been a much lonelier place. My thanks also go out to Dora Biro and Tim Guildford, who assessed my transfer of status and gave me a much-needed wakeup call in a difficult time.

I am indebted to John and Jason Hogg and the rest of the building services staff of the Department, who kept the lights on and things running smoothly even in the most extreme of circumstances, both in the Tinbergen building, and later at the John Krebs Field Station. I would especially like to thank Christine Soper for her assistance with animal husbandry at the field station labs, and without whom the labs would likely have collapsed.

Within Zoology I would like to extend my gratitude to Ollie Padget and Tom Hart for their statistics support, and to Lynn Lewis-Bevan and Em Dale who offered a sympathetic ear during trying times. In the University at large, I am indebted to my academic mentor Bella Jackson for her endless patience and uncanny insight.

To all my friends outside the Department, who kept me (mostly) sane through the long years of this project and supported me through it all – from the bottom of my heart, thank you. To the many members of *Taruithorn*, who provided a safe harbour for my many eccentricities and provided a wonderful and supportive community – thank you and remember: *Valaraukar hínilmar mantër!*

Of all my friends, I cannot thank Aym Maidment enough, who has been my rock and a constant source of joy and inspiration in my life. Thanks for being there, pal.

Most of all, I am grateful to my wonderful parents Jane and Jamie for their unconditional love and support on both sides of the pond, for raising me, and for standing by me through thick and thin. My dad in particular, who has reviewed more versions of this manuscript than I can count, and always had an insightful comment to bring to bear.

Last but not least, I would like to recognise all the animal subjects who ever participated or otherwise contributed to this thesis. Firstly, to the many unnamed elephantnose fish whose care characterised my first year, and who taught me the value of empathy – you are not forgotten. But most of all, to the 37 Picasso triggerfish who were the sturdy backbone of this thesis.

| | | |
|---|---|---|
| Aquila | Fishy McFishface | Pumpkin |
| Armitage | Foamrider | Randir |
| Aubrey | Gil-Galad | Reverend Gill |
| Beefcake | Gog | Rheia |
| Clint | Hester | Rhys |
| Connie | Lapis lazuli | Sardine |
| Duck | Little Spinner | Scar |
| EOR | Magog | Skybiter |
| Eugenia | Megaton Mary | The Maw |
| Fëanor | Miranda II | Turgon |
| Fingolfin | Nettles | Vetch |
| Fingon | Nicholas | |
| Finwë | Peridot | |

Thank you all.

# Abstract

Sensory information necessarily informs decision-making, but both sensory information and the biological structures that receive and process it are subject to inherent limitations and uncertainty. Both these factors degrade an animal's capacity to make informed, adaptive decisions in a complex and ever-changing world. How is it that animals manage these constraints?

In this thesis, I set out to examine multiple cue use and decision-making in a highly derived marine teleost: the Picasso or lagoon triggerfish, *Rhinecanthus aculeatus*. The use of multiple cues represents an integrative solution to the problem of informational uncertainty; however, it remains poorly understood outside of a narrow range of species and behaviours, largely mammalian. This thesis begins to address the question of how the processes of cue integration are accomplished in other vertebrate taxa.

I approached this through two complementary sets of experiments, addressing related questions of decision-making in relation to information use. In Chapter 2 I examined whether triggerfish are capable of sensory updating, i.e., weighting their prior estimates with incoming information to inform their decision-making. In Chapter 3, I re-examined the same data set and performed a modified replicate to dissect multiple factors contributing to decision-making, focusing on task learning and how subjects might have taken advantage of low-level heuristics to solve the problems presented to them. I placed this discussion in an integrated associative learning framework. For Chapter 4, I was curious as to how aquatic environments, which are highly labile and may favour different sensory channels in different contexts, might affect learning style in triggerfish. To examine this, I tested how triggerfish learn to discriminate objects with multiple competing cues in an environment where cue quality is subject to varying environmental transmission quality. I approached all these questions through the framework of associative learning, using conditioning techniques to train triggerfish to discriminate between sets of stimuli with different properties and examine how they make decisions under different conditions.

Fishes are not only taxonomically distant to more typically studied models of decision-making such as mammals, but also represent an ecological contrast, moving with relative freedom in a three-dimensional environment that exerts radically different pressures on their movement, sensory, and decision ecology. The findings presented in this thesis increase our understanding of information use within fishes and expand our understanding of the evolution and ecology of decision-making within the vertebrates.

# TABLE OF CONTENTS


**GENERAL INTRODUCTION** ..... 1

Background ..... 2
Informational integration ..... 2
Drawbacks of integration ..... 4
Sensory updating ..... 5
Taxonomic distribution of research ..... 6
Study System ..... 7
Aims ..... 9
Significance ..... 11
References ..... 12

**UPDATING IN *RHINECANTHUS ACULEATUS*** ..... 17

Abstract ..... 18
Introduction ..... 19
Methods ..... 24
*Subjects and Husbandry* ..... 24
*General Procedure* ..... 25
Experiment 1: Pilot ..... 27
*Methods* ..... 27
*Results* ..... 30
Experiments 2 and 3 ..... 35
*Flaws with the Pilot experiment* ..... 35
*Methods* ..... 36
*Results* ..... 39
Discussion ..... 50
*Prediction 1* ..... 50
*Prediction 2* ..... 51
*Models of updating* ..... 52
*Caveats* ..... 53
*Conclusions* ..... 54
References ..... 55

**RESPONSE BIASES IN *RHINECANTHUS ACULEATUS*** ..... 59

Abstract ..... 60
Background ..... 61
1.0 Sequence effects ..... 63
*1.1 Introduction* ..... 63
*1.2 Hypotheses* ..... 66
*1.3 Null models* ..... 67
*1.4 The inter-trial interval* ..... 69
*1.5 Methods* ..... 70
*1.6 Results* ..... 71
2.0 Response bias and lateralization ..... 77
*2.1 Introduction* ..... 77
*2.2 Methods* ..... 77
*2.3 Results* ..... 78
3.0 Task acquisition ..... 82
*3.1 Introduction* ..... 82
*3.2 Methods* ..... 83
*3.3 Results* ..... 84
General Discussion ..... 86
*Task acquisition* ..... 86

*Side bias* ... *88*
*Sequence effects* ... *89*
*Response biases as decision strategies* ... *91*
**Concluding remarks** ... **96**
**References** ... **97**

**ASSESSMENT OF CUE RELIABILITY IN *R. ACULEATUS*** ... **101**

**Abstract** ... **102**
**Background** ... **104**
**Aims** ... **107**
**Significance** ... **108**
**Methods** ... **109**
*Subjects & husbandry* ... *109*
*Experimental paradigm* ... *110*
*Stimulus quality and design* ... *110*
*Pre-training* ... *112*
*General procedure* ... *113*
*Minimising inadvertent biases* ... *114*
*Training* ... *115*
*Testing* ... *116*
*Managing misbehaviour* ... *117*
*Data analysis* ... *118*
**Results** ... **121**
**Discussion** ... **124**
*Health* ... *126*
*Sensory Factors* ... *126*
*Cognitive Factors* ... *127*
**Concluding remarks** ... **131**
**References** ... **131**

**GENERAL DISCUSSION** ... **136**

**Introduction** ... **137**
**Aims** ... **138**
*Chapter 2* ... *138*
*Chapter 3* ... *139*
*Chapter 4* ... *139*
**Summary of results** ... **140**
*Chapter 2* ... *140*
*Chapter 3* ... *143*
*Chapter 4* ... *145*
**Significance** ... **146**
*Background* ... *146*
*My findings* ... *147*
**Unanswered Questions** ... **151**
*Bayesian modelling* ... *151*
*Speed-accuracy trade-offs* ... *151*
*Anomalous results* ... *153*
**Final observations and future directions** ... **154**
*Frustration* ... *154*
*The caudal apparatus* ... *159*
*Ecology, evolutionary history, and cognition* ... *161*
*Metacognition and uncertainty monitoring* ... *162*
**Final remarks** ... **163**
**References** ... **164**

# CHAPTER 1
# General Introduction

# General Introduction

## Background

In order to behave adaptively, animals must both sense and respond to their environment. Since natural environments are complex, heterogeneous, and variable, having an accurate and up-to-date model of the environment is essential to adaptive behaviour [1]–[3]. Commonly, this involves making decisions between competing choices – whether due to competing priorities (e.g., between current and future reproductive investment), or uncertainty in the consequences of alternative courses of action (e.g., imperfect assessment of mate quality). Such decisions occur across all domains of behaviour, including mate choice [4], nest site choice [5], parental investment [6], group association [7], and foraging [8]. Optimally, making such decisions requires animals to have as accurate a picture of their environment as possible. However, individual sensory systems are subject to inherent limitations and flaws – noise, ambiguity, and limited bandwidth – that place limitations on this capacity; e.g. [9]. These limitations to sensory acquisition pose a challenge to the production of adaptive behaviour. Animals can reduce their sensory uncertainty by incorporating multiple sources of information about the state of the environment [10]. While extensively documented in some areas of study (e.g., foraging theory), research into informational integration remains limited in both taxonomic and behavioural scope. In this thesis, I will examine how and why triggerfish use different sources information to make decisions, using operant conditioning as an experimental framework.

## Informational integration

One type of strategy to reduce sensory uncertainty is to combine multiple sources of information to produce a more robust posterior estimate. This includes combining information from multiple spatially separated receptors (e.g., binocular vision), from

multiple distinct sensory modalities (e.g., vision and touch), or through time (i.e., memory and learning). Combining information in this manner can improve the accuracy of an animal's internal model in several ways. In this thesis, I focus on the use of prior information (i.e., integrating sensory information through time), but briefly discuss other forms of integration here, as they provide a valuable perspective.

Combining information across modalities may compensate for any limitations or drawbacks of a single modality, such as spatial range or temporal resolution, taking over in different behavioural or environmental contexts. For example, dwarf scorpionfish (*Scorpaena papillosus*) shift between using vision and the mechanosensory lateral line system to locate prey depending on ambient light levels [11], while largemouth bass (*Micropterus salmoides*) rely on vision for initial prey detection and orientation, but shift to lateral line cues to guide the final strike [12]. Alternatively, animals may combine information from multiple modalities simultaneously. This can provide multiple, statistically independent estimates of a feature or object, reducing uncertainty compared to multiple unimodal stimuli [13]. Electrically sensitive elephantnose fish (*Gnathonemus petersii*) can more accurately track a moving shelter with all senses intact than when restricted to only a single sense [14]. In addition to reducing uncertainty, combining multimodal cues can also enhance the speed of response, either by simple probability (i.e. the race model [15]) or through direct enhancement. For example, barn owls orient to a visual stimulus both more precisely and more rapidly when an auditory component is also present [16].

Within a single modality, integration across multiple spatially separated receptors may produce information that could not be acquired by a single stationary receptor. The iconic example of this phenomenon is the emergence of depth perception arising from binocular stereoscopic vision [17]. By contrast, integrating distinct signals within a sensory modality may also enhance perception. For example, in human vision different components such as

colour, shape, and motion are initially processed separately before being re-integrated at higher cortical levels. The integration of these distinct cues can contextually enhance both speed and accuracy of perception in various contexts, e.g. [18]. The extent to which these organizational principles extend to non-primates is currently unclear. However, some recent research has found that some fish species respond to various illusions in like manner to humans, indicating that visual integration shares some key features such as the separate processing of colour and shape [19]–[21]. This suggests that fish might benefit from visual integration in a similar manner to mammals.

Like binocular stereopsis, a single moving visual receptor can generate depth information through motion parallax if information is integrated through time [22]. For example, the characteristic head-bobbing of various owl species is suggested to be an adaptation to improve distance estimation during hunting [23]. Many different organisms exhibit stereotyped sensorimotor routines to improve the acquisition of sensory information. For example, the "probing motor acts" displayed by weakly electric fishes when interacting with novel objects have recently been characterised as active motor strategies to enhance perception using the active electric sense [24], [25].

## Drawbacks of integration

These strategies are not without potential costs and drawbacks. Environmental sampling may expose an animal to the risk of predation, and requires investment of both time and energy, especially if sampling requires active sensorimotor strategies (such as tactile exploration, echolocation, or active electroreception [26], [27]). Integrating may also result in inaccurate estimates if there is a mismatch with the dynamic scale of the environment – an animal that integrates foraging information on seasonal scales may not be able to respond appropriately to diel changes.

Furthermore, animals have limited attention [28], and employing multiple sensory modalities in a single task may distract from other priorities. In a classic example, female wolf spiders presented with both the vibratory and visual components of a male's courtship signal were less likely to detect a synthetic predator and were hence more vulnerable to predation than females presented with only the visual component of the courtship ritual [29]. Limited attention may also slow down processing in demanding multisensory tasks. For example, rats trained to a cross-modal auditory/visual discrimination task exhibit greater response latencies than individuals trained to an equivalent unimodal task [30]. Given these drawbacks, we should expect an animal's integration strategies to be tuned to its environment.

## Sensory updating

The use of prior information is a broad subject, which includes everything from the moment-by-moment summation that produces sensory flow [31], to the decades-long spatial learning of landscape features by African elephants [32]. On a macro scale, the integration of prior and current sensory information (updating) in decision-making has been most thoroughly investigated in the context of foraging theory; e.g. [8], [33]–[35], also see [36] for a historical overview. Animals searching for food in a patchy environment must develop decision rules to determine when to leave a partially depleted patch in favour of a new one. While an inflexible rule (such as a fixed giving-up time) can be adaptive in highly static and predictable environments, updating a prior expectation of patch richness with new information during foraging can potentially greatly improve an animal's provisioning rate [37]. While numerous statistical models for incorporating prior information have been proposed (e.g. [38], [39]), by far the most successful is Bayesian updating [36], in which prior knowledge is rendered as a probability distribution which is combined with new information to form a more robust posterior estimate. While there are abundant models

and field studies of updating in the context of foraging, comparatively few studies have explicitly manipulated the prior information animals have access to [10]. This evidence gap makes distinguishing Bayesian updating from other mathematical solutions problematic, as prior information cannot be readily measured and quantified. Furthermore, uncontrolled priors may be the result of evolutionary adaptation rather than individual learning [40]. As such, our understanding of how animals incorporate prior information into their decision-making remains relatively poor. Finally, there have been only a scattering of studies examining updating in other behavioural contexts (mate selection [41], [42], dispersal [43], nesting [44]). In short, our understanding of how animals make use of prior information is largely restricted to a narrow range of species and behaviours.

## Taxonomic distribution of research

The overwhelming bulk of research on updating in non-human animals has focused on mammals and, to a lesser degree, birds. The rest of the vertebrate clade has received comparatively little attention. This is especially true of the ray-finned fishes (Actinopterygii), which form an outgroup to the Tetrapod lineage [45], having diverged approximately 450mya [46]. As a result, the degree to which such cognitive abilities are conserved across the vertebrate clade remains unclear. In addition, the taxonomic bias towards mammals means research has focused on species adapted to a limited range of principally terrestrial sensory environments. Expanding the breadth of research to include piscine systems holds the added benefit of exploring information use in organisms adapted to a radically different – i.e., aquatic – physical environment.

In mammals, multisensory integration in particular is most often associated with brain structures lacking clear equivalents in fish, such as the parietal and prefrontal cortex of the cerebrum [47]. It is worth noting, though, that evidence does indicate that some multisensory processes may occur at lower levels directly in the sensory cortices [48], and

the superior colliculus, the latter of which does have a homolog in the piscine optic tectum [49].

If informational integration abilities are shared between mammals and ray-finned fishes it would indicate that integration can take place without the derived cortical structures that are present in mammals. Such a finding could suggest either that sensory integration evolved independently in the two lineages, or that integration is a conserved trait (synapomorphy) across the ray-finned fishes and tetrapods, and consequently that the characteristic mammalian cortical structures were only later co-opted for this function. In either case, this would require a drastic revaluation of the evolution of sensory integration.

## Study System

The bulk of the ray-finned fishes fall into the infraclass Teleostei, which are not only the most ecologically and morphologically diverse group of fishes, but at 26,000 described species numerically dominate the entire vertebrate lineage, containing as many species as the other Vertebrata combined, [50]. Behavioural research on this group is sparse compared to that in mammals and birds, and it tends to be restricted to a few model species which do not represent the phylogenetic or behavioural diversity of the clade. Despite this, in recent decades research has uncovered a wealth of previously unrecognised sophisticated cognitive abilities in teleosts, including numeracy [51], [52], cooperative hunting [53], and even tool use [54]. Furthermore, recent research has indicated that many cognitive abilities are surprisingly flexible – for example, archerfish can be conditioned to discriminate between different human faces irrespective of presentation angle, despite lacking any evolutionary history of exposure to human faces [55]. A particularly burgeoning area of research concerns visual processing, in which some teleosts exhibit certain fascinating commonalities with humans. For example, multiple species are be able to perceive various visual illusions [20], [21], [56], [57], recognise 3D-

rotated objects [58], and form object categories [58]. Elephantnose fish even spontaneously exhibit cross-modal recognition of objects, an ability only previously recorded in mammals [59].

The Tetraodontiformes are an order of highly derived, principally marine teleosts characterised by extreme modification of the typical teleost bodyplan [50], including reduced pelvic fins, fewer vertebrae, cranial bones fused into a beak, and defensive adaptations including bony plates or thick leathery skin. The group is only distantly related to other teleosts [60]. A comparative lack of behavioural research in this taxon, coupled with their taxonomic distance from other more commonly studied teleosts such as the goldfish (Cyprinidae: *Carassius auratus*), zebrafish (Cyprinidae: *Danio rerio*), and livebearers (Poecilidae spp.); make them a promising taxon for comparative behavioural research.

The lagoon or Picasso triggerfish (Balistidae: *Rhinecanthus aculeatus*) is a small member of the order native to shallow and especially intertidal reef environments throughout the Indo-Pacific and eastern Atlantic [61], [62]. The species' vulnerability has not been formally assessed (e.g. IUCN, CITES); however, Fishbase [62] assigns it a low vulnerability score based on life-history and distribution data [63]. Being reef-dwelling fishes, triggerfish inhabit complex habitats with extensive three-dimensional structure, and their motor apparatus is adapted accordingly for slow but efficient and precise movement. Their principal form of propulsion is balistiform locomotion, a form of median paired fin (MPF) swimming dependent upon undulations of the dorsal and anal fins [64]. This affords triggerfish extraordinarily versatile motion, including the ability to swim backwards, and both roll and yaw without significant translation of the body in space ([65], personal observations). This extreme manoeuvrability has attracted attention as a model for more mobile aquatic robots, e.g. [66]. While omnivorous, the Picasso triggerfish is principally a visually guided predator of armoured benthic invertebrates such as sea urchins, shrimps,

and molluscs [61], [67], [68]. Unusually for a fish, over the past decade a large body of experimental work has been published establishing the capabilities of the species' visual system at both a behavioural [69]–[72] and physiological level [73], [74]. However, there has been comparatively little work on its cognitive abilities. The species is known to take well to captivity; being a popular species in the aquarium trade, and is receptive to operant conditiong, e.g. [70]. Taken together, these features make it a promising species for comparative sensory and behavioural research.

## Aims

In this thesis, I set out to examine the role of information weighting and updating in decision making in the Picasso triggerfish. Informational updating remains understudied both in the range of species and behaviours studied, and at an experimental level has principally been restricted to field studies of certain natural behaviours (e.g., foraging). This leaves open the question of whether the use of prior information is a general cognitive ability or is restricted to particular behaviours.

I approached this question through the framework of operant conditioning. Conditioning is a set of techniques in which novel or existing behaviours can be reinforced, extinguished, or modified through pairing with a system of rewards and punishments. It represents a highly flexible and powerful set of tools to interrogate sensory and cognitive abilities of animals behaviourally. While behaviours generated in laboratory conditions may lack for ecological validity, the advantage of the approach is that it permits precise control of an animal's sensory environment during learning, allowing the experimenter to dissect precisely which stimuli contribute to any particular behaviour. In this context, it permits direct control of the prior information an animal has available to complete a set behavioural task.

Two major sets of experiments constitute the core of this thesis, addressing complimentary aspects of information integration. In the first, triggerfish were trained to complete a series of object discrimination tasks using unfamiliar stimuli. While a control group were presented with a neutral stimulus placement, a treatment group had a positionally biased training history. I then observed how this biased information was incorporated into subjects' decisions when the reliability of the primary visual cue was manipulated. After a pilot produced mixed results, I carried out two further modified versions of the experiment. In Chapter 2, I describe the details of these experiments, and demonstrate that triggerfish can both use prior information in the context of a visual discrimination task, and that this information is weighted against the reliability of the visual cue.

In Chapter 3 I extend this analysis to assess multiple alternative low-level strategies that the triggerfish could have used to solve the task, using both the original and new data to determine how differences in the initial training schedule affect how subjects learned and attended to the different stimuli. I then attempt to place these results within a larger decision-making framework, informed by associative learning theory and with reference to the animals' ecology.

In the second set of experiments, triggerfish were trained to complete a similar discrimination task with multiple competing visual cues. These cues were presented at either high quality, low quality, or variable quality to simulate the effect of a shifting signalling environment which favours different channels in different contexts. I planned to examine how this training history impacted which cues subjects learned and used to complete the discrimination task. Chapter 4 of this thesis discusses the results of these experiments, in particular the failure of the subjects to learn the exigencies of the task, and the implications of these results for our understanding of triggerfish sensory processing and cognition.

## Significance

At present, our understanding of the evolution of information integration and processing in vertebrates is highly biased towards more derived vertebrates, e.g., mammals. Expanding the taxonomic breadth of study to fish will help develop our understanding of the evolution of key cognitive abilities in vertebrates, and potentially raise interesting questions as to the underlying neural structures that facilitate those abilities. Fish also represent an interesting sensory and ecological contrast in being primarily aquatic, compared to the principally terrestrial and surface-bound mammals. The aquatic environment supports movement with lower energetic expenditure and more degrees of freedom [75] than the terrestrial, and tends to vary on differing spatial and temporal scales. This may influence how aquatic organisms process and makes use of information.

Human activities have a disproportionate effect on freshwater and coastal marine environments [76], with increasing sedimentation and turbidity being among the most frequently-cited factors [77], [78]. Among the most common effects of increasing turbidity is a breakdown in visually mediated behaviours in affected species, especially in the contexts of foraging [79], [80], group cohesion [81], and sexual signalling [82], [83]. These may cascade through communities, with unpredictable consequences. As a result, most research has focused on community- and population-level responses to turbidity, despite these responses often being mediated by individual-level behavioural changes [84]. By contrast, there is a dearth of experimental work examining how individuals of species respond to changing sensory conditions. Improving our understanding of how aquatic animals process and use sensory information may aid our understanding of how those animals will respond to drastic environmental change.

Given the above considerations, I believe that there is a gap in our knowledge that this thesis is uniquely suited to fill, integrating concepts from multiple overlapping disciplines to improve our understanding of the cognitive, taxonomic, and ecological basis of behaviour.

## References

[1] S. R. X. Dall and R. A. Johnstone, "Managing uncertainty: Information and insurance under the risk of starvation," *Philos. Trans. R. Soc. B Biol. Sci.*, vol. 357, no. 1427, pp. 1519–1526, 2002.

[2] S. R. X. Dall, L.-A. Giraldeau, O. Olsson, J. M. McNamara, and D. W. Stephens, "Information and its use by animals in evolutionary ecology," *Trends Ecol. Evol.*, vol. 20, no. 4, pp. 187–193, 2005.

[3] N. E. Munoz and D. T. Blumstein, "Multisensory perception in uncertain environments," *Behav. Ecol.*, vol. 23, no. 3, pp. 457–462, May 2012.

[4] U. Candolin, "The use of multiple cues in mate choice," *Biol. Rev.*, vol. 78, no. 4, p. S1464793103006158, Nov. 2003.

[5] G. P. Jones, "Spawning-site choice by female Pseudolabrus celidotus (Pisces: Labridae) and its influence on the mating system," *Behav. Ecol. Sociobiol.*, vol. 8, no. 2, pp. 129–142, Apr. 1981.

[6] R. J. Lavery and M. H. A. Keenleyside, "Parental investment of a biparental cichlid fish, Cichlasoma nigrofasciatum, in relation to brood size and past investment," *Anim. Behav.*, vol. 40, no. 6, pp. 1128–1137, 1990.

[7] C. Agrillo and M. Dadda, "Discrimination of the larger shoal in the poeciliid fish Girardinus falcatus," *Ethol. Ecol. Evol.*, vol. 19, no. 2, pp. 145–157, 2007.

[8] J. M. McNamara and A. I. Houston, "The application of statistical decision theory to animal behaviour," *J. Theor. Biol.*, vol. 85, no. 4, pp. 673–690, Aug. 1980.

[9] W. Bialek, "Physical limits to sensation and perception.," *Annual review of biophysics and biophysical chemistry*, vol. 16. pp. 455–478, 1987.

[10] T. J. Valone, "Are Animals Capable of Bayesian Updating ? An Empirical Review," *Oikos*, vol. 112, no. 2, pp. 252–259, Feb. 2006.

[11] J. C. Montgomery and A. R. Hamilton, "Sensory contributions to nocturnal prey capture in the dwarf scorpion fish (Scorpaena papillosus)," *Mar. Freshw. Behav. Physiol.*, vol. 30, no. 4, pp. 209–223, 1997.

[12] J. G. New, "Multimodal integration in the feeding behaviors of predatory teleost fishes.," *Brain. Behav. Evol.*, vol. 59, no. 4, pp. 177–89, 2002.

[13] M. O. Ernst and H. H. Bülthoff, "Merging the senses into a robust percept," *Trends Cogn. Sci.*, vol. 8, no. 4, pp. 162–169, 2004.

[14] R. Rojas and P. Moller, "Multisensory contributions to the shelter-seeking behavior of a mormyrid fish, Gnathonemus petersii Günther (Mormyridae, Teleostei): the role of vision, and the passive and active electrosenses.," *Brain. Behav. Evol.*, vol. 59, no. 4, pp. 211–21, 2002.

[15] D. H. Raab, "Statistical facilitation of simple reaction times.," *Trans. N. Y. Acad. Sci.*, vol. 24, no. 5, pp. 574–590, 1962.

[16] E. A. Whitchurch and T. T. Takahashi, "Combined Auditory and Visual Stimuli Facilitate Head Saccades in the Barn Owl ( Tyto alba )," *J. Neurophysiol.*, vol. 96, no. 2, pp. 730–745, Aug. 2006.

[17] J. I. Nelson, "Globality and stereoscopic fusion in binocular vision," *J. Theor. Biol.*, vol. 49, no. 1, pp. 1–88, 1975.

[18] M. W. Self and S. Zeki, "The integration of colour and motion by the human visual brain," *Cereb. Cortex*, vol. 15, no. 8, pp. 1270–1279, Aug. 2005.

[19] L. Mitchell, K. L. Cheney, F. Cortesi, N. J. Marshall, and M. Vorobyev, "Triggerfish uses chromaticity and lightness for object segregation," *R. Soc. Open Sci.*, vol. 4, no. 12, p. 171440, Dec. 2017.

[20] V. A. Sovrano, L. Albertazzi, and O. Rosa Salva, "The Ebbinghaus illusion in a fish (Xenotoca eiseni)," *Anim. Cogn.*, vol. 18, no. 2, pp. 533–542, Mar. 2015.

[21] E. E. Simpson, N. J. Marshall, and K. L. Cheney, "Coral reef fish perceive lightness illusions," *Sci. Rep.*, vol. 6, Oct. 2016.

[22] S. H. Ferris, "Motion parallax and absolute distance," *J. Exp. Psychol.*, vol. 95, no. 2, pp. 258–263, Oct. 1972.

[23] S. Ohayon, R. F. Van Der Willigen, H. Wagner, I. Katsman, and E. Rivlin, "On the barn owl's visual pre-attack behavior: I. Structure of head movements and motion patterns," *J. Comp. Physiol. A Neuroethol. Sensory, Neural, Behav. Physiol.*, vol. 192, no. 9, pp. 927–940, Sep. 2006.

[24] M. J. Toerring and P. Moller, "Locomotor and electric displays associated with electrolocation during exploratory behavior in mormyrid fish," *Behav. Brain Res.*, vol. 12, no. 3, pp. 291–306, 1984.

[25] V. Hofmann *et al.*, "Sensory flow shaped by active sensing: sensorimotor strategies in electric fish.," *J. Exp. Biol.*, vol. 216, no. Pt 13, pp. 2487–500, Jul. 2013.

[26] J. R. Speakman, M. E. Anderson, and P. A. Racey, "The energy cost of echolocation in pipistrelle bats (Pipistrettus pipistrellus)," *J. Comp. Physiol. A*, vol. 165, no. 5, pp. 679–685, Sep. 1989.

[27] V. L. Salazar, R. Krahe, and J. E. Lewis, "The energetics of electric organ discharge generation in gymnotiform weakly electric fish," *Journal of Experimental Biology*, vol. 216, no. 13. pp. 2459–2468, Jul-2013.

[28] R. Dukas, "Behavioural and ecological consequences of limited attention," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 357, no. 1427. pp. 1539–1547, 2002.

[29] E. A. Hebets, "Attention-altering signal interactions in the multimodal courtship display of the wolf spider Schizocosa uetzi," *Behav. Ecol.*, vol. 16, no. 1, pp. 75–82, Jan. 2005.

[30] J. McGaughy, J. Turchi, and M. Sarter, "Crossmodal divided attention in rats: effects of chlordiazepoxide and scopolamine," *Psychopharmacology (Berl).*, vol. 115, no. 1–2, pp. 213–220, 1994.

[31] V. Hofmann, J. I. Sanguinetti-Scheck, and L. Gómez-Sena, "From static electric images to electric flow: Towards dynamic perceptual cues in active electroreception," *J. Physiol.*, vol. 107, no. 1–2, pp. 95–106, Jan. 2013.

[32] L. Polansky, W. Kilian, and G. Wittemyer, "Elucidating the significance of spatial memory on movement decisions by African savannah elephants using state–space models," *Proc. R. Soc. B Biol. Sci.*, vol. 282, no. 1805, 2015.

[33] A. S. Dunlap, D. R. Papaj, and A. Dornhaus, "Sampling and tracking a changing environment: persistence and reward in the foraging decisions of bumblebees," *Interface Focus*, vol. 7, no. 3, p. 20160149, Jun. 2017.

[34] H. H. Marshall, A. J. Carter, A. Ashford, J. M. Rowcliffe, and G. Cowlishaw, "How do foragers decide when to leave a patch? A test of alternative models under natural and experimental conditions," *J. Anim. Ecol.*, vol. 82, no. 4, pp. 894–902, Jul. 2013.

[35] P. Nonacs and J. L. Soriano, "Patch sampling behaviour and future foraging expectations in Argentine ants, Linepithema humile," *Anim. Behav.*, vol. 55, no. 3, pp. 519–527, Mar. 1998.

[36] J. M. McNamara, R. F. Green, and O. Olsson, "Bayes' theorem and its applications in animal behaviour," *Oikos*, vol. 112, no. 2, pp. 243–251, Feb. 2006.

[37] J. C. Alonso, J. A. Alonso, L. M. Baustita, and R. Munoz-Pulido, "Patch use in cranes: a field test of optimal foraging predictions," *Anim. Behav.*, vol. 49, no. 5, pp. 1367–1379, May 1995.

[38] L. D. Devenport and J. A. Devenport, "Time-dependent averaging of foraging information in least chipmunks and golden-mantled ground squirrels," *Anim. Behav.*, vol. 47, no. 4, pp. 787–802, 1994.

[39] A. Lange and R. Dukas, "Bayesian approximations and extensions: Optimal decisions for small brains and possibly big ones too," *J. Theor. Biol.*, vol. 259, no. 3, pp. 503–516, Aug. 2009.

[40] J. C. Ramírez and J. A. R. Marshall, "Can natural selection encode Bayesian priors?," *J. Theor. Biol.*, vol. 426, pp. 57–66, Aug. 2017.

[41] B. Luttbeg, "A Comparative Bayes tactic for mate assessment and choice," *Behav. Ecol.*, vol. 7, no. 4, pp. 451–460, Dec. 1996.

[42] B. Luttbeg and R. R. Warner, "Reproductive decision-making by female peacock wrasses: flexible versus fixed behavioral rules in variable environments," *Behav. Ecol.*, vol. 10, no. 6, pp. 666–674, Nov. 1999.

[43] V. Selonen and I. K. Hanski, "Decision making in dispersing Siberian flying squirrels," *Behav. Ecol.*, vol. 21, no. 2, pp. 219–225, Mar. 2010.

[44] K. A. Schmidt and C. J. Whelan, "Nesting in an uncertain world: Information and sampling the future," *Oikos*, vol. 119, no. 2, pp. 245–253, Feb. 2010.

[45] K. Kikugawa *et al.*, "Basal jawed vertebrate phylogeny inferred from multiple nuclear DNA-coded genes," *BMC Biol.*, vol. 2, no. 1, p. 3, 2004.

[46] S. Kumar and S. B. Hedges, "A molecular timescale for vertebrate evolution," *Nature*, vol. 392, no. 6679, pp. 917–920, Apr. 1998.

[47] C. Chandrasekaran, "Computational principles and models of multisensory integration," *Curr. Opin. Neurobiol.*, vol. 43, pp. 25–34, Apr. 2017.

[48] J. K. Bizley, G. P. Jones, and S. M. Town, "Where are multisensory signals combined for perceptual decision-making?," *Current Opinion in Neurobiology*, vol. 40. Elsevier Ltd, pp. 31–37, 01-Oct-2016.

[49] T. R. Stanford, S. Quessy, and B. E. Stein, "Evaluating the Operations Underlying Multisensory Integration in the Cat Superior Colliculus," *J. Neurosci.*, vol. 25, no. 28, pp. 6499–6508, Jul. 2005.

[50] J. S. Nelson, T. C. Grande, and M. V. H. Wilson, *Fishes of the World*, 5th ed., vol. 2. Wiley, 2016.

[51] M. E. Miletto Petrazzini, C. Agrillo, V. Izard, and A. Bisazza, "Relative versus absolute numerical representation in fish: Can guppies represent 'fourness'?," *Anim. Cogn.*, vol. 18, no. 5, 2015.

[52] C. Agrillo, M. Dadda, G. Serena, and A. Bisazza, "Use of number by fish," *PLoS One*, vol. 4, no. 3, p. e4786, Mar. 2009.

[53] M. E. Arnegard and B. A. Carlson, "Electric organ discharge patterns during group hunting by a mormyrid fish," *Proc. R. Soc. B Biol. Sci.*, vol. 272, no. 1570, pp. 1305–1314, Jul. 2005.

[54] C. Brown, "Tool use in fishes," *Fish Fish.*, vol. 13, no. 1, pp. 105–115, Mar. 2012.

[55] C. Newport, G. M. Wallis, and U. E. Siebeck, "Object recognition in fish: accurate discrimination across novel views of an unfamiliar object category (human faces)," *Anim. Behav.*, vol. 145, pp. 39–49, Nov. 2018.

[56] K. Wyzisk and C. Neumeyer, "Perception of illusory surfaces and contours in goldfish," *Vis. Neurosci.*, vol. 24, no. 03, pp. 291–298, May 2007.

[57] C. Agrillo, M. E. Miletto Petrazzini, and M. Dadda, "Illusory patterns are fishy for fish, too," *Front. Neural Circuits*, vol. 7, no. AUG, 2013.

[58] V. Schluessel, G. Fricke, and H. Bleckmann, "Visual discrimination and object categorization in the cichlid Pseudotropheus sp.," *Anim. Cogn.*, vol. 15, no. 4, pp. 525–537, Jul. 2012.

[59] S. Schumacher, T. Burt de Perera, J. Thenert, and G. von der Emde, "Cross-modal object recognition and dynamic weighting of sensory inputs in a fish.," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 113, no. 27, pp. 7638–43, Jul. 2016.

[60] Y. Yamanoue *et al.*, "Phylogenetic position of tetraodontiform fishes within the higher teleosts: Bayesian inferences based on 44 whole mitochondrial genome sequences," *Mol. Phylogenet. Evol.*, vol. 45, no. 1, pp. 89–101, Oct. 2007.

[61] T.-C. Chen, R. F. G. Ormond, and H.-K. Mok, "Feeding and territorial behaviour in juveniles of three co-existing triggerfishes," *J. Fish Biol.*, vol. 59, no. 3, pp. 524–532, Sep. 2001.

[62] R. Froese and D. Pauly, "Fishbase," 2000. [Online]. Available: https://www.fishbase.de. [Accessed: 09-Oct-2018].

[63] W. W. L. Cheung, T. J. Pitcher, and D. Pauly, "A fuzzy logic expert system to estimate intrinsic extinction vulnerabilities of marine fishes to fishing," *Biol. Conserv.*, vol. 124, no. 1, pp. 97–111, Jul. 2005.

[64] R. W. Blake, "On balistiform locomotion," *J. Mar. Biol. Assoc. United Kingdom*, vol. 58, no. 1, pp. 73–80, 1978.

[65] H. Loofbourrow, "Hydrodynamics of balistiform swimming in the picasso triggerfish, Rhinecanthus aculeatus," University of British Columbia, 2006.

[66] T. Hu, G. Wang, L. Shen, and F. Li, "Bionic inspirations of fish-like robots from Rhinecanthus aculeatus," in *2006 IEEE International Conference on Mechatronics and Automation, ICMA 2006*, 2006, vol. 2006, pp. 639–643.

[67] T. Kuwamura, "Habitat segregation, coexistence or interspecific territoriality between two triggerfishes, Rhinecanthus aculeatus and Sufflamen chrysopterus, with notes on distribution of other balistids at Sesoko Island, Okinawa," *Galaxea*, vol. 10, pp. 65–78, 1991.

[68] T. R. McClanahan and S. H. Shafir, "Causes and consequences of sea urchin abundance and diversity in Kenyan coral reef lagoons," *Oecologia*, vol. 83, no. 3, pp. 362–370, Jun. 1990.

[69] C. M. Champ, M. Vorobyev, and N. J. Marshall, "Colour thresholds in a coral reef fish," *R. Soc. Open Sci.*, vol. 3, no. 9, Sep. 2016.

[70] C. Newport *et al.*, "Fish use colour to learn compound visual signals," *Anim. Behav.*, vol. 125, pp. 93–100, Mar. 2017.

[71] N. F. Green, H. H. Urquhart, C. P. Van Den Berg, N. J. Marshall, and K. L. Cheney, "Pattern edges improve predator learning of aposematic signals," *Behav. Ecol.*, vol. 29, no. 6, pp. 1481–1486, Jul. 2018.

[72] K. L. Cheney *et al.*, "An Ishihara-style test of animal colour vision.," *J. Exp. Biol.*, vol. 222, no. Pt 1, p. jeb189787, Jan. 2019.

[73] C. Champ, G. M. Wallis, M. Vorobyev, U. E. Siebeck, and N. J. Marshall, "Visual acuity in a species of coral reef fish: Rhinecanthus aculeatus.," *Brain. Behav. Evol.*, vol. 83, no. 1, pp. 31–42, 2014.

[74] V. Pignatelli, C. Champ, N. J. Marshall, and M. Vorobyev, "Double cones are used for colour discrimination in the reef fish, Rhinecanthus aculeatus.," *Biol. Lett.*, vol. 6, no. 4, pp. 537–9, Aug. 2010.

[75] R. I. Holbrook and T. Burt de Perera, "Separate encoding of vertical and horizontal components of space during orientation in fish," *Anim. Behav.*, vol. 78, no. 2, pp. 241–245, Aug. 2009.

[76] R. K. Kopf, C. M. Finlayson, P. Humphries, N. C. Sims, and S. Hladyz, "Anthropocene Baselines: Assessing Change and Managing Biodiversity in Human-Dominated Aquatic Ecosystems," *BioScience*, vol. 65, no. 8. pp. 798–811, 01-Aug-2015.

[77] D. H. Wilber and D. G. Clarke, “Biological Effects of Suspended Sediments: A Review of Suspended Sediment Impacts on Fish and Shellfish with Relation to Dredging Activities in Estuaries,” *North Am. J. Fish. Manag.*, vol. 21, no. 4, pp. 855–875, Nov. 2001.

[78] W. Berry, N. Rubenstein, B. Melzian, and B. Hill, “The Biological Effects of Suspended and Bedded Sediment (SABS) in Aquatic Systems: A Review,” 2015.

[79] J. C. Barrett, G. D. Grossman, and J. Rosenfeld, “Turbidity-Induced Changes in Reactive Distance of Rainbow Trout,” *Trans. Am. Fish. Soc.*, vol. 121, no. 4, pp. 437–443, Jul. 1992.

[80] A. De Robertis, C. H. Ryer, A. Veloza, and R. D. Brodeur, “Differential effects of turbidity on prey consumption of piscivorous and planktivorous fish,” *Can. J. Fish. Aquat. Sci.*, vol. 60, no. 12, pp. 1517–1526, Dec. 2003.

[81] J. L. Kelley, B. Phillips, G. H. Cummins, and J. Shand, “Changes in the visual environment affect colour signal brightness and shoaling behaviour in a freshwater fish,” *Anim. Behav.*, vol. 83, no. 3, pp. 783–791, Mar. 2012.

[82] U. Candolin, T. Salesto, and M. Evers, “Changed environmental conditions weaken sexual selection in sticklebacks,” *J. Evol. Biol.*, vol. 20, no. 1, pp. 233–239, 2007.

[83] O. Seehausen, J. J. M. Van Alphen, and F. Witte, “Cichlid Fish Diversity Threatened by Eutrophication that Curbs Sexual Selection,” *Science (80-. ).*, vol. 277, no. 5333, pp. 1808–1811, 1997.

[84] D. P. Chivers, F. Al-Batati, G. E. Brown, and M. C. O. Ferrari, “The effect of turbidity on recognition and generalization of predators and non-predators in aquatic ecosystems,” *Ecol. Evol.*, vol. 3, no. 2, pp. 268–277, Feb. 2013.

# CHAPTER 2
## Updating in *Rhinecanthus aculeatus*

# UPDATING IN *RHINECANTHUS ACULEATUS*

## Abstract

This chapter describes testing of the capacity of Picasso triggerfish (*R. aculeatus*) to assess and incorporate prior information regarding statistical distributions into their decisions in the context of a visual discrimination task and examined the role of sensory discriminability in determining how triggers use and weight different information sources.

Subjects were trained to complete a simple object discrimination protocol using operant conditioning. The experiment proceeded in three stages. Firstly, subjects were trained to select a single positive stimulus from three distractors, based on colour. Once subjects had acquired this task, a controlled prior was introduced by manipulating the spatial distribution of the positive stimulus, so that it was over-represented in one position (henceforth, the "bias position"). Having completed bias training, subjects were presented with a series of probes in which the colour difference between the distractors and positive stimulus was reduced, making discrimination more challenging, thereby enhancing the relative value of the prior. As visual reliability decreased, subjects became less accurate in their choices, but also increasingly selected the bias position. Hence, when confronted with reduced visual discriminability, the fish adaptively adjusted their decision strategy by increasing their reliance on the prior positional information.

This study thus not only demonstrates that triggerfish are able to dynamically shift their decision strategies to cope with complex and changing information but provides preliminary evidence that they can dynamically integrate and weight different sources of information to optimise their decision-making, consonant with Bayesian integration theory. The use of dynamic, flexible decision strategies has not previously been demonstrated in fish, and this research may have dramatic implications both for our understanding of the evolution of these cognitive mechanisms, and of foraging ecology in fishes.

## Introduction

Animals select among alternatives in a variety of behavioural contexts, such as mate selection, foraging, habitat selection, parental investment, predator avoidance, and group living. To make these decisions adaptively, animals must acquire, interpret and assess environmental information using their senses [1]. Much work on animal behaviour has assumed that sensory information is reliable [2], [3]. For example, foraging theory has frequently modelled animals as prescient actors with perfect knowledge of the quality of food patches [4], [5]. While useful, this assumption misrepresents the process of perception; in that all sensory modalities are fundamentally probabilistic, and all sensory information is to varying degrees uncertain and incomplete [6]. Uncertainty and incompleteness impact animals' behavioural strategies in domains as diverse as navigation [7], the timing of reproduction [8], and foraging [9].

Sensory uncertainty hence presents a major barrier to the production of adaptive behaviour. The barrier can be partially overcome, however, by using information from multiple, statistically-independent samples – be they spatially, temporally, or qualitatively separated. With multiple samples, animals may either use information redundantly (relying on whichever source of information is currently most reliable) or may combine sources of information into a unified sensory estimate. The latter is potentially a very powerful tool for reducing sensory uncertainty [10], [11], but is also more cognitively demanding. A number of different models to describe this process have been proposed, but arguably the most successful is that of Bayesian statistics [1].

Bayesian statistics constitutes a robust and systematic approach to resolving conflicting information in the context of uncertainty. Each source of information is conceptualized as a probability distribution (a prior estimate), in which the spread of the function can be interpreted as the confidence in that estimate. Bayes theorem is used to produce a

combined posterior estimate, which is weighted towards the prior estimate with higher confidence (Figure 1). This represents a mathematically optimal solution to the problem of integrating uncertain information, minimising the mean squared error in the estimate [12]–[14]. In the context of animal behaviour and perception, confidence in an estimate may be a function of, for example, the variance or noise in a sensory estimate, or the associative strength of a conditioned stimulus.

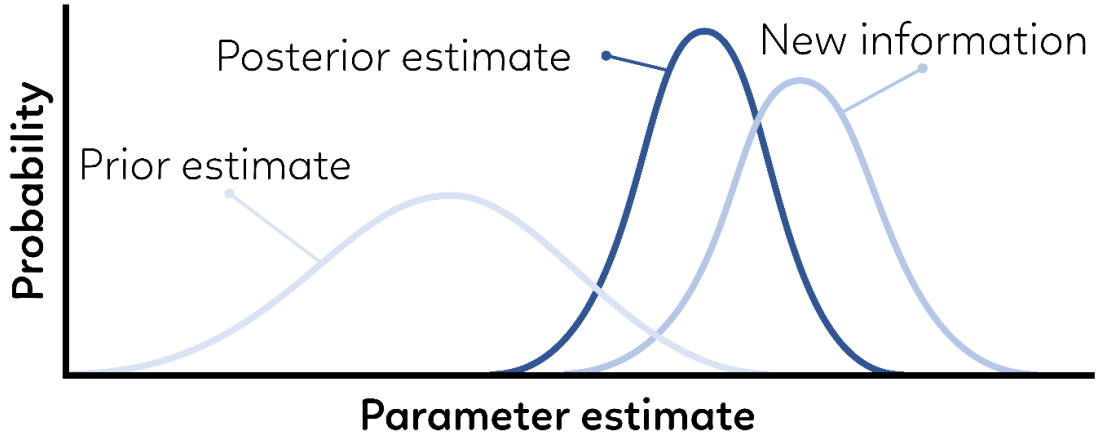



*Figure 1:* Illustration of principle of Bayesian inference, demonstrating how two estimates are combined to produce a posterior estimate.

Bayesian models are extensively employed to tackle problems of perception. For example, multimodal cue integration within and between modalities has been shown to match Bayesian predictions in humans (e.g. [15]–[19] and several other species such as macaques [20], cats [21], [22], rats [23], barn owls [24], and *Drosophila* fruit flies [25].

Within the fields of animal learning and decision making, the most pertinent aspect of Bayesian statistics is the use of Bayes Theorem to update learned prior distributions with incoming sensory information, an idea known as Bayesian inference [26]. Bayesian inference has been widely applied in the fields of animal learning and behaviour, to great success. It has its longest pedigree in foraging theory [27], where the model has been

extensively explored as an alternative to models featuring prescient actors or fixed decision strategies [28]–[32]. Bayesian Inference models have also been applied to dispersal [33] and mate selection [2], [34], [35]. Finally, some aspects of human sensorimotor control [36], [37] and sensory estimation [38] correspond to Bayesian predictions. The bulk of work in animal behaviour has been theoretical [2], [3]. while the comparatively small body of experimental work is taxonomically and ecologically biased towards avian species and foraging theory respectively. Additionally, few studies have explicitly manipulated and controlled for prior experience, a prerequisite for demonstrating updating [39] , though see [32], [40], [41]. As such, it is often impossible to disentangle the effect of learned prior distributions from innate and evolved prior distributions [13], [42]. Finally, few models explicitly test Bayesian approaches against other models which incorporate both prior and sample information.

In this chapter, I set out to test for Bayesian updating in a taxonomic outgroup of the vertebrates, while using an operant conditioning framework to explicitly control for subject's learned priors. The teleost fishes (Actinopterygii: Teleostei) – of which the triggerfishes are members – represent not only the single largest and most morphologically-diverse clade of fishes, but numerically dominate the entire vertebrate lineage, containing as many species as the other Vertebrata combined [43]. Despite this, research on teleost learning and cognition lags far behind that about other vertebrates, especially mammals and birds [44], and their cognitive capabilities are only recently becoming an object of study [45], [46]. The Picasso triggerfish (Balistidae: *Rhinecanthus aculeatus*) is a small, highly derived reef-dwelling teleost species originating from the indo-pacific. It is a generalist omnivore [47], with armoured invertebrates such as shrimps and molluscs making up the bulk of its natural diet [48], [49]. Native to highly-cluttered reef environments, the species is a visually-guided forager with high visual acuity and

developed colour vision [50], [51]. It is known to take readily both to captivity and to operant conditioning [52], [53]. The species' vulnerability has not been directly assessed (IUCN, CITES); however Fishbase [54] assigns it a low vulnerability score based on life-history and distributional data [55]. Taken together, these features make it a highly suitable subject for this research.

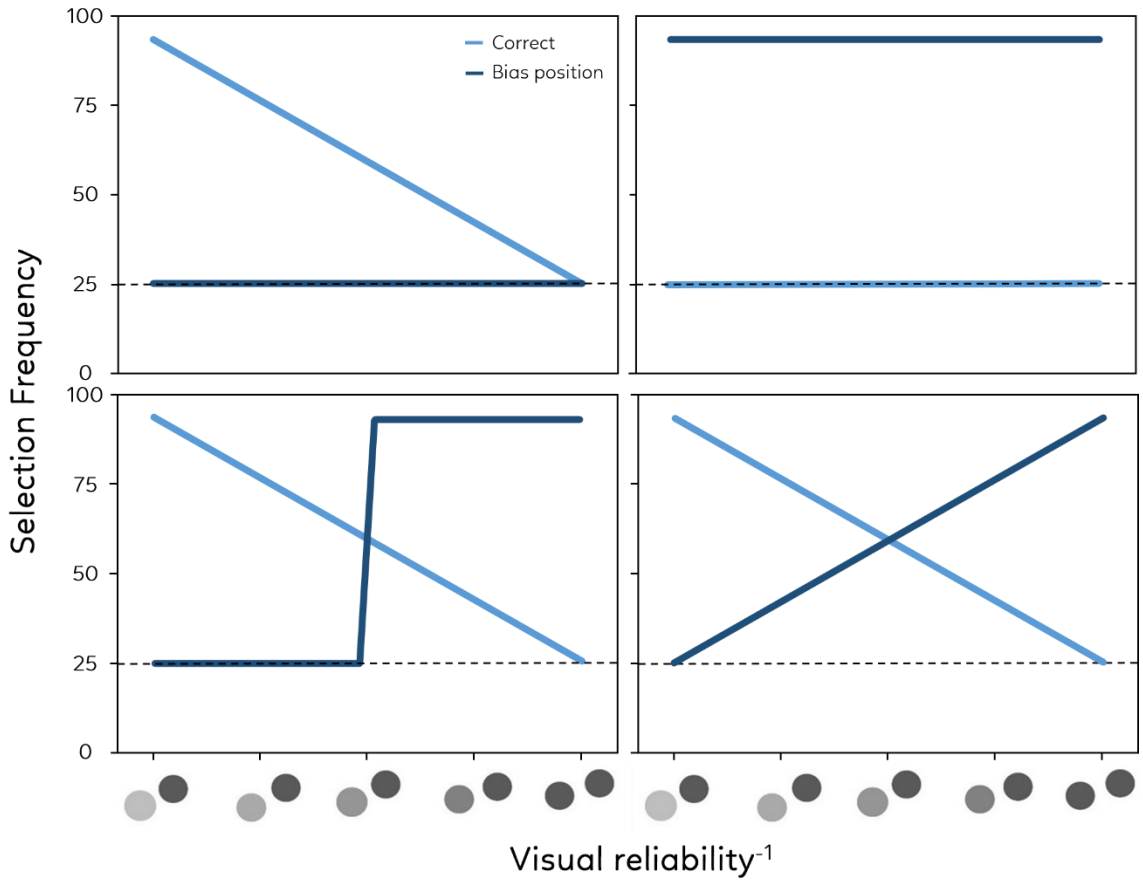



*Figure 2:* Four alternative models of subject behaviour under conditions of reduced visual reliability. "Correct" indicates percentage of choices allocated to the CS+, while "Bias Position" is the same in refence to the conditioned bias position. Dotted line indicates chance (25%). Left to right, top to bottom: (a) Null model 1: subjects to not attend to prior positional information, and exclusively use visual information. (b) Null model 2: subjects do not attend to visual information, and exclusively use prior positional information. (c) Redundant model: subjects attend exclusively to either positional or visual information, and switch depending on the reliability of the visual stimulus. (d) Bayes model: subjects attend to both sources of information, and dynamically weight the relative weight given to either stimulus depending on the reliability of each source of information.

I used an operant conditioning approach to tackle the question of prior information use and Bayesian Inference in triggerfish. Conditioning is a technique in which certain behaviours are reinforced, extinguished, or modified through pairing with a system of positive and negative reinforcement (i.e., rewards, punishments). It represents a highly flexible and powerful set of tools to interrogate the sensory and cognitive abilities of animals behaviourally. While behaviours generated under artificial conditions may lack for ecological validity, the advantage of this approach is that it permits precise control of an animal's sensory environment during learning, allowing the experimenter to dissect precisely what information is necessary for any particular behaviour.

Triggerfish were trained to complete simple visual object-discrimination tasks in which the spatial distribution of the positive conditioned stimulus (CS+) was either random or was systematically biased towards one position. This provided two possible sources of information which could be used to solve the discrimination task – visual contrast between the stimuli, and a positional cue derived from the statistical distribution of the positive stimulus. Those subjects that received a fully randomised CS+ served as a control group. Subjects were then tested on trials with reduced visual information. Under a null model, subjects would exclusively attend to the visual information available, and while accuracy would decline as visual reliability decreased, they would distribute their choices randomly, resulting in a flat response curve (Figure 2a). I hypothesized that subjects would learn to attend to the bias in the position of the stimulus and would rely on this cue when visual information was reduced. This can be further split into redundant and Bayesian models. Under a redundant model, any single trial will be exclusively informed by one source of information, and subjects will rapidly switch depending upon the reliability of the visual cue: resulting in a sigmoidal response curve (Figure 2c). By contrast the Bayesian model is dynamic, with both sources of information being incorporated on each trial but their

relative weights being determined by the reliability of the visual cue, resulting in a smooth response curve (Figure 2d). I predict that subjects would principally rely on visual information during training but would dynamically weight positional information as visual reliability declined. An initial pilot study (experiment 1) produced mixed results, with some subjects exhibiting updating but no population-level effect. Using the pilot to inform the methodologies, two modified experiments were then run in parallel (experiments 2 and 3).

# Methods

## Subjects and Husbandry

All individuals used in these experiments were wild-caught sub-adults (standard length ca. 10-20cm), acquired from a local aquarium supplier (The Goldfish Bowl, 118-122 Magdalen Road, Oxford OX4 1RQ). Subjects were cared for in accordance with the local code of practice for the care of animals (Ethical Review Committee, Department of Zoology, University of Oxford). Once data collection was complete, subjects were either retained for further experiments, or returned to the supplier. Twenty-nine subjects were used across all experiments, with some subjects being used in more than one experiment.

Subjects were housed individually in glass aquaria (35.5 x 60.0 x 31.5 cm), in two tiers on metal racks (Figure 3a). A shared external sump system filtered and pumped water through aquaria, and opaque PVC sheets were used to create visual barriers between aquaria. Aquaria were illuminated both by laboratory lighting and a dedicated lighting system, both on a closed-circuit 12:12h day: night cycle, with ramp-up and ramp-down to mitigate the stress of rapid lighting changes. Aquaria were furnished with a coral gravel substrate, coral rock caves provided both shelter and enrichment. Outside of experiments, subjects were fed twice daily six days a week, using a mixture of standard commercial pellet food (Formula One Marine Pellet, Ocean Nutrition, USA) and frozen foods (lancefish and cockle).

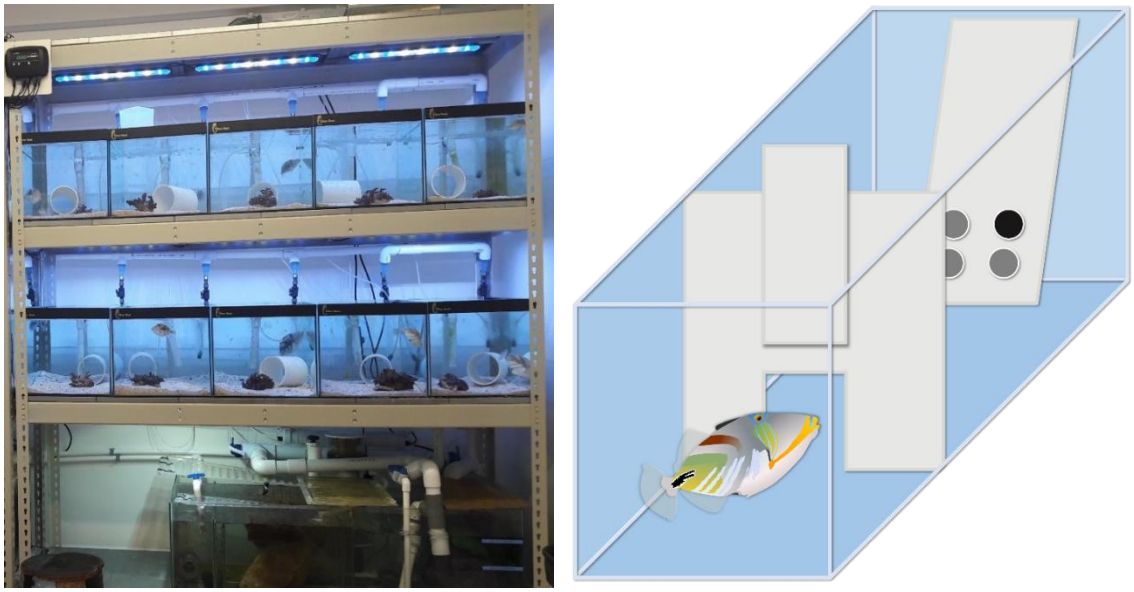


*Figure 3:* (a) photograph of husbandry setup, showing individual aquaria subjects are housed in, and conjoined sump filtration system, (b) schematic illustration of experimental arena.

## General Procedure

I used a simultaneous four-alternate forced-choice (4AFC) task for all experiments discussed in this chapter, consisting of a single positive conditioned stimulus (the CS+) and three identical distractors (CS-). This design affords greater statistical power than a similar two-alternate (2AFC) task given an equal sample size.

Experiments were carried out in the subjects' home aquaria. At the start of a session, all décor (excluding substrate) was removed, and an opaque PVC wall was inserted 18cm from the front. This divided the aquarium into an anterior "start box", and a posterior "choice chamber" where stimuli were presented. The subject could move between these chambers by means of a guillotine door operated by the experimenter. A trial began with a subject in the anterior start box and was marked by the experimenter opening the guillotine door. Subjects entered the choice chamber, after which the guillotine door was closed, and they were able to interact with the stimuli placed at the posterior end of the chamber. After assessing the stimuli, subjects would indicate a choice by biting a stimulus. Correct choices were promptly reinforced with a food reward, presented using forceps. Once a choice was

made, the presentation board was removed, and door re-opened. Subjects could return to the start box of their own volition. No negative reinforcement was used for incorrect responses. If a subject failed to make a detectable choice within 300 seconds, then the session was terminated, and training was resumed in the next training slot.

### Presentation of Stimuli

Stimuli were presented on an opaque PVC board, which was placed at the posterior end of the choice chamber (Figure 3b). The board was white PVC (40 x 20cm), with four Velcro disks for the attachment of stimuli. Stimuli were placed at the vertices of an imaginary square 9cm on a side, positioned such that the centre of the square should be at subjects' eye level as they enter the choice chamber. The four stimulus positions are henceforth referred to as upper-left, upper-right, lower-left, and lower-right (UL, UR, LL, and LR). Stimuli consisted of 2.5cm paper disks (Canon matte photo paper) printed with varying values of greyscale (Canon PIXMA iP7250 printer) and laminated to prevent waterlogging. The darkest of these (luminance = 90) was used as the CS+ in all experiments.

### Mitigating confounding factors

Several precautions were employed against experimenter behaviour biasing outcomes. Firstly, the experimenter remained motionless during trials, and kept their behaviour consistent surrounding this phase. Secondly, the experimenter delayed reaching for the food reward until a definitive response was produced. Thirdly, the experimenter attended to the motion of the fish (following the subject's rostrum); and ignored the stimuli, to avoid their gaze acting as a cue.

# Experiment 1: Pilot

## Methods

### Pre-training

Prior to the start of the experiment, all subjects were pre-conditioned to the CS+ presented at various points around their home aquarium using a transparent Perspex rod. This consisted of rewarding any contact the subject made with the stimulus, which was matched to the CS+ used in the main experiment. Pre-conditioned sessions were composed of sessions of 10 trials each, conducted twice daily for two days.

### Initial training

Subjects were trained to complete a four-alternative forced-choice (4-AFC) task, selecting a single positive stimulus from three distractors. Initial training consisted of sessions of 10 trials each, at 1 to 3 sessions per day, and continued until the fish had at least reached the pre-assigned learning criterion (3 consecutive sessions of >80% correct choices). In multi-trial procedures, the order in which trials are presented may affect subjects' responses to those trials - such disturbances being known as schedule effects. Subjects may learn to rely on simple heuristics such as repeating or alternating their choices, thereby rendering the results inadmissible. Truly random schedules often contain long trains of repeats (among other patterns), and thereby may inadvertently reinforce such biases. For this reason, the structure of schedules is constrained to mitigate such biases – the results being commonly known as a Gellerman-like schedules in honour of the originator of the concept. During training, the position of the CS+ followed a pseudo-random Gellerman-like schedule [56]–[58] that was generated using a MATLAB code (Mathworks, Natick, MA, USA).

### Bias training

Once subjects had reached the learning criterion, they were pseudo-randomly assigned to either a control group, or one of four treatment groups (UL, UR, LL, LR). During bias training, the treatment groups received a modified schedule in which the CS+ was overrepresented

in one of these four positions at 60% of trials. The bias training period was fixed at 6 sessions of 10 trials each. The schedule was generated using a modified MATLAB code.

## Testing

Once subjects had completed bias training, a series of probe sessions was carried out in which the luminance contrast of the visual cue was varied. In probe trials the luminance value of the distractors was decreased, increasing their resemblance to the CS+, and thus increasing the difficulty of visual discrimination. Four probe trials were used in which the distractors were visually intermediate between the training distractor and CS+. In a final probe trial, the distractors were identical to the CS+, hence no visual information was available (Figure 4). To minimise the potential for further feedback learning to occur, probe (Treatment, "T") trials were not reinforced. This necessitated interspersing T trials with equal numbers of rewarded (Reinforcement, "R") trials to prevent task extinction. These R trials maintained the original stimulus pairing (S+ and S-) and positional bias used during bias training. Testing consisted of 24 sessions of 10 trials each, at 1-2 sessions per day, with a pseudorandom schedule.

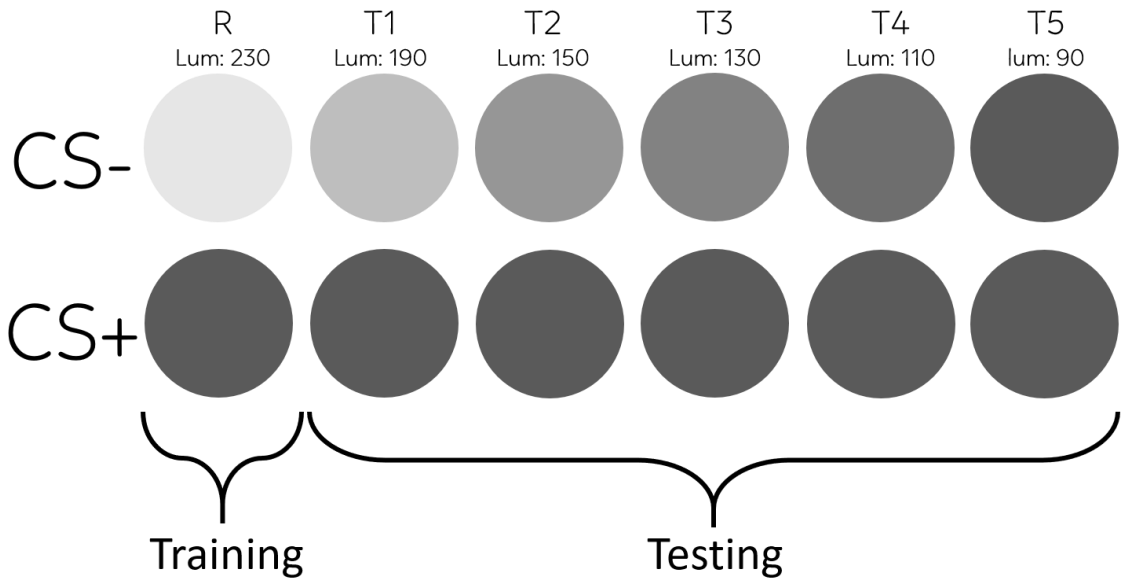



*Figure 4:* schematic of the stimulus combinations used during training and testing. Top row indicated distractor; bottom row indicates positive stimulus. Leftmost pairing (R) represents the training condition, subsequent pairings (T1-T5) indicate probe pairings.

## Statistical analysis

The principal statistical test used to analyse choice data was the binomial test (one-sample, two-tailed) to determine the minimum and maximum number of trials where a particular choice that would be expected by chance with an alpha of 0.05. In cases with values falling outside this range (i.e., $p \neq 0.25$), the null hypothesis that choices do not differ from a chance distribution can be rejected. This is indicated as significance bounds on plots. For plots of subjects' performance through time, this is calculated for a single session ($n = 10$, $p = 0.25$, range = 0-5 trials, 0.00-50.00%), whereas for plots broken down by trial type this was calculated for the total number of trials per trial type for a single subject ($n = 24$, $p = 0.25$, range = 2-10 trials, 8.33-41.67%). In addition, a generalised linear mixed effects model was used to test the relationship of bias position choice to visual reliability versus other predictors.

## Results

### Training

All subjects acquired the basic discrimination task, reaching the learning criterion in 4-14 sessions (figure 5). Once subjects reached the learning criterion in the discrimination task, a statistical bias was introduced for six sessions, in which the CS+ appeared in one position for 60% of trials. During bias training performance remained high (93.3% accuracy, ±4.5%) (Figure 6).

### Testing

During testing, subjects were presented with 24 sessions of unreinforced probe trials, in which the discriminability of the visual cue was reduced. All subjects completed this phase. During testing subject's accuracy declined as a function of luminance contrast, as predicted (figure 7). When no visual information was available (i.e., during T5 trials) accuracy did not statistically differ from chance, confirming that the subjects did not have access to uncontrolled sources of information to complete the task (e.g., experimenter bias). Accuracy was comparable between treatment and control groups (figure 8) despite highly varying sample sizes, and overall accuracy in R trials was comparable to that found during the bias training phase (94.2% ±4.1%, compared to 93.3% ±4.5%).

There was extensive individual variation in the rate at which subjects selected the biased position during testing (figure 9). Two subjects (6-Nicholas and 7-Eugenia) exhibited a strong positive relationship with trial type, exhibiting increasing positional bias as visual reliability decreased (e.g., figure 10a). However, the remaining subjects either did not develop consistent conditioned biases, or developed spontaneous, unreinforced biases towards other positions (e.g., figure 10b,c). The population as a whole did not significantly bias their choices towards the prior biased position (figure 11).

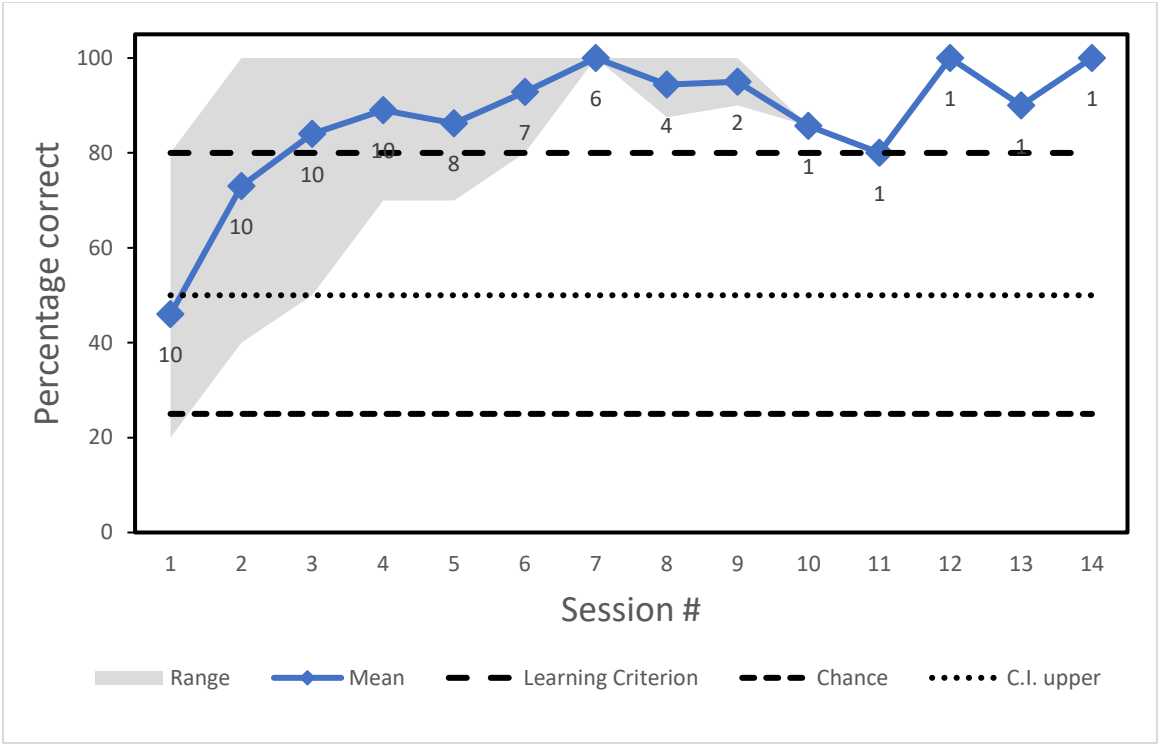



*Figure 5:* Pilot experiment, mean subject discrimination performance during initial training. Task learning follows a typical asymptotic relationship. Point values indicate the number of subjects in training at that point in time – as subjects reached the learning criterion, they were moved to bias training. Shaded area indicates accuracy range.

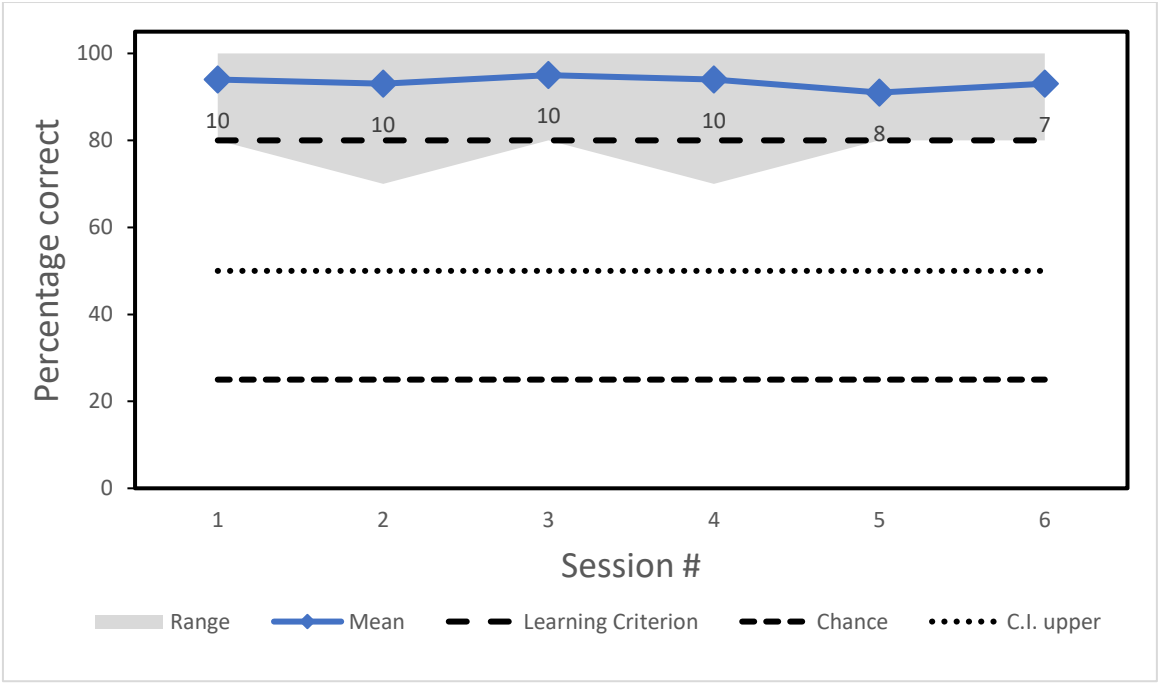



*Figure 6:* Pilot experiment, mean subject discrimination performance during bias training. Shaded area indicates accuracy range.

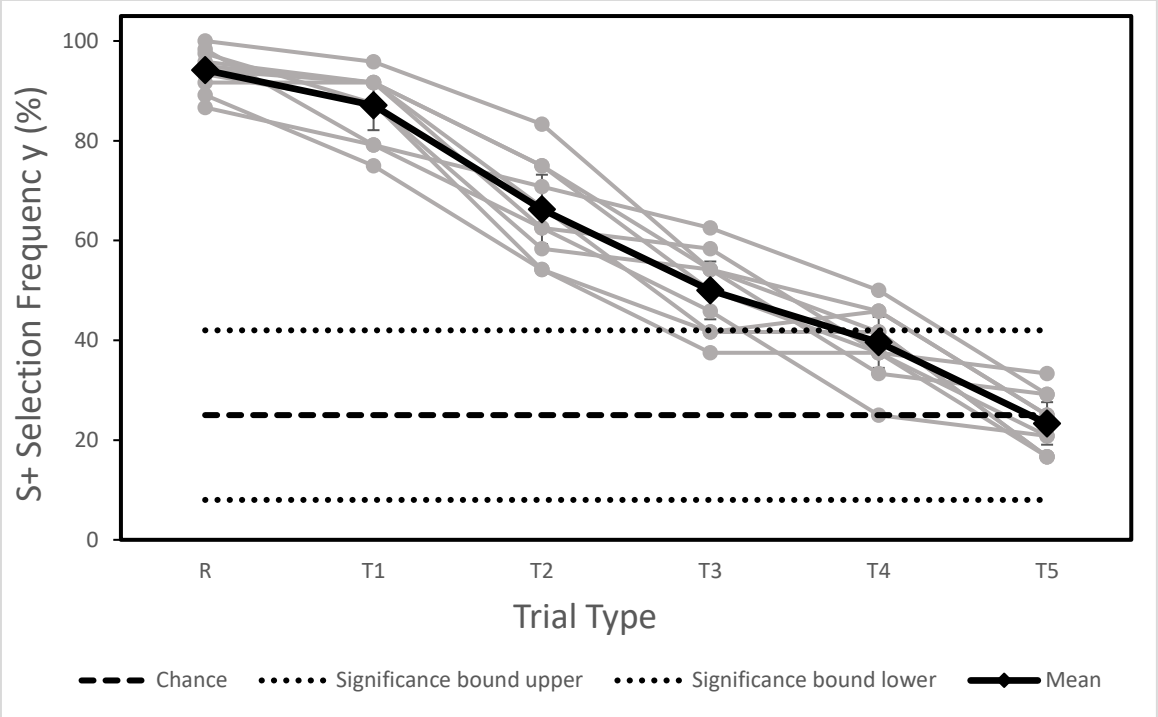



*Figure 7:* Accuracy during testing in the pilot, by trial type, in decreasing order of visual discriminability. Error bars from 95% C.I. of the mean.

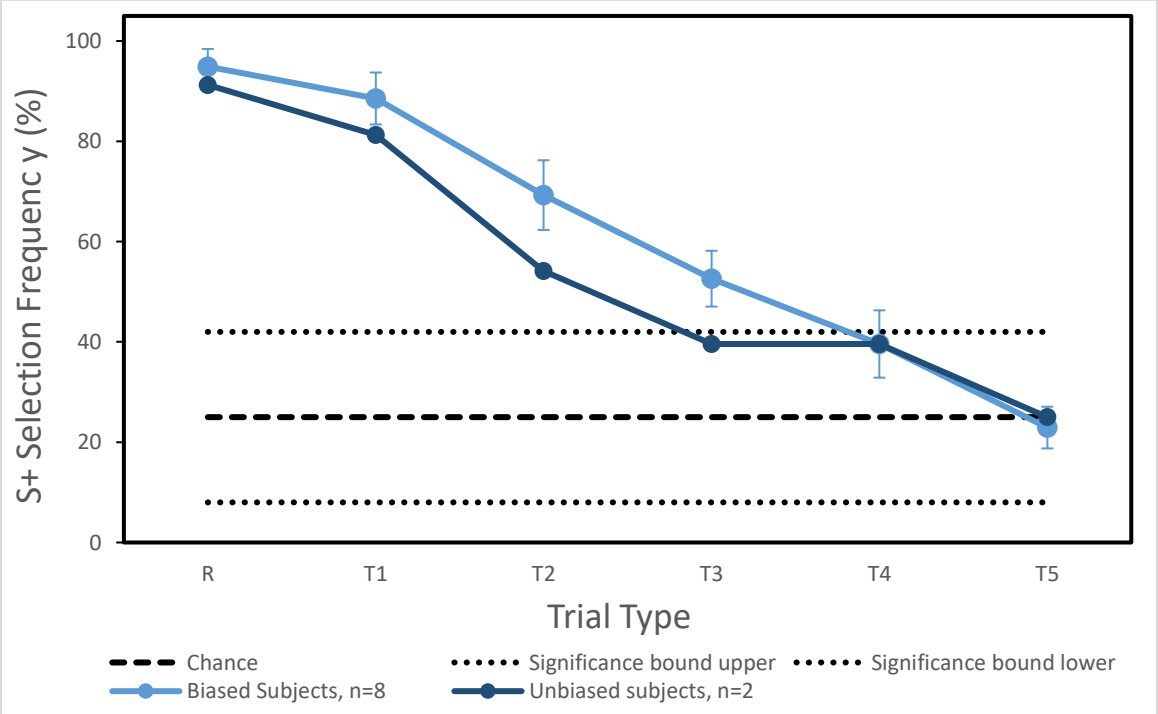



*Figure 8:* Comparison of accuracy between the control and test groups in the pilot experiment. Error bars from 95% C.I of the mean. Appropriate error bars could not be calculated for the control group due to small sample size (n=2), and so are omitted.

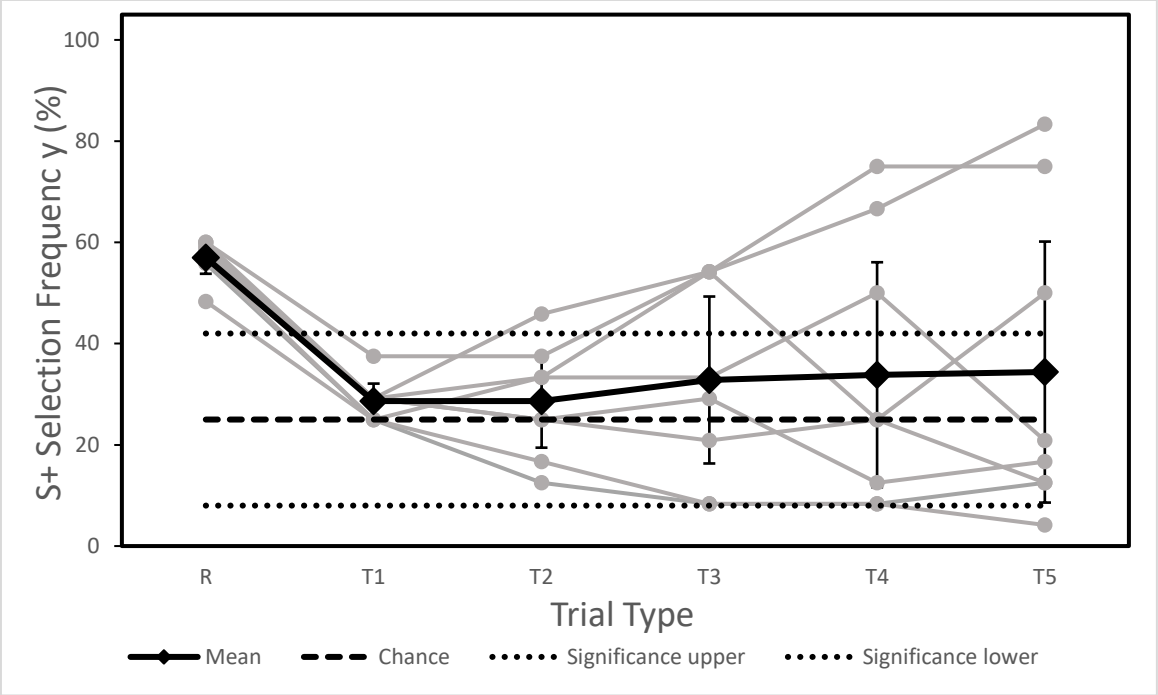



*Figure 9:* Individual variation in bias position selection frequency. As a population mean, subjects did not exhibit a significant preference for the bias position in any trial type. Error bars from 95% C.I of the mean.

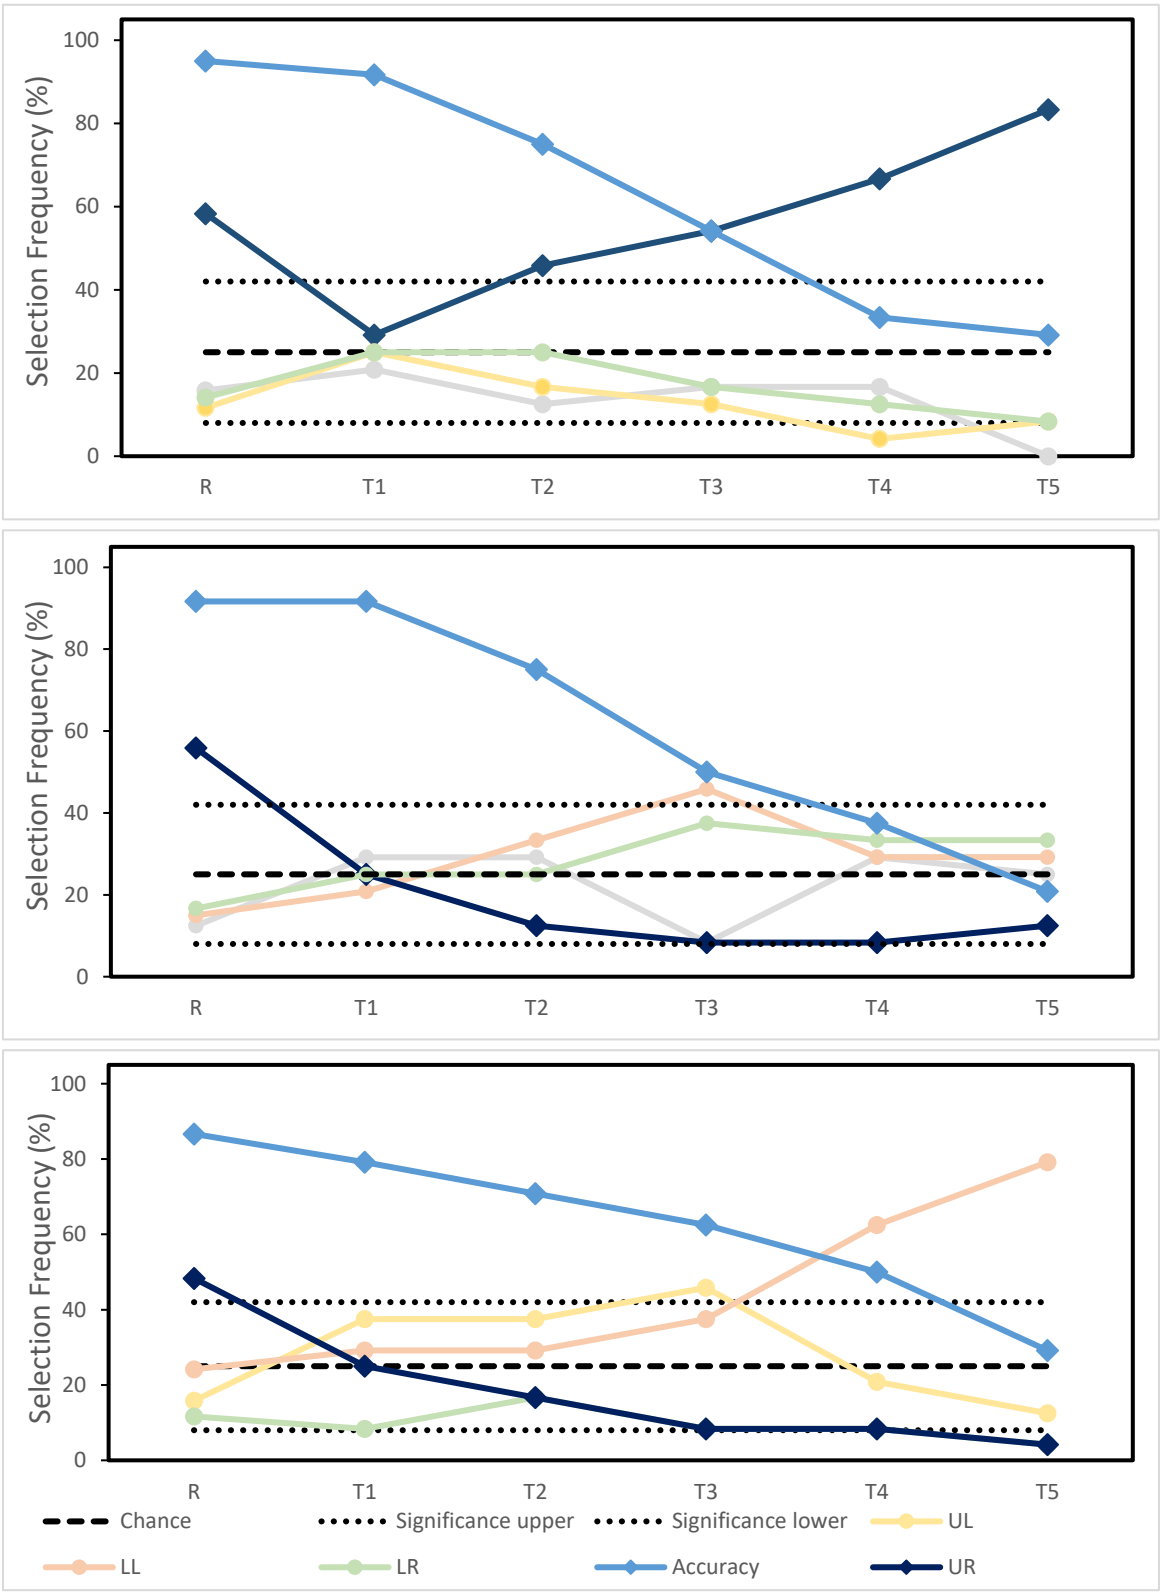



*Figure 10:* Examples of individual choice distributions during testing. From top to bottom; (a) 7-Eugenia, (b) 4-Armitage, (c) 8-Gog. Dark blue line indicates the percentage of choices allocated to the bias position. (a) exhibits the hypothesized shift, biasing choices towards the bias position as visual reliability declines; (b) exhibits no significant positional bias; (c) exhibits a bias towards a non-reinforced position (LL). UL = upper-left, UR = upper-right, LL = lower-left, LR = lower-right.

# Experiments 2 and 3

## Flaws with the Pilot experiment

The initial pilot produced mixed results. While some subjects did prefer the bias position in trials with low visual reliability, the magnitude of the effect in those individuals was comparatively small, and six of the eight subjects in the treatment group either failed to develop positional biases or developed unreinforced biases towards other positions. In brief, most subjects failed to learn the positional bias or failed to use the information when visual cues were lacking. I hypothesized that this occurred due to a failure to form strong associative links between the positional stimulus and the reinforcer (i.e., food) compared to the associative link between food and the visual stimulus. Assuming this to be the case, I hypothesized four factors that could contribute to this:

- *Lower salience* compared to the visual CS+
- *Insufficient time* to form a robust positional CS-US association
- *Overshadowing* by the visual CS
- *Blocking* by the visual CS

Low salience compared to the visual CS would result in only a slow growth of associative strength, contributing to insufficient time to form a robust association. Salience is typically regarded as an inherent property of the CS (Pearce & Bouton 2001), thus the solution to this problem is to increase the time available by increasing the duration of bias training.

The Rescorla-Wagner model [59] holds that any single unconditioned stimulus (US) can only sustain a limited total associative strength, irrespective of the number of conditioned stimuli (CS) it is paired with. This limitation explains the apparently maladaptive Blocking Effect [60], [61], wherein subjects conditioned to a compound stimulus “AB”, having previously been conditioned to “A” fail to develop associations with “B”, since “A” has already taken up all the available associative strength possible with the US. In the context of our experiment, the introduction of the positional cue after reinforcing the visual cue in

isolation may have resulted in associative blocking of position by vision. An alternative is Overshadowing. When presented with two stimuli in a compound (in our case, a positional and visual stimulus), neither may reach the maximum possible associative strength, as each are competing for associations with the same US – thereby overshadowing each other. Ultimately, both models are compatible with the results presented in the pilot. Consequently, I formulated two variants of the original experiment to address different potential problems. These were carried out on two separate populations of fish, both incorporating several individuals that had participated in the pilot experiment. Experiments 2 and 3 were carried out concurrently, with an initial sample size of 4 in each experiment (cohort 1), which was increased with a second round of data collection (cohort 2). There were minimal methodological differences between cohorts, outlined below.

It should be noted that the Rescorla-Wagner framework was formulated for simple conditioning, and to the best of my knowledge has not been directly applied to choice-limited simultaneous discrimination tasks such as 2- and 4AFC. As such, the deductions made here are my own inferences regarding the model.

## Methods

### Experiment 2

This variant closely follows the design on the pilot (experiment 1) but addresses the issues of time and salience. To address time the duration of bias training was doubled (from 6 to 12 sessions of 10 trials). To address lower salience the positional cue (i.e., CS+ presented in the bias position) was paired with a preferred reward (a slice of lancefish). Correct responses in other positions continued to be reinforced using the standard pellet reward used in experiment 1 (the pilot). In this manner, the positional cue was explicitly and distinctly reinforced, contrary to the previous experiment. A short assay was carried out to determine that all subjects preferred lancefish to dry pellet. A control group with no

positional bias received the lancefish at the same frequency (60% of trials), at parity across positions.

## Experiment 3

Experiment 3 was a more radical modification of the original experiment, designed to address the factors of overshadowing and blocking. To avoid these factors arising from the presentation of a compound stimulus, the positional and visual cues were modified to become separate trial types, to which subjects were conditioned in tandem. The "Object" (O) task was identical to the visual task found in the pilot, with no positional bias. In the "Position" (P) task, subjects were reinforced for picking a single position when presented with four identical stimuli (luminance = 90). Subjects were assigned to one of three conditions; (UR), lower left (LL), or a control that appeared in all positions (UL, UR, LL, LR) with equal frequency. Training consisted of sessions of 8 trials (4 Object, 4 Position). A more stringent learning criterion of 75% correct in both tasks over four consecutive sessions was used, to ensure subjects had robustly acquired both tasks. In the first cohort of subjects, subjects that failed to reach this criterion within 50 sessions (i.e., 400 trials) were removed from the experiment. In the second cohort, all subjects (meeting and not meeting the 50-session criterion) were advanced to the testing phase.

The "testing" phase was correspondingly modified, and the same system of O and P trials were interspersed between probe trials.

## Calibrating visual reliability

In order to test whether the cue integration strategy used by the fish matches an optimal Bayesian Inference model, it was first necessary to establish the uncertainty associated with each cue (visual sample and prior) in isolation. Two complementary methods were used to measure these.

Firstly, a behavioural assay was used to determine a subject's expected performance with access to only (i) prior positional bias, or (ii) only current sample visual information. In the former case this is simply the extent of the statistical bias during training, 0.6. However, the latter depends upon the subjective discriminability of the S+ in each test condition (see figure 4). This property cannot be objectively measured at a behavioural level in a single individual – that would require a direct measure of sensory uncertainty before any other sources of information were incorporated, and subjects will always have access to unknown and uncontrolled natural priors [13], [36]. However, it is possible to derive an aggregate estimate across a sample population, which will likely differ in their natural priors. To derive this estimate, a replicate of experiment 2 was carried out, without a reinforced positional bias. This experiment was as described in experiment 2, with the following differences: (i) no positional bias was introduced (ii) only pellet food was used to reinforce correct choices (iii) correct choices in probe trials were also reinforced. Six subjects were enrolled in this assay. Since this method was functionally identical to the control group of the experiment 2, data from these individuals was included in the analysis, for a total population of 8 subjects.

A second, accessory approach was also used. A photometer was used to measure the luminance of the distractor in each test condition, with reference to the positive stimulus. Measurements were taken *in-situ* through the glass front of the aquarium under standard lighting conditions. Measurements were taken by a single observer and the mean calculated over multiple repeated observations in all four positions on the stimulus presentation board. This data was collected by C. Newport and is included here with her permission.

# Results

## Experiment 2

### *Training*

Four subjects were enrolled in cohort 1 for this experiment, followed by eight more in the second cohort for an experimental population of twelve. All learned the basic object discrimination task, reaching the learning criterion in 4-11 training sessions (figure 12). Once subjects had reached the learning criterion, a statistical bias was introduced into the position of the positive stimulus (S+), so it appeared in one position for 60% of trials. During bias training (12 sessions), subject performance remained high (mean = 96.1 ±7.9%) (figure 13). However, one subject died of unknown causes during this phase, such that the testing population was 11 individuals. A subsequent experiment (see "Calibrating visual reliability") was methodologically identical to the control group, so this population (n=6) was included in subsequent analyses for a total population of 17 individuals.

### *Testing*

Of the population of 11, one subject frequently failed to complete testing sessions within a single sitting and was consequently removed from the experiment for a total testing population of 10 individuals.

As in the pilot, accuracy declined as the quality of visual information decreased (figure 14). When no visual information was available (i.e., during T5 trials) accuracy did not statistically differ from chance. Accuracy did not differ significantly between control and treatment groups (figure 15). In Reinforcement (R) trials, accuracy was comparable to that found during bias training (96.9% accuracy, ±5.6%). The extent of positional bias in both individual fish and in aggregate (Figure 16) was analysed using a binomial test. All subjects in the bias treatments exhibited a statistically significant preference for the bias position in probe (T) trials, while the control group did not exhibit any biases that were consistent across treatment conditions (Figure 17).

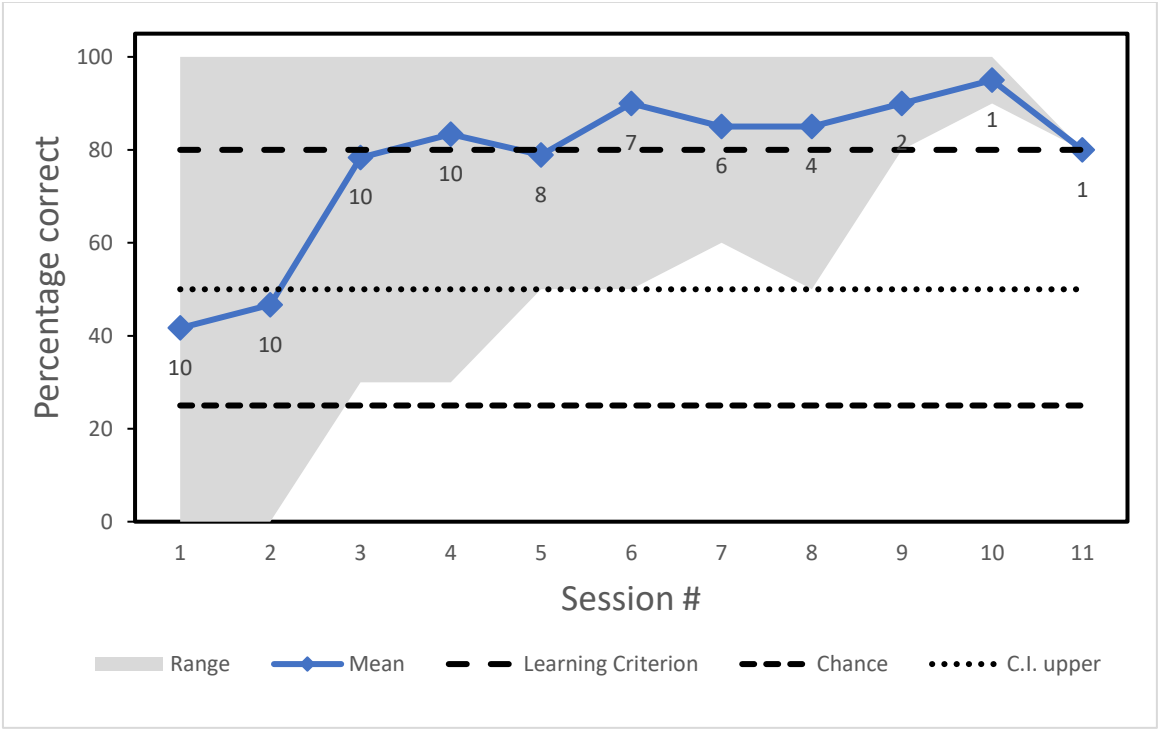



*Figure 11:* Experiment 2, mean subject discrimination performance during initial training. Point values indicate the number of subjects in training at that point in time – as subjects reached the learning criterion, they were moved to bias training. Shaded area indicates range.

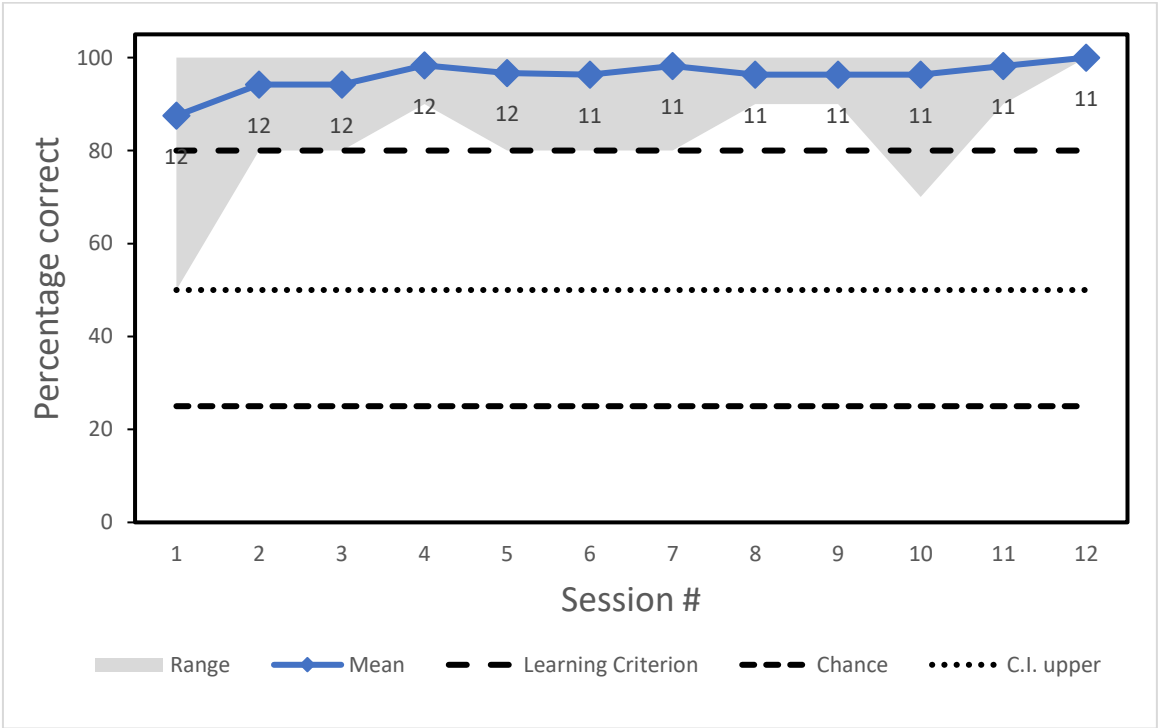



*Figure 12:* Experiment 2, mean subject discrimination performance during bias training. Point values indicate the number of subjects in training at that point in time – one subject died before training could be completed. Shaded area indicates range.

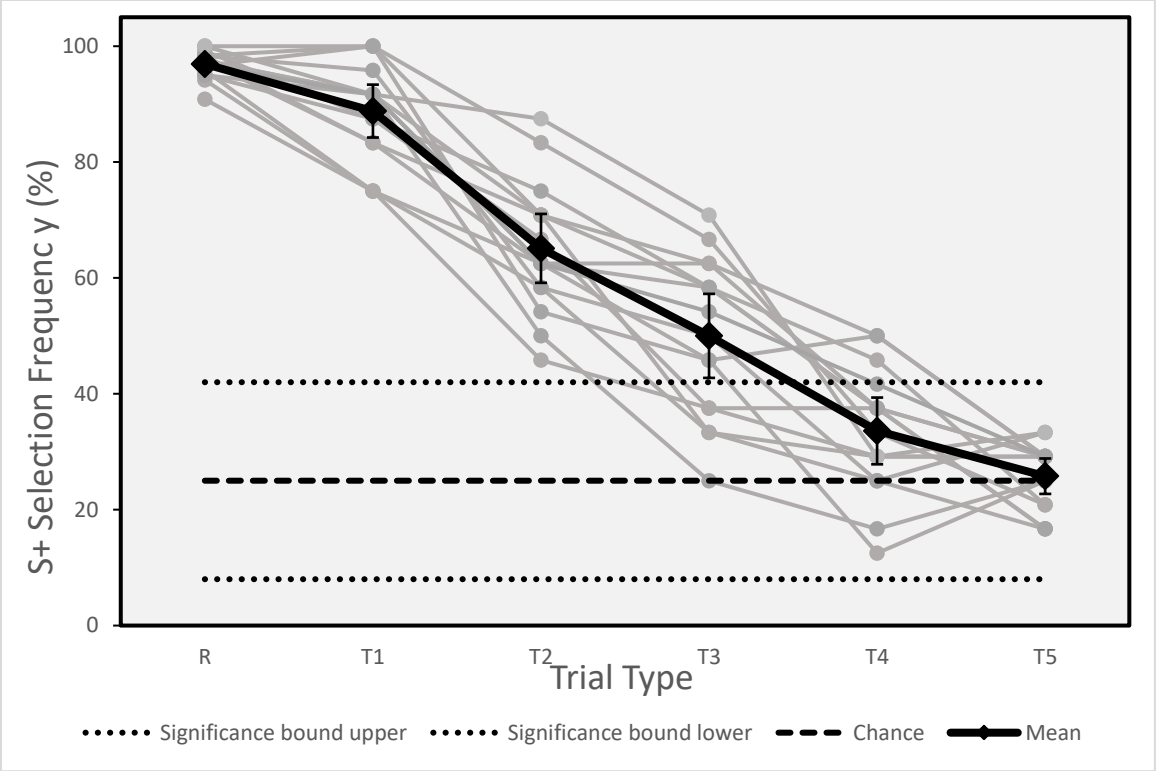



*Figure 13:* Variation in accuracy during testing in the experiment 2, by trial type, in decreasing order of visual discriminability.

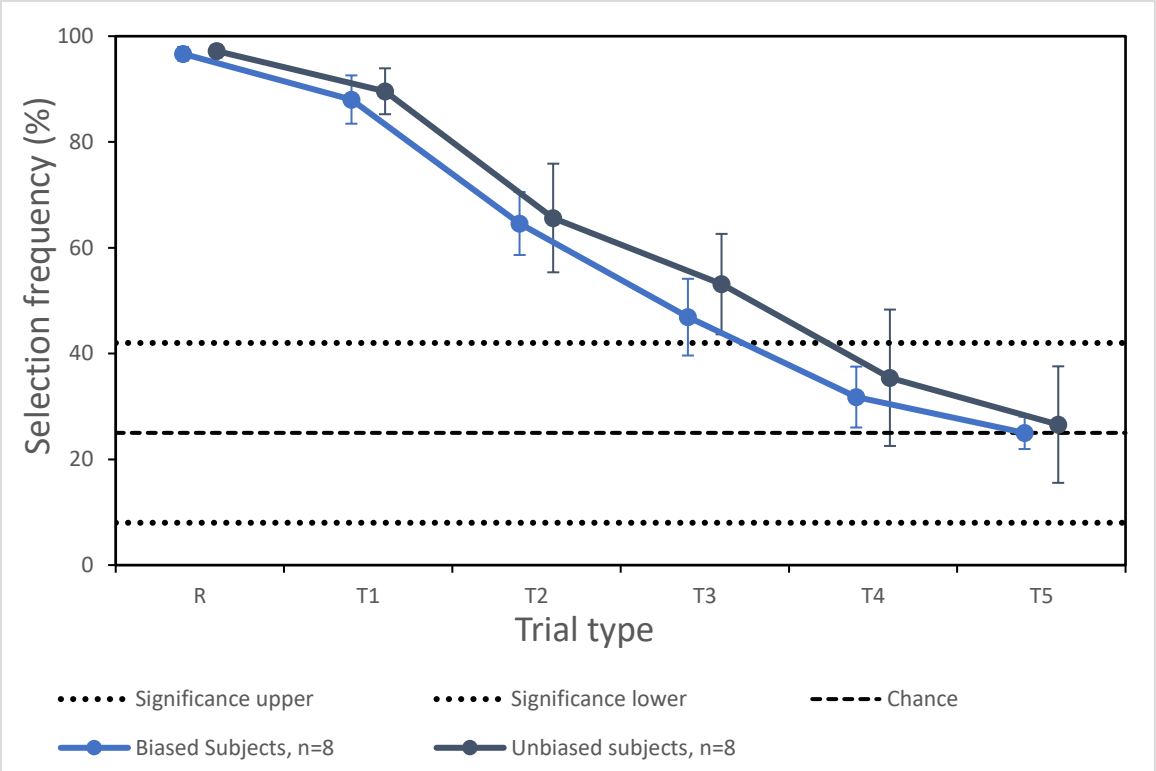



*Figure 14:* Comparison of accuracy between the control and test groups in the pilot experiment.

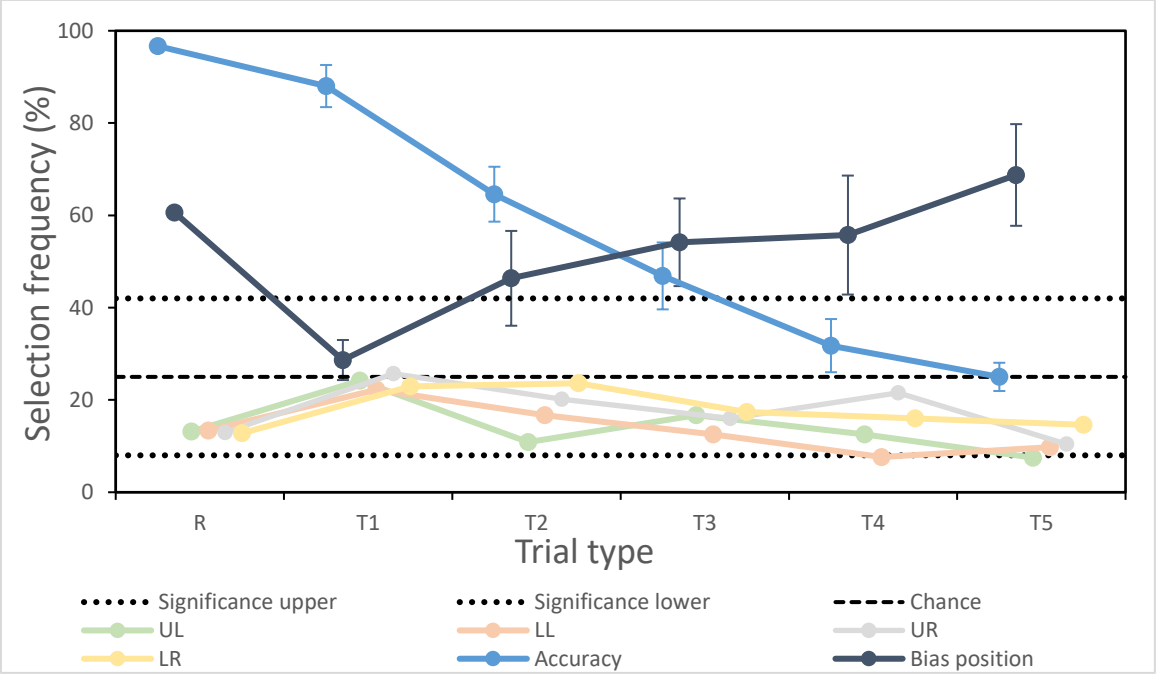



*Figure 15:* Experiment 2, mean distribution of choices during testing, broken down by trial type (bias treatment group). As visual discrimination became more challenging (R→T5), and accuracy declined (light blue) subjects increasingly distributed their choices towards the bias position (dark blue). Correspondingly, fewer choices were allocated to other stimulus positions (UL, UR, LL, LR). Curves for non-bias positions were approximately parallel, indicating that subjects did not have secondary biases to other stimulus positions. Error bars from 95% C.I. of the mean.

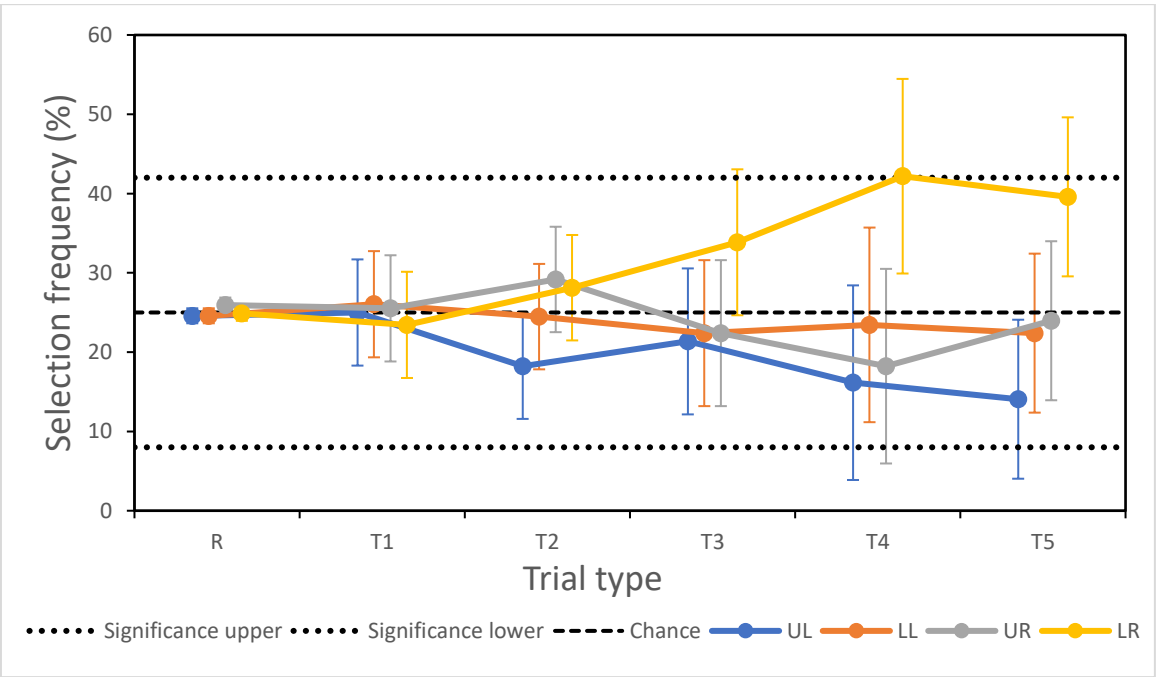



*Figure 16:* Experiment 2, mean distribution of choices during testing, broken down by trial type (control treatment group). Subjects did not exhibit statistically significant preferences for any stimulus position. Error bars from 95% C.I. of the mean.

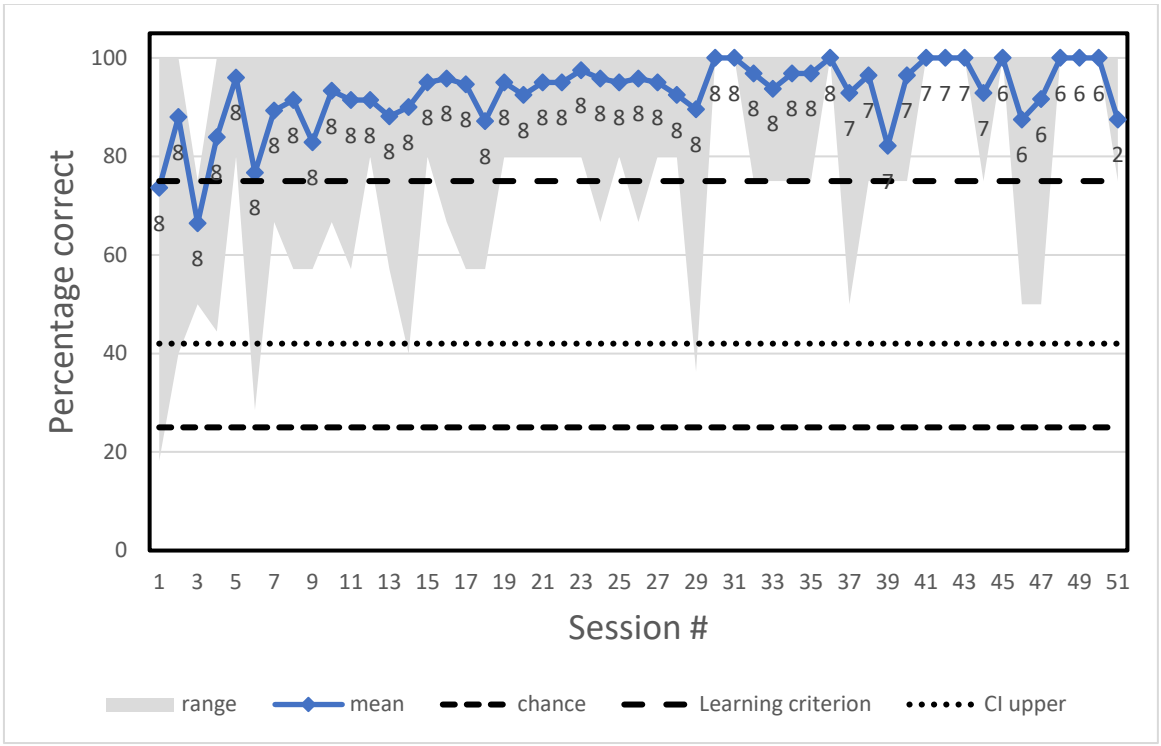
## Experiment 3

### *Training*

Four subjects were enrolled in cohort 1 for this experiment, followed by four more in the second cohort for an experimental population of eight. Subjects in both cohorts acquired the visual task (O) rapidly (within 6 sessions), and performance remained above learning criterion for the remaining duration of training (Figure 18a). Initial training with the first cohort of subjects found that accuracy in the Positional (P) task was very low and did not improve over the course of the initial 5 sessions (Figure 18b). In addition, several subjects developed positional biases contrary to the reinforced bias (i.e., lateralized behaviour / side bias) (Figure 20). To overcome this, a battery of sessions was carried out in which trials were repeated until a correct answer was reported, at which point the subject was rewarded, and the next trial began. This lasted for 24 sessions. The second cohort experienced similar challenges to task acquisition. Accuracy was especially low in these individuals, resulting in low subject motivation (assessed by increasing response delay). Several temporary, modified training techniques were used to overcome these problems, as follows (i) alternating sessions of all-P and all-O trials, (ii), placing the reward directly on the stimulus, (iii) *ad-libidum* trials, and (iv) iterated trials as in the first cohort.

### *Testing*

Of the initial population of eight, six subjects were progressed to the testing phase. Results from experiment 3 were broadly similar to those found in experiment 2. Of the testing population of 6 subjects, all successfully completed the testing phase. Accuracy was closely correlated with the reliability of the visual cue, and when no visual information was available subjects performed at chance (figure 21). Accuracy in the Object Reinforcement ("RP") task was similar to that found in "R" trials in experiment 2 (96.5% accuracy, ±3.1%).



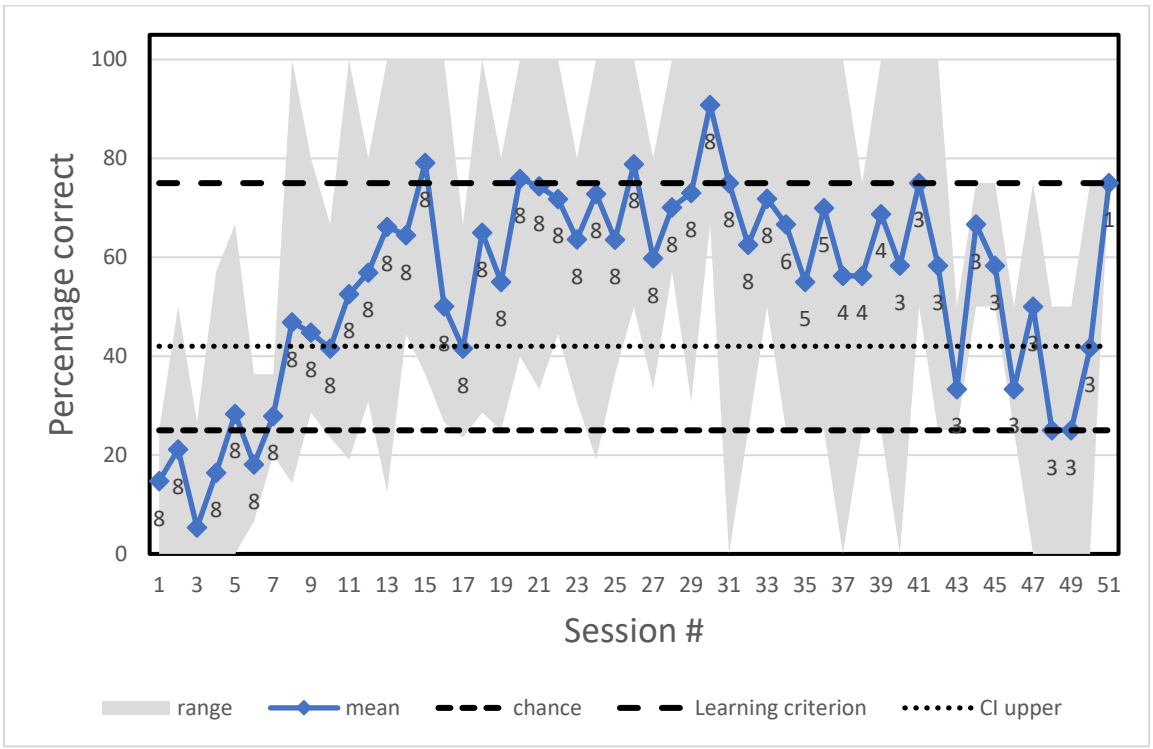




*Figure 18:* Experiment 3, mean subject discrimination performance during training. From top to bottom; (a) Object ["O"] task, (b) Position ["P"] task. Point values indicate the number of subjects in training at that point in time – as subjects reached the learning criterion, they were moved to testing. Shaded area range. While subjects rapidly acquired the object discrimination task following a steep asymptotic curve (b), acquisition of the position task was much slower and less consistent (a).

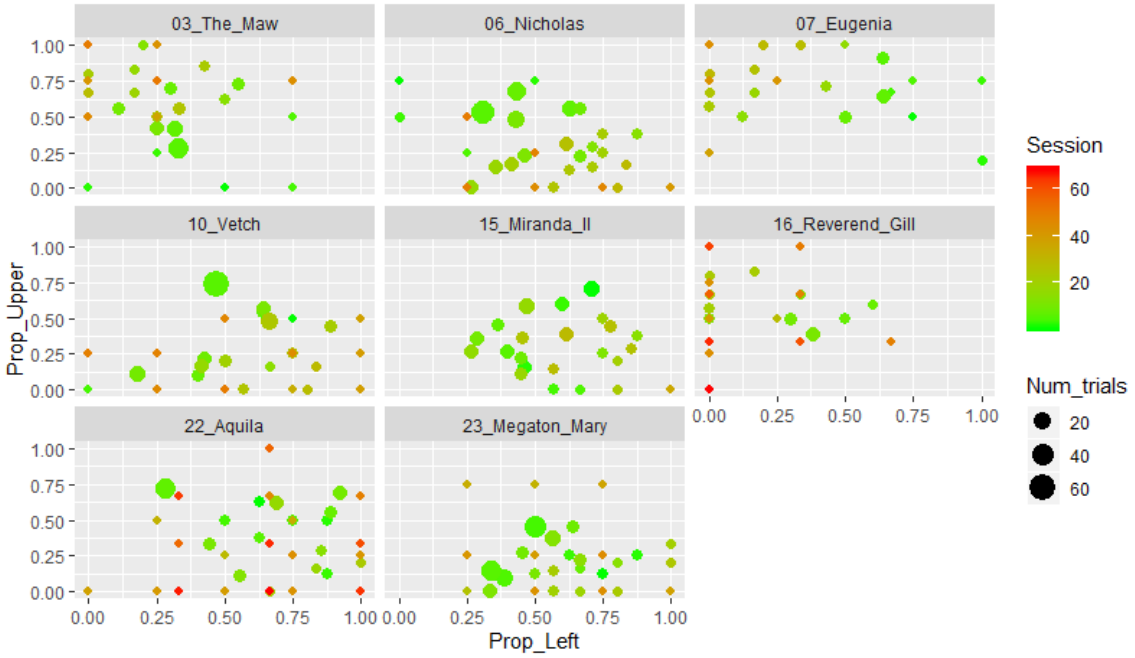



*Figure 20:* Change in degree of positional bias through training, by subject. Colour gradient corresponds to session #, point size corresponds to the number of trials per session. A value of 1.0 on either axis indicates all choices were in upper positions (UL or UR) or left positions (UL or LL) respectively, and vice versa. Sessions are widely distributed across the plots, and similarly coloured points (i.e., chronologically adjacent sessions) are not clustered, indicating a lack of consistency in positional bias between sessions.

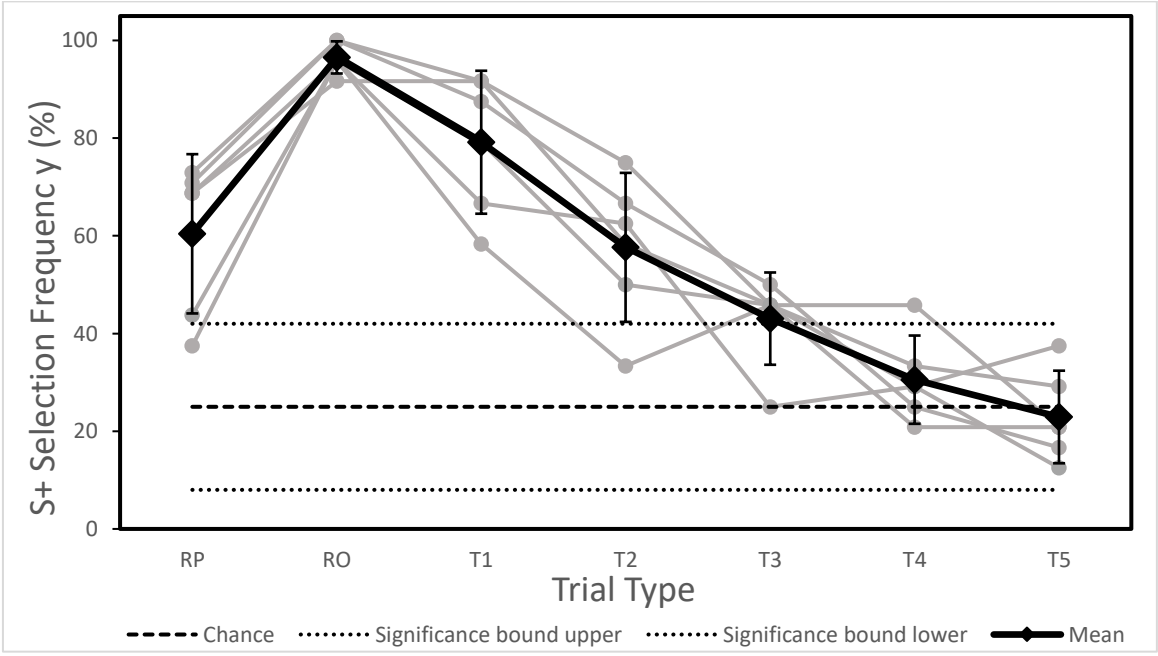



*Figure 21:* Accuracy during testing in the experiment 3, by trial type, in decreasing order of visual discriminability. Error bars from 95% C.I. of the mean.

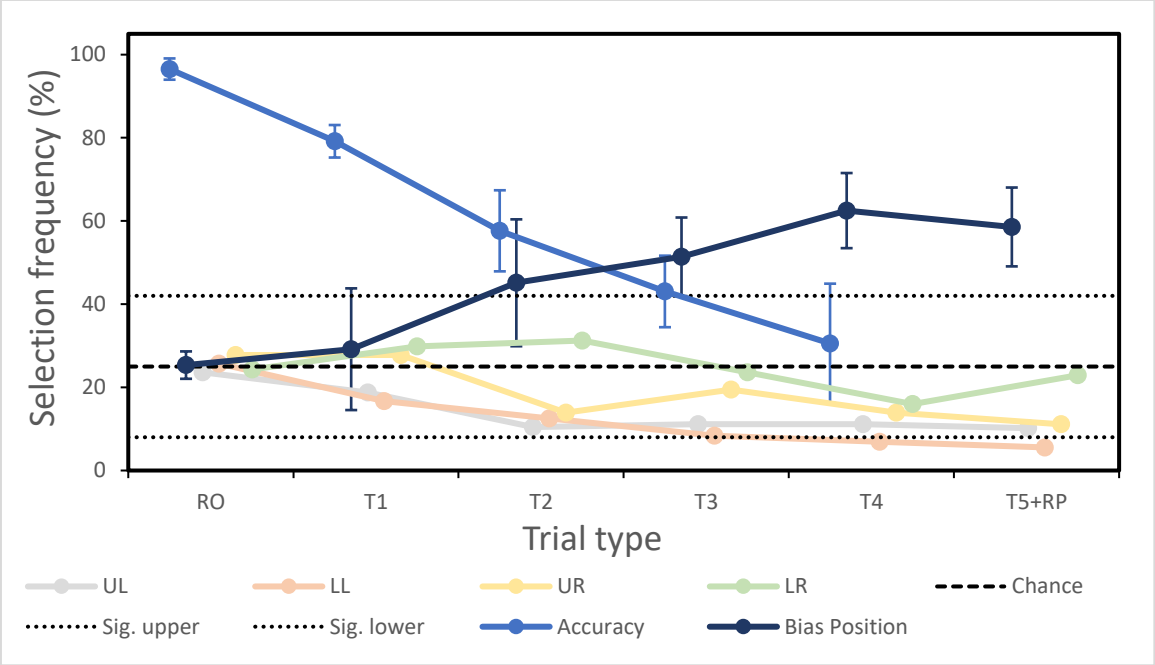



*Figure 22:* Experiment 3, mean distribution of choices during testing, broken down by trial type (bias treatment group). As visual discrimination became more challenging (R→T5), and accuracy declined (light blue) subjects increasingly distributed their choices towards the bias position (dark blue). Correspondingly, fewer choices were allocated to other stimulus positions (UL, UR, LL, LR). Curves for non-bias positions were approximately parallel, but subjects consistently chose the LR position at a higher rate than predicted. Error bars from 95% C.I. Accuracy is not given for T5+RP as there was no objectively "correct" choice for these trials.

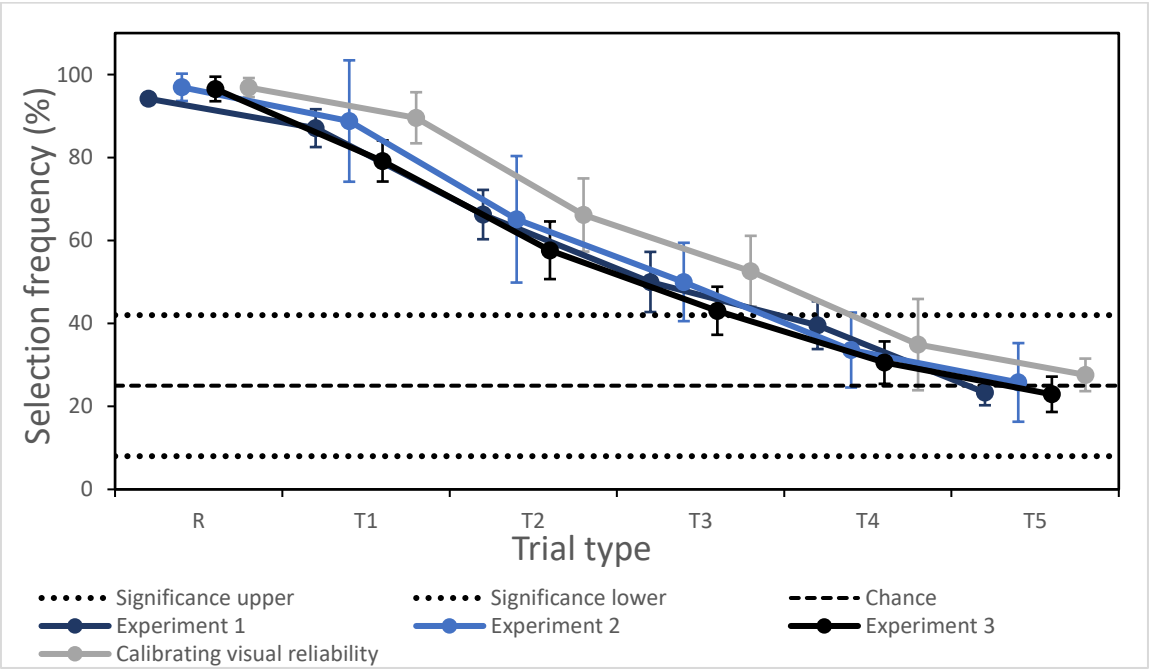



*Figure 23:* Comparison of mean subject accuracy across experiments. Error bars from 95% C.I. Accuracy not given for T5+RP as there was no objectively "correct" choice for these trials.

## Bias position preference

While the learning curves differed radically, the results from probe trials across experiments 2 and 3 were similar both in the degree of positional bias (Figures 15,22), and in overall accuracy (figure 23). As such, the two data sets were aggregated to evaluate whether trial type (i.e., visual discriminability) predicted the rate at which subjects selected the bias position. Since the response variable (selection frequency of the bias position) is binomial (each trial scoring either 0 or 1), a generalised linear mixed effects model was computed with a "logit" (binomial) link function using the "glmer" function from the lme4 R package [62] to test for an effect of discriminability on bias position preference. A GLMM permits testing for the effect of a treatment (trial type) while accounting for repeated measures from individual subjects (coded as a random, intercept-only predictor "subject" in the verbal equation below). The effect of inter-experiment differences was included as a predictor (experiment_identity), and to account for differing experience levels of a two-level factor ("with" and "without" prior experience) was coded as an interaction with trial type.

*Bias position selection frequency ~ treatment*experience + experiment_identity + (1|subject)*

P-values were calculated using a z test on the one full model computed. I chose the parameters based on the contingencies of the experimental setup, and the biological plausibility of the system rather than use a backward deletion process. The z test therefore is akin to measuring each parameter's contribution to the full model.

There was no effect of the interaction between the effect of treatment and the experiment ($z = 1.2$, $p = 0.23$), suggesting that subjects responded similarly during testing regardless of the precise training method used. Prior experience had a non-significant positive relationship with bias, experienced subjects being more likely to select the bias ($z = 1.8$, $p = 0.08$). Having accounted for repeated measures from individual subjects, trial type accounted for the bulk of variance in bias position selection, with a significant positive

relationship between choice bias and trial type (R→T5) ($z = 7.72$, $p < 0.0001$). This confirms the interpretation (figures 16, 22) that subjects biased their choices towards the previously biased position as visual reliability declined.

## Calibrating visual reliability

To test whether the cue integration strategy used by the fish matches an optimal Bayesian Inference model, it was first necessary to establish the uncertainty associated with each cue (visual sample and prior) in isolation. The mean accuracy of subjects in the absence of conditioned prior was comparable to that found in experiments 1-3 (figure 23). While the subjects were not conditioned to any biased position, there was some suggestion of a population-level preference for the lower-right (LR) position, though this did not reach statistical significance (figure 24). A photometer was used to measure the relative luminance of the distractors in each test condition, with reference to the positive stimulus. The results suggest that the measured relative luminances of the stimuli *in-situ* are a good match for the values coded while printing (Figure 25), and that these correspond closely to the values for subject accuracy derived above (Figure 26).

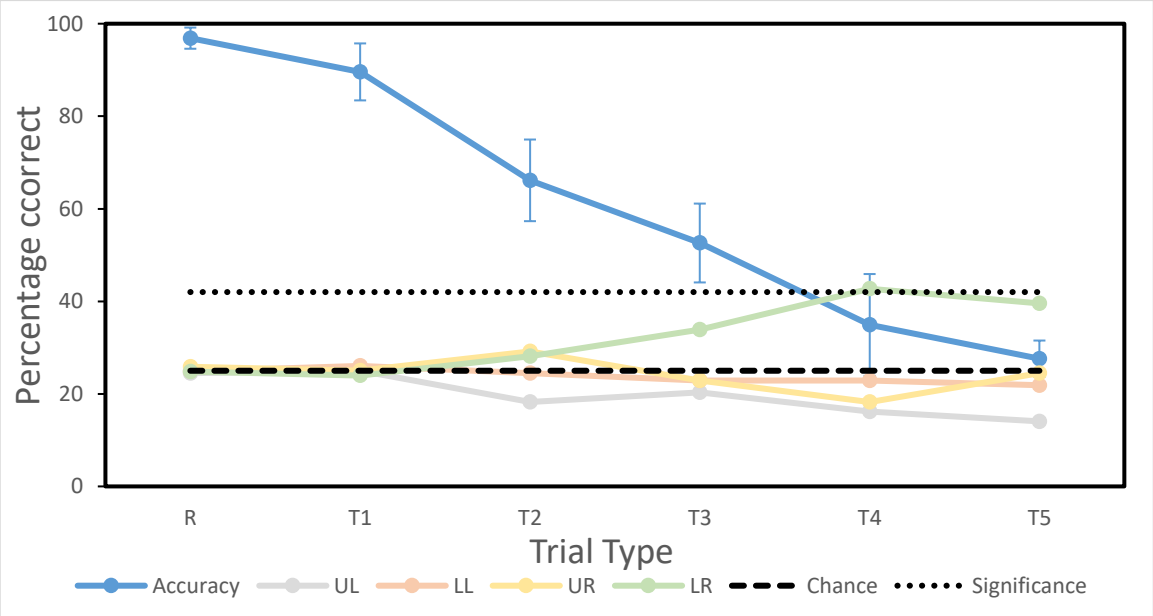



*Figure 24:* Distribution of choices in subjects conditioned without a positional bias, by trial type. In trials T4-T5 the population exhibited a noticeable but non-significant

preference for the LR position. Error bars from 95% C.I. Accuracy not given for T5+RP as there was no objectively "correct" choice for these trials.

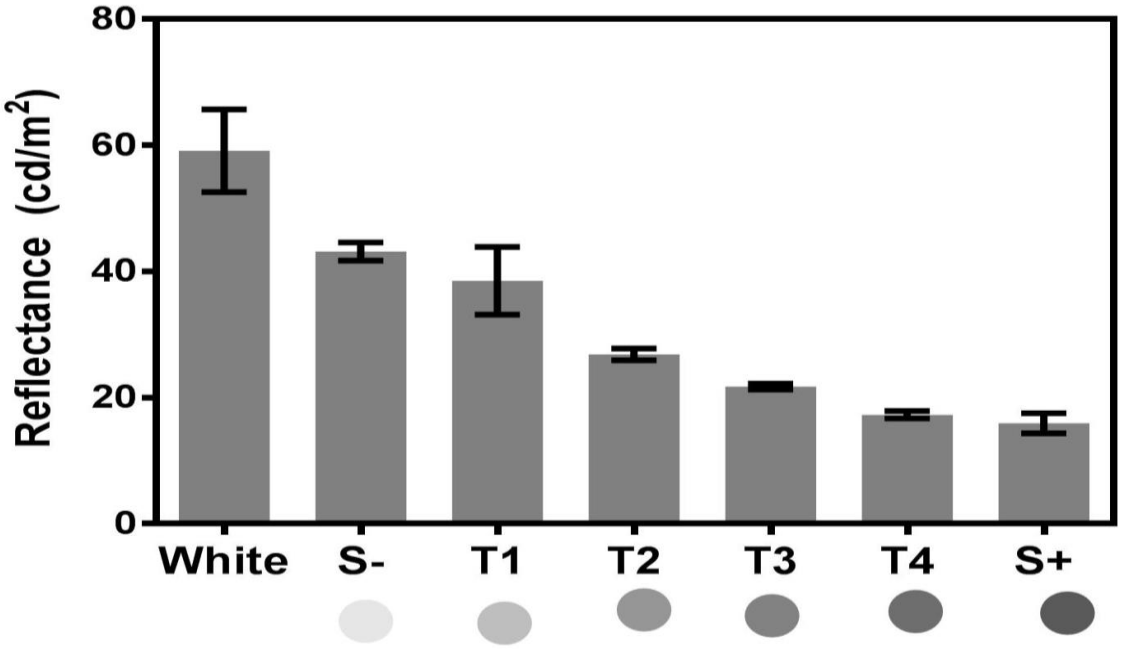



*Figure 25:* Stimuli reflectance measurements. Data collected and plotted by C. Newport 2017, used with permission.

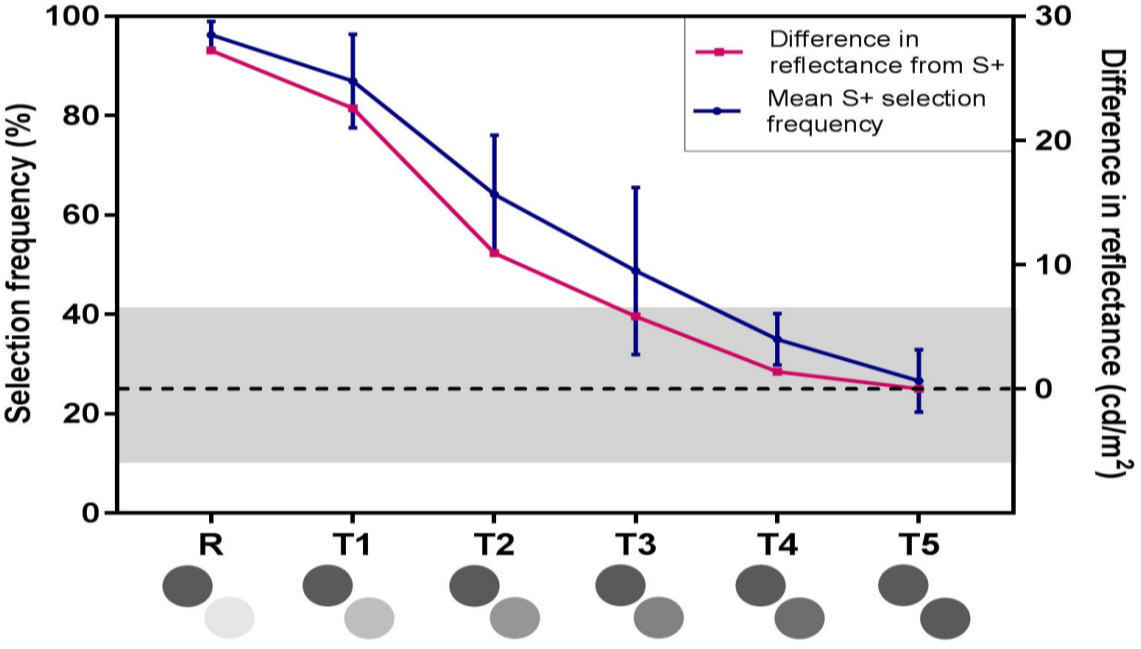



*Figure 26:* Comparison between subject accuracy (blue) and measured reflectance difference between CS+ and distractor (magenta). Dotted line indicates chance discrimination performance (25%). error bars from 95% C.I. Data plotted by C. Newport 2017, used with permission.

## Discussion

In this chapter, I present the results of three core experiments designed to assess the ability of Picasso triggerfish to incorporate multiple sources of information – including priors – into a decision-making task. To achieve this, I presented subjects with an object discrimination task with two potential cues; visual discrimination of the positive stimulus, and prior positional information regarding the statistical distribution of said positive stimulus. Having been conditioned to both these cues, I observed subject's decisions when the quality of visual information was reduced. I predicted that triggerfish would principally use visual information to complete the task but would shift to relying on prior information when the visual cue was less discriminable. I also distinguished two possible paradigms of prior information use – a weighted integrative approach (Bayes), or a simple alternating exclusive approach (Redundant) – see figure 2. This framework made two major predictions about the subject's behaviour – first; regarding task accuracy, and second; regarding positional biases.

### Prediction 1

Subject accuracy would be negatively correlated with luminance contrast. Since the positional bias was not present in probe trials, subjects' accuracy was solely a function of their ability to discriminate the CS+ from the distractors. As such, subject performance should be negatively correlated with stimulus similarity; and when no visual information was available (i.e., when CS+ and distractors were identical, T5 condition), subject performance should not differ from chance. Deviation from this assumption would have indicated that the subjects had access to uncontrolled sources of information, such as cues in the experimenter's behaviour. Confirming the prediction, accuracy was strongly correlated with the measured luminance contrasts between the S+ and distractors in all experiments. Where no visual information was available (i.e., in the T5 testing condition) performance did not differ from chance in all cases. This demonstrates that subjects

principally relied on visual information and perceived more similar stimuli as less discriminable.

## Prediction 2

Subjects would linearly increase their reliance on prior information as visual reliability declined. I distinguished four separate cases of prior information uses, dubbed the null models, redundant model, and weighted model (figure 2). The null models represent rigid decision strategies, depending upon only a single source of information – either visual or prior information. The latter was rendered extremely unlikely by our training schedule, since the visual stimulus was introduced first, and potentially resulted in a greater reward frequency and reliably than the positional cue. In a redundant model, subjects would attend to both sources of information, but any single trial would be informed exclusively by one or another, depending on the reliability of either cue. In this model, we would predict a sharp sigmoidal curve to the bias position response, reflecting a rapid shift in information use. Finally, in the weighted model, each trial would be informed by both sources of information, with the relative weighting according to stimulus reliability. In this model, we would expect a roughly linear shift in the selection of the bias position with changing visual reliability. If subjects did use such a computation, an optimal cue combination strategy would correspond to a Bayesian model [12], [13]. I found that subjects in experiments 2 and 3 did indeed use a dynamic decision strategy, with a roughly linear response to changing visual reliability. During training and reinforcement trials (when the stimuli were readily discriminable) subject's decisions were dominated by the visual cue. However, reduced discriminability of the visual cue was correlated with a shift towards choosing the bias position, indicating the subjects' estimates were dominated by prior information (figure 2d). Furthermore, when no visual information was available, subjects selected the bias position at a frequency comparable to the degree of positional bias they experienced during training. I infer that the subjects can learn statistical prior information to inform

decisions and weight this information against sensory information to produce a posterior estimate that incorporates both sources of information.

## Models of updating

An outstanding question is how triggerfish are weighting information in the experiment. I have proposed a Bayesian inference model, but there are other models which incorporate both current sample and prior information. Lower-level approximations, such as linear approximants, may require less memory or knowledge while performing similarly under certain conditions [63]. For example, [64] elucidated a simple model describing foraging behaviour of ground squirrels, in which a linear operator is used to control the relative weighting of prior information according to recency. One possibility – which remains ambiguous due to the binary output of our study design – is that subjects did attend to both cues, but that the prior informational cue was only integrated over very small time periods. For example, consider a subject's decisions at the end of the current experiment. Is its internal representation of the prior distribution informed by its entire prior experience in the task (400+ trials over 3-4 weeks), by only the past week, the past day, or even only the last few trials? Drawing parallels to the foraging theory literature, one finds that timescales for initial learning of patch distributions are typically on the order of 1-2 weeks and a few hundred behavioural iterations, both in birds [65], [66] and invertebrates [67]. However, this does not necessarily indicate that prior information was consistently integrated over such a period. One possibility is that there is a strong effect of recency (i.e., decision dominated by either the most recent rewarded, or the most recent readily discriminable trial). When confronted with a discrimination, a subject may prefer to repeat its last choice than to consult an internalised statistical distribution built up over days – after all, such an approach would be much less cognitively demanding. This model greatly resembles (and may be compatible with) the notion of the sequential effect [68], in which subjects spuriously detect patterns in random sequences of stimuli. Such a decision strategy might

be adaptive if the natural environment is consistent on short temporal scales. This is a subject that I explore in greater detail in subsequent chapters.

## Caveats

Once they had completed training, subjects in experiments 2 and 3 performed very similarly in the testing phase. However, the learning curves for the two groups were drastically different, as was the proportion of subjects that reached the learning criterion. Given the comparatively small sample sizes (i.e. 12 vs 8), this could have arisen as a result of individual differences, such as varying sensory acuity, memory, or personality [69]. However, the consistency within each experiment, and the large disparity with how quickly subjects in experiment 3 learned the positional versus visual tasks suggests that methodological differences were the driving factor. This is consonant with field studies of animal foraging, which have found that small changes in the environment and context may result in animals using quite different strategies [41]. During the acquisition phase, subjects in experiment 3 appeared to exhibit a range of non-random behavioural strategies, such as lateralization. This suggests that, in addition to different sources of information, subjects may weight different decision strategies in the context of such tasks.

Additionally, several individuals assigned to the control groups developed spontaneous biases, and there was a degree of consistency across the population, with a near-significant population bias towards the lower-right stimulus position. These did not correlate with experience in subjects that had been used in prior experiments. While the bias did not exceed significance in the test population aggregate, there was evidence of significant bias in some individuals. For example, In T5 trials, with no sample information available, subject 17 chose the LR position in 54% of trials, greatly exceeding statistical significance (42%, calculated from binomial test). This pattern resulted in a less frequent reward schedule for the individual concerned. Such Skinnerian “superstitious” behaviour

[70], may be maladaptive in experimental setting, where stimuli are randomised, but adaptive in natural environment where events are rarely truly random. The consistency of the positional bias through time and across test population is more difficult to explain, however. One possibility is that these subjects opted for a robust non-optimal strategy (i.e., lateralised choices, reward ration ¼) in a situation with an initially low reward ratio, and by reaching a local decision optimum failed to reach the optimal response which maximised reward rate. Alternatively, bias may have arisen from uncontrolled priors due to common experiences outside the experiment. Comparable side preferences in randomised discrimination tasks have been found by other investigators working with similar systems [71]. Side bias and task acquisition are further explored in subsequent chapters.

## Conclusions

With this study, I have provided a robust demonstration that not only are triggerfish capable of attending to and incorporating prior information into their decision making, but also that they are capable of integrating this with sensory information and weighting each source of information according to reliability. To my knowledge, this constitutes the first demonstration that a fish species is capable of using and weighting prior information in a conditioned task. This demonstrated that fish are capable of using complex, dynamic decision strategies in conditioned behaviours, and suggests that they may be able to apply these cognitive tools to more abstract aspects of decision-making such as problem solving, mate selection and foraging - potentially changing our understanding of when these advanced cognitive abilities evolved within the vertebrate lineage.

# References

[1] S. R. X. Dall, L.-A. Giraldeau, O. Olsson, J. M. McNamara, and D. W. Stephens, "Information and its use by animals in evolutionary ecology," *Trends Ecol. Evol.*, vol. 20, no. 4, pp. 187–193, 2005.

[2] B. Luttbeg, "A Comparative Bayes tactic for mate assessment and choice," *Behav. Ecol.*, vol. 7, no. 4, pp. 451–460, Dec. 1996.

[3] M. Mangel, "Dynamic information in uncertain and changing worlds," *J. Theor. Biol.*, vol. 146, no. 3, pp. 317–332, Oct. 1990.

[4] L.-A. Giraldeau, "Solitary foraging strategies," in *Behavioural Ecology*, E. Danchin, L.-A. Giraldeau, and F. Cézilly, Eds. Oxford, UK: Oxford University Press, 2008, pp. 233–256.

[5] G. Perry and E. R. Pianka, "Animal foraging: past, present and future," *Trends Ecol. Evol.*, vol. 12, no. 9, pp. 360–364, Aug. 1997.

[6] D. C. Knill and A. Pouget, "The Bayesian brain: The role of uncertainty in neural coding and computation," *Trends Neurosci.*, vol. 27, no. 12, pp. 712–719, 2004.

[7] T. Merkle and R. Wehner, "Desert ants use foraging distance to adapt the nest search to the uncertainty of the path integrator," *Behav. Ecol.*, vol. 21, no. 2, pp. 349–355, 2010.

[8] K. A. Schmidt and C. J. Whelan, "Nesting in an uncertain world: Information and sampling the future," *Oikos*, vol. 119, no. 2, pp. 245–253, Feb. 2010.

[9] C. W. Clark and M. Mangel, "Foraging and flocking strategies: information in an uncertain environment.," *Am. Nat.*, vol. 123, no. 5, pp. 626–641, 1984.

[10] M. O. Ernst and M. S. Banks, "Humans integrate visual and haptic information in a statistically optimal fashion," *Nature*, vol. 415, no. 6870, pp. 429–433, 2002.

[11] R. J. van Beers, A. C. Sittig, and J. J. Gon, "Integration of proprioceptive and visual position-information: An experimentally supported model.," *J. Neurophysiol.*, vol. 81, no. 3, pp. 1355–64, Mar. 1999.

[12] K. P. Körding, "Decision theory: What 'should' the nervous system do?," *Science*, vol. 318, no. 5850. pp. 606–610, 26-Oct-2007.

[13] J. M. McNamara, R. F. Green, and O. Olsson, "Bayes' theorem and its applications in animal behaviour," *Oikos*, vol. 112, no. 2, pp. 243–251, Feb. 2006.

[14] M. Ernst, "A Bayesian view on multimodal cue integration," *Perception*, vol. 131, no. Chapter 6, pp. 105–131, 2006.

[15] D. Alais and D. Burr, "The Ventriloquist Effect Results from Near-Optimal Bimodal Integration," *Curr. Biol.*, vol. 14, no. 3, pp. 257–262, Feb. 2004.

[16] M. O. Ernst and H. H. Bülthoff, "Merging the senses into a robust percept," *Trends Cogn. Sci.*, vol. 8, no. 4, pp. 162–169, 2004.

[17] S. Gepshtein and M. S. Banks, "Viewing Geometry Determines How Vision and Haptics Combine in Size Perception," *Curr. Biol.*, vol. 13, no. 6, pp. 483–488, Mar. 2003.

[18] D. C. Knill, "Robust cue integration: A Bayesian model and evidence from cue-conflict studies with stereoscopic and figure cues to slant," *J. Vis.*, vol. 7, no. 7, p. 5, May 2007.

[19] J. Lewald, W. H. Ehrenstein, and R. Guski, "Spatio-temporal constraints for auditory-visual integration," *Behav. Brain Res.*, vol. 121, no. 1–2, pp. 69–79, Jun. 2001.

[20] Y. Gu, D. E. Angelaki, and G. C. DeAngelis, "Neural correlates of multisensory cue integration in macaque MSTd," *Nat. Neurosci.*, vol. 11, no. 10, pp. 1201–1210, Oct. 2008.

[21] M. A. Meredith and B. E. Stein, "Spatial determinants of multisensory integration in cat superior colliculus neurons," *J. Neurophysiol.*, vol. 75, no. 5, pp. 1843–1857, 1996.

[22] T. R. Stanford, S. Quessy, and B. E. Stein, "Evaluating the Operations Underlying Multisensory Integration in the Cat Superior Colliculus," *J. Neurosci.*, vol. 25, no. 28, pp. 6499–6508, Jul. 2005.

[23] J. D. Cushman *et al.*, "Multisensory Control of Multimodal Behavior: Do the Legs Know What the Tongue Is Doing?," *PLoS One*, vol. 8, no. 11, p. e80465, Nov. 2013.

[24] E. A. Whitchurch and T. T. Takahashi, "Combined Auditory and Visual Stimuli Facilitate Head Saccades in the Barn Owl ( Tyto alba )," *J. Neurophysiol.*, vol. 96, no. 2, pp. 730–745, Aug. 2006.

[25] F. J. Stewart, D. A. Baker, and B. Webb, "A model of visual-olfactory integration for odour localisation in free-flying fruit flies.," *J. Exp. Biol.*, vol. 213, no. 11, pp. 1886–900, Jun. 2010.

[26] J. M. McNamara and A. I. Houston, "The application of statistical decision theory to animal behaviour," *J. Theor. Biol.*, vol. 85, no. 4, pp. 673–690, Aug. 1980.

[27] A. Oaten, "Optimal foraging in patches: A case for stochasticity," *Theor. Popul. Biol.*, vol. 12, no. 3, pp. 263–285, Dec. 1977.

[28] J. M. Biernaskie, S. C. Walker, and R. J. Gegear, "Bumblebees learn to forage like bayesians," *Am. Nat.*, vol. 174, no. 3, pp. 413–423, Sep. 2009.

[29] B. Luttbeg and T. A. Langen, "Comparing alternative models to empirical data: cognitive models of Western Scrub-Jay foraging behaviour," *Am. Nat.*, vol. 163, no. 2, pp. 263–276, Feb. 2004.

[30] O. Olsson, U. Wiktander, N. M. A. Holmgren, and S. G. Nilsson, "Gaining Ecological Information about Bayesian Foragers through Their Behaviour. II. A Field Test with Woodpeckers," *Oikos*, vol. 87, no. 2, p. 264, Nov. 1999.

[31] T. J. Valone and J. S. Brown, "Measuring Patch Assessment Abilities of Desert Granivores," *Ecology*, vol. 70, no. 6, pp. 1800–1810, Dec. 1989.

[32] J. A. van Gils, I. W. Schenk, O. Bos, and T. Piersma, "Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches.," *Am. Nat.*, vol. 161, no. 5, pp. 777–93, May 2003.

[33] V. Selonen and I. K. Hanski, "Decision making in dispersing Siberian flying squirrels," *Behav. Ecol.*, vol. 21, no. 2, pp. 219–225, Mar. 2010.

[34] E. J. Collins, J. M. McNamara, and D. M. Ramsey, "Learning rules for optimal selection in a varying environment: mate choice revisited," *Behav. Ecol.*, vol. 17, no. 5, pp. 799–809, Sep. 2006.

[35] B. Luttbeg and R. R. Warner, "Reproductive decision-making by female peacock wrasses: flexible versus fixed behavioral rules in variable environments," *Behav. Ecol.*, vol. 10, no. 6, pp. 666–674, Nov. 1999.

[36] K. P. Körding and D. M. Wolpert, "Bayesian integration in sensorimotor learning," *Nature*, vol. 427, no. 6971, pp. 244–247, Jan. 2004.

[37] M. Berniker and K. Kording, "Bayesian approaches to sensory integration for motor control," *Wiley Interdiscip. Rev. Cogn. Sci.*, vol. 2, no. 4, pp. 419–428, Jul. 2011.

[38] Y. Weiss, E. P. Simoncelli, and E. H. Adelson, "Motion illusions as optimal percepts," *Nat. Neurosci.*, vol. 5, no. 6, pp. 598–604, Jun. 2002.

[39] T. J. Valone, "Are Animals Capable of Bayesian Updating ? An Empirical Review," *Oikos*, vol. 112, no. 2, pp. 252–259, Feb. 2006.

[40] J. C. Alonso, J. A. Alonso, L. M. Baustita, and R. Munoz-Pulido, "Patch use in cranes: a field test of optimal foraging predictions," *Anim. Behav.*, vol. 49, no. 5, pp. 1367–1379, May 1995.

[41] H. H. Marshall, A. J. Carter, A. Ashford, J. M. Rowcliffe, and G. Cowlishaw, "How do foragers decide when to leave a patch? A test of alternative models under natural and experimental conditions," *J. Anim. Ecol.*, vol. 82, no. 4, pp. 894–902, Jul. 2013.

[42] J. C. Ramírez and J. A. R. Marshall, "Can natural selection encode Bayesian priors?," *J. Theor. Biol.*, vol. 426, pp. 57–66, Aug. 2017.

[43] J. S. Nelson, T. C. Grande, and M. V. H. Wilson, *Fishes of the World*, 5th ed., vol. 2. Wiley, 2016.

[44] R. O. Salva, V. A. Sovrano, and G. Vallortigara, "What can fish brains tell us about visual perception?," *Front. Neural Circuits*, vol. 8, p. 119, 2014.

[45] V. A. Braithwaite, "Cognitive ability in Fish," in *Fish Physiology, volume 24*, vol. 24, no. C, Sloman Katherine A., R. W. Wilson, and S. Balshine, Eds. Academic Press, 2005, pp. 1–37.

[46] C. Brown, "Fish intelligence, sentience and ethics," *Anim. Cogn.*, vol. 18, no. 1, pp. 1–17, Jan. 2015.

[47] C. L. McCord and M. W. Westneat, "Evolutionary patterns of shape and functional diversification in the skull and jaw musculature of triggerfishes (Teleostei: Balistidae)," *J. Morphol.*, vol. 277, no. 6, pp. 737–752, Jun. 2016.

[48] T. Kuwamura, "Habitat segregation, coexistence or interspecific territoriality between two triggerfishes, Rhinecanthus aculeatus and Sufflamen chrysopterus, with notes on distribution of other balistids at Sesoko Island, Okinawa," *Galaxea*, vol. 10, pp. 65–78, 1991.

[49] T.-C. Chen, R. F. G. Ormond, and H.-K. Mok, "Feeding and territorial behaviour in juveniles of three co-existing triggerfishes," *J. Fish Biol.*, vol. 59, no. 3, pp. 524–532, Sep. 2001.

[50] C. Champ, G. M. Wallis, M. Vorobyev, U. E. Siebeck, and N. J. Marshall, "Visual acuity in a species of coral reef fish: Rhinecanthus aculeatus.," *Brain. Behav. Evol.*, vol. 83, no. 1, pp. 31–42, 2014.

[51] K. L. Cheney, C. Newport, E. C. McClure, and N. J. Marshall, "Colour vision and response bias in a coral reef fish.," *J. Exp. Biol.*, vol. 216, no. Pt 15, pp. 2967–73, Aug. 2013.

[52] C. Newport *et al.*, "Fish use colour to learn compound visual signals," *Anim. Behav.*, vol. 125, pp. 93–100, Mar. 2017.

[53] U. E. Siebeck, L. Litherland, and G. M. Wallis, "Shape learning and discrimination in reef fish," *J. Exp. Biol.*, vol. 212, no. 13, pp. 2113–2119, Jul. 2009.

[54] R. Froese and D. Pauly, "Fishbase," 2000. [Online]. Available: https://www.fishbase.de. [Accessed: 09-Oct-2018].

[55] W. W. L. Cheung, T. J. Pitcher, and D. Pauly, "A fuzzy logic expert system to estimate intrinsic extinction vulnerabilities of marine fishes to fishing," *Biol. Conserv.*, vol. 124, no. 1, pp. 97–111, Jul. 2005.

[56] L. W. Gellermann, "Chance Orders of Alternating Stimuli in Visual Discrimination Experiments," *Pedagog. Semin. J. Genet. Psychol.*, vol. 42, no. 1, pp. 206–208, Mar. 1933.

[57] B. J. Fellows, "Chance Stimulus Sequences for Discrimination Tasks," *Psychol. Bull.*, vol. 67, no. 2, pp. 87–92, 1967.

[58] C. J. Gerard, H. A. Mackay, B. Thompson, and W. J. McIlvane, "Rapid generation of balanced trial distributions for discrimination learning procedures: a technical note.," *J. Exp. Anal. Behav.*, vol. 101, no. 1, pp. 171–8, Jan. 2014.

[59] R. A. Rescorla and A. R. Wagner, "A theory of classical conditioning: Variations in the effectiveness of reinforcement and nonreinforcement.," *Class. Cond. II Curr. Res. Theory*, vol. 21, no. 6, pp. 64–99, 1972.

[60] L. J. Kamin, "Predictability, Surprise, Attention and Conditioning," in *Punishment and Aversive Behavior*, B. A. Campbell and R. M. Church, Eds. New York: Appleton-Century-Crofts, 1969, pp. 279–296.

[61] S. J. Shettleworth, *Cognition, evolution, and behavior*. Oxford University Press, 2010.

[62] D. Bates, M. Mächler, B. M. Bolker, and S. C. Walker, "Fitting linear mixed-effects models using lme4," *J. Stat. Softw.*, vol. 67, no. 1, Oct. 2015.

[63] A. Lange and R. Dukas, "Bayesian approximations and extensions: Optimal decisions for small brains and possibly big ones too," *J. Theor. Biol.*, vol. 259, no. 3, pp. 503–516, Aug. 2009.

[64] L. D. Devenport and J. A. Devenport, "Time-dependent averaging of foraging information in least chipmunks and golden-mantled ground squirrels," *Anim. Behav.*, vol. 47, no. 4, pp. 787–802, 1994.

[65] S. L. Lima, "Downy Woodpecker Foraging Behavior: Efficient Sampling in Simple Stochastic Environments," *Ecology*, vol. 65, no. 1, pp. 166–174, Feb. 1984.

[66] T. J. Valone and L.-A. Giraldeau, "Patch estimation by group foragers: what information is used?," *Anim. Behav.*, vol. 45, no. 4, pp. 721–728, Apr. 1993.

[67] P. Nonacs and J. L. Soriano, "Patch sampling behaviour and future foraging expectations in Argentine ants, Linepithema humile," *Anim. Behav.*, vol. 55, no. 3, pp. 519–527, Mar. 1998.

[68] A. J. Yu and J. D. Cohen, "Sequential effects: Superstition or rational behavior?," *Adv. Neural Inf. Process. Syst.*, vol. 21, pp. 1873–1880, 2009.

[69] F. Dubois, L.-A. Giraldeau, and D. Réale, "Frequency-dependent payoffs and sequential decision-making favour consistent tactic use.," *Proceedings. Biol. Sci.*, vol. 279, no. 1735, pp. 1977–85, May 2012.

[70] B. F. Skinner, "'Superstition' in the pigeon," *Journal of Experimental Psychology*, vol. 38, no. 2. American Psychological Association, US, pp. 168–172, 1948.

[71] S. Dangelmayer, J. Benda, and J. Grewe, "Weakly electric fish learn both visual and electrosensory cues in a multisensory object discrimination task," *J. Physiol.*, vol. 110, no. 3, pp. 182–189, Oct. 2016.

# CHAPTER 3

## Response biases in *Rhinecanthus aculeatus*

# RESPONSE BIASES IN *RHINECANTHUS ACULEATUS*

## Abstract

This chapter focuses on the emergence of behavioural strategies in triggerfish that were not directly or deliberately reinforced and attempts to place these within a larger decision-making framework.

I previously demonstrated that triggerfish integrate both prior and visual information to form weighted posterior estimates. Subjects behaved as if they integrated positional information across many trials to achieve this. However, it is possible that simpler strategies could produce similar behaviour. In this chapter, I examined three possible deviations from behavioural integration – response biases, sequence effects, and anomalies in learning curves. The experimental population exhibited a small positional response bias that was consistent across subjects but not correlated with reinforcement, suggesting either early-life effects or population-wide lateralization. I found that while subjects did exhibit sequence effects in certain contexts, this could not account for updating performance. Furthermore, counter-intuitively sequence effects were negatively correlated with both prior reinforcement and prior stimulus quality, counter to theoretical predictions. The context-specific use of sequence effects suggests that subjects employed them as a backup strategy when other information was poor. I examined differential task acquisition rates and determined that slow learning was likely partially due to task interference, but that inadvertent counterconditioning may have contributed to this. I finally attempted to place these observations within a larger decision-making framework with reference to the ecological context in which animal behaviour occurs.

## Background

In the previous chapter, I examined how triggerfish weigh multiple sources of information in a visual discrimination task. Triggerfish were conditioned to a 4AFC task with a statistical bias in the distribution of stimuli, and then presented with trials featuring reduced signal quality. I found that when the quality of the visual cue declined, subjects increasingly favoured the previously biased position. This response was dynamic; with the degree of positional bias being negatively correlated with visual reliability. This demonstrated that subjects could consider both the current visual information and the past statistical distribution in the position of the stimulus and were able to dynamically weight those two sources of information when making decisions.

However, the results raised several follow-up questions which I set out to explore here. These focus on anomalous behaviour exhibited by subjects during training (response bias and different learning rates between experiments) and the possibility that different behavioural strategies could have contributed to subjects' performance. I attempt to integrate these into a larger framework of decision making and learning, informed by the Rescorla-Wagner model of associative learning [1].

My original analysis did not examine the timeframe over which subjects integrated prior information. The degree of positional bias was analysed in aggregate over the entirety of testing following a fixed period of exposure to the conditioned bias. Given that the degree of positional bias remained constant throughout both training and testing, a subject could viably exhibit very similar bias response curves whether integrating information across the entire experiment, or only the few dozen trials. Particularly, it is important to distinguish true updating from low-order heuristics that reproduce similar results. For example, a subject that simply repeats its previous choice when current sample information is poor could conceivably exhibit very similar behaviour to a subject that weights current sample

information with a statistical distribution built up over many trials. Such a strategy would be less cognitively demanding and not require ongoing retention and integration of information. Such heuristics produce phenomena known as sequence effects, which can be detected and measured within trial sequences. To test for the use of low-order strategies I re-analysed the data set at the level of individual trial-to-trial transitions and observed the frequency of sequence effects in the data.

While subjects across experiments exhibited very similar behaviour during testing, initial acquisition rates differed radically across experiments, implying that variation in the structure of conditioning had large effects on subjects' capacity to attend to and learn different stimuli. Subjects trained to a concurrent schedule of alternating position-discrimination and colour-discrimination training (i.e., experiment 3) were very slow to learn the position task, and only a subset of subjects reached the learning criterion. I hypothesized that this was partially a result of interference between the two tasks inhibiting associations forming between the reinforcer and the positional stimulus. To test this possibility, I carried out a replicate of this experiment in which subjects were conditioned to complete only the positional task and compared learning curves across the two experiments.

Over the course of experiments during both training and testing, I noted two trends in the distribution of choices that subjects made. Firstly, there appeared to be a weak population-wide directional preference downwards and rightwards (i.e., to the lower and right-ward positions). This pattern was consistent across multiple experiments, experimenters, tank positions, rooms, and subjects. Secondly, during the testing phase several subjects developed large positional biases counter to their conditioning. Both these observations potentially affect interpretation of the results. Consequently, I set out to test whether this overall bias was statistically significant and examine both whether the schedule history

impacted the patterns, and how this affected subject's reinforcement to determine if they affect my interpretation of the results.

In summary, here I set out to examine questions of sequence effects, response biases and task acquisition. These three primary research questions are set out sequentially in this chapter, and the outcomes are brought together and integrated into a framework in the discussion.

# 1.0 Sequence effects

## 1.1 Introduction

Research on the role of prior information use in decision-making has largely focused on broad statistical trends – prior distributions built up over dozens to thousands of trials over extended periods of time. This is true across disciplines, from foraging theory [2]–[5] (reviewed by [6]) to reproductive decision-making [7], [8], human sensorimotor control [9] and decision-making [10]–[12].

However, this approach assumes that subjects integrate information on similar timescales. If the statistical contingencies of the task can be captured in a short sequence of trials, then a subject could equally produce similar performance by only integrating information over that sequence. To illustrate the problem in the context of a discrimination task, consider the example given by Cho *et al.* [13]. A subject is presented with a 2AFC task with the sequence 11121_, in which "1" and "2" stand for alternative choices. Let us compare two imaginary experimental subjects, one of which considers the entire training history, and another which considers only the immediately preceding trial. In the final trial, both would have a similar predisposition to choose "1" over "2".

Recent work in humans has demonstrated that in the context of 2AFC tasks, humans appear to "assume" that statistical trends in the data persist for around 4 trials [13], [14].

Comparably, in the context of a matching task where the statistical contingencies of the task periodically shift, rhesus macaques (*Macacca mulatta*) appear to discount previous information at an approximately linear rate, thereby rapidly adjusting to new statistical contingencies [15]. In some cases, animals are known to adjust their sampling (and hence the rate at which they update their priors) according to environmental variability, e.g. [16]; but may not be able to adjust optimally when presented with novel contingencies under experimental conditions [5].

Sequence effects – being described as early as 1953 by Garner [17] – have been variously defined in different fields of research, e.g. [14], [18], [19]; but are typically framed as differences in subject behaviour depending on the order in which experimental trials are presented. For example, when estimating the position of an object, subjects tend to bias their estimates towards the first-presented position [20]. In the context of *n*-AFC tasks, sequence effects generally refer to biases in response towards perceived local patterns in trial order. Subjects tend to respond both more rapidly and more accurately when a trial presentation confirms a perceived local pattern, compared to when the pattern is violated [14], [21]–[23]. This effect intensifies as the sequence grows in length. The patterns that sequence effects refer to are low-order relationships between trials. In the context of 2AFC tasks – where the literature on sequence effects principally developed – the two common patterns are repeats (e.g., LLLLLL) and alternations (e.g., LRLRLR). When translating to n-dimensional tasks where n>2, alternations are less common, but may appear as “A” or “not-A” – however, these are difficult to test for in a 4AFC task. My training schedule specifically excluded prolonged trains of repeats and alternations, and even single train repeats are infrequent. This is a deliberate design feature intended to minimise the formation of such associations [24]. However, first-order effects are still possible in this system.

In the fields of comparative psychology and psychophysics (among others), the emergence of sequence effects and other phenomena such as lateralization – often dubbed "superstitious" behaviours [24] – are generally regarded as a methodological problem to be solved, rather than a behavioural or cognitive phenomenon in their own right [18], [19], [25]–[27]. Furthermore, they are typically conceptualised as a consequence solely of poor trial structure and interaction between experimental conditions; and typically, little consideration is given to the light the phenomena may shed on the animal's own cognition, e.g. [25], [28]. However, the frequency with which sequence effects and other low-level heuristics emerge in both human and non-human experiments suggests that they are adaptive. There is some evidence to support this claim. Building on experimental demonstrations of sequence effects in humans working on 2AFC tasks, Cohen *et al.* [13], [14] constructed a model demonstrating that sequential effects arise from subjects discounting information linearly with decaying coefficients into the past, and that such an approach approximates optimal Bayesian inference, assuming that the task is statistically non-stationary, i.e. it undergoes periodic statistical changes. Behavioural lateralization has also been demonstrated to be an adaptive strategy for mice in a 2AFC water maze task [29] (reviewed in [30]).

Truly random contingencies are rare in nature, and many phenomena (the distribution of prey, etcetera) are either meta-stable through time, or clustered in space such that the detection of small-scale patterns may be highly adaptive in a variety of contexts – even if they are not adaptive in a laboratory context [31]. In this context, it may be more adaptive for an animal to selectively attend only to recent information, and discount older information that may be out of date. If such heuristics operate in parallel with other quantifiable strategies (such as that shown in Chapter 2), this suggests a deeper level of weighting – that of behavioural strategies, rather than sources of information. Given these

considerations, I posited that testing for sequence effects in my data was both necessary to test updating, and to determine whether triggerfish may have operated multiple strategies in parallel. To approach this, I aggregated the three main data sets collected in Chapter 2, with several further experiments, into a single data set. This was then subjected to several analyses examining relationships between sequential trials.

## 1.2 Hypotheses

I hypothesized that subjects might use low-order heuristics to solve the discrimination tasks. While the dimensionality of the experiments (i.e., 4AFC) naturally discouraged most such strategies (e.g., fixed side bias, alternation), position repeats were still plausible. Specifically, if subjects repeated their choices at a rate greater than chance, this could result both in above-chance reinforcement, and the observed bias preference observed in the data. Given this, I hypothesize that subjects would have repeated choices at a rate greater than chance. Further, I expected that this effect would be enhanced under the following conditions:

- When current sample visual information was poor
- When the previous choice was reinforced
- When the previous choice was reinforced with a preferred reward

Regarding the first postulate, as a simple heuristic, sequence effects should be more frequent when more reliable strategies are less effective, such as when stimulus quality for visual discrimination is poor.

Regarding the second and third postulates, if sequence effects are a consequence of a strategy that considers only the immediately preceding trial, then the reinforcement of that trial is also a valid source of information. Extended in this way, sequence effects may be equivalent to a win-stay/lose-shift heuristic. This effect should also be enhanced when the potential benefits are greater, such as when a greater magnitude reinforcer was used.

In an associative RW framing, this would correspond to employing an unconditioned stimulus (US) with a larger value for β and/or λ, wherein:

$$\Delta V_A = \alpha\beta(\lambda - V_T)$$

Where ΔVA is the change in associative strength of a single stimulus A, α is the salience of the CS, β is the strength of the reinforcer, λ is the asymptotic maximum associative strength the US can support, and $V_T$ is the sum of the associative strengths of all stimuli present on the trial.

## 1.3 Null models

To determine if subjects repeated their choices at a rate above chance, it was necessary to have a robust model of chance behaviour. I considered three alternative null models, reflecting different hypothesis subjects might have about the statistical features of the task.

- *Structure matching:* match the rate at which positions repeated during training
- *Distribution matching:* match the overall distribution of positions during training
- *No matching:* repeats are a function of the overall distribution of choices

### Structure matching

Subjects may have attended to the actual frequency of position repeats during training and adjusted their own repeat rate to match. The overall frequency of repeats was not a specifically controlled factor in schedule design – rather, the length of individual repeat trains was curtailed in a Gellerman-like schedule. Consequently, rate of position repeats in the schedule was varied by training phase and experiment (table 1). This inconsistency makes creating a null model uncertain, and highly contingent on how quickly subjects responded to changes in reinforcement. Consequently, it was discounted as non-predictive.

| Treatment | Initial Training | Bias Training | Testing |
|---|---|---|---|
| Experiments 1, 2; control | 0.0% | 0.1% | 6.2% |
| Experiments 1, 2; bias | 1.3% | 38.3% | 12.7% |
| Experiment 3; bias | 46.3% | | 14.5% |

*Table 1:* Actual mean rate of CS+ position repeats as a percentage of total trial-to-trial transitions broken down by experiment, treatment, and testing phase.

### Distribution matching

Subjects could have attended to the overall distribution of stimulus presentations, and thereby produced a position repeat rate determined by treatment condition (i.e., control or bias treatments) (table 2). Overall, this would correspond to a repeat rate of 0.413 (biased treatment) or 0.25 (control treatment). This null model does assume that subjects learned the positional cue and adjusted their responses accordingly. To validate the computed values, I also employed a simulation approach (described in the methods section).

### No matching

Finally, subjects may not have attended to either the structure or overall distribution of trials. The most parsimonious model of such behaviour is a parameter-free model, informed only by subjects' actual choices. In this case, the rate of repeats should be entirely determined by the overall distribution of their choices. An example of this for a single subject is given in table 3, with an expected overall repeat rate of 0.288. This is a highly parsimonious model and makes no assumptions regarding the extent to which subjects have learned the positional bias. The parsimony of this model allows all data, including the pilot experiment, to be analysed. Consequently, this is the model I used for most subsequent analyses.

| Unbiased Subjects | | | Choice 1 | | | |
|---|---|---|---|---|---|---|
| | | | P1 *bias* | P2 | P3 | P4 |
| | | | 0.60 | 0.13 | 0.13 | 0.13 |
| Choice 2 | P1 *bias* | 0.60 | 36.00 | 8.00 | 8.00 | 8.00 |
| | P2 | 0.13 | 8.00 | 1.78 | 1.78 | 1.78 |
| | P3 | 0.13 | 8.00 | 1.78 | 1.78 | 1.78 |
| | P4 | 0.13 | 8.00 | 1.78 | 1.78 | 1.78 |

| Biased Subjects | | | Choice 1 | | | |
|---|---|---|---|---|---|---|
| | | | P1 | P2 | P3 | P4 |
| | | | 0.25 | 0.25 | 0.25 | 0.25 |
| Choice 2 | P1 | 0.3 | 6.25 | 6.25 | 6.25 | 6.25 |
| | P2 | 0.3 | 6.25 | 6.25 | 6.25 | 6.25 |
| | P3 | 0.3 | 6.25 | 6.25 | 6.25 | 6.25 |
| | P4 | 0.3 | 6.25 | 6.25 | 6.25 | 6.25 |

*Table 2 (above):* Markov transition table of a distribution-matching model for both the bias treatment (a) and control treatment (b). The model produces an overall rate of position repeats of 41.3% and 25.0% respectively (from the sum of diagonal). P1-4 correspond to nominal stimulus positions, equivalent or UR, UL, etc. Colours included as a visual aid, correlate to percentage scalar (RED-low, GREEN-high).

| Exp1 1.Hester UL | | | Choice 1 | | | |
|---|---|---|---|---|---|---|
| | | | UL | UR | LL | LR |
| | | | 0.40 | 0.24 | 0.13 | 0.23 |
| Choice 2 | UL | 0.40 | 16.00 | 9.50 | 5.17 | 9.33 |
| | UR | 0.24 | 9.50 | 5.64 | 3.07 | 5.54 |
| | LL | 0.13 | 5.17 | 3.07 | 1.67 | 3.01 |
| | LR | 0.23 | 9.33 | 5.54 | 3.01 | 5.44 |

*Table 3 (left):* Example Markov transition table from a single subject (01_Hester), calculated from the subject's overall distribution of choices across stimulus positions. UL = upper-left, UR = upper-right, LL = lower-left, LR = lower-right. Subject has an overall model position repeat rate of 0.288 (sum of diagonal). Colours included as a visual aid, correlate to percentage scalar (RED-low, GREEN-high).

## 1.4 The inter-trial interval

The inter-trial interval is generally thought to be of great importance in determining the extent of sequence effects, with longer intervals decreasing their magnitude; e.g. [32], presumably due to trace memory decay. In the experiments discussed, the inter-trial interval was not controlled, but was determined by the time it took the experimenter to record a trial outcome and re-set the apparatus. I estimate the interval as between 30 and

180 seconds. The interval between sessions was much larger and considerably more variable, between 2 and 72 hours. As such, sessions are treated as independent. This allowed n-1 inter-trial transition comparisons per session where n was the number of trials per session.

## 1.5 Methods

### Simulation

The null expected position repeat rate is contingent on the treatment the subject is exposed to, i.e., control or bias treatments. The relate selection frequencies expected under each treatment in turn biases the degree of expected position repeats. To determine whether the observed data did deviate from a distribution matching model, a simulation of each treatment was carried out using sample with replacement (fish null v1). The simulations were carried out according to the number of trials within the experiment. There were four choices, as described previously (UL, UR, LL, LR).

```
sesn<-24 #number of sessions
```

For the simulations, the number of sessions was specified, the number of fish as per the experimental setup and the number of repeat simulations to estimate variance.

```
for (i in 1:sesn){
thissession<-sample(options, nrep, prob=(c(0.6,0.1333,0.1333,0.1333)), replace = TRUE)
#thissession<-sample(options, nrep, replace = TRUE)
p <- matrix(nrow = 4, ncol = 4, 0)
for (t in 1:(length(thissession) - 1)) p[thissession[t], thissession[t + 1]] <- p[thissession[t], thissession[t + 1]] + 1
p<-p/(nrep-1)
```

Each simulation consisted of a single loop, with the simulation means being composed into a matrix of position probabilities and transition probabilities between each trial and the next for each subject in each treatment. Finally, the diagonals of each transition matrix (i.e.,

repeated position) versus non-diagonals (i.e., non-repeated positions) were averaged to produce a mean, broken down by treatment and position.

### Markov tables

Testing-phase data was aggregated across the four component experiments outlined in Chapter 2, including all subjects that completed testing. I constructed Markov tables for each subject breaking down the proportion of inter-trial transitions by position. These were matched with null Markov tables for each subject based on individual's actual distribution of choices across stimulus positions. From this the population mean model deviation was calculated, with treatment (bias versus control) as a factor.

## 1.6 Results

### Analytic approach

I created Markov tables of each subject's deviation from a null model of their behaviour. Subjects were compared to both a distribution-matching null model accounting only for their individual distribution of choices, and a no-matching parameter free model. The mean across individual tables was then calculated for the treatment group, and the mean model deviation of each position transition calculated. Subjects in the biased group exhibited a very poor overall match to the distribution-matching model (table 4). This was accounted for by a reduced bias compared to the training distribution and is expected given that subjects primarily used visual discrimination. Consequently, only the non-matching null was employed for the bias treatment (table 5b).

Contrary to my hypothesis, the biased treatment group made repeat position choices significantly less frequently in the no-matching model compared to non-repeat choices (Mann-Whitney U test, $U = 0$, $n_1=4$, $n_2=12$, $P=0.001$, two-tailed), suggesting they actually avoided positional repeats. The relationship in the control group was more equivocal, with no significant difference from the no-matching null ($U = 11$, $n_1 = 4$, $n_2 = 12$, $P= 0.133$, two-

tailed) but a very significant deviation from the distribution-matching null, also repeating choices less frequently than expected (U = 3.5, $n_1$ = 4, $n_2$ = 12, P = 0.007, two-tailed).

| UL n=4 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 19.4 | 11.1 | 7.8 | 7.5 |
| UR | 13.2 | 4.1 | 3.2 | 2.4 |
| LL | 8.6 | 2.8 | 1.5 | 2.3 |
| LR | 7.3 | 3.8 | 2.8 | 2.2 |

| UR n=7 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 2.4 | 9.2 | 3.1 | 3.2 |
| UR | 6.1 | 19.1 | 8.8 | 9.6 |
| LL | 4.2 | 6.9 | 2.7 | 3.4 |
| LR | 3.0 | 10.4 | 2.9 | 5.2 |

| LL n=7 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 1.8 | 2.2 | 7.2 | 3.6 |
| UR | 2.2 | 1.8 | 7.8 | 3.9 |
| LL | 7.2 | 7.3 | 24.0 | 9.1 |
| LR | 2.7 | 3.7 | 11.0 | 4.7 |

| LR n=4 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 1.9 | 1.5 | 2.8 | 8.2 |
| UR | 2.7 | 2.3 | 2.1 | 11.0 |
| LL | 2.1 | 2.4 | 1.6 | 9.4 |
| LR | 6.8 | 10.8 | 8.3 | 26.2 |

| UL n=4 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 36.0 | 8.0 | 8.0 | 8.0 |
| UR | 8.0 | 1.8 | 1.8 | 1.8 |
| LL | 8.0 | 1.8 | 1.8 | 1.8 |
| LR | 8.0 | 1.8 | 1.8 | 1.8 |

| UR n=7 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 1.8 | 8.0 | 1.8 | 1.8 |
| UR | 8.0 | 36.0 | 8.0 | 8.0 |
| LL | 1.8 | 8.0 | 1.8 | 1.8 |
| LR | 1.8 | 8.0 | 1.8 | 1.8 |

| LL n=7 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 1.8 | 1.8 | 8.0 | 1.8 |
| UR | 1.8 | 1.8 | 8.0 | 1.8 |
| LL | 8.0 | 8.0 | 36.0 | 8.0 |
| LR | 1.8 | 1.8 | 8.0 | 1.8 |

| LR n=4 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | 1.8 | 1.8 | 1.8 | 8.0 |
| UR | 1.8 | 1.8 | 1.8 | 8.0 |
| LL | 1.8 | 1.8 | 1.8 | 8.0 |
| LR | 8.0 | 8.0 | 8.0 | 36.0 |

*Table 4:* Markov transition tables of mean actual (top row) and distribution-matching model (bottom row) repeat frequencies in the treatment condition, broken down by position transition. From left to right, (a) upper-left bias, (b) upper-right bias, (c) lower-left bias, (d) lower-right bias. While the overall distribution of repeated choices matches expectations, subjects consistently repeat bias-position choices at a lower rate than predicted.

| Unbiased subjects, n=10 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | -1.75 | 0.69 | 0.44 | -0.75 |
| UR | 0.13 | -1.36 | 0.05 | 1.47 |
| LL | 0.51 | -0.87 | -1.47 | 0.75 |
| LR | 0.15 | 1.29 | 0.59 | 0.13 |

| Unbiased subjects, n=10 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | -2.11 | 1.43 | -0.59 | 0.73 |
| UR | -0.39 | -0.70 | -1.05 | 2.68 |
| LL | 0.13 | 0.01 | -2.35 | 1.95 |
| LR | -0.96 | 1.30 | -1.05 | 0.98 |

| Biased subjects, n=22 | Choice 1 | | | |
|---|---|---|---|---|
| Choice 2 | UL | UR | LL | LR |
| UL | -1.11 | 0.64 | 0.26 | 0.53 |
| UR | 0.23 | -1.00 | 0.63 | 0.44 |
| LL | 0.61 | -0.14 | -1.03 | -0.04 |
| LR | -0.15 | 0.84 | 0.30 | -0.99 |

*Table 5:* Markov transition tables, showing mean deviation from different null models by position transition. From left to right, (a) unbiased treatment, distribution-matching null model, (b) unbiased treatment, no-matching null model, (c) biased treatment, no-matching null model. The diagonal of choice transitions (i.e., repeat choices, UL-UL, etc.) exhibit consistently negative model deviation, indicating that subjects repeated positions at a lower rate than predicted by either model.

## Simulation approach

To complement the analytic approach described in the previous section, the behaviour of subjects was simulated according to a distribution-matching model. In this analysis, both bias and control groups repeated their choices at a rate less than simulated. However, this manifested differently in each treatment. In the control group, effect size was small but consistent across positions. In the bias group, subjects repeated choices in the bias position less than expected but repeats in other positions did not differ from simulation expectations (figure 1). This broadly aligns with the analytic results, and further suggests that biased subjects specifically avoided repeating their conditioned position. However, due to the poor match between biased subject data and the distribution-matching model, these results must be considered with caution.

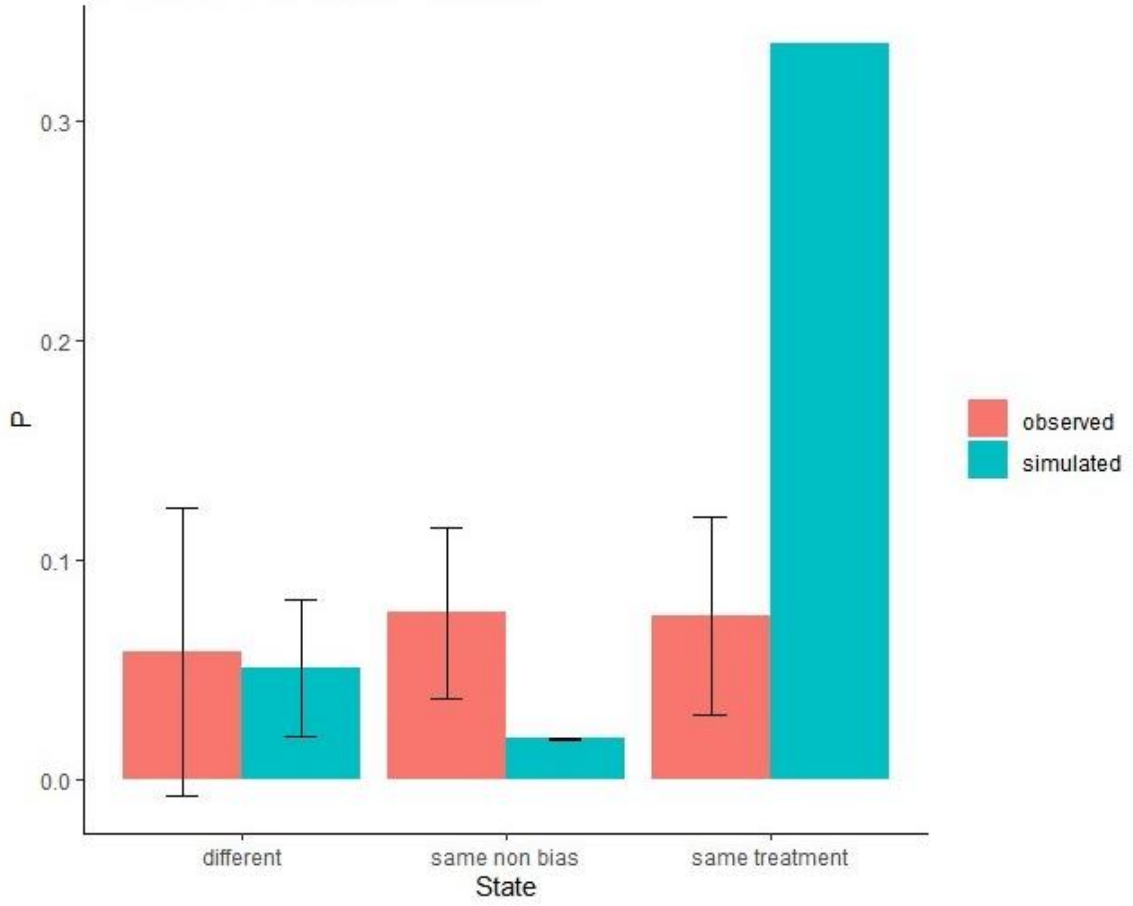



*Figure 1:* mean frequency of different transition states (repeats and non-repeats), broken down by treatment and position type. Different = non-repeats, same non bias = repeats in the control group, same treatment = repeats in the bias group. Error bars derived from standard deviation.

## Role of prior reinforcement

The mean position repeat rate exhibited a complex relationship with prior reinforcement (figure 2), with different treatments exhibiting different responses. Excluding trials with a preferred (lancefish) reinforcer (due to a subset, see figure 2), the data did not deviate from assumption of normality (Shapiro-Wilks, $p = 0.20$) but did violate homogeneity of variance (Levene test, $p = 0.0000164$). ANOVAs are robust to minor violations of homoscedasticity, but not with unequal sample sizes [33]. Consequently, I employed an aligned rank transform (ART) on a repeated measures ANOVA, with one within subject factor (*Reinforcement* - pellet or nothing) and one between-subjects factor (*Treatment* – control or bias), and subject identity as an error term. Prior reinforcement had a significant effect on the rate of repeated choices with a large effect size ($F(1,30) = 31.870$, $p = 1.25 \times 10^{-5}$), with higher repeat rates following non-reinforcement. Treatment did not have a significant effect alone ($F(1,30) = 2.896$, $p = 0.099$), but did in interaction with prior reinforcement $F(1,30) = 5.235$, $p = 0.029$), with significant differences emerging only in the control group.

In experiment 2, a two-tier reinforcement system was employed. Subjects were always reinforced for correct responses, but correct responses in the bias position were rewarded with a preferred food, a slice of lancefish ("fish" in figure 2). In the control group, the preferred reinforcement was distributed across all positions. Subjects in this experiment exhibited marked response differences to reinforcement with lancefish – biased subjects greatly increased their repeat rate, while control group did not (figure 2). The sample size of this comparison was too small for a meaningful statistical test (2 vs 8). However, the results are tentatively striking, with the mean difference between groups accounted for 59% of the total variance, and no overlap between treatments.

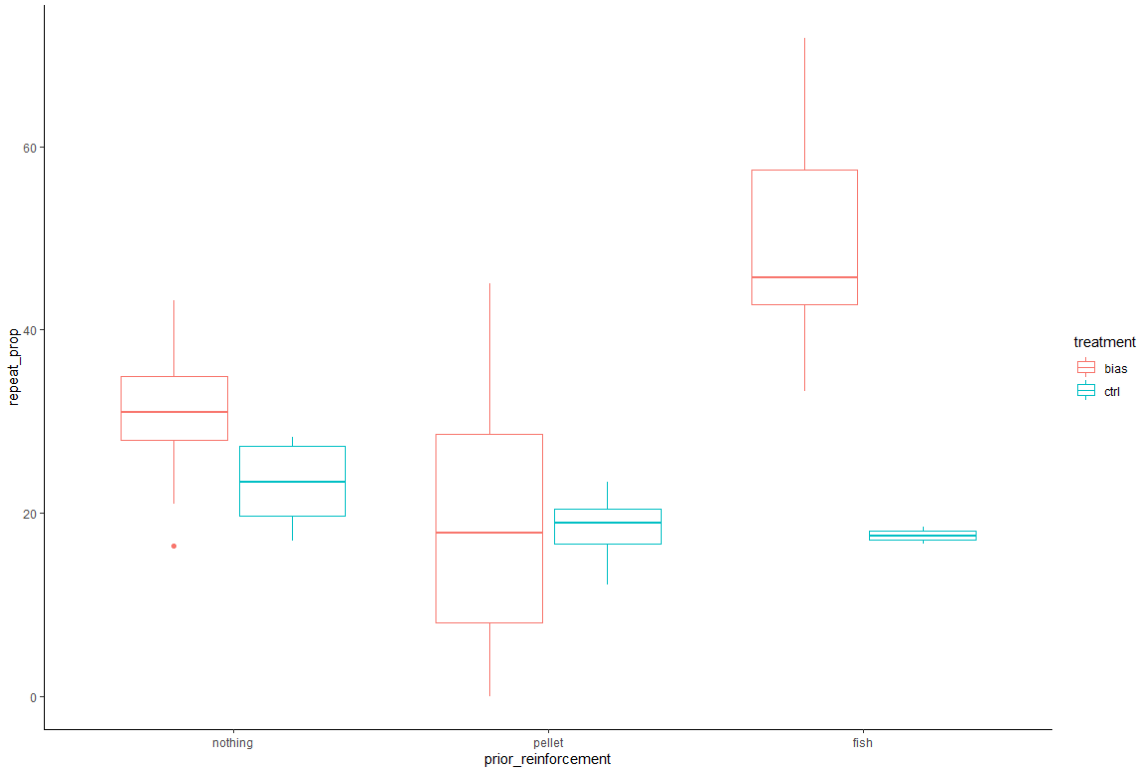



*Figure 2:* Percentage of repeat choices by prior reinforcement. "Fish" (i.e., lancefish reward) includes data from only experiment 2, as only this experiment used a two-tier reward system. Sample size varies by group, as follows: "nothing" (bias n=22, ctrl n=10); "pellet" (bias n=22, ctrl n=10); "nothing" (bias n=10, ctrl n=2)

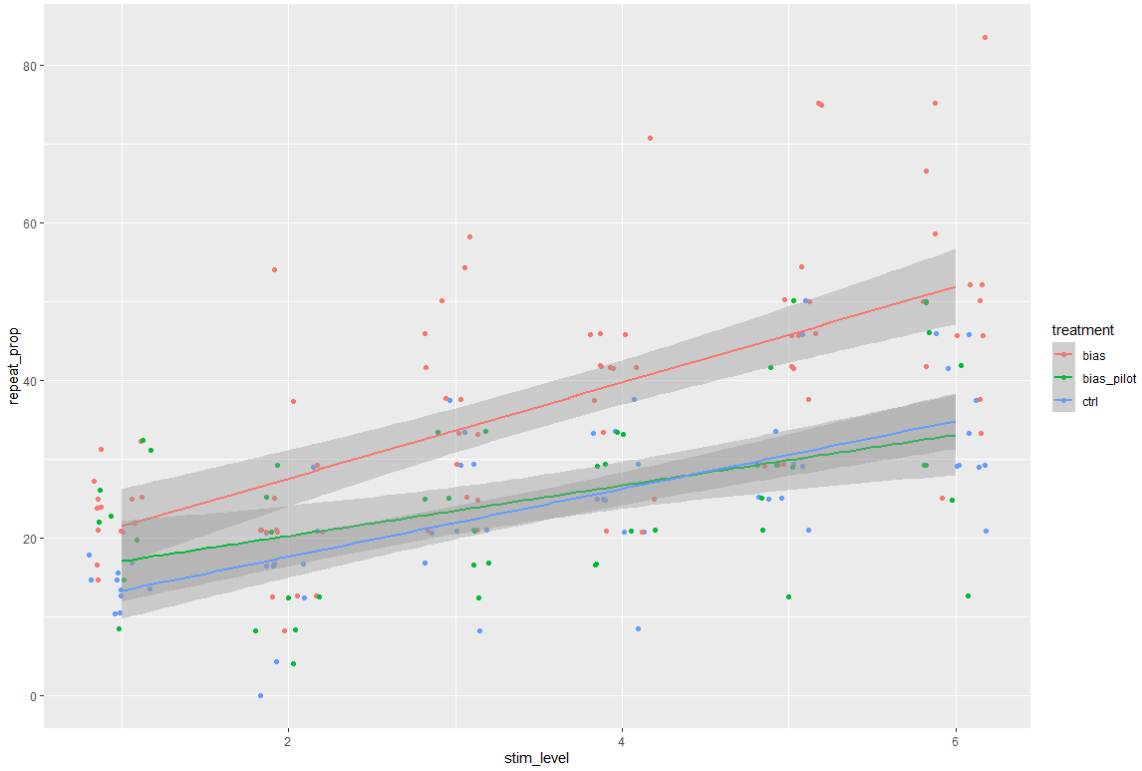



*Figure 3:* Frequency with which each subject repeated positional choice (testing phase only), plotted against trial type ("stim_level"; 1 = R, 2 = T1, etc.)

## Role of stimulus quality

To analyse the relationship between the repeat rate and stimulus quality, the repeat frequencies for each subject were broken down by trial type. These data violated the assumption of normality (Shapiro-Wilk test, w = 0.945, p = $2.73 \times 10^{-6}$). To assess the relationship, a simple linear regression was constructed considering three explanatory variables, having discarded subject identity as non-significant:

$$repeat\ frequency \sim trial\ type + treatment + experiment$$

Due to previously discussed poor learning of the positional cue, the biased treatment in the pilot experiment was considered as a distinct treatment condition (figure 3). This model had good explanatory power ($R^2$ = 0.487). Trial type (β = 4.807, p < $5.43 \times 10^{-16}$) and treatment (ctrl β = -13.889, p = $2.32 \times 10^{-6}$) were significant predictors with trial type having the largest contribution, but experiment was not a significant predictor (largest effect experiment 3 replicate, β = 3.241, p = 0.45). Treatment (i.e., bias versus control) affected the intercept (with the control group having a lower overall repeat rate), but not the overall slope of the relationship, indicating that there was a negative correlation between repeat rate and stimulus quality in both treatments.

# 2.0 Response bias and lateralization

## 2.1 Introduction

Over the course of experiments during both training and testing, I noted two trends in the distribution of choices that subjects made. Firstly, there appeared to be a weak population-wide directional preference downwards and rightwards (i.e., to the lower and right positions). This pattern was consistent across multiple experiments, experimenters, tank positions and laboratories. Secondly, during testing several subjects developed large positional biases counter to their conditioning. Both observations appear to represent significant deviations from basic models of subject behaviour, i.e., the assumption that subjects will optimise their responses to maximise reward frequency. If subjects maintained such biases without cost to their reinforcement rate, this would constitute a flaw in the method. By contrast, if triggerfish maintained bias contrary to reinforcement this would suggest that they were subject to uncontrolled biases at perceptual, decisional, or motor levels that caused them to deviate from theoretically optimal behaviour.

## 2.2 Methods

### Population bias

To formally test the former hypothesis, I carried out a simple post-hoc analysis of subject choices in the testing phase. I pooled position choice data from probe trials across all experiments (excluding the pilot as the results from that experiment were highly inconsistent, and subjects consequently could not be fitted to any simple model). I constructed two model subjects, NULL and TREATMENT, corresponding to the control and test treatments. NULL was assumed to make positional choices randomly, thereby choosing each position with equal frequency. TREATMENT was assumed to prefer one position, otherwise distributing choices equally among the remaining positions. The degree of preference in TREATMENT was derived by averaging the bias position selection frequency across all biased subjects (n=14) (figure 4).

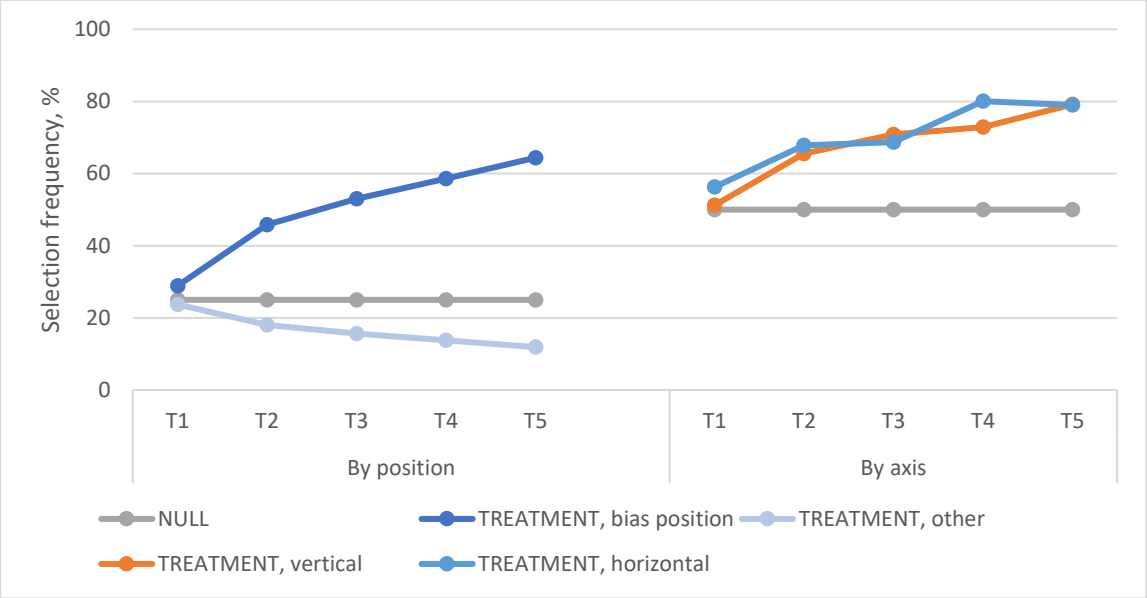



*Figure 4:* Model choice distributions by trial type.

Each subject's deviation from their respective model was calculated, and the mean population values calculated (control n=8; test n=14, all n=22). This calculation was made both for the raw positional choice data, and with choices decomposed into their cartesian coordinates (i.e., up/down, left/right).

### Individual bias

Subject that selected an unconditioned position beyond the calculated confidence interval in two or more trial types was judged to have developed unconditioned biases. This criterion selected four subjects from the pilot (n=3) and experiment 2 (n=1).

## 2.3 Results

### Population bias

Subjects' distribution of choices was subtracted from a model to extract individual deviation. There were significant differences in mean model deviation between stimulus positions (single-factor ANOVA $F(3, 84) = 6.556$, $MSE = 0.546$, $p = 0.00049$). A post-hoc Tukey HSD test found that there were significant differences between LR ($M = 6.514$, $SD = 10.17$) and both UL ($M = -5.612$, $SD = 5.611$) and LL ($M = -2.185$, $SD = 6.698$) positions. Overall, subjects chose LR on 6.5%±1.8SE more trials than expected.

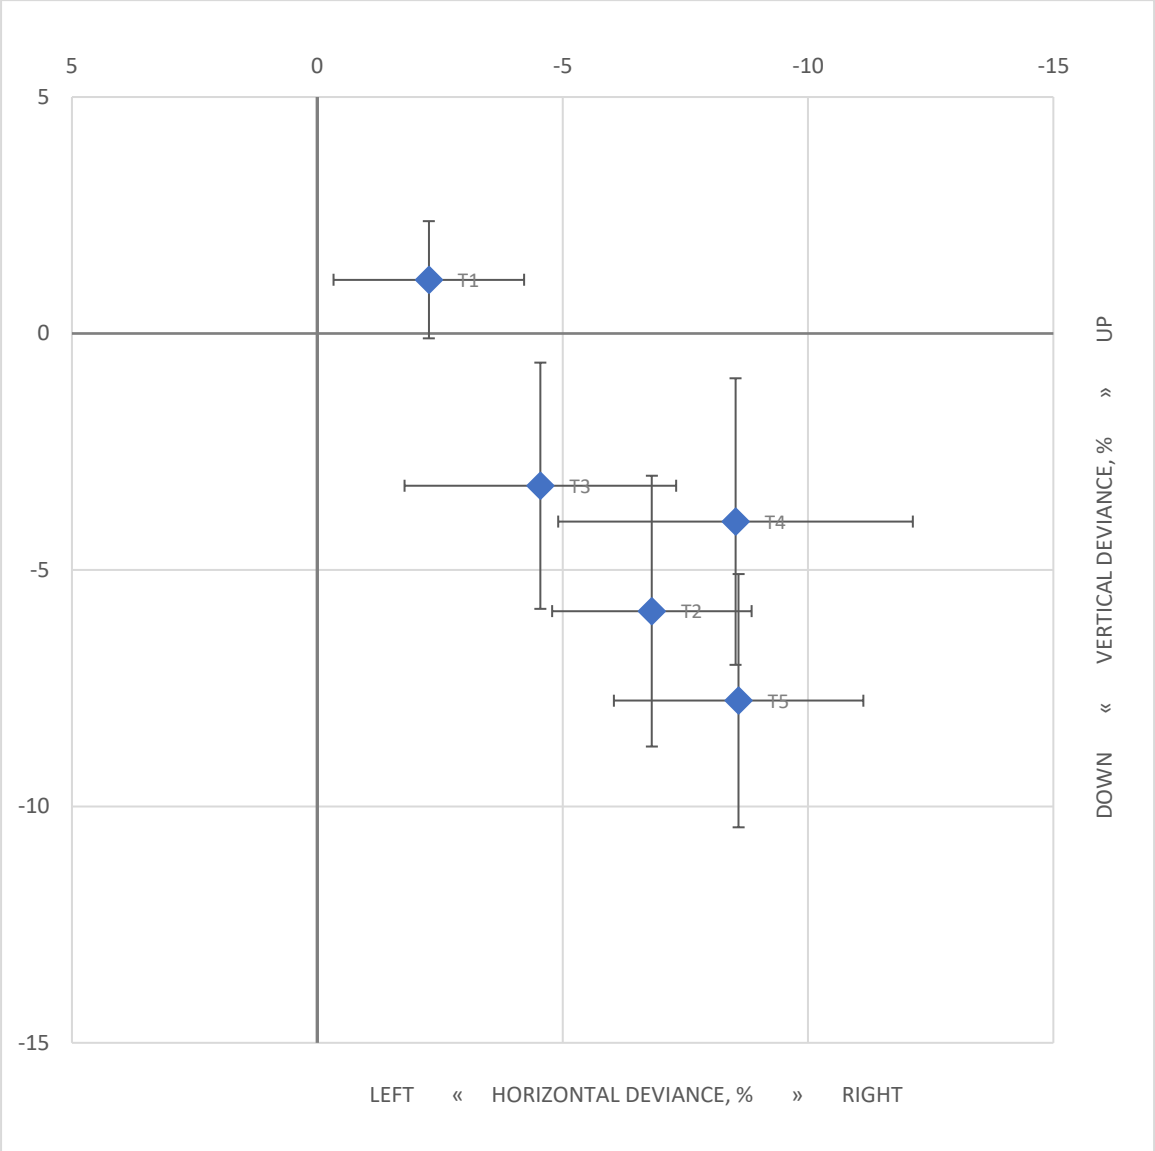



*Figure 5:* Mean deviation from modelled position repeat frequencies by trial type, with choices broken down into their cartesian components. Errors bars from standard error. Data labels correspond to trial types, with contrast decreasing with increasing number.

This pattern was conserved when subject choices were decomposed into their coordinate values. Mean horizontal and vertical model deviation were both negative (i.e., right-biased, and down-biased) and significant (two-tailed one sample t-test; horizontal $t(21) = -2.611$, $p = 0.016$; vertical $t(21) = -2.134$, $p = 0.045$). This overall pattern was correlated with stimulus quality, with the degree of bias increasing as visual information declined (figures 6,7).

This bias was not detected by an initial side-bias test conducted prior to the pilot. This might have been a result of the comparatively small sample size of the assay – only 4 subjects at 40 trials each. To test whether the observed bias was an unintentional artefact of conditioning, I examined the distribution of experimentally naïve subjects' choices during the initial training phase before the introduction of a conditioned bias. During initial training, naïve subjects did exhibit a right-wards lateral bias in the distribution of their errors, but no significant overall vertical bias. These data suggest the population-level bias was not a consequence of the training schedule but was pre-existing.

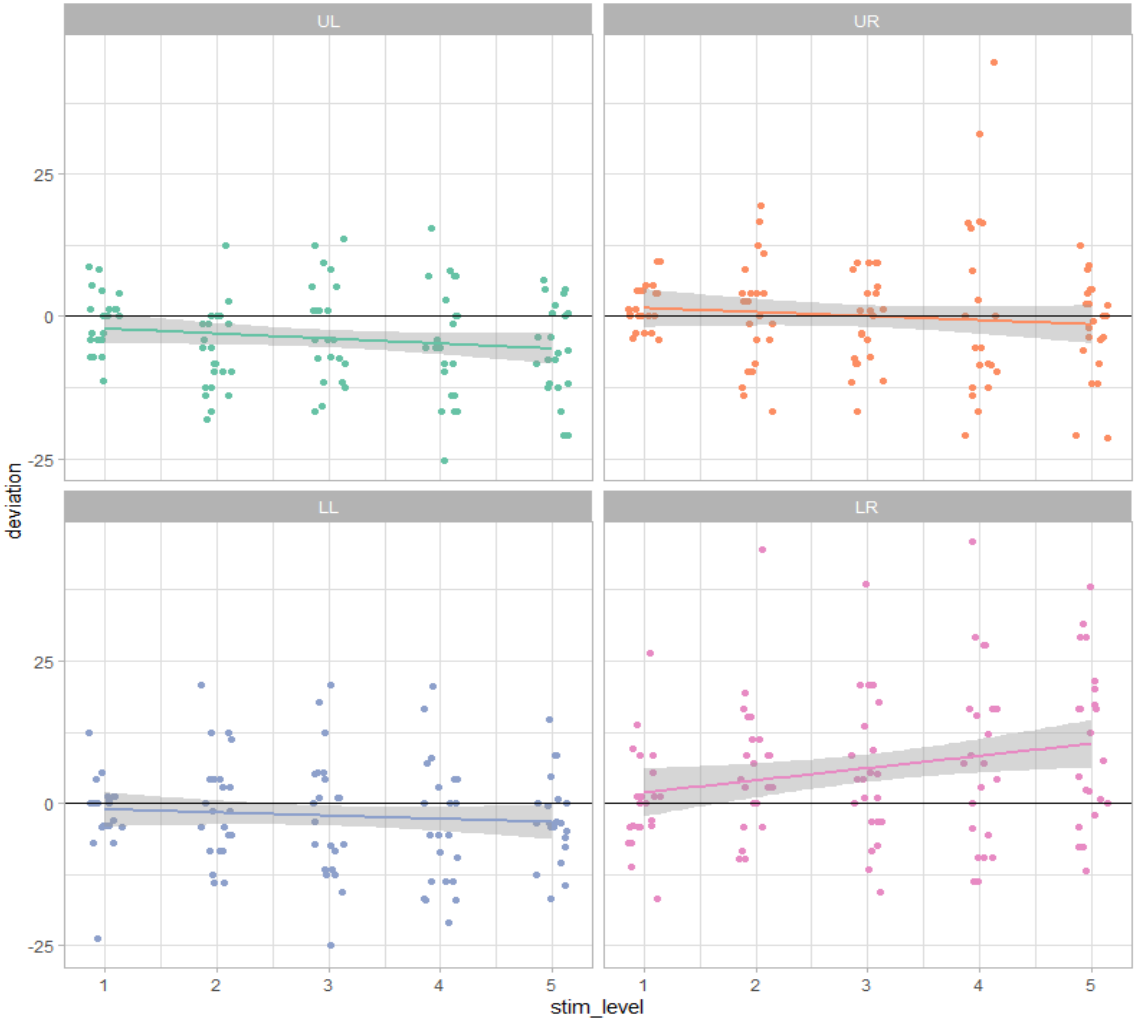



*Figure 6:* Deviation from modelled position repeat frequencies plotted against stimulus intensity. "Stim_level" corresponds to trial type, i.e., 1=T1, 2=T2, etcetera. Regression lines derived from linear model. Model deviation is correlated with stimulus quality across all positions, being negative for UL, UR, LL, and positive for LR, indicating that subjects increasingly allocated more choices than predicted to the LR position as visual quality declined.

## Individual bias

Several subjects exhibited consistent positional biases during the testing phase that were counter to conditioning. This appeared largely in the pilot, but there was one case in experiment 2 (figure 7). In the pilot, subjects with inconsistent biases were no less accurate overall than consistent subjects (One-tailed Wilcoxon rank-sum test. W = 16.5, r=0.433, p=0.1), but were significantly less accurate in R trials (One-tailed Wilcoxon rank-sum test. W = 19, r=0.613, p=0.033), resulting in a lower overall reinforcement ratio. This demonstrates that the unreinforced bias was an inferior strategy for those subjects. These biases were not present during either training phase. Consequently, it is likely that these unconditioned biases developed during testing.

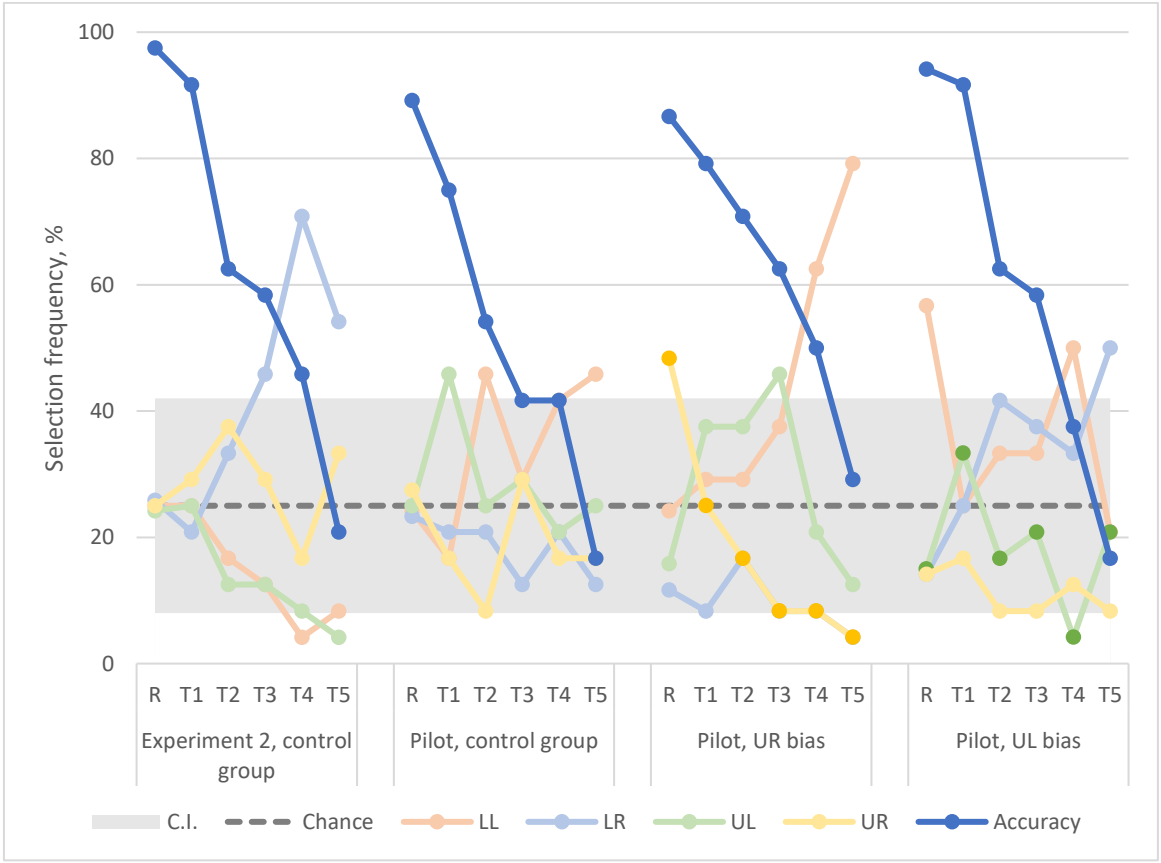



*Figure 7:* (left to right) subjects Scar (control), Rheia (control), Gog (UR bias), and Vetch (UL bias), demonstrating unconditioned biases towards LR, LL, LL, and LR respectively.

# 3.0 Task acquisition

## 3.1 Introduction

Comparing across the experiments described in the previous chapter, I found that different training methods resulted in radically different performances – in experiment 3 particularly, nearly half of subjects (3/8) failed to acquire the task within the allotted 50 sessions, compared to only 2/20 with the original method (experiments 1 and 2). In addition, the mean rates of task acquisition were radically different between approaches, with naïve subjects in experiments 1 and 2 reaching the learning criterion in 4-14 sessions (M=7.3±2.4SD) (figure 8).

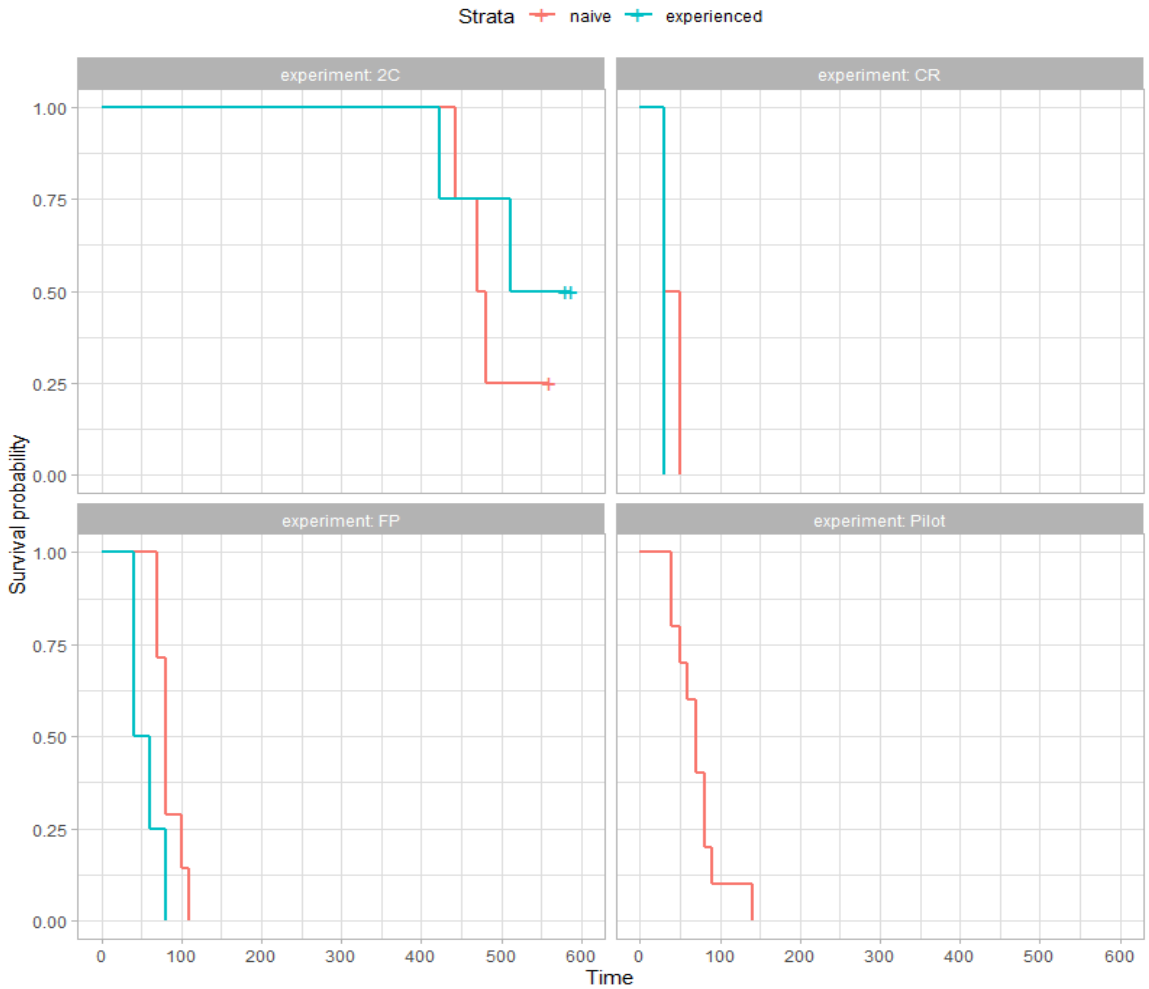



*Figure 8:* Kaplan-Meier survival plot of task acquisition rates by experiment, split by degree of prior experience. 2C = exp. 3, CR = visual reliability assay, FP = exp. 2, Pilot = exp. 1. The x-axis (time) indicates the number of trials taken to reach the learning criterion, while the y-axis indicates the proportion of subjects in the experiment. Crosses indicate where a subject was disenrolled without reaching the learning criterion.

While there were differences in method between these experiments, these were comparatively minor, and all experiments used the same basic schedule structure and stimuli (refer to Chapter 2 for details). This begs the question of how small methodological differences so radically affected subjects' ability to learn different components of the task. I considered that the alternating presentation of two different tasks with similar context and exigencies may have resulted in interference, thereby inhibiting learning. The formation of associations in learning is a process of refinement, as associative value is assigned to those features of the task that are both unique and relevant. During this process, the subject may initially associate reward with irrelevant features of the context. Where two tasks are simultaneously conditioned, features that are reinforced in one task are not reinforced in the other, forcing an extended period of simultaneous associative refinement that may slow learning [1].

## 3.2 Methods

I hypothesized that learning the two similar tasks in tandem had a confounding effect on the formation of US-CS associations. To test this hypothesis, an experiment was carried out to determine how rapidly the positional task might be acquired when reinforced in isolation. Five experienced and 3 naïve fish were used for this experiment, including some (n=3) that were previously used in experiment 3. Training was identical to that described in Chapter 2: Methods: Experiment 3, with the exception that all 8 trials per session were "Position" (P) trials. Subjects were conditioned to discriminate between four identical stimuli (dark grey circles, 90,90,90) based on their position. Subjects were assigned pseudo-randomly to one of the four positions, and were always reinforced for selecting their assigned position, and never for selecting any other position. As previously, training was considered complete once fish had reached a pre-assigned learning criterion: a minimum of 75% correct responses over three training days. To ensure a valid comparison, incorrect

trials were initially iterated until a correct response was elicited (see Chapter 2: Results: Experiment 3). Iteration was applied consistently in sessions 1-10 across the population.

## 3.3 Results

All subjects (n=8) completed training, reaching the learning criterion in 12-17 sessions (figure 9). In comparison to experiment 3, subjects conditioned to a single task reached the learning criterion in significantly fewer trials (Cox proportional hazards model, likelihood ratio = 7.63, n=16, z=2.473, p = 0.01). Neither prior experience nor subject identity had a significant effect on learning rate (Cox proportional hazard; experience z = -0.08, p = 0.938; subject z = -0.62, p = 0.533) (figure 10a,b).

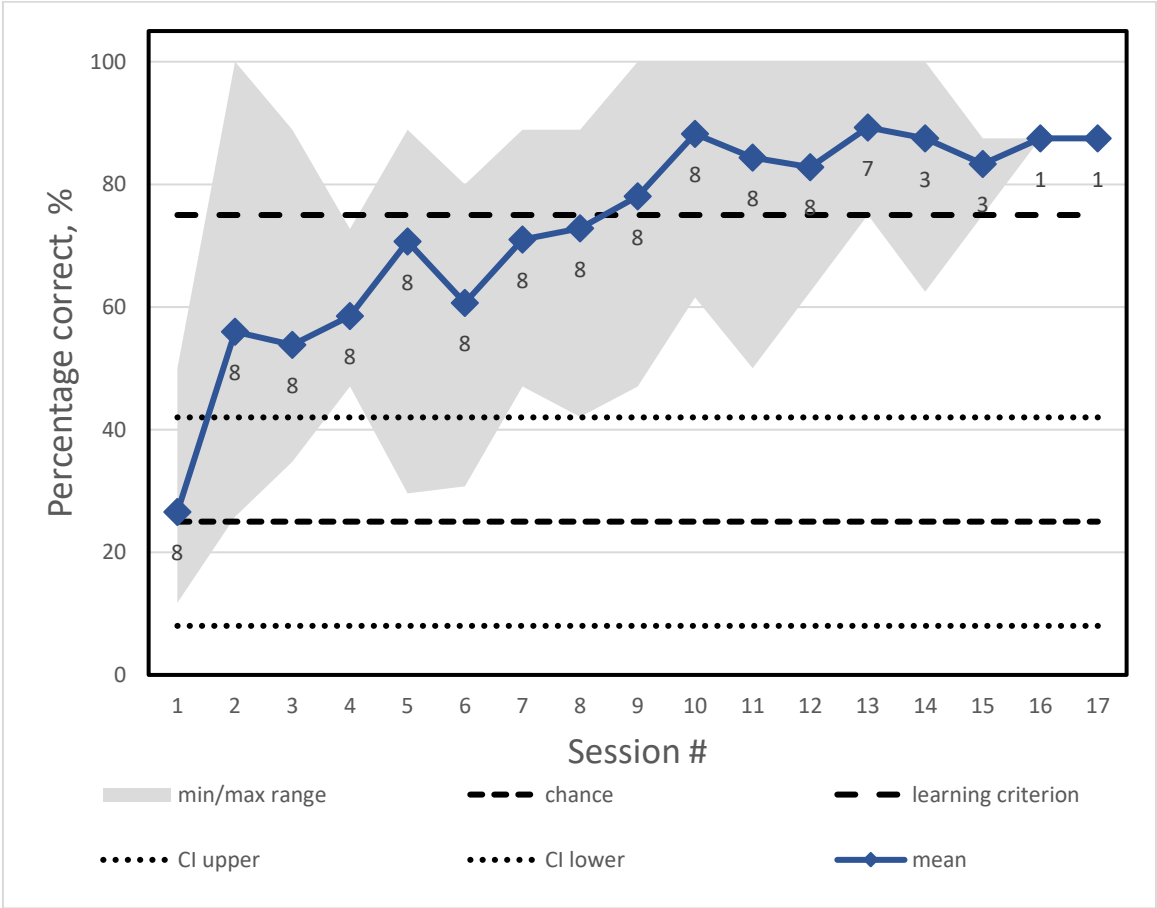



*Figure 9:* Mean subject discrimination performance during training. Point values indicate the number of subjects in training at that point in time – subjects were removed once they reached the learning criterion. Shaded area indicates maximum and minimum subject accuracy.

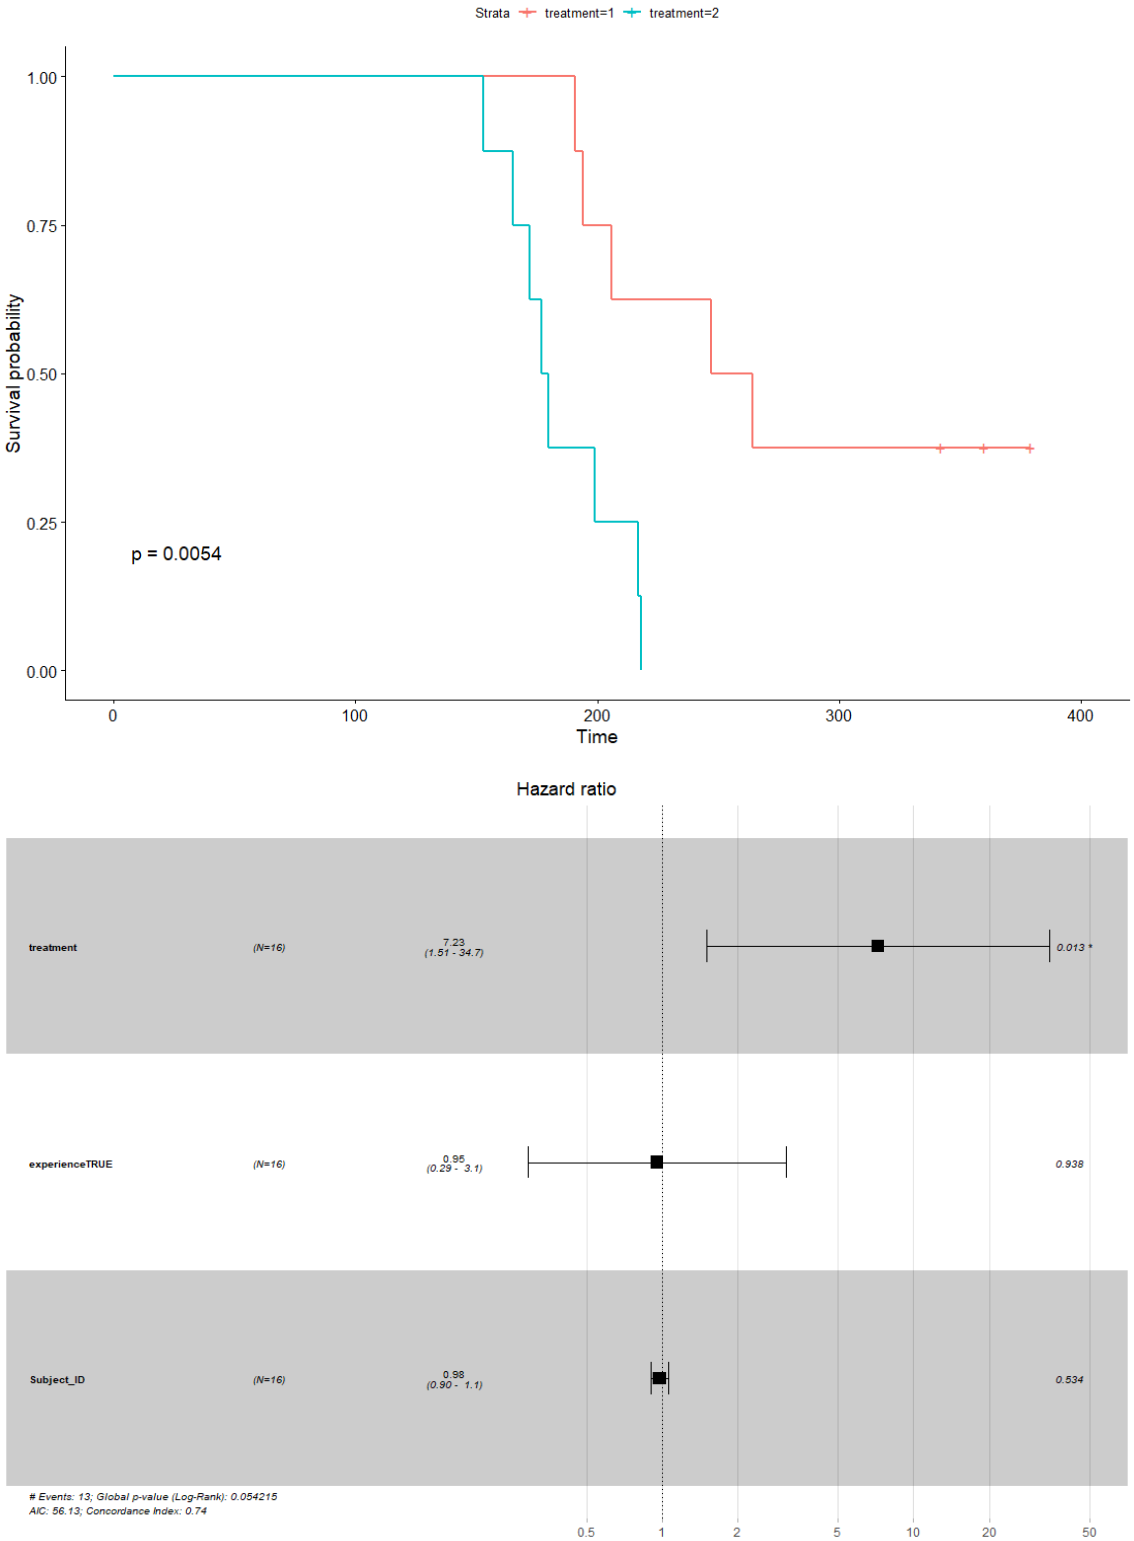



Figure 10: (a) Kaplan-Meier survival plot comparing the number of trials to reach the learning criterion between experiment 3 and the replicate. "Time" stands for the number of trials to reach criterion. "treatment 1" = experiment 3, "treatment 2" = replicate. Crosses indicate where a subject was disenrolled without reaching the learning criterion (b) Forest plot showing hazard ratios of each covariate in the Cox proportional hazard model.

# General Discussion

## Task acquisition

During the experiments discussed in Chapter 2, I observed that subjects assigned to experiment 3 exhibited dramatically poorer task acquisition performance than their peers in other treatments, regardless of prior experience. I hypothesized that the structure of this experiment (i.e., conditioning subjects to two distinct tasks in alternating trials) might have produced interference between the two tasks, thereby slowing learning. To test this hypothesis, I replicated the experiment conditioning subjects only to the positional tasks and compared performance between the two groups. Subjects conditioned to only the positional task acquired the task much more quickly than subjects conditioned to both, suggesting that the concurrent conditioning to both rules did have a detrimental effect on learning. However, their performance was still poorer than subjects with comparable experience in either experiment 1 or 2, suggesting interference was not the sole factor. Here, I consider some factors that may have contributed to this differential performance.

### Inadvertent counterconditioning

If a subject remained in the experimental compartment for too long at the end of a trial, a net was used to gently usher it back to the start box. This was more common following incorrect responses, as subjects would often wait for reinforcement. Typically, subjects simply avoided the net, returning to the start box in good order. During training in experiment 3, two subjects reacted rather differently – vigorously attacking the net, including biting and tail beating. This behaviour is reminiscent of territorial defence in wild triggerfish. It is possible that – counterintuitively – this interaction served as positive reinforcement for incorrect responses. The opportunity to engage in agonistic behaviour can be appetitive to an organism, and act as a positive reinforcer. For example, the opportunity to engage in territorial display with a (simulated) competitor acted as a positive reinforcer for male fighting fish (*Betta splendens*) [34]. By analogy, it is possible

that agonistic interactions with the net may have acted as a positive reinforcer of incorrect responses. This phenomenon is examined in greater detail in the general discussion (Chapter 5, section on "frustration").

### Exploratory behaviour

One factor that may have affected acquisition was the reinforcement schedule. The transition from pre-training to testing proper was accompanied by a shift from a ≤100% reinforcement schedule to a much-reduced effective reinforcement schedule, due to low initial subject accuracy. Classical operant conditioning theory states that the shift from a total to a partial reinforcement schedule accelerates task extinction [35]. It is thus likely that subjects experienced partial task extinction following the start of testing. An increase in explorative behaviour would also be predicted as an adaptive response during this transition. Trimmer *et al.* [36] showed that, when actors can expect to make multiple sequential choices, there is a beneficial trade-off to selecting a more ambiguous option that reduces uncertainty in future choices, rather than selecting an option with a known beneficial outcome. Since testing introduced potentially novel and ambiguous stimulus combinations, one might expect an increase in exploratory behaviour at this point in the experiment. Houston *et al.* [37] similarly provided an analytical demonstration that it may be beneficial for actors to deviate from currently optimal choices when certain conditions are met.

### Robust and optimal strategies

Subjects may have become mired in a behavioural strategy that was locally efficient but globally sub-optimal. Subjects can be (approximately) assumed to maximise their reinforcement rate, but this is not the same as maximising accuracy. Reinforcement rate is also contingent on factors such as the reaction time, inter-trial interval, and the expected number of trials per day. Particularly, object discrimination is a process of evidence

accumulation, resulting in a necessary trade-off between accuracy and speed (the speed-accuracy trade-off, or SAT) [38].

If the number of trials per day is not constrained and the risks of failure are small; then it may be favourable to choose quickly and with minimal inspection. In low-order tasks where the probability of making a correct response by chance is high (e.g., 50%), this strategy may maximise provisioning rate. However, this does not match the contingencies of the experiments discussed, which – with a low success rate for random choices (0.25) and a limited number of trials per day – present a very high opportunity cost for incorrect responses. Subjects typically exhibited rapid reaction times (<2 seconds) with correspondingly short assessment periods. It is possible then that subjects were labouring under a maladjusted SAT and decision strategy. However, this model has uncertain predictions, and so is not directly testable using the existing data set.

## Side bias

I anecdotally observed two phenomena of positional biases across the experiments discussed in the previous chapter – a consistent population-wide preference for the lower-right choices, and the development of idiosyncratic biases during testing. Re-analysing the dataset confirmed my observations that participants displayed a significant preference for the lower and right choices, with a larger effect in the horizontal dimension. This bias was larger when visual quality was poor. Of the individuals exhibiting idiosyncratic biases, I found that while their overall accuracy was comparable to the rest of the cohort, they received a lower reinforcement ratio due to lower accuracy in reinforcement trials. Furthermore, these biases only developed during testing.

Neither phenomenon disrupts the overall interpretation of the data, due to the factorial assignment of conditions. However, the question of why both phenomena appeared remains significant for our understanding of decision-making.

Regarding individual bias, it was primarily found in the pilot (experiment 1), in which subjects did not form as strong associations with the positional cue. The transition from bias training to testing was associated with a reduction in the reinforcement ratio from 100% to at most 50% assuming perfect accuracy. As discussed above, such transitions cause partial task extinction, and given the weaker positional association may have resulted in subjects exhibiting more exploratory behaviour or falling back on pre-existing strategies.

Regarding the population bias, while extraneous factors cannot be excluded, I find it unlikely that these influenced the results. The LR bias was consistent across multiple laboratories, experimenters, tank positions, and experiment phases, and (of course) subjects. Consequently, I hypothesize that this is a by-product of behavioural lateralization. Such spontaneous bias is commonly observed in behavioural experiments in a range of species, including fishes; e.g. [39]. This interpretation is discussed in more detail below in in the section “Response biases as decision strategies”.

## Sequence effects

In order to test whether subjects relied on low-level heuristics to solve the positional bias aspect of the experiments, I examined the testing data for sequence effects between trials. The results ran counter to expectations. Overall, triggerfish repeated positional choices at a rate lower than predicted by any models of behaviour. Furthermore, position repeat behaviour was further suppressed in the biased treatment group, further indicating that sequence effects could not account for biased behaviour. However, more complex responses were revealed upon breaking down the data. While overall repeat rates were below expectations, subjects linearly increased the rate of repeat responses as stimulus quality decreased. If subjects only repeated their choices when visual information was poor, this suggests that repeating choices was an active strategy chosen only in certain

contexts. This idea is examined in further detail in the following section (*Response biases as decision strategies*).

Subjects' repeat rates exhibited a complex relationship with prior reinforcement. Both biased and control subjects repeated their choices more frequently following non-reinforcement (i.e., win-shift, lose-stay), though this effect was larger in the bias group. This may represent a genuine difference in behaviour between the two treatment groups, or an artefact of the reduced sample size of the control group. There is no obvious adaptive explanation for this pattern, which appears to correlate to an inverse win-stay/lose-shift rule. Further research to determine if this is a real effect is warranted.

A further inflection on this pattern was observed in experiment 2. This method involved a two-tier reward system, in which correct responses in the biased position were rewarded with a preferred food (lancefish). In the control treatment, this was rewarded randomly. Despite the small sample size, there was a marked difference in behaviour between treatments, with only biased subjects dramatically increasing their repeat rate following reinforcement with lancefish. The increase in sequence effects in the treatment group follows associative learning predictions, but the behaviour of the treatment group is anomalous. A possible explanation for this lies in the differential outcomes procedure (DOP) [40]. In this model, subjects are concurrently conditioned to two distinct tasks, each with a distinct reinforcer. This has the effect of enhancing learning in both tasks. Alternatively, it could suggest that subjects in the control treatment learned that fish was a non-predictive reinforcer, and selectively ignored it. This would be a highly unusual finding – the notion that an unconditioned stimulus (a reinforcer, the lancefish) could have its associative value altered by how it is paired with a conditioned stimulus is novel, and as far as I am aware has but recently been considered [41]. While these results are informal, they are interesting enough to merit further investigation.

## Response biases as decision strategies

The emergence of response biases in animals has been documented in the literature on animal behaviour since its earliest development [42], [43], being included under the banner of "superstitious" behaviour by B.F. Skinner [24]. In the context of discrimination tasks, animals respond strongly to minor differences in the ordering of stimulus presentations, even when the number of presentations in each position remains at parity [44]. Animals may also be influenced by factors extraneous to the experiment, such as inadvertent cues by the experimenter, the positioning of objects in the laboratory, and etcetera. Sensitivity to stimulus ordering is especially high early in conditioning, when subjects have not developed strong associations to the exigencies of the stimuli [28]. These factors may induce a range of biases, including lateralized responses (i.e. side bias), sequence effects, alternation, and win-stay/lose-shift responses [42].

The emergence of these phenomena has been most extensively discussed in the domain of psychophysics, since deviations in subject's behaviour radically affect estimates of psychometric functions [28]. In this context response biases are typically treated as non-informative "nuisance" behaviours to be solved, rather than a potentially informative aspect of decision-making [28], [45]. Accordingly, conditioning architecture is specifically structured to minimise the emergence of response biases, through pseudo-random or Gellerman-like schedules [42], [43], [46], [47], or more recently dynamic schedules that adapt to subject choices to counter-condition emerging biases (e.g. [48]). The unifying assumption underlying all these approaches is that output (i.e. response) biases can be eliminated by removing input (i.e. schedule) biases (typified in [49]).

While these methods can minimise the emergence of response biases, some degree of bias typically remains after these factors are accounted for. With positional biases, "fixed" biases typically do not respond to reinforcement schedule and may not be ameliorated by

counterconditioning [50], and consequentially are typically compensated for via statistical correction without investigating their root cause [51]. In Fellows' nomenclature [42], fixed side bias is a response set hypothesis, i.e. a strategy that is independent of response reinforcement. These he contrasted with predictive hypotheses such as win-stay/lose-shift, in which the subsequent response is contingent on the outcome of the previous response. Fellows' "response set" is identifiable with the "global choice biases" of Linares *et al.* [52], response biases which are idiosyncratic and independent of reinforcement history. Notably, these two papers frame the phenomenon in deeply contrasting terms. Fellows characterises "preservationism" (side bias) as a behavioural strategy based on an inferred hypothesis of the subject about the rules of the system. By contrast, Linares *et al.*'s work on human behaviour (where, notably, the rules of the system are unambiguous to the subject) characterise response bias as an inadvertent artefact of either perceptual or behavioural tuning. While both approaches have merit, they likely describe quite distinct processes.

## Response bias as inherent bias

While adaptive strategies (i.e., in which the subject is following an internal model) may account for side bias in some contexts, it fails in the context of human studies in which subjects are aware of the rules of the task, and in equivalent cases with animals in which the subject has been conditioned to some idealized accuracy criterion, and yet exhibits side bias in their residual incorrect choices. Such residual fixed biases may arise from asymmetries at a perceptual, decisional, or a post-decisional (i.e., effectuator) level. For example, Lengyal *et al.* [45] demonstrated that in a 2AFC amplitude discrimination task, rat's degree of lateral bias modulated with signal intensity in a fashion that better matched a model with a post-decisional bias term than perceptual. Proximally, such asymmetries may arise from physiological asymmetries in receptors or due to the stochastic nature of neural dynamics without gross physiological asymmetry [50], [52].

However, distinguishing these alternatives is extremely challenging, and conclusions typically cannot be derived from accuracy data alone [52].

## Response bias as adaptive strategy

Laboratory animals do not share our priorities and have no interest in finding the "correct" solution to a discrimination task [53], instead operating under their own motivations. Indeed, the extent to which an animal "understands" the rules of the system as the experimenter conceptualises them is an epistemic question far beyond the scope of this thesis. Subjects are typically modelled as rational actors that aim to optimise some function, typically their reinforcement ratio, e.g. [28]. This too is an oversimplification, however. Animals – like humans – are not rational actors and exhibit various deviations from this model. For example, animals' perception of a reward (and change in associative value) is frequently modulated by their affective state [54], and reinforcement of one trial may also retroactively affect subject's valuation of previous trials [44], [55]. Furthermore, animal's behaviour in choice tasks is better predicted by matching frequency of behaviours to the proportion of reinforcement they elicit (i.e. melioration) [56], a strategy which is frequently outperformed by other approaches [31]. Crucially, we do not know what assumptions the subject brings to the laboratory. Natural environments are statistically non-stationary in their properties [57], and consequently animals may not adapt to the static conditions of many experiments.

Nevertheless, optimisation is a useful and functional approximation [58] of behaviour. While this process can be framed in purely associative terms, it is increasingly typical to frame decision-making in terms of strategies. Fawcett *et al.* (2013) [31] propose that animals should have an evolved suite of strategies that can be adapted to a range of circumstances, encompassing both fixed and flexible learning rules. The apparent universality of response biases (side bias, win-stay/lose-shift, alternation, etc.) across taxa

suggests they may be part of this suite of strategies. The use of such approaches should be expected when subjects have access to little information, and rapid decisions with a low but consistent payoff are favoured. As Louis Gellerman noted nearly a century ago [43], under truly random stimulus presentations, side bias can result in performance in excess of 70% in two-choice tasks - far above chance. Even under Gellerman-like schedules, the efficacy of side bias can be greater than chance or choosing randomly [59].

"Value set" strategies (i.e. those that are insensitive to the outcomes of an action, such as side bias) especially may be favoured when signal noise is high (i.e. the outcome of single decision not very informative) [57]. Supporting this hypothesis, Trevino *et al.* (2013) [59] documented that mice were more likely to develop side bias in a discrimination experiment when visual information was poor. This model would also fit the sequence effect results reported here, with the observation that the frequency of sequence effects increased when visual information was lacking.

But do these strategies have relevance outside of a limited experimental context? If they are only artefacts of laboratory conditions, the observed behaviours would lack any ecological significance. Work on behavioural lateralization argues otherwise. Beyond the context of decision-making experiments, cognitive and behavioural lateralization have been extensively documented in a wide range of animal taxa, including multiple fish species [60]. Many natural behaviours are lateralized, including predator avoidance [61] and novel object exploration [62]. While the principal focus has been on visually mediated behaviours, sensory lateralization has been observed across multiple modalities. For example, when encountering novel objects blind cave fish *(Astyanax fasciatus)* preferentially use the right flank of their lateral line organ [63]. The frequency of the phenomenon demands an adaptive explanation, and several have been proposed. Principally, evidence suggests that it may reduce redundancy and thereby increase neural

capacity by eliminating duplication of processes across cerebral hemispheres. As an extension, this division of functions may also enhance multitasking by preventing behavioural "jamming" occurring when incompatible responses are initiated by input from distinct receptors [30]. The benefits of this form of lateralization may be especially pronounced when there is little binocular overlap (such that each optic lobe receives separate inputs) and in organisms that lack a flexible neck. Fish exhibit both traits [60].

Even within an adaptive framework, the use of low-level heuristics is necessarily less effective for a subject that attending to the exigencies of the task – in the experiments discussed here, neither subjects' global biases nor their repeat rate matched the actual schedule of reinforcement, demonstrating that they were sub-optimal strategies.

A speed-accuracy trade-off (SAT) could possibly account for this pattern of behaviour, depending on the function that subjects were maximising. A local reinforcement frequency maximising function could be viable assuming visual discrimination were a more time- or energy-consuming strategy than position-repeating. However, this would come at a cost of overall reinforcement, due to the limited number of trials per day. Furthermore, the long inter-trial intervals compared to decision latencies argue against this (mean of ~3s versus ~180s). An untapped question is whether there are other invisible trade-offs at work. Different strategies may differ in their energetic cost (such as due to different motor approaches) or subjects may face cognitive or processing limitations. While the experimenter knows that predators are absent in laboratory settings, the subject certainly does not, and may still allocate attention to vigilance, thereby sabotaging discrimination performance through split attention [64].

A more parsimonious explanation relates to incomplete training. In the original formulation (experiment 1) most subjects failed to acquire the positional task – however, when the duration of bias training was doubled in experiment 2 performance dramatically improved.

By extension, subjects enrolled in experiments 2 and 3 may still not have formed complete associations with either cue for the task, allowing variance in their behaviour to be accounted for by other factors.

Fundamentally, the analyses presented are represent an *ad-hoc* reanalysis of data collected to examine a separate question, coupled with modest data supplementation. As such, the results are necessarily somewhat tentative and require dedicated replication to fully test my interpretation. Here, triggerfish have exhibited sub-optimal but complex decision strategies, the variance in which can only partially be accounted for by the strategies investigated here. Further experiments with this species may help to further decode the nature of decision making in non-cortical vertebrates.

## Concluding remarks

Throughout the experiments discussed in this thesis, subjects developed a range of different strategies to "solve" the discrimination tasks. Typically, subjects rapidly approached optimal solutions to solve the discriminations, i.e., attending to the intended exigencies of the task, and the specific conditioned stimulus features which predicted the unconditioned stimulus. However, this process of refining discrimination strategies was subject to ongoing bias both at an individual and population level. Triggerfish readily defaulted to simple positional rules to solve the tasks when contingencies shifted, and the population as a whole exhibited consistent behavioural bias in their choices. I suggest that these deviations from expected behaviour represent the emergence of flexible and adaptive behavioural strategies under unusual (i.e., experimental) conditions where they may not be generally adaptive. Particularly, the context-specific appearance of sequence effects strongly suggests that it represents a fallback strategy under conditions of poor visual information – a low-level equivalent of updating. This implies triggerfish are capable

of weighting not just different sources of information, but the relative merit of different behavioural strategies themselves.

## References

[1] R. A. Rescorla and A. R. Wagner, "A theory of classical conditioning: Variations in the effectiveness of reinforcement and nonreinforcement.," *Class. Cond. II Curr. Res. Theory*, vol. 21, no. 6, pp. 64–99, 1972.

[2] T. J. Valone, "Information for patch assessment: A field investigation with black-chinned hummingbirds," *Behav. Ecol.*, vol. 3, no. 3, pp. 211–222, 1992.

[3] H. H. Marshall, A. J. Carter, A. Ashford, J. M. Rowcliffe, and G. Cowlishaw, "How do foragers decide when to leave a patch? A test of alternative models under natural and experimental conditions," *J. Anim. Ecol.*, vol. 82, no. 4, pp. 894–902, Jul. 2013.

[4] J. A. van Gils, I. W. Schenk, O. Bos, and T. Piersma, "Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches.," *Am. Nat.*, vol. 161, no. 5, pp. 777–93, May 2003.

[5] A. S. Dunlap, D. R. Papaj, and A. Dornhaus, "Sampling and tracking a changing environment: persistence and reward in the foraging decisions of bumblebees," *Interface Focus*, vol. 7, no. 3, p. 20160149, Jun. 2017.

[6] T. J. Valone, "Are Animals Capable of Bayesian Updating ? An Empirical Review," *Oikos*, vol. 112, no. 2, pp. 252–259, Feb. 2006.

[7] B. Luttbeg and R. R. Warner, "Reproductive decision-making by female peacock wrasses: flexible versus fixed behavioral rules in variable environments," *Behav. Ecol.*, vol. 10, no. 6, pp. 666–674, Nov. 1999.

[8] W. Hunte, R. A. Myers, and R. W. Doyle, "Bayesian mating decisions in an amphipod, Gammarus lawrencianus Bousfield," *Anim. Behav.*, vol. 33, no. 2, pp. 366–372, 1985.

[9] K. P. Körding and D. M. Wolpert, "Bayesian integration in sensorimotor learning," *Nature*, vol. 427, no. 6971, pp. 244–247, Jan. 2004.

[10] M. Moutoussis, P. Fearon, W. El-Deredy, R. J. Dolan, and K. J. Friston, "Bayesian inferences about the self (and others): A review," *Consciousness and Cognition*, vol. 25, no. 1. Academic Press Inc., pp. 67–76, 2014.

[11] C. A. Holt and A. M. Smith, "An update on Bayesian updating," *J. Econ. Behav. Organ.*, vol. 69, no. 2, pp. 125–134, Feb. 2009.

[12] G. Charness and D. Levin, "When optimal choices feel wrong: A laboratory study of bayesian updating, complexity, and affect," *American Economic Review*, vol. 95, no. 4. Kahneman and Tversky, pp. 1300–1309, 2005.

[13] R. Y. Cho *et al.*, "Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task," *Cogn. Affect. Behav. Neurosci.*, vol. 2, no. 4, pp. 283–299, 2002.

[14] A. J. Yu and J. D. Cohen, "Sequential effects: Superstition or rational behavior?," *Adv. Neural Inf. Process. Syst.*, vol. 21, pp. 1873–1880, 2009.

[15] L. P. Sugrue, G. S. Corrado, and W. T. Newsome, "Matching behavior and the representation of value in the parietal cortex," *Science (80-. ).*, vol. 304, no. 5678, pp. 1782–1787, 2004.

[16] M. L. Wildhaber, R. F. Green, and L. B. Crowder, "Bluegills continuously update patch giving-up times based on foraging experience," *Anim. Behav.*, vol. 47, no. 3, pp. 501–513, Mar. 1994.

[17] W. R. Garner, "An informational analysis of absolute judgments of loudness," *J. Exp. Psychol.*, vol. 46, no. 5, pp. 373–380, Nov. 1953.

[18] M. E. Putt, "Carryover and Sequence Effects," in *Encyclopedia of Statistics in Behavioral Science*, B. S. Everitt and D. C. Howell, Eds. Chichester, UK: John Wiley & Sons, Ltd, 2005, p. Volume 1, A-D.

[19] A. Acheson, "Sequence Effects." 2010.

[20] E. L. Crawford and S. Duffy, "Sequence effects in estimating spatial location," *Psychon. Bull. Rev.*, vol. 17, no. 5, pp. 725–730, Oct. 2010.

[21] E. Soetens, L. C. Boer, and J. E. Hueting, "Expectancy or Automatic Facilitation? Separating Sequential Effects in Two-Choice Reaction time," *J. Exp. Psychol. Hum. Percept. Perform.*, vol. 11, no. 5, p. 598, 1985.

[22] N. Stewart, G. D. A. Brown, and N. Chater, "Sequence Effects in Categorization of Simple Perceptual Stimuli," *J. Exp. Psychol. Learn. Mem. Cogn.*, vol. 28, no. 1, pp. 3–11, Jan. 2002.

[23] C. L. Ecott and T. S. Critchfield, "Noncontingent reinforcement, alternative reinforcement, and the Matching Law: a laboratory demonstration," *J. Appl. Behav. Anal.*, vol. 37, no. 3, pp. 249–265, Sep. 2004.

[24] B. F. Skinner, "'Superstition' in the pigeon," *Journal of Experimental Psychology*, vol. 38, no. 2. American Psychological Association, US, pp. 168–172, 1948.

[25] M. Milinski, "How To Avoid Seven Deadly Sins in the Study of Behavior," in *Advances in the Study of Behavior*, vol. 26, no. C, 1997, pp. 159–180.

[26] G. R. Lockhead and J. Hinson, "Range and sequence effects in judgment," *Percept. Psychophys.*, vol. 40, no. 1, pp. 53–61, 1986.

[27] R. Diaz-Uriarte, *The analysis of cross-over trials in animal behavior experiments : review and guide to the statistical literature*. 2001.

[28] D. Herrera and M. Treviño, "Undesirable choice biases with small differences in the spatial structure of chance stimulus sequences," *PLoS One*, vol. 10, no. 8, pp. 1–13, Aug. 2015.

[29] M. Treviño, "Stimulus similarity determines the prevalence of behavioral laterality in a visual discrimination task for mice," *Sci. Rep.*, vol. 4, no. 1, p. 7569, May 2014.

[30] G. Vallortigara and L. J. Rogers, "Survival with an asymmetrical brain: Advantages and disadvantages of cerebral lateralization," *Behav. Brain Sci.*, vol. 28, no. 4, pp. 575–633, 2005.

[31] T. W. Fawcett, S. Hamblin, and L.-A. Giraldeau, "Exposing the behavioral gambit: the evolution of learning and decision rules," *Behav. Ecol.*, vol. 24, no. 1, pp. 2–11, Jan. 2013.

[32] W. J. Matthews and N. Stewart, "The effect of interstimulus interval on sequential effects in absolute identification," *Q. J. Exp. Psychol.*, vol. 62, no. 10, pp. 2014–2029, Oct. 2009.

[33] M. M. A. Ananda and S. Weerahandi, "Two-Way ANOVA with unequal cell frequencies and unequal variances," *Stat. Sin.*, vol. 7, no. 3, pp. 631–646, 1997.

[34] T. I. Thompson, "Visual reinforcement in Siamese fighting fish," *Science (80-. ).*, vol. 141, no. 3575, pp. 55–57, 1963.

[35] N. J. Mackintosh, *The psychology of animal learning*. Oxford, England: Academic Press, 1974.

[36] P. C. Trimmer, A. I. Houston, J. A. R. Marshall, M. T. Mendl, E. S. Paul, and J. M. McNamara, "Decision-making under uncertainty: biases and Bayesians," *Anim. Cogn.*, vol. 14, no. 4, pp. 465–476, Jul. 2011.

[37] A. I. Houston, J. M. McNamara, and M. D. Steer, "Violations of transitivity under fitness maximization.," *Biol. Lett.*, vol. 3, no. 4, pp. 365–7, Aug. 2007.

[38] L. Chittka, P. Skorupski, and N. E. Raine, *Speed-accuracy tradeoffs in animal decision making*, vol. 24, no. 7. 2009, pp. 400–407.

[39] C. Newport, G. Wallis, and U. E. Siebeck, "Concept learning and the use of three common psychophysical paradigms in the archerfish (Toxotes chatareus)," *Front. Neural Circuits*, vol. 8, no. APR, p. 39, Apr. 2014.

[40] R. Ramos and L. M. Savage, "The Differential Outcomes Procedure Can Interfere or Enhance Operant Rule Learning," *Integrative Physiological and Behavioral Science*, vol. 38, no. 2. Transaction Publishers, pp. 17–35, 2003.

[41] M. Treviño, "Associative Learning Through Acquired Salience," *Front. Behav. Neurosci.*, vol. 9, p. 353, Jan. 2016.

[42] B. J. Fellows, "Chance Stimulus Sequences for Discrimination Tasks," *Psychol. Bull.*, vol. 67, no. 2, pp. 87–92, 1967.

[43] L. W. Gellermann, "Chance Orders of Alternating Stimuli in Visual Discrimination Experiments," *Pedagog. Semin. J. Genet. Psychol.*, vol. 42, no. 1, pp. 206–208, Mar. 1933.

[44] C. A. Catania, "Reinforcement Schedules: The Role of Responses Preceding the One that Produces the Reinforcer," *J. Exp. Anal. Behav.*, vol. 15, no. 3, pp. 271–287, 1971.

[45] G. Lengyel, A. Hyafil, J. Fiser, and J. de la Rocha, "The origin of fixed, history-independent choice biases of rodents in perceptual decision making tasks," in *Conference on Cognitive Computational Neuroscience*, 2019.

[46] C. J. Gerard, H. A. Mackay, B. Thompson, and W. J. McIlvane, "Rapid generation of balanced trial distributions for discrimination learning procedures: a technical note.," *J. Exp. Anal. Behav.*, vol. 101, no. 1, pp. 171–8, Jan. 2014.

[47] M. van Casteren and M. H. Davis, "Mix, a program for pseudorandomization," *Behav. Res. Methods*, vol. 38, no. 4, pp. 584–589, Nov. 2006.

[48] A. T. Piet, A. El Hady, and C. D. Brody, "Rats adopt the optimal timescale for evidence integration in a dynamic environment," *Nat. Commun.*, vol. 9, no. 1, pp. 1–12, Dec. 2018.

[49] W. M. Baum, "On two types of deviation from the matching law: bias and undermatching.," *J. Exp. Anal. Behav.*, vol. 22, no. 1, pp. 231–242, 1974.

[50] L. Lebovich, R. Darshan, Y. Lavi, D. Hansel, and Y. Loewenstein, "Idiosyncratic choice bias naturally emerges from intrinsic stochasticity in neuronal dynamics," *Nat. Hum. Behav.*, vol. 3, no. 11, pp. 1190–1202, Nov. 2019.

[51] K. Vancleef, J. C. A. Read, W. Herbert, N. Goodship, M. Woodhouse, and I. Serrano-Pedraza, "Two choices good, four choices better: For measuring stereoacuity in children, a four-alternative forced-choice paradigm is more efficient than two," *PLoS One*, vol. 13, no. 7, p. e0201366, Jul. 2018.

[52] D. Linares, D. Aguilar-Lleyda, and J. López-Moliner, "Decoupling sensory from decisional choice biases in perceptual decision making," *Elife*, vol. 8, Mar. 2019.

[53] A. I. Houston, "Flying in the face of nature," *Behavioural Processes*, vol. 80, no. 3. pp. 295–305, 2009.

[54] J. M. Aw, R. I. Holbrook, T. Burt de Perera, and A. Kacelnik, "State-dependent valuation learning in fish: Banded tetras prefer stimuli associated with greater past deprivation," *Behav. Processes*, vol. 81, no. 2, pp. 333–336, Jun. 2009.

[55] P. R. Killeen, "Mathematical principles of reinforcement," *Behavioral and Brain Sciences*, vol. 17, no. 01. p. 105, 1994.

[56] R. J. Herrnstein and W. Vaughan, "Melioration and Behavioral Allocation," in *Limits to Action*, J. E. R. Staddon, Ed. Academic Press, 1980, pp. 143–176.

[57] M. R. Gross, A. I. Houston, E. J. Collins, J. M. McNamara, F.-X. Dechaume-Moncharmont, and N. R. Franks, "Simple learning rules to cope with changing environments," *J. R. Soc. Interface*, vol. 5, no. 27, pp. 1193–1202, Oct. 2008.

[58] F. Balci *et al.*, "Acquisition of decision making criteria: Reward rate ultimately beats accuracy," *Attention, Perception, Psychophys.*, vol. 73, no. 2, pp. 640–657, Feb. 2011.

[59] M. Treviño, T. Oviedo, P. Jendritza, S. Bin Li, G. Köhr, and R. J. De Marco, "Controlled variations in stimulus similarity during learning determine visual discrimination capacity in freely moving mice," *Sci. Rep.*, vol. 3, no. 1, pp. 1–13, Jan. 2013.

[60] A. Bisazza and C. Brown, "Lateralization of Cognitive Functions in Fish," in *Fish Cognition and Behavior*, 2nd ed., vol. 2, C. Brown, K. Laland, and J. Krause, Eds. Oxford: Wiley-Blackwell, 2011, pp. 298–324.

[61] G. Lippolis, A. Bisazza, L. J. Rogers, and G. Vallortigara, "Lateralisation of predator avoidance responses in three species of toads," *Laterality*, vol. 7, no. 2, pp. 163–183, 2002.

[62] Á. Miklósi and R. J. Andrew, "Right eye use associated with decision to bite in zebrafish," *Behav. Brain Res.*, vol. 105, no. 2, pp. 199–205, Nov. 1999.

[63] T. Burt De Perera and V. A. Braithwaite, "Laterality in a non-visual sensory modality - The lateral line of fish [1]," *Current Biology*, vol. 15, no. 7. Cell Press, pp. R241–R242, 12-Apr-2005.

[64] R. Dukas, "Behavioural and ecological consequences of limited attention," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 357, no. 1427. pp. 1539–1547, 2002.

# CHAPTER 4

## Assessment of Cue reliability in *R. aculeatus*

# ASSESSMENT OF CUE RELIABILITY IN *RHINECANTHUS ACULEATUS*

## Abstract

This chapter describes an attempt to examine how triggerfish learn to recognise and discriminate objects in the context of a shifting signalling environment. In previous chapters, I have demonstrated that triggerfish can use both prior experience and sensory information in tandem, and that they shift the relative reliance on visual information according to its quality. This raises the question of how learned recognition proceeds when multiple competing sensory cues are present. The conspicuousness hypothesis proposes that predators attend to prey signals that strongly contrast with the background, thereby driving the evolution of prey crypsis. Considering the highly labile nature of aquatic environments, different signals may be favoured at different times and in different contexts. Thus, I hypothesized that triggerfish would selectively attend to those signals that are consistently reliable, and shift signal use depending on context.

To address this question, I trained triggerfish to complete a compound object discrimination task, in which two distinct cues could be used to distinguish the stimuli. During training one of the cues was consistently conspicuous, while the other either varied, being either consistently poor quality, or consistently high quality, depending on the treatment group. This simulated the effect of a dynamic sensory environment changing the reliability of different visual channels. Once subjects had learned to reliably discriminate the two stimuli, they were to be presented with a series of probe trials in which the cues were presented in isolation, in opposition, or with a dimensional shift. This was intended to quantify the degree to which subjects had learned and used each cue under different sensory regimes.

No subject successfully learned to discriminate the stimuli, despite the success of the training paradigm in previous experiments, and multiple sequential modifications made to the method in order to overcome possible sensory or cognitive barriers to learning. Herein, I discuss several hypotheses as to the failure of the experiment, particularly focusing on possible cognitive and sensory barriers, and their implications for our understanding of triggerfish cognition and decision-making.

## Background

A central question in the study of animal behaviour is how animals learn to recognise objects and events in their environment – be they food, threat, a potential mate, or otherwise. Animals are equipped with multiple senses, each of which can provide a wealth of potentially redundant or complimentary information about their environment [1]. However, detecting any individual signal represents a challenge due to individual sensory limitations [2], signal degradation and ambiguity, and limited attention and processing ability [3]. Perception, then, is a process characterised at once by great redundancy and terrible scarcity; both of information and attention. It follows that animals must be selective in the signals they attend to – the question is, which are these?

How a signal is perceived is heavily influenced not only by the properties of the object or event of interest, but also the context in which it occurs. For example, a visual signal is determined not only by the reflectance of the object or event of interest, but by the intensity and spectral composition of light illuminating it, the way it contrasts with the scene background, and the optical properties of the transmission medium coupled with the distance to the receiver. Once the signal reaches the eye, it will further be perceived differently depending upon factors such as light adaptation, which is primarily determined by the overall illumination of the receiver. Hence, the physical environment not only provides a vital context for understanding signals but will also strongly affects how effectively different signals are transmitted and received by observers.

It follows that different environments favour different sensory channels, and consequently both signals and sensory systems evolve to suit the peculiar conditions organisms find themselves in. Take, for example, the case of the golden- (*Manacus vitellinus*) and white-collared (*Manacus candei*) manakins. These species are largely allopatric but overlap in a narrow geographic area wherein hybridisation occurs, resulting in some *M. candei* males

exhibiting an intermediate yellow collar. These individuals enjoy a large mating advantage over white-collared males, but despite this the overlap zone appears stable, and yellow-collared males have not spread further into the white-collared manakin's range. Uy and Stein (2007) demonstrated that this is the result of the physical environment in the hybrid zone producing a peculiar light environment in which yellow-collared males are more conspicuous to the females' visual system [4]. In the absence of a specific female preference for yellow collars, a minor difference in the light environment preferentially transmitting yellow-collared signals is sufficient to a large mating advantage for those males bearing that signal.

The medium through which a signal is transmitted can also have a major effect on how it is perceived [5] – a factor which is especially prominent in aquatic environments. Light is far more strongly attenuated in water than air, and the degree of attenuation varies strongly with wavelength. This process is exacerbated by both the presence of dissolved matter and suspended particles, and by the properties of incoming light [6]–[9]. The effects of these interactions on visual perception are complex and not consistently predictable. Increasing turbidity is associated with a reduction in spectral bandwidth due to scattering and absorption, and reduction in object brightness and background contrast due to scattering. However, some research has indicated that low levels of turbidity can actually increase background contrast, depending on the size of the suspended particles and other factors [10], [11]. For example, cuttlefish *(Sepia prashadi & S. pharaonis)* are better able to detect a pattern through polarisation vision in low-turbidity conditions than in fully clear water.

These factors result in different aquatic environments favouring radically different visual channels. For example, female guppies *(Poecilia reticulata)* respond differently to the chromatic components of male nuptial colours when placed in different water colours [12]. The consequences of this have been most extensively explored in the context of sexual

signals, with evidence of evolutionary adaptation at the levels of signal expression [13], [14], receptor design [15], and behaviour [16]. The degree of sensitivity to small differences in photic regime in some species is remarkable. For example, expression of male nuptial colouration in three-spine sticklebacks (*Gasterosteus aculeatus*) varies considerably with depth on spatial scales as small as 1.5 metres [17], while small differences in depth distribution radically affect the degree to which male colouration in two related Lake Victoria cichlids (*Haplochromis pundamilia* and *H. nyererei*) varies with turbidity [18].

In addition to evolutionary adaptation, phenotypic plasticity plays a major role in responses to changing sensory environments. Photoreceptor gene expression in three-spined sticklebacks [19], guppies [20], and bluefin killifish (*Lucania goodei*) [21] varies with photic regime, while expression of nuptial colouration has been found to be sensitive to light conditions in a range of species including Haplochromine cichlids [22] and Nile tilapia (*Oreochromis niloticus*) [23].

While there has been extensive research into both evolutionary and developmental adaptation to the sensory environment, the role of individual learning has gone comparatively unexamined, and while behavioural responses to changing signalling conditions have been widely documented [8], [24]–[26], the cognitive basis of these differences has rarely been interrogated. This absence is despite the fact that much environmental variation occurs on spatial and temporal scales that are experienced by individual organisms. Aquatic ecosystems in particular are highly labile on a range of scales [27], [28]. Given these observations, the capacity to for behavioural accommodation should be adaptive, and we should expect animals to shift which signals they attend to as the signalling environment changes.

## Aims

In this chapter, I set out to examine how triggerfish might adjust their learning strategy in the context of a shifting sensory environment. In a dynamic sensory environment different sensory channels may be favoured at different times, and this should affect how animals attend to and learn different signals. This represents a continuation and elaboration of the experiments discussed in Chapters 2 and 3, in which I found that triggerfish weight their reliance on visual information versus prior information according to its reliability. Here I extended this to examine how triggerfish learn to recognise complex objects, and if they can attend to both statis and dynamic differences in signal reliability to maximise their performance when multiple signals are present. The purpose of this feasibility study was simply to determine whether the baseline cognitive capacity to assess relative cue reliability is present in triggerfish. This would have informed the design of subsequent experiments investigating if triggerfish could apply this facility in an ecologically relevant context, i.e., object recognition in shifting sensory environments.

To tackle this question, I trained triggerfish to complete an object discrimination task with compound stimuli that differed in two features, which were distinct but matched in form and magnitude. A subject could potentially use either of these features – henceforth "cues" – to successfully discriminate the stimuli. Depending upon the treatment group, one of the cues was presented at high quality, and the other either at constant high quality, constant medium quality, or variable quality. This simulated a situation where multiple cues are present, but a varying signal environment favoured different cues at different points in time. Once subjects had learned to complete the discrimination task, they were to be presented with a series of probe trials that variously broke down the task into its component elements; presenting only a single cue, placing the cues in opposition, or dimensionally shifted the cues.

By presenting the conditioned cues in isolation or opposed, subjects' choices would indicate the degree to which they had learned each individual cue, and by comparing between treatment groups it would have been possible to examine how the presentation of a cue (as opposed to its inherent properties) affected how triggerfish responded to it. I hypothesized that subjects would learn to attend to the cue that was consistently reliable over one that was intermittently reliable, even if the mean reliability of the two was matched. All subjects enrolled in the initial pilot group failed to reach the pre-established learning criterion or indeed perform consistently above chance, and so I was not able to carry out any probe trials. Several sequential modifications were made to the method to ameliorate this problem, with similar results in each case. The results of this training including the modified experiments are presented below, accompanied by a discussion of the outcomes and proposals for future research directions.

## Significance

Human activities are an ongoing driver of increasing turbidity in aquatic ecosystems globally, with the degree of disturbance is likely to increase in the future [29]. Primary drivers of increasing turbidity include increased sediment loading from mining activities and erosion, eutrophication due to agricultural runoff, and sediment resuspension from trawling [30]–[33]. The adverse consequences of these activities on ecosystems have been widely documented in a variety of contexts, reviewed in [34]. To focus on visual perception, it is accepted that increased turbidity generally has a detrimental effect on most visually mediated behaviours, including foraging [8], [10], [35]–[37], group cohesion [26], [38], and sexual signalling [14], [39]–[42]. The literature on the relationship between turbidity and sexual selection is especially extensive, with sticklebacks [43]–[45] being a case study in the phenomenon. At a population level, anthropogenic turbidity degrades specificity in visual sexual signals, undermining the processes that produce and maintain diversity – as

exemplified by the collapse of the Lake Victoria haplochromine cichlid radiation [46]–[48]. Meanwhile, adverse effects on foraging and movement ecology can have drastic consequences for both recruitment and population structure [49].

Nevertheless, how individual species, communities and populations respond to disturbances remains unpredictable. A factor contributing to this may be signal redundancy – wherein multiple cues may contribute to a behaviour. If a behaviour can be mediated by multiple sensory cues, individuals may shift their attention to alternative cues if turbidity disrupts cues on other visual channels. The cognitive capacity for this attention shifting is the subject of this chapter. Such individual behavioural variance can affect community-level responses to disturbance [24], [34], and therefore improving our understanding of animals' cognitive and behavioural flexibility may be essential to assess how those species and communities will respond to drastic environmental change.

# Methods

## Subjects & husbandry

As in previous chapters, all triggerfish used in these experiments were wild-caught sub-adults (standard length ca. 10-20cm), acquired from a local aquarium supplier (The Goldfish Bowl, 118-120 Magdalen Road, Oxford OX41RQ). Subjects were cared for in accordance with the local code of practice for the care of animals (Ethical Review Committee, Department of Zoology, University of Oxford). Once data collection was complete, subjects were either retained for further experiments, or returned to the supplier. Eight subjects were used across all experiments, all with no prior experimental experience. Several changes were made to the husbandry and housing of experimental subjects from previous work, focusing on an increased swimming volume, improved filtration capacity, and the inclusion of more enrichment objects for improved welfare.

Subjects were individually housed in glass aquaria (105.0 x 37.2 x 31.5cm), arrayed on two tiers of metal racks. A shared external sump system provided filtration and water flow, and opaque PVC sheets were used to create visual barriers between compartments. Aquaria were illuminated both by laboratory lighting and a dedicated lighting system, both on a closed-circuit 12:12h day: night cycle, with ramp-up and ramp-down to mitigate the stress of rapid lighting changes. Aquaria were furnished with a coral gravel substrate, while coral rock caves provided both shelter and enrichment. Outside of experiments, subjects were fed twice daily six days a week, using a mixture of standard commercial pellet food (Formula One Marine Pellet, Ocean Nutrition, USA) and frozen foods (krill, lancefish, and whole cockles).

## Experimental paradigm

Subjects were presented with a discrimination task, in which they were conditioned to select a positive stimulus (CS+) while ignoring a distractor (CS-). The stimuli were similar but could be distinguished based on two individual features ("cues"), either of which was theoretically sufficient to discriminate the stimuli. A subject could attend to either one or both in order to complete the task. During training, one cue would be presented at high quality, while the other would either be low quality, or vary in quality. The former simulated a consistent visual environment in which one cue was always preferable, while the latter simulated a varying environment in which the optimal cue was contingent and unpredictable. In addition, a control group was presented with both cues at high quality.

## Stimulus quality and design

There were multiple possible approaches available to control the quality of a cue. For example, the chromatic composition or intensity of illumination could be changed, or the turbidity of the water altered. However, these approaches come with significant drawbacks. In addition to being difficult to maintain consistency, turbidity can act as a

physiological stressor for fish [50], [51], and would require much finer control of stimulus viewing distance. I considered using coloured filters to alter the light environment. However, triggerfish are known to have strong colour preferences and aversions [52] which would introduce problematic motivational artefacts with chromatic stimuli. Instead, for these experiments I opted to adjust the physical properties of the stimuli rather than the visual environment itself and use only achromatic stimuli.

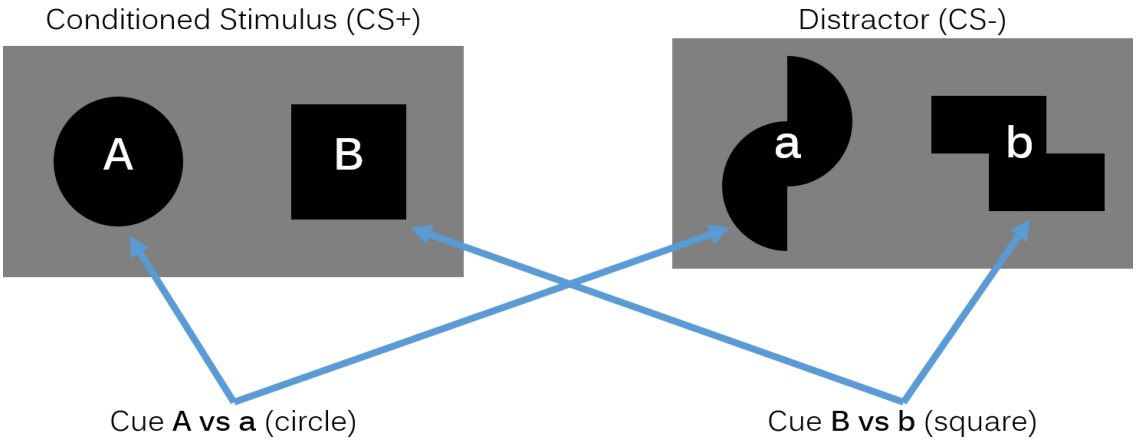



*Figure 1:* **Comparison of stimuli used in the experiment, not actual size.**

Each stimulus consisted of a tile featuring two elements – "A" and "B". Each element could exist in one of two configurations ("A" vs "a", and "B" vs "b"), and the difference between these configurations constituted a cue. Thus, a subject would be conditioned to choose (for example) stimulus "aB", while ignoring the distractor "Ab". The subject could successfully complete the task by attending to both cues, to only a/A, or *vice versa*. To ensure that any preference for a cue was a result of the *quality* with which it was presented, it was essential that each should be – insofar as possible – equivalent, and of equal psychophysical salience to the subjects. As such, the elements were matched by area, and each cue was dependent on an offset of equal magnitude, albeit in different spatial dimensions (horizontal versus vertical). Furthermore, each element was highly abstract and monochromatic to mitigate motivational effects. In a pilot where subjects were conditioned

to a single cue in isolation, there was no observed difference in learning speed or accuracy. Finally, the experiment was symmetric and fully factorial, with each element being used as either positive or negative stimulus for different subjects to account for any remaining differences.

The quality of each cue was controlled by adjusting the contrast between the element and its background. The background was presented at a fixed luminance, while the elements could be made lighter or darker, thereby changing the quality of the cue. Before beginning the experiment, a behavioural assay was carried out with a separate population of subjects to determine an appropriate range of lightness differences, using a variant staircase protocol.

Stimuli were generated in PowerPoint (Microsoft, etc.), then printed (Canon matte photo paper MP-101, 170gm$^{-2}$, Canon PIXMA iP7250) and laminated to prevent water damage.

## Pre-training

In previous work, experiments took place in subjects' home aquaria. However, I found that this introduced too many uncontrolled variables. As such, a dedicated experimental tank was used for this project. This was a bare glass tank (35.5 x 60.0 x 31.5 cm), the long sides masked with opaque white PVC panels. When not in use, filtration was provided by an internal filter, and daily 20% water changes were carried out on experiment days. Between sessions, any debris or uneaten food was removed using a net.

To minimise stress of transport between aquaria, subjects were moved in a large (2l) beaker rather than a net. Other members of our lab group have successfully used this approach and have found that subjects rapidly learn to approach and enter the beaker of their own volition (H.C. Karlsson, personal communication). Before experiments began, three steps of pre-training took place – forceps training, beaker training, and experiment pre-training. Subjects were initially conditioned to feed from forceps in their home tank.

Once this was achieved, subjects were accustomed to take food adjacent to an immersed beaker, and then to enter and accept being moved in the beaker without exhibiting stress responses. Finally, once subjects could be reliably moved back and forth between their home and experimental tanks, they were conditioned to swim between compartments of the experimental tank and peck at a blank stimulus presented on a transparent Perspex rod to receive a food reward. Once subjects could reliably complete this entire sequence, they were enrolled in the main experiment.

## General procedure

As previously, I used a simultaneous four-alternate forced-choice (4-AFC) task for this experiment, consisting of a single positive conditioned stimulus (CS+) and three identical distractors (CS-).

The experimental tank was divided into two compartments by an opaque PVC wall placed 20cm from the front glass pane. This separated the aquarium into an anterior “start box”, and a posterior “choice chamber” where stimuli were presented on an opaque white PVC board. Subjects could navigate between the two compartments by means of a guillotine door operated by the experimenter. A trial began with a subject in the anterior start box and was marked by the experimenter opening the guillotine door (figure 2). Subjects entered the choice chamber, after which the guillotine door was closed, and they were able to interact with the stimuli placed at the posterior end of the chamber. After assessing the stimuli, subjects would indicate a choice by biting a stimulus. Correct choices were promptly reinforced with a food reward, presented using forceps. Once a choice was made, the board was removed, and door re-opened. Subjects could return to the start box of their own volition. No negative reinforcement was used for incorrect responses. If a subject failed to make a detectable choice within 300 seconds, then the session was terminated, and training was resumed in the next training slot.

## Minimising inadvertent biases

Subjects indicted a choice by biting at a stimulus. This behaviour mimics the natural feeding of triggerfish and is thus easily conditioned onto new stimuli. However, given the large area of each stimulus compared to a triggerfish's mouth, we considered that subjects might develop self-reinforcing biases towards one part of a stimulus, and thereby towards one or another cue. To minimise the chance of this, a target consisting of a 10mm red star was added to each stimulus, placed between each element (figure 3). Red was chosen since previous research has found that Picasso triggerfish have a preference for this colour [53]. Only responses that directly targeted this spot were reinforced, whereas bites elsewhere on the stimulus were treated as incorrect responses.

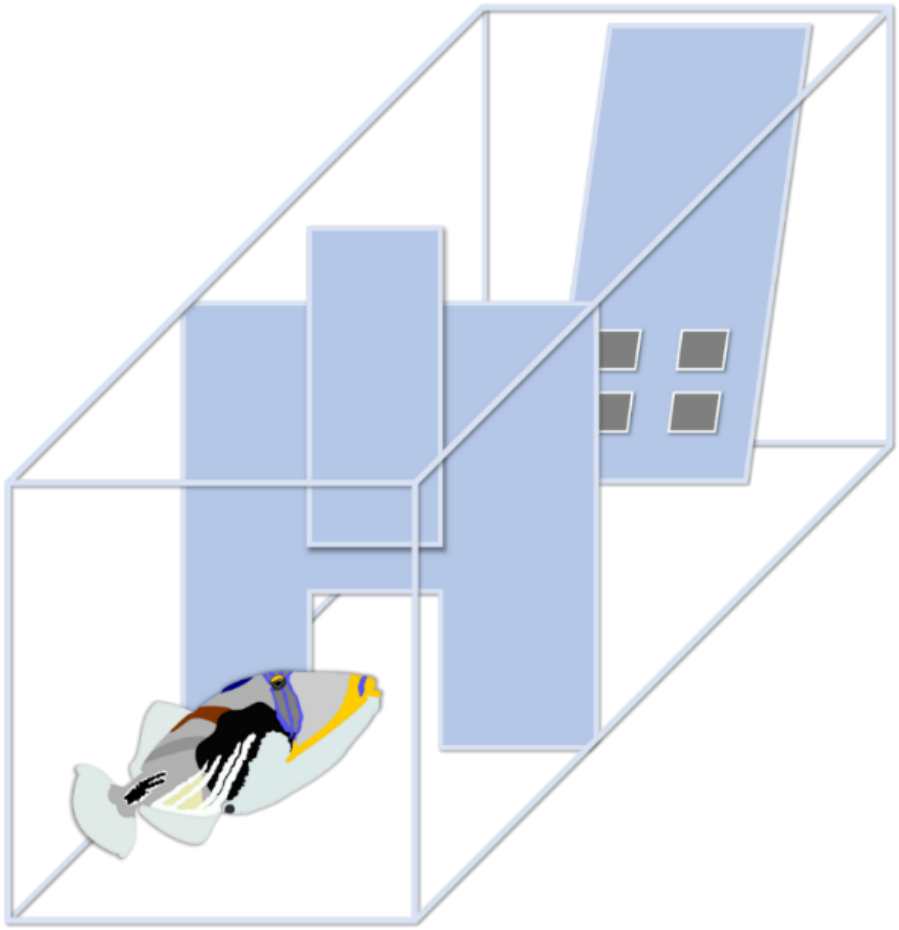


*Figure 2:* schematic illustration of experimental arena, showing anterior start box divided from the choice chamber by a guillotine door, and presentation board featuring stimuli. Subject approximately to scale.

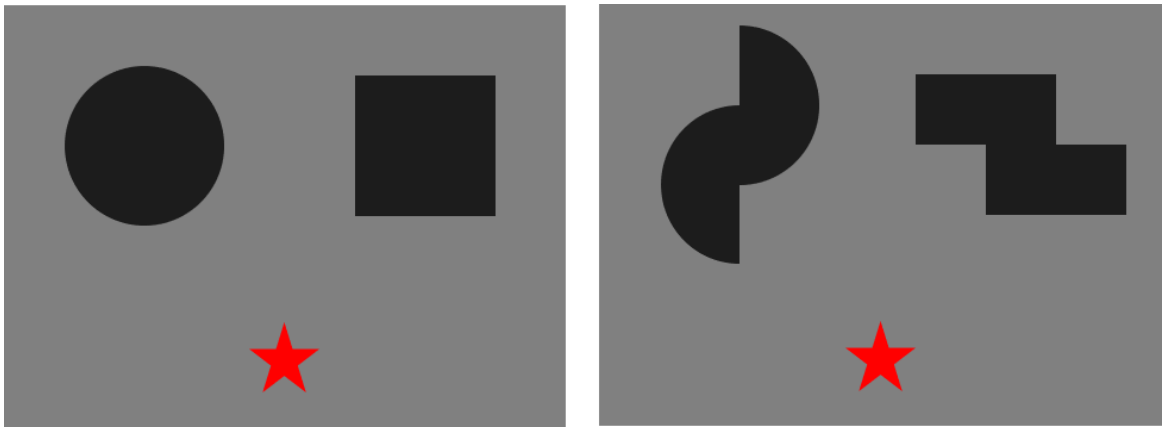


*Figure 3:* Illustration of the stimuli employed in the experiment, showing the addition of the appetitive element – a red star. Approximately actual size. Only bites directed to the red star were reinforced.

Several precautions were employed against experimenter behaviour biasing the subject's choices. Firstly, the experimenter remained motionless during trials, and kept their behaviour consistent surrounding this phase. Secondly, the experimenter delayed reaching for the food reward until a definitive response was produced. Thirdly, the experimenter attended to the motion of the fish; and ignored the stimuli to prevent gaze acting as a cue.

## Training

Five subjects were initially enrolled in the experiment, and assigned randomly to one of three treatment groups using a true random number generator [54]. This initial group was planned as a pilot, with more subjects to be included later, with balanced test conditioned across subjects. In the Control group (Group 0), both the Sq and Ci elements were presented at high contrast with the background (n=1). In the Medium group (Group 1), one element was presented at high contrast, and one at intermediate contrast (n=2). Finally, in the Variable group (Group 2), one element was presented with high contrast, and the other presented variably. The range of contrasts was informed by a pilot assay, in which an intensity-accuracy curve was generated for comparable stimuli. The mean log contrast in Group 2 was matched to the intermediate contrast in Group 1. The training schedule was sessions of 12 trials at 2 sessions per day, with trial order following a pseudo-random schedule generated using Mix [55].

In previous experiments, I employed a simple accuracy criterion to determine when a subject had acquired the task (typically 75% correct over 3 consecutive sessions). However multiple subjects also spontaneously developed positional biases which this criterion did not account for. As such, I applied a more robust set of criteria here to account for this possibility, considered over four consecutive sessions:

1. Accuracy equal to or greater than 75% correct (i.e., ≥36/48 trials correct)
2. Positional bias (single position) not exceeding 50% (i.e., ≤24/48 chose same position)
3. Directional bias (vertical or horizontal) not exceeding 75% (i.e., ≤36/48 chose same side)

Once a subject satisfied these criteria, it would be moved to the testing phase. Subjects that failed to reach these criteria within 40 sessions were removed from the experiment.

## Testing

Once the criteria were reached, subjects would have been presented with a series of probe trials to determine the extent to which they had attended to and learned different elements of the stimuli. To minimise further learning biasing the results, probe trials would have been unreinforced. These probe trials would have been interspersed with an equal number of training trials to prevent task extinction. Two batteries of testing sessions would have been carried out in sequence. The first battery would have tested the relative ability of subjects to recognise and discriminate individual elements of the stimuli. The second battery would have investigated whether subjects could generalise the learned discrimination rules across different shapes (figure 4). In both cases the training schedule was sessions of 12 trials at 2 sessions per day, with trial order following a pseudo-random schedule generated using Mix.

In the conflict test (Ab/aB), there is not an objectively “correct” choice. As such, the 4-AFC task featuring a single CS+ and three distractors was inappropriate – subjects potentially could have used simple heuristics e.g., an odd-one-out to solve the task. Instead, a 2-AFC

task was designed, with modified positions. Sessions would have alternated either left/right or up/down.

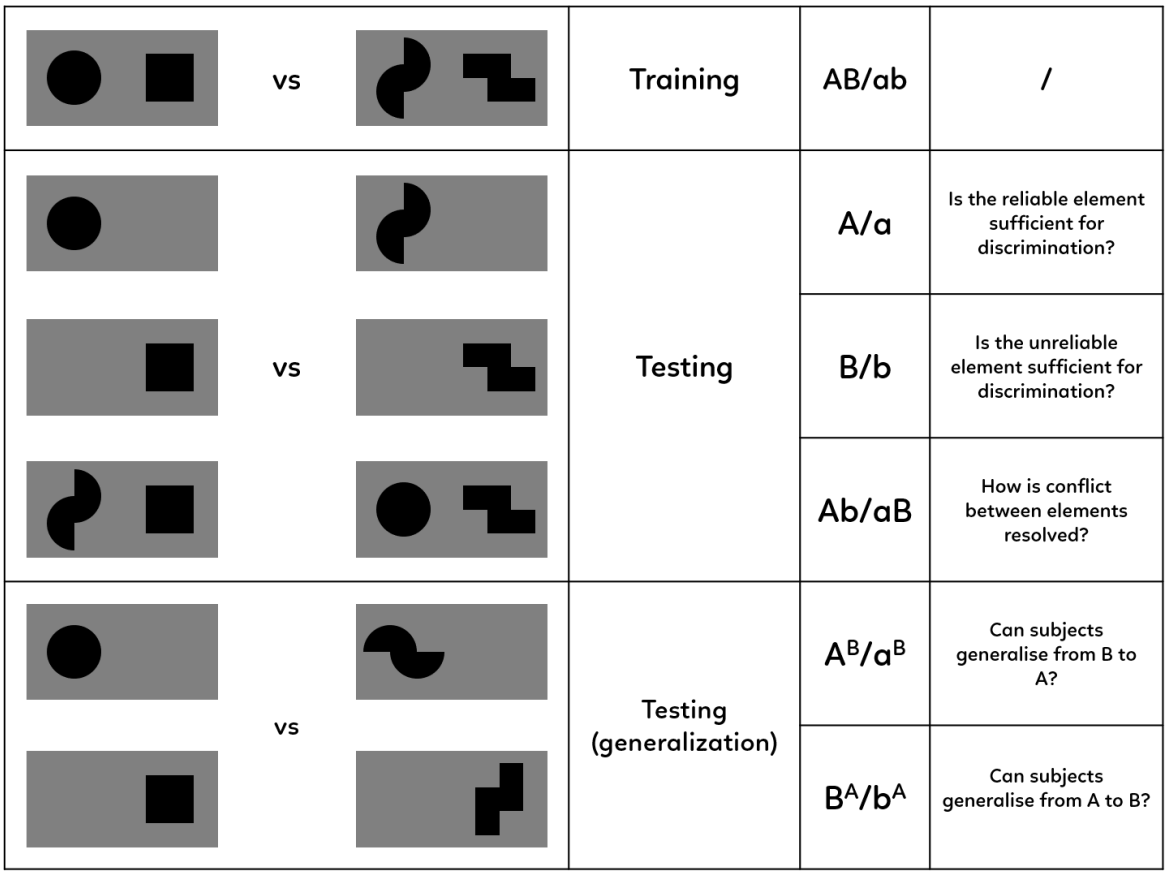



*Figure 4:* Illustration of the stimulus pairings employed during each phase of the experiment, with the interpretation of each discrimination outcome.

## Managing misbehaviour

A rolling set of criteria were planned to assess the development of positional biases, and to provide a formalised and consistent method of handling such deviations from expected behaviour. If a subject failed to leave the start box within 300 seconds, the training session would be suspended and resumed on the same trial at the next available training slot. If a subject left the start box but failed to make a detectable choice within 300 seconds, the trial would be marked "FAIL", and training continued immediately on the next trial. Trials scored "FAIL" would not be included in subsequent data analyses. If subject performance fell below certain criteria during the testing phase; then data collection would be

terminated, the affected session discontinued, and the subject returned to the training phase. Once the subject once again reached the learning criteria, testing would be resumed as previously. The criteria for being removed from testing were as follows:

1. Accuracy (R trials) ≤58% (i.e., ≤7/12 trials correct)
2. Positional bias (single position, all trials) ≥66% (i.e., ≥16/24 trials chose same position)
3. Directional bias (coordinate, all trials) ≥83% (i.e., ≥20/24 trials chose same coordinate)
4. Trials scored "FAIL" ≥12.5% (i.e., ≥3/24 trials incomplete)

## Data analysis

In this section I lay out the general scheme for data analysis that was planned, had the experiment been completed. Planned statistical methods are included in bullet points in each section.

### 1.0 Were the core assumptions of the experiment valid?

This experiment aimed to examine how triggerfish learn to use visual cues which vary in quality. However, in order the experiment to be valid, certain core assumptions must be verified. Firstly, did the population exhibit a strong preference for one or the other cue? If subject attention was driven by some inherent property of a cue, this would complicate analysis (though not invalidate the results, as the experiment was counterbalanced). Secondly, did the quality of a cue affect accuracy? If reducing cue quality did not impair subject performance, then subjects would have no basis for preference. Expanding on this, was mean accuracy matched in the Med and Var treatments?

#### *1.1 Was there a population-level preference for either cue?*

If most subjects exhibited a preference for either the Sq or Ci cues, this would invalidate the comparison.

- 2-sample, 2-sided equality test: compare accuracy across entire population in P1a and P1b trials

#### *1.2 Did cue quality predict subject accuracy?*

The experiment was predicated on the model that low-quality cues impair discrimination.

- Linear regression of accuracy x trial type

*1.3 Was the range of stimulus intensities used in the Var treatment appropriate?*

The range of intensities used in the Var group was informed by a pilot assay and intended that the Var and Med treatments should have matched accuracy in order to compare preferences for cues that are consistent versus variable cues.

- 2-sample, 2-sided equality test: compare accuracy in R trials between Var and Med training conditions

## 2.0 How did treatment affect how subjects learned and used cues?

*2.1 Did treatment affect how quickly subjects learned the task?*

- Graphical representation as Kaplan-Meier curves
- Survival Analysis, Cox-Mantel: Did the # trials to differ between treatments?

*2.2 Did treatment affect overall accuracy once subjects had learned the task?*

- ANOVA / Kruskal-Wallis test comparing mean accuracy of R trials between treatments

## 3.0 Relative learning of each cue

*3.1 Did subjects employ an elemental or configural approach to learning the stimuli?*

In discussing the design of this experiment, I have implicitly assumed that subjects would employ an elemental associative structure in learning the contingencies of the task. That is to say, they would form associations between each cue (the conditioned stimuli) and the unconditioned stimulus (i.e., food) separately. An alternative is the configural model, in which the two stimuli are associatively linked by a third and independent configural representation as they co-occur [56]. In this model, subjects conditioned only to the compound stimulus would principally have learned the configural representation, and so would perform much more poorly when presented with either cue in isolation

- 2-sample, 2-sided equality test: Does accuracy in single-cue trials differ from accuracy in R trials

### *3.2 Did subjects more effectively learn the reliable cue?*

3.2a When presented with cues in isolation, was accuracy higher with the high-quality cue?

- 2-sample, 2-sided equality test: Did accuracy in P1a trials differ from accuracy from P1b trials

3.2b When the cues were presented in conflict, which cue did subjects prefer?

- 1-sample, 2-sided equality test: Did the ratio of Ab: aB choices in P2 trials differ significantly from 1:1

3.2c If differences were present, did the magnitude differ between treatment groups?

## 4.0 Were subjects able to generalize learning rules across shapes?

Subjects were presented with two cues, consisting of the difference in the vertical displacement of a circle, and horizontal displacement of a square. Assuming subjects can learn this discrimination, the obvious question follows whether they can generalize the learning rules across shapes – e.g., horizontal displacement of a circle. Furthermore, if subjects can generalize rules across shapes, is this affected by conditioning history with the cue?

### *4.1 Could subjects successfully complete generalization trials?*

Were there differences in generalization ability between treatment groups?

- 1-sample, 2-sided equality test: did population accuracy in P3 trials differ from chance (25%)?

### *4.2 Was generalization symmetric?*

Did performance differ between $A^B/a^B$ and $B^A/b^A$ trials?

- 2-sample, 2-sided equality test, did average accuracy differ between $A^B/a^B$ and $B^A/b^A$ trials?

## Results

Five subjects were initially enrolled in the experiment pilot. Two subjects consistently failed to complete the pre-training procedure, as they could not consistently complete 10 trials of moving between start box and choice chamber. Three subjects completed pre-training and were assigned to training conditions.

These subjects failed to reach the pre-assigned learning criterion within the allotted period (40 sessions) (figure 6). As a result of this, several sequential modifications were made to the procedure to improve subject performance. Initially, to address motivation as a possible confounding factor, the reward was changed to a preferred food (brine shrimp). Concurrently, the magnitude of each cue was increased to overcome the possibility that the resolution of the subject's visual system was a limiting factor. The degree of offset in the figures was increased from 50% of their diameter to 100% of their diameter (figure 5).

As previously, only three of the five initial subjects successfully completed pre-training (same subjects). Of these subjects, none substantially improved their performance over the course of 14 sessions (figure 7). Considering this, I hypothesized that the number of alternative choices in the task, coupled with the complexity of the individual figures; might make the task too cognitively demanding for the subjects. To test this, a transitional stage of training was introduced in which only one distractor was present (2-AFC). This experiment was terminated after 24 sessions – at this point, there was no sign of improving performance beyond chance, and the allotted time to carry out the experiments had been exhausted (figure 8).

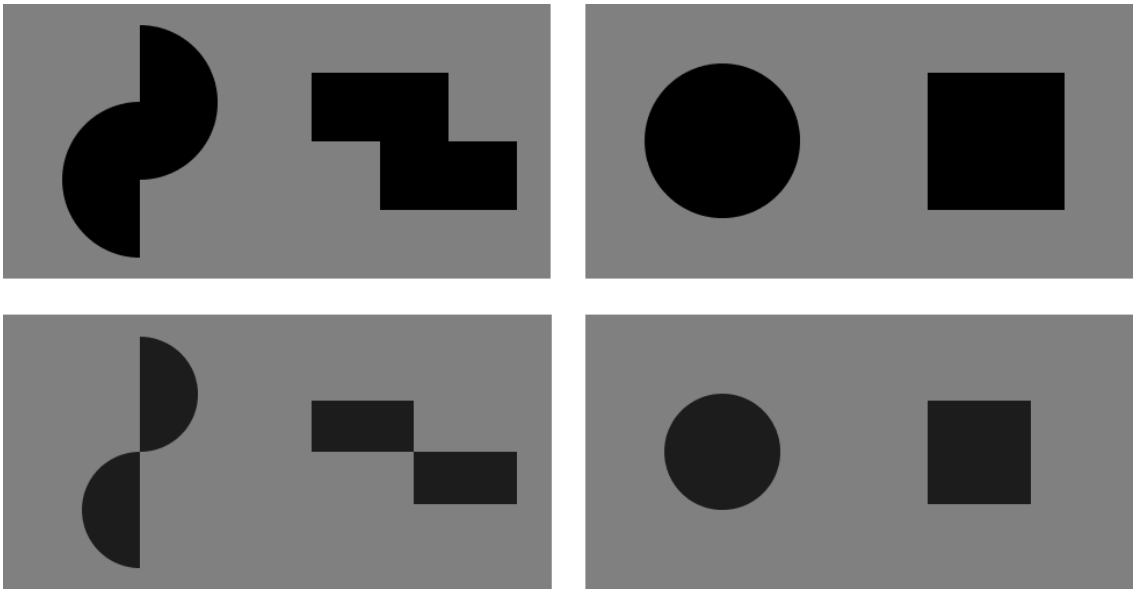


*Figure 5:* Original stimulus pairing (top) and modified stimulus pairing (bottom). In the modified design, the elements are offset by their full diameter.

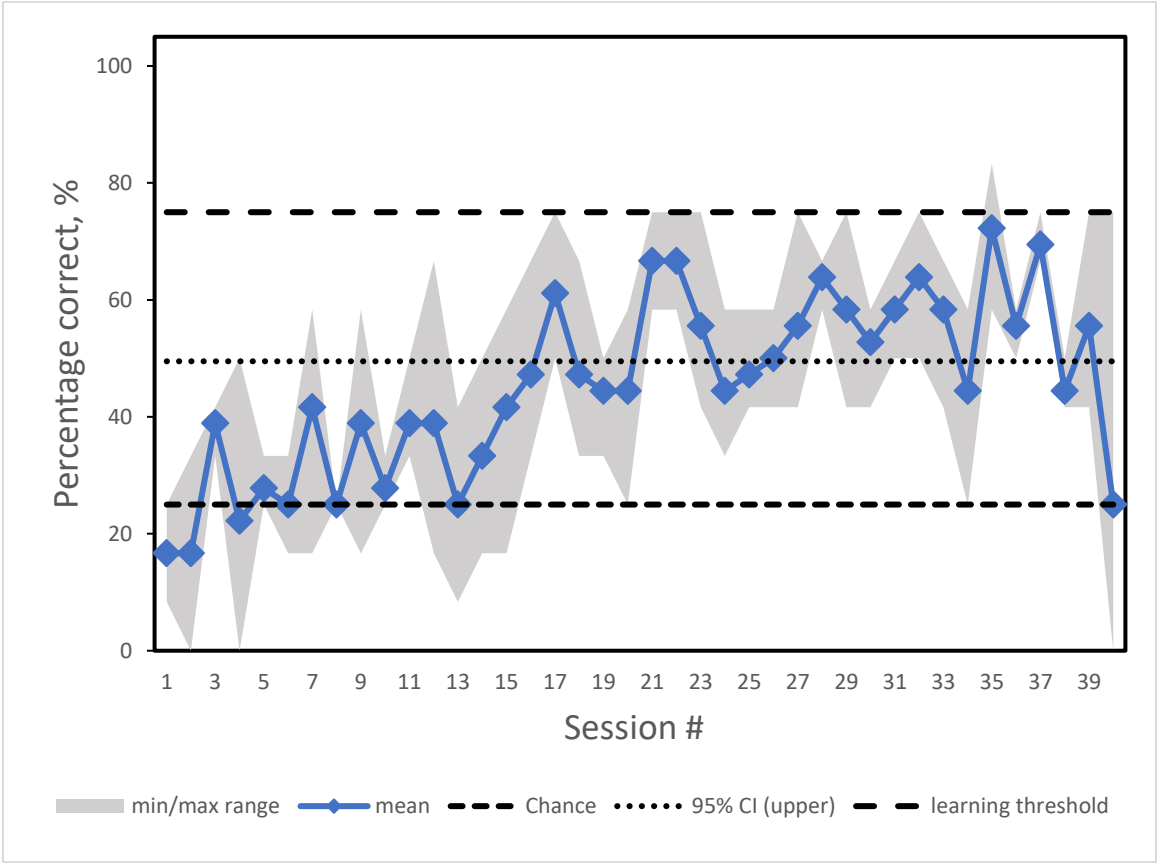



*Figure 6:* Mean subject discrimination performance during initial training, original formulation, n=3. Shaded area indicates maximum and minimum subject accuracy.

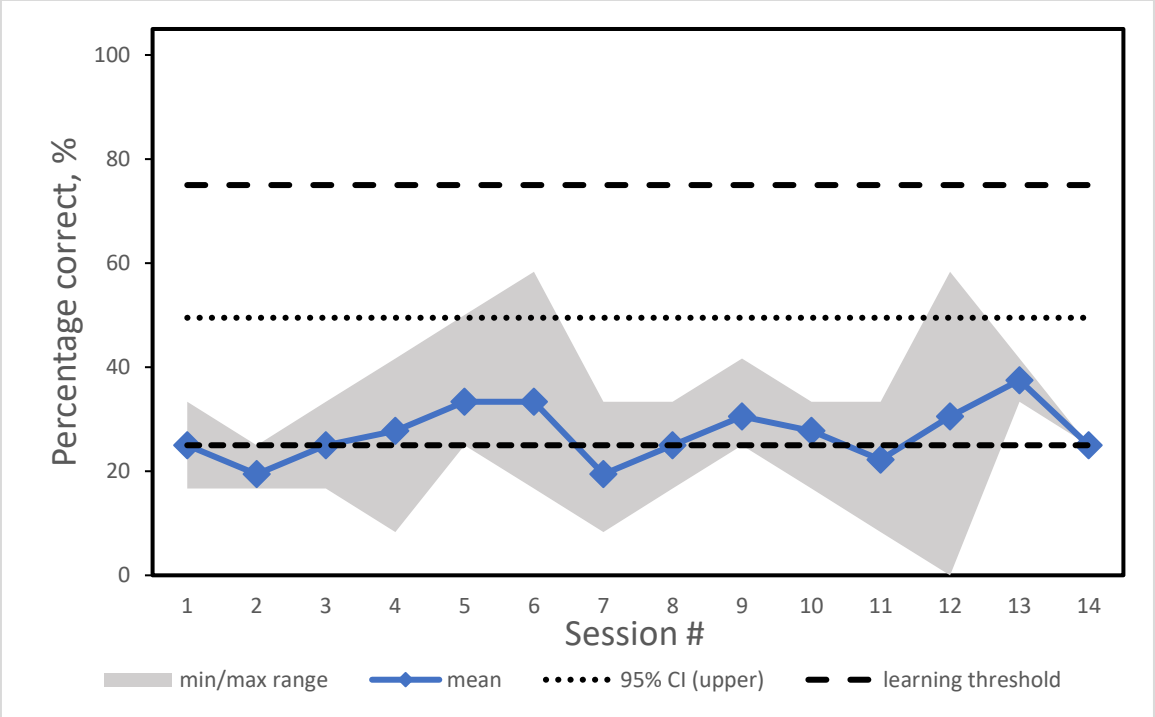



*Figure 7:* Mean subject discrimination performance during initial training, second formulation (increase in stimulus and reinforcer magnitude), n=3. Shaded area indicates maximum and minimum subject accuracy.

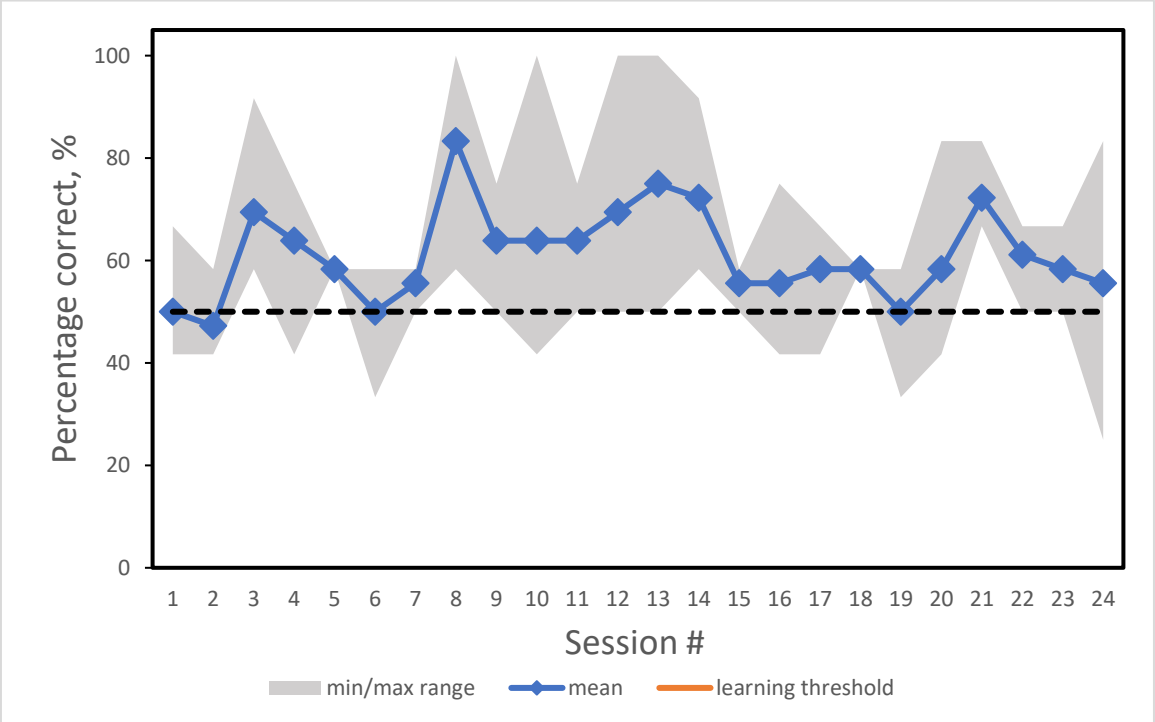



*Figure 8:* Mean subject discrimination performance during initial training, third formulation (2AFC), n=3. Shaded area indicates maximum and minimum subject accuracy.

## Discussion

In this chapter, I set out to examine how triggerfish might learn different signals in an environment where changing signal transmission results in different signals being reliable at different times. To test this, I set out an experiment in which triggerfish were conditioned to recognise an object on the basis of two distinct cues (signals), in which the reliability of either cue was manipulated independently. During training, subjects experienced either all cues at high intensity (control), one cue at high intensity and one at medium intensity (intermediate group) or one cue at high intensity and one at variable intensity. Having been conditioned to this task, subjects were to be presented with a series of deconstructed stimuli featuring the cues singly, transposed or placed in opposition; to determine how cue quality affected how the cue was learned and used. I hypothesized that subjects would develop stronger associations with the consistently high-quality cue, and that this effect would be larger in magnitude when the less-reliable cue was also variable. The probes were designed such that any response subjects produced would be informative. However, I did not anticipate that subjects would fail to acquire the basic object discrimination task.

Triggerfish are visually guided predators. Previous research has demonstrated that their visual system is capable of fine-scale spatial resolution (to ~1-2°) [57], and are capable of discriminating both chromatic and lightness differences [52], [58]. Furthermore, triggerfish can learn to attend to separate elements of a stimulus such as colour and pattern in a similarly constructed experiment [59]. Finally, Chapter 2 of this thesis robustly established that triggerfish are quite sensitive to lightness differences in achromatic stimuli. Triggerfish are also readily conditioned to simple discrimination tasks with both 2 and 4 stimuli. Previous research with triggerfish has not typically presented data on task acquisition [57]–[61]. In an exception, Green *et al.* [62] report results of conditioning *R. aculeatus* to discriminate pattern shapes in a 2AFC task, using the same learning criterion I applied in

this chapter. Green *et al.* found that 50/57 subjects learned the discrimination within 30 sessions (120 trials), with the experimental population exhibiting a typical learning curve. These observations provided a baseline expectation of success for this experiment. I found that the initial experimental population (n=5) universally failed to reach a pre-assigned learning criterion within the allocated time (40 sessions), and the experiment could not be concluded. This outcome contrasts with previous experiments, in which triggerfish typically acquired simple discrimination tasks in a 4AFC framework within 11 sessions.

As a result of this failure, several modifications were made to the experimental procedure to improve subject performance. In the first instance, two approaches were used. To address possible motivational issues, the food reward (pellet food) was changed to a preferred food (krill). This was a factor in re-designing the experiments discussed in Chapter 2 and succeeded in that context. Also, I considered that – while significantly larger than the species' established spatial resolution at the distances observed – the differences between the CS+ and distractors might be too slight for the subjects to perceive. As such, the stimuli were modified such that the magnitude of offset in the distractors was doubled, increasing the magnitude of the cue. When these modifications were found to be ineffective, I considered that, given the complexity of the stimuli, the number of alternative choices (four) might have an effect of "informational overload". To compensate for this, a transitional stage of training was introduced in which only one distractor was present (2-AFC). This approach also failed, with subjects failing to improve their performance beyond chance. At this juncture, experiments had to be concluded as the allotted time set aside for the experiments had concluded.

Given the structural similarity to the experiments outlined in Chapter 2, I estimated that data collection should take a similar amount of time. The bulk of data collection for Chapter 2 (12 subjects across two experiments) was carried out in two months (February –

March 2017). This includes experiment 2, in which subjects were also anomalously slow to learn, but which nevertheless only increased the duration of data collection by two weeks. Given this, the entire period of October – December 2018 (3 months) was set aside for data collection, with the option to extend through March 2019 in the event of unforeseen circumstances. Nevertheless, this time was insufficient in which to complete data collection. This anomalous finding begs the question of why subjects failed to acquire the task.

## Health

In between data collection for Chapter 2 and the experiments described here several members of the experimental population began to exhibit a suite of behavioural symptoms suggesting pathogen infection. This was initially characterised by behavioural quirks and a loss of physical coordination, followed by lassitude, loss of appetite, and finally death. This has subsequently been diagnosed as a previously unknown parasitic infection affecting the brain, as diagnosed by an Oxford University veterinarian following investigation. The experiments discussed in this chapter were carried out in a clean laboratory with quarantined subjects, and no diagnostic symptoms of the infection were observed. Nevertheless, the infection was present in other adjacent laboratories, and it is also possible the subjects were infected in the wild, or during handling by suppliers. The loss of motivation observed in the late stages of the experiments could interpreted as symptoms of this disease. Consequently, it is possible that the subjects experienced a pathogen-related cognitive impairment that impacted their behaviour.

## Sensory Factors

One possibility is that the cue differences were beyond the subjects' visual capabilities to resolve. *R. aculeatus'* maximum sensory resolution has been measured at approximately 1° [57]. In this experiment, the magnitude offset in the revised stimuli was equal to the element diameter, 1.5-1.7cm (as the Sq and Ci elements were matched by area, not

diameter). Subjects would initially perceive the stimuli at ~35cm upon opening of the door and would typically make a secondary inspection at a range of 2-10cm. At these ranges, the elements would correspond to 2.5° in the former instance, and 8.6-41.1° in the latter. Thus, visual resolution was unlikely to be a limiting factor.

## Cognitive Factors

### Geometrical factors

The CS+/- employed in the first version of the psychophysical assay consisted of mirror-images pairings. Subjects were extremely slow to acquire this task. This is surprising, as numerous other basal vertebrates have been shown to be able to learn mirror-image discriminations – for example, bamboo sharks (*Chiloscyllium griseum*) and Malawi cichlids (*Pseudotropheus* sp.) [70], [72] – as well as discriminate and recognise 3D rotated objects [70], [73]. Consequently, the stimuli were modified to differ in the vertical or horizontal translation of an element (figure 9). While acquisition was still slow compared to previous experiments (mean 31.7 sessions), subjects were able to learn this discrimination. While the element pairings used in the first experiment iteration did not possess rotational or mirror-image ambiguity, it is possible that subjects preferentially relied upon other rules (such as area or hue) for discrimination.

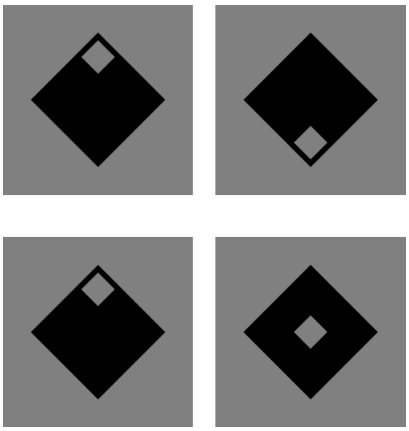


*Figure 9:* Top row: original stimulus pairing used in assay. Bottom row: modified stimuli

The initial failure of subjects to discriminate mirror-image stimuli presents another surprising result, given that related taxa are capable of it. One possibility is this phenomenon is a consequence of their movement ecology. While most fish species are capable of free movement within the 3D aquatic environment, each species has evolved distinct movement and motor ecologies, affecting

their interactions with their environment. Demersal, benthic, and surface-bound species will necessarily interact with their spatial environment differently to pelagic species.

For example, Davis *et al.* [74] found that, while bronze cory catfish (*Corydoras aeneus*) were challenging to condition to move vertically in a y-maze task, they subsequently encoded vertical information more accurately than horizontal information. Triggerfish not only inhabit environments with extensive 3D structure, but their balistiform locomotion affords exceptional freedom of movement, and they regularly approach objects from a range of angles (personal observation). It may be that the ability to distinguish object rotations in a foraging context is not advantageous for this species.

### Eco-cognitive factors

It is possible that the highly abstracted form of the experiment impeded subjects' ability to engage with the task. Animals often exhibit remarkable cognitive abilities only within certain narrow ecological or behavioural contexts. For example, when presented with a number discrimination task guppies exhibited abilities superior to many mammals if the task mimicked their natural foraging activities [63], but not when a more standardised format was employed [64]. This phenomenon of selective association is typically presumed to be an evolved feature of cognition to favour biologically significant associations. Comparably, inanga fish (*Galaxias maculatus*) were able to associate a specific location and time with a food reward, but were not able to associate predation risk with a time-location in a similar experiment [65]. However, since this experiment engages the same sensory and behavioural mechanisms as previous experiments carried out both by me and other researchers, this factor seems unlikely to be significant.

### In-attentional blindness

I included a red star – the "appetitive element" – in each stimulus to provide a target for subjects to bite at (see method for details). This was included to prevent potential self-

reinforcing biases occurring due to subjects selectively biting one or another element of the stimulus. While subjects did consistently direct their responses towards this element, it is possible that they treated this as the most salient feature of the stimulus and ignored the other elements. This idea could be identified with the concept on inattentional blindness [66], in which tasks that monopolise attention can prevent perception of other environmental cues. While widely-documented in humans, to my knowledge, only a single paper has been published on this phenomenon in a non-human species, and that only last year [67]. Gabor *et al.* demonstrated that horses conditioned to approach a food reward exhibited reduced startle responses to a novel stimulus than did a control group.

Similarly, it is possible that the presence of a colour associated with food in triggerfish had a similar effect. Something analogous in concept may be rendered in an associative learning framework, similar to Kamin's [71] associative blocking effect. In short, it is possible that the appetitive element was sufficiently distracting that it monopolised attention away from the salient features of the task. In terms of the Rescorla-Wagner model, this would correspond to the relative salience and quality of the appetitive element ($\alpha$ and $\lambda$) sufficiently outweighing those of the intended cues that no increase in associative value could occur. If this is correct, then to proceed the salience of the appetitive element would need to be reduced (for example, by choosing a different colour) or be removed entirely. However, the latter would require an alternative solution to the problem of self-reinforcing positional biases. One possibility would be to change the mechanism by which subjects indicated a choice. A commonly used modality with fish is association; requiring a subject to spend a certain amount of time near one or another stimulus (e.g. [68]–[70]). However, the current approach has the advantage that it engages triggerfish's natural haptic exploration of novel objects.

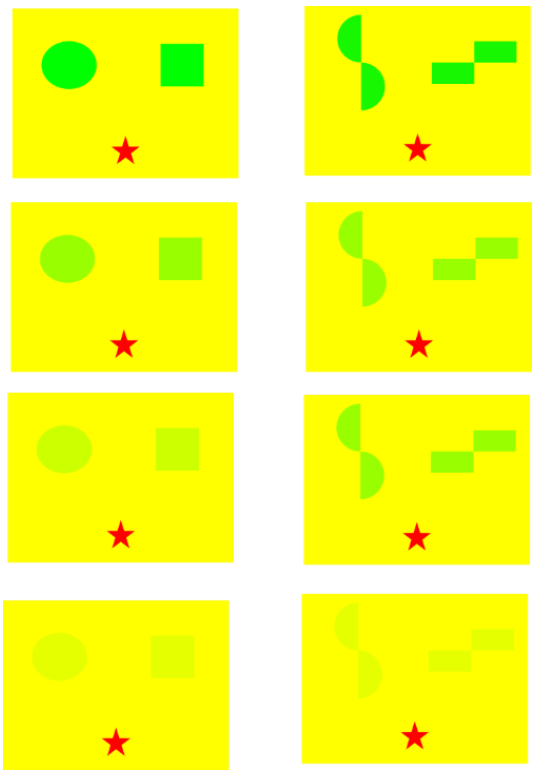


*Figure 10:* Potential modified stimulus pairings, in which contrast between figures and background is based on chroma rather than luminance.

### Cognitive loading

It is possible that the large number of elements and comparably large number of shared features between stimuli simply made the act of discrimination too cognitively taxing for the subjects to complete. This is argued against to some degree by the successful completion of the psychophysical assay by a separate population of triggerfish – the assay also employed a 4AFC procedure and used stimuli that were comparable to individual elements in the final experiment. Furthermore, when the full task was reduced to a 2AFC protocol, subjects showed no improvement in performance. Nevertheless, this could be tested in a modified task in which subjects were conditioned to discriminate simplified stimuli featuring only one or the other element.

### Sensory bias

A recently-published paper [58] found that this species performs poorly at identifying shapes that are defined purely by achromatic contrast (i.e. being lighter or darker than the background, but not differing in chroma). By comparison, triggerfish were adept at identifying objects defined by chromatic contrast. I judge that of the possibilities discussed above, this is the most likely to underly my findings, being demonstrated in the same species in an extremely similar experimental context. A test of this hypothesis would be to replicate the original stimulus pairings, but with element/background contrast defined by hue rather than luminance. A possible mock-up of this is shown in figure 10.

## Concluding remarks

While I was unable to address the original questions under consideration, the experiments described in this chapter present a fascinating contrast to previous results. Coupled with my previous findings, the data presented here allowed me to analyse various conditions which triggerfish were able to be conditioned to different stimuli. The failure of subjects to acquire the discrimination tasks presented represents a significant deviation from previous work in this species and may yet shed light on as yet unrecognised aspects of triggerfish cognition. The Picasso triggerfish is highly appealing model species for examining cognitive and sensory abilities in basal vertebrates, for reasons outlined in Chapter 2. However, the results presented here emphasize that future research should take into consideration nuances of both sensory tuning and of ecology in doing so.

## References

[1] S. R. X. Dall, L.-A. Giraldeau, O. Olsson, J. M. McNamara, and D. W. Stephens, "Information and its use by animals in evolutionary ecology," *Trends Ecol. Evol.*, vol. 20, no. 4, pp. 187–193, 2005.

[2] D. C. Knill and A. Pouget, "The Bayesian brain: The role of uncertainty in neural coding and computation," *Trends Neurosci.*, vol. 27, no. 12, pp. 712–719, 2004.

[3] R. Dukas, "Behavioural and ecological consequences of limited attention," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 357, no. 1427. pp. 1539–1547, 2002.

[4] J. A. C. Uy and A. C. Stein, "Variable visual habitats may influence the spread of colourful plumage across an avian hybrid zone," *J. Evol. Biol.*, vol. 20, no. 5, pp. 1847–1858, Sep. 2007.

[5] J. A. Endler, "Signals, signal conditions, and the direction of evolution," *Am. Nat.*, vol. 139, no. Suppl., Oct. 1992.

[6] G. L. Cole, "Lost in Translation: Adaptation of Mating Signals in Changing Environments," *Springer Sci. Rev.*, vol. 1, no. 1–2, pp. 25–40, Dec. 2013.

[7] C. J. Lorenzen, "Extinction of light in the ocean by phytoplankton," *ICES J. Mar. Sci.*, vol. 34, no. 2, pp. 262–267, Mar. 1972.

[8] M. Jönsson, L. Ranåker, P. A. Nilsson, and C. Brönmark, "Foraging efficiency and prey selectivity in a visual predator: differential effects of turbid and humic water," *Can. J. Fish. Aquat. Sci.*, vol. 70, no. 12, pp. 1685–1690, Dec. 2013.

[9] M. P. F. Costa, E. M. L. M. Novo, and K. H. Telmer, "Spatial and temporal variability of light attenuation in large rivers of the Amazon," *Hydrobiologia*, vol. 702, no. 1, pp. 171–190, Oct. 2013.

[10] J. M. Hinshaw, "Effects of illumination and prey contrast on survival and growth of larval perch *Perca flavescens*," *Trans. Am. Fish. Soc.*, vol. 114, no. 4, pp. 540–545, Jul. 1985.

[11] A. C. Utne-Palm, "Visual feeding of fish in a turbid environment: Physical and behavioural aspects," *Mar. Freshw. Behav. Physiol.*, vol. 35, no. 1–2, pp. 111–128, Jan. 2002.

[12] G. L. Cole and J. A. Endler, "Variable Environmental Effects on a Multicomponent Sexually Selected Trait," *Am. Nat.*, vol. 185, no. 4, pp. 452–468, Apr. 2015.

[13] J. R. Morrongiello, N. R. Bond, D. A. Crook, and B. B. M. Wong, "Nuptial coloration varies with ambient light environment in a freshwater fish," *J. Evol. Biol.*, vol. 23, no. 12, pp. 2718–2725, Dec. 2010.

[14] S. T. Giery, C. A. Layman, and R. B. Langerhans, "Anthropogenic ecosystem fragmentation drives shared and unique patterns of sexual signal divergence among three species of Bahamian mosquitofish," *Evol. Appl.*, vol. 8, no. 7, pp. 679–691, Aug. 2015.

[15] K. E. Körner, I. Schlupp, M. Plath, and E. R. Loew, "Spectral sensitivity of mollies: Comparing surface- and cave-dwelling Atlantic mollies, Poecilia mexicana," *J. Fish Biol.*, vol. 69, no. 1, pp. 54–65, Jul. 2006.

[16] K. U. Heubel and I. Schlupp, "Turbidity affects association behaviour in male Poecilia latipinna," *J. Fish Biol.*, vol. 68, no. 2, pp. 555–568, Feb. 2006.

[17] C. D. Brock, M. E. Cummings, and D. I. Bolnick, "Phenotypic plasticity drives a depth gradient in male conspicuousness in threespine stickleback, Gasterosteus aculeatus," *Evolution (N. Y).*, vol. 71, no. 8, pp. 2022–2036, Aug. 2017.

[18] R. F. Castillo Cajas, O. M. Selz, E. A. P. Ripmeester, O. Seehausen, and M. E. Maan, "Species-Specific Relationships between Water Transparency and Male Coloration within and between Two Closely Related Lake Victoria Cichlid Species," *Int. J. Evol. Biol.*, vol. 2012, pp. 1–12, Jul. 2012.

[19] C. G. McDonald and C. W. Hawryshyn, "Intraspecific variation of spectral sensitivity in threespine stickleback (Gasterosteus aculeatus) from different photic regimes," *J. Comp. Physiol. A*, vol. 176, no. 2, pp. 255–260, 1995.

[20] Y. Sakai, H. Ohtsuki, S. Kasagi, S. Kawamura, and M. Kawata, "Effects of light environment during growth on the expression of cone opsin genes and behavioral spectral sensitivities in guppies (Poecilia reticulata)," *BMC Evol. Biol.*, vol. 16, no. 1, p. 106, Dec. 2016.

[21] R. C. Fuller, K. L. Carleton, J. M. Fadool, T. C. Spady, and J. Travis, "Genetic and environmental variation in the visual properties of bluefin killifish, Lucania goodei," *J. Evol. Biol.*, vol. 18, no. 3, pp. 516–523, Mar. 2005.

[22] D. S. Wright, E. Rietveld, and M. E. Maan, "Developmental effects of environmental light on male nuptial coloration in Lake Victoria cichlid fish," *PeerJ*, vol. 6, p. e4209, Jan. 2018.

[23] M. A. W. Hornsby, S. Sabbah, R. M. Robertson, and C. W. Hawryshyn, "Modulation of environmental light alters reception and production of visual signals in Nile tilapia," *J. Exp. Biol.*, vol. 216, no. 16, pp. 3110–3122, Aug. 2013.

[24] D. P. Chivers, F. Al-Batati, G. E. Brown, and M. C. O. Ferrari, "The effect of turbidity on recognition and generalization of predators and non-predators in aquatic ecosystems," *Ecol. Evol.*, vol. 3, no. 2, pp. 268–277, Feb. 2013.

[25] L. Jacobsen, S. Berg, H. Baktoft, and C. Skov, "Behavioural strategy of large perch Perca fluviatilis varies between a mesotrophic and a hypereutrophic lake," *J. Fish Biol.*, vol. 86, no. 3, pp. 1016–1029, Mar. 2015.

[26] J. L. Kelley, B. Phillips, G. H. Cummins, and J. Shand, "Changes in the visual environment affect colour signal brightness and shoaling behaviour in a freshwater fish," *Anim. Behav.*, vol. 83, no. 3, pp. 783–791, Mar. 2012.

[27] J. T. O. Kirk, "The vertical attenuation of irradiance as a function of the optical properties of the water," *Limnol. Oceanogr.*, vol. 48, no. 1, pp. 9–17, Jan. 2003.

[28] J. T. O. Kirk, "Effects of suspensoids (turbidity) on penetration of solar radiation in aquatic ecosystems," *Hydrobiologia*, vol. 125, no. 1, pp. 195–208, Jun. 1985.

[29] B. S. Halpern *et al.*, "A global map of human impact on marine ecosystems," *Science (80-. ).*, vol. 319, no. 5865, pp. 948–952, Feb. 2008.

[30] V. H. Smith, G. D. Tilman, and J. C. Nekola, "Eutrophication: Impacts of excess nutrient inputs on freshwater, marine, and terrestrial ecosystems," in *Environmental Pollution*, 1998, vol. 100, no. 1–3, pp. 179–196.

[31] M. S. Islam and M. Tanaka, "Impacts of pollution on coastal and marine ecosystems including coastal and marine fisheries and approach for management: A review and synthesis," *Marine Pollution Bulletin*, vol. 48, no. 7–8. Pergamon, pp. 624–649, 01-Apr-2004.

[32] W. F. Henley, M. A. Patterson, R. J. Neves, and A. D. Lemly, "Effects of Sedimentation and Turbidity on Lotic Food Webs: A Concise Review for Natural Resource Managers," *Rev. Fish. Sci.*, vol. 8, no. 2, pp. 125–139, Apr. 2000.

[33] T. Linders, P. Nilsson, A. Wikström, and M. Sköld, "Distribution and fate of trawling-induced suspension of sediments in a marine protected area," *ICES J. Mar. Sci.*, vol. 75, no. 2, pp. 785–795, Mar. 2018.

[34] S. P. Collin and N. S. Hart, "Vision and photoentrainment in fishes: The effects of natural and anthropogenic perturbation," *Integrative Zoology*, vol. 10, no. 1. Wiley-Blackwell, pp. 15–28, 01-Jan-2015.

[35] A. C. Utne-Palm, "The effect of prey mobility, prey contrast, turbidity and spectral composition on the reaction distance of Gobiusculus flavescens to its planktonic prey," *J. Fish Biol.*, vol. 54, no. 6, pp. 1244–1258, Jun. 1999.

[36] A. De Robertis, C. H. Ryer, A. Veloza, and R. D. Brodeur, "Differential effects of turbidity on prey consumption of piscivorous and planktivorous fish," *Can. J. Fish. Aquat. Sci.*, vol. 60, no. 12, pp. 1517–1526, Dec. 2003.

[37] J. C. Barrett, G. D. Grossman, and J. Rosenfeld, "Turbidity-Induced Changes in Reactive Distance of Rainbow Trout," *Trans. Am. Fish. Soc.*, vol. 121, no. 4, pp. 437–443, Jul. 1992.

[38] L. Eaton and K. A. Sloman, "Subordinate brown trout exaggerate social signalling in turbid conditions," *Anim. Behav.*, vol. 81, no. 3, pp. 603–608, Mar. 2011.

[39] A. K. Schwartz and A. P. Hendry, "Testing the influence of local forest canopy clearing on phenotypic variation in Trinidadian guppies," *Funct. Ecol.*, vol. 24, no. 2, pp. 354–364, Feb. 2010.

[40] M. Michelangeli, U. Tuomainen, U. Candolin, and B. B. M. Wong, "Habitat alteration influences male signalling effort in the Australian desert goby," *Behav. Ecol.*, vol. 26, no. 4, pp. 1164–1169, 2015.

[41] J. Sundin, T. Aronsen, G. Rosenqvist, and A. Berglund, "Sex in murky waters: algal-induced turbidity increases sexual selection in pipefish," *Behav. Ecol. Sociobiol.*, vol. 71, no. 5, p. 78, May 2017.

[42] M. Järvenpää and K. Lindström, "Water turbidity by algal blooms causes mating system breakdown in a shallow-water fish, the sand goby Pomatoschistus minutus," *Proc. R. Soc. B Biol. Sci.*, vol. 271, no. 1555, pp. 2361–2365, Nov. 2004.

[43] U. Candolin, T. Salesto, and M. Evers, "Changed environmental conditions weaken sexual selection in sticklebacks," *J. Evol. Biol.*, vol. 20, no. 1, pp. 233–239, 2007.

[44] J. Engström-Öst and U. Candolin, "Human-induced water turbidity alters selection on sexual displays in sticklebacks," *Behav. Ecol.*, vol. 18, no. 2, pp. 393–398, Mar. 2007.

[45] B. B. M. Wong, U. Candolin, and K. Lindström, "Environmental Deterioration Compromises Socially Enforced Signals of Male Quality in Three-Spined Sticklebacks," *Am. Nat.*, vol. 170, no. 2, pp. 184–189, Aug. 2007.

[46] O. Seehausen, J. J. M. Van Alphen, and F. Witte, "Cichlid Fish Diversity Threatened by Eutrophication that Curbs Sexual Selection," *Science (80-. ).*, vol. 277, no. 5333, pp. 1808–1811, 1997.

[47] M. E. Maan, O. Seehausen, and J. J. M. Van Alphen, "Female mating preferences and male coloration covary with water transparency in a Lake Victoria cichlid fish," *Biol. J. Linn. Soc.*, vol. 99, no. 2, pp. 398–406, Jan. 2010.

[48] F. Witte, O. Seehausen, J. H. Wanink, M. A. Kishe-Machumu, M. Rensing, and T. Goldschmidt, "Cichlid species diversity in naturally and anthropogenically turbid habitats of Lake Victoria, East Africa," *Aquat. Sci.*, vol. 75, no. 2, pp. 169–183, 2013.

[49] J. Engström-Öst and J. Mattila, "Foraging, growth and habitat choice in turbid water: An experimental study with fish larvae in the Baltic Sea," *Mar. Ecol. Prog. Ser.*, vol. 359, pp. 275–281, May 2008.

[50] M. Hasenbein, L. M. Komoroske, R. E. Connon, J. Geist, and N. A. Fangue, "Turbidity and salinity affect feeding performance and physiological stress in the endangered delta smelt," *Integr. Comp. Biol.*, vol. 53, no. 4, pp. 620–634, Oct. 2013.

[51] J. W. Sigler, T. C. Bjornn, and F. H. Everest, "Effects of Chronic Turbidity on Density and Growth of Steelheads and Coho Salmon," *Trans. Am. Fish. Soc.*, vol. 113, no. 2, pp. 142–150, Mar. 1984.

[52] C. M. Champ, M. Vorobyev, and N. J. Marshall, "Colour thresholds in a coral reef fish," *R. Soc. Open Sci.*, vol. 3, no. 9, Sep. 2016.

[53] K. L. Cheney, C. Newport, E. C. McClure, and N. J. Marshall, "Colour vision and response bias in a coral reef fish.," *J. Exp. Biol.*, vol. 216, no. Pt 15, pp. 2967–73, Aug. 2013.

[54] M. Haahr, "Random.org," *Randomness and Integrity Services Ltd.*, 1998. [Online]. Available: https://www.random.org/. [Accessed: 15-May-2019].

[55] M. van Casteren and M. H. Davis, "Mix, a program for pseudorandomization," *Behav. Res. Methods*, vol. 38, no. 4, pp. 584–589, Nov. 2006.

[56] R. C. Honey, M. D. Iordanova, and M. Good, "Associative structures in animal learning: Dissociating elemental and configural processes," *Neurobiol. Learn. Mem.*, vol. 108, pp. 96–103, Feb. 2014.

[57] C. Champ, G. M. Wallis, M. Vorobyev, U. E. Siebeck, and N. J. Marshall, "Visual acuity in a species of coral reef fish: Rhinecanthus aculeatus.," *Brain. Behav. Evol.*, vol. 83, no. 1, pp. 31–42, 2014.

[58] L. Mitchell, K. L. Cheney, F. Cortesi, N. J. Marshall, and M. Vorobyev, "Triggerfish uses chromaticity and lightness for object segregation," *R. Soc. Open Sci.*, vol. 4, no. 12, p. 171440, Dec. 2017.

[59] C. Newport *et al.*, "Fish use colour to learn compound visual signals," *Anim. Behav.*, vol. 125, pp. 93–100, Mar. 2017.

[60] E. E. Simpson, N. J. Marshall, and K. L. Cheney, "Coral reef fish perceive lightness illusions," *Sci. Rep.*, vol. 6, Oct. 2016.

[61] V. Pignatelli, C. Champ, N. J. Marshall, and M. Vorobyev, "Double cones are used for colour discrimination in the reef fish, Rhinecanthus aculeatus.," *Biol. Lett.*, vol. 6, no. 4, pp. 537–9, Aug. 2010.

[62] N. F. Green, H. H. Urquhart, C. P. Van Den Berg, N. J. Marshall, and K. L. Cheney, "Pattern edges improve predator learning of aposematic signals," *Behav. Ecol.*, vol. 29, no. 6, pp. 1481–1486, Jul. 2018.

[63] A. Bisazza, C. Agrillo, and T. Lucon-Xiccato, "Extensive training extends numerical abilities of guppies," *Anim. Cogn.*, vol. 17, no. 6, pp. 1413–1419, 2014.

[64] E. Gatto, T. Lucon-Xiccato, B. B. Savaşçı, M. Dadda, and A. Bisazza, "Experimental setting affects the performance of guppies in a numerical discrimination task," *Anim. Cogn.*, vol. 20, no. 2, pp. 187–198, Mar. 2017.

[65] S. G. Reebs, "Time-place learning based on food but not on predation risk in a fish, the inanga (Galaxias maculatus)," *Ethology*, vol. 105, no. 4, pp. 361–371, Apr. 1999.

[66] A. Mack and I. Rock, *Inattentional Blindness*. MIT Press, 1998.

[67] V. Gabor, S. Wall, M. Gerken, and L. Brinkmann, "Does inattentional blindness exist in horses (Equus caballus)?," *Appl. Anim. Behav. Sci.*, vol. 215, pp. 45–51, Jun. 2019.

[68] J. Mackintosh and N. S. Sutherland, "Visual discrimination by the goldfish: The orientation of rectangles," *Anim. Behav.*, vol. 11, no. 1, pp. 135–141, Jan. 1963.

[69] C. Agrillo, M. Dadda, G. Serena, and A. Bisazza, "Use of number by fish," *PLoS One*, vol. 4, no. 3, p. e4786, Mar. 2009.

[70] V. Schluessel, H. Kraniotakes, and H. Bleckmann, "Visual discrimination of rotated 3D objects in Malawi cichlids (Pseudotropheus sp.): A first indication for form constancy in fishes," *Anim. Cogn.*, vol. 17, no. 2, pp. 359–371, Mar. 2014.

[71] L. J. Kamin, "Predictability, Surprise, Attention and Conditioning," in *Punishment and Aversive Behavior*, B. A. Campbell and R. M. Church, Eds. New York: Appleton-Century-Crofts, 1969, pp. 279–296.

[72] S. Gierszewski, H. Bleckmann, and V. Schluessel, "Cognitive Abilities in Malawi Cichlids (Pseudotropheus sp.): Matching-to-Sample and Image/Mirror-Image Discriminations," *PLoS One*, vol. 8, no. 2, p. e57363, Feb. 2013.

[73] C. Newport, G. Wallis, and U. E. Siebeck, "Object recognition in fish: accurate discrimination across novel views of an unfamiliar object category (human faces)," *Anim. Behav.*, vol. 145, pp. 39–49, Nov. 2018.

[74] V. A. Davis, R. I. Holbrook, S. Schumacher, T. Guilford, and T. B. de Perera, "Three-dimensional spatial cognition in a benthic fish, corydoras aeneus," *Behav. Processes*, vol. 109, no. PB, pp. 151–156, Nov. 2014.

[75] R. A. Rescorla and A. R. Wagner, "A theory of classical conditioning: Variations in the effectiveness of reinforcement and nonreinforcement.," *Class. Cond. II Curr. Res. Theory*, vol. 21, no. 6, pp. 64–99, 1972.

# CHAPTER 5
## General Discussion

# General Discussion

## Introduction

In this thesis, I set out to investigate information use and informational weighting in a fish, with complementary experiments examining both prior information use (i.e., updating) and the weighting of multiple concurrent sensory cues. Multiple cue use is a broad category, including the integration of information through time, across multiple modalities, across multiple receptors, or across multiple sub-systems within one modality; such as the integration or motion and colour in vision [1]. Integrating multiple sources of information to mediate behaviour is a nigh-ubiquitous phenomenon in the natural world, and provides a vital tool in overcoming the inherent limitations of an animal's individual senses [2]. However, how information is acquired, integrated, and used to inform behaviour remains poorly understood outside narrow range of taxa and behaviours [3]. For example, work on informational updating (the synthesis of prior knowledge with incoming sensory information to produce a weighted posterior estimate) has been biased in both taxonomy (focusing on mammals and birds) and methodology (focusing on field studies and natural foraging behaviours) [4]. As such, the taxonomic distribution and functional flexibility of this cognitive toolset remains unresolved. This thesis represents an attempt to partially bridge this knowledge gap, focusing on an understudied taxon and using conditioned behaviours as a theoretical and practical framework.

In Chapter 2, I demonstrated firstly that the Picasso triggerfish (*Rhinecanthus aculeatus*) incorporate an implied prior cue alongside visual information in an object discrimination task, and secondly that this information is weighted according to reliability. In Chapter 3, I further investigated the temporal scale over which prior information is acquired and incorporated, as well as the different strategies that subjects employed throughout training. In Chapter 4, I aimed to interrogate how triggerfish solve discrimination tasks in

the context of multiple competing cues that differ in their reliability, with an eye to investigating how triggerfish may cope with shifting sensory environments. Triggerfish enrolled in these experiments exhibited highly unusual learning curves which form the basis of certain interpretations of learning and decision-making in this species.

In this final chapter, I summarise the results and significance of this work in the context of current research, before discussing potential further avenues of investigation suggested by the results. Finally, I end with a discussion of how both this thesis in particular and work on piscine cognition in general may be important for furthering our understanding of cognition across a wide range of species, including our own.

# Aims

## Chapter 2

In addition to having a distinct taxonomic bias towards derived vertebrates, previous work on updating and prior information use has principally focused on field studies of naturally occurring behaviours; especially foraging [3], [5]–[7]. Furthermore, only a few studies have explicitly manipulated the prior information that animals have access to [3]. This renders quantifying the exact role of prior information use problematic and opens the possibility of evolved (rather than learned) priors being involved. Here, I set out to bridge this knowledge gap by testing updating using conditioned behaviours while controlling subjects' access to prior information under laboratory conditions.

Pursuant of two aims, I carried out a series of experiments in which triggerfish were conditioned to complete a 4AFC object discrimination with two potential cues – visual (colour difference between stimulus and distractors) and positional (a statistical bias in the stimulus towards one position). Subjects were then presented with a series of probe trials in which the reliability visual cue was reduced, and their responses recorded. A test group

with a biased stimulus distribution were contrasted against a control group with no such bias.

I hypothesized that, in the presence of a biased stimulus distribution, triggerfish would learn to attend to this cue, and use it to inform their choices when visual cue information was compromised. Furthermore, I expected that triggerfish would use a dynamic cue weighting strategy, using both visual and prior information, and weighting each according to their respective reliability.

## Chapter 3

In Chapter 2 I established that triggerfish could learn and incorporate prior distributions into their decision-making. However, this left ambiguous the timescale over which evidence was accumulated, with the possibility of low-order strategies explaining some variance in trigger behaviour. To test this possibility, I re-analysed the previously collected to data to analyse the frequency with which subjects repeated their choices. I hypothesized that subjects could have exhibited a positional bias simply by repeating their choices at a rate greater than chance, and that this would appear as sequence effects in the data.

In addition, I anecdotally observed several deviations from expected decision-making across the experimental population, including unconditioned biases and highly divergent task acquisition depending on minor changes to the conditioning schedule. I set out to quantify these observations, and to test whether they impacted my interpretation of results.

## Chapter 4

Aquatic environments are highly labile both in space and time, presenting aquatic organisms with particular challenges to visual perception. Not only do the appearance of stimuli vary depending on the environment, but different visual channels are favoured in different contexts, with well-documented ecological effects [8]. While both evolutionary and

ontogenetic adaptations to the sensory environment are well-documented, the possible role of behavioural accommodation has been comparatively under-examined. I was curious as to whether triggerfish can attend to the shifting reliability of individual cues and change their behaviour accordingly in order to compensate for this factor.

To examine this question, I carried out a series of experiments in which triggerfish were presented with compound stimuli differing in two features, with one feature being presented at either consistent reduced quality, or with variable quality. This simulated the effects of a consistent and variable visual environment, respectively.

# Summary of results

## Chapter 2

### Protocol acquisition

Triggerfish were initially trained to discriminate a single positive stimulus from three distractors in a 4AFC task. Once subjects had reached a pre-assigned criterion, the position of the positive stimulus was systematically biased towards one position. A block of training sessions was carried out under these conditions before subjects were presented with a battery of probe trials with manipulated visual cue reliability.

Out of 28 subjects, only two failed to reach the assigned learning criterion; being dis-enrolled due to low motivation. Of the remaining experimental population, naïve subjects (i.e., those without prior experience in similar experiments) took a mean of 7.3±2.6 sessions (n=19) to reach criterion, while experienced fish took only 3.7±1.1 sessions (n=7), 3 being the minimum possible. Accuracy remained extremely high both during bias training (94.0±5.1%, n=26) and in reinforcement trials during testing (95.9±3.4%, n=26), with no effect of experience. The large effect of experience on initial task acquisition suggests experienced individuals retained some memory of their prior conditioning.

Given the large gap between experiments (2-6 months), this retention is notable. The study of memory in fishes is comparatively underdeveloped, but this observation compares favourably with Ingraham and colleagues' [9] work on retention in cichlids (*Labidochromis caereuleus*), which demonstrated that cichlids retained a significant preference for a reinforced stimulus after 12 days of non-reinforcement.

The Picasso triggerfish has recently come to greater scientific attention thanks to work carried out in the last decade by the Sensory Neurobiology Group at the University of Queensland, and collaborators. Their research has focused on empirically establishing the species' spectral sensitivity and visual acuity at both a physiological and behavioural level [10]–[13], as well as their colour response bias, and mechanisms underlying colour and pattern discrimination [14]–[18]. So far, a limited range of conditioning procedures have been attempted with this species, principally two-alternate forced choice (2AFC) (e.g. [12]) and recently a variation on a simultaneous oddity from sample procedure [18]. Furthermore, most of this work has been carried out at field stations where subjects are held in captivity for only short periods before being released, which is neither ethically nor financially feasible for laboratories not located at an appropriate field site.

Two-alternate forced-choice (2AFC) is a more commonly used discrimination paradigm; however, the increased statistical power and informativeness of each trial in 4AFC can radically reduce experiment duration, which is doubly valuable when working with subjects that have limited motivation (briefly reviewed in [19]). Previous work with triggerfish has struggled with the issue of time. For example, Champ *et al.* [12] took 4 months to condition triggerfish to a 2AFC discrimination task, and fewer than half of the subjects reached the required threshold. Save for dimensionality (i.e., stimuli differing in two dimensions rather than 1), this experiment was very similar to those outlined in Chapter 2. In contrast, across multiple experiments naïve subjects reached learning criterion within 21 calendar days, and

a mean of only 4.9±1.9 calendar days (n=19), while experienced subjects reached the learning criterion within 4 calendar days, (mean = 2.6±0.8 days, n=7). These results unambiguously demonstrate that not only are triggerfish capable of being conditioned to a higher-order choice task, but also that doing so drastically increases efficiency. However, as experiment 3 demonstrated, minor changes in the structure of the conditioning procedure can lead to radically divergent acquisition curves. Finally, the success of these experiments demonstrates that Picasso triggerfish can be successfully maintained long-term in captivity for repeated experiments.

### Multiple cue use

Having been successfully conditioned to discriminate the stimuli visually, subjects were provided with a second source of information – position. I predicted that when deprived of visual information, subjects would instead rely on the positional cue, and bias their choices accordingly. As predicted, when the quality of the visual cue was degraded, subjects increasingly biased their choices towards the bias position. The relationship between the two factors was approximately linear. By contrast, subjects in the control group (trained without a positional bias) exhibited no consistent bias independent of visual cue quality. These results indicate that not only did subjects learn the statistical distribution of the stimulus, but that they used this information to complete the task when visual information was lacking.

Furthermore, the linear correlation observed suggests that subjects weighted the two sources of information when making decisions. If subjects attended to visual cue quality and thereby decided on each trial whether to use (exclusively) prior or current sample (visual) information, this would produce a sigmoid response curve, in contrast to the linear response observed. Consequently, these results provide evidence of updating in triggerfish.

## Chapter 3

### Task acquisition

I had previously observed that there were large differences in learning rates between experiments. Particularly, subjects in experiment 3 were much slower to acquire the positional aspect of the task, irrespective of prior experience. I hypothesized that this was due to interference between the two tasks inhibiting the formation of associations between the reinforcer and the relevant aspects of the conditioned stimuli (i.e., position). To test this hypothesis, I replicated the experiment using only the positional task, and then compared task acquisition rates. All subjects (n=8) acquired the task, compared to 6/8 for the original experiment, and in significantly fewer trials (mean = 185±23.9SD vs. 506±63.3), confirming the hypothesis. In both cases, prior experience reduced acquisition time. However, subjects were still slower to learn the positional than visual tasks, suggesting other factors also contributed to the learning disparity.

### Sequence effects

Sequence effects are a common phenomenon in comparative psychology and psychophysics, being responses by subjects to low-level patterns in a series of stimulus presentations. Both human and animal subjects rend to respond more quickly and accurately when a trial reinforces a (perceived) pattern in past presentations, such as a train of repeats or alternations. While typically discussed as experimental artefacts, these must necessarily reflect a strategy by the subject to solve the task using simple, undemanding rules. While the number of stimulus positions discouraged most simple rules, repeats were plausible in my experiments, and further, given the biased input could potentially produce the bias response curve exhibited in the bias treatment. To test this hypothesis, I analysed the rate at which subjects repeated their choices, factored by treatment, prior reinforcement, and stimulus quality. Actual subject choices were matched to null models to assess choice deviance, using both analytic and simulation approaches.

Overall, the mean repeat rate was significantly lower than in any null model, rejecting the hypothesis that subject's ability to incorporate prior information could be accounted for by low-order computations. However, analysing interaction with other factors revealed more complex patterns in behaviour. While overall repeat rates were below model expectations, the frequency of position repeats was positively correlated with task difficulty, suggesting that subjects selectively repeated their choices when visual information was poor. Repeat rates also exhibited a counter-intuitive relationship with prior reinforcement. I expected that subjects would exhibit a win-stay/lose-shift relationship with reinforcement, increasing repeat rates after a reward. This would align with associative learning predictions. However, the opposite was true – subjects were more likely to repeat their choices following non-reinforcement. The experiment 2 population exhibited a final inflection in this pattern. This experiment used a two-tier reward system, with preferred food being awarded to correct bias choices in the bias treatment, and randomly in the control treatment. Subjects in the bias treatment greatly increased their repeat rate following enhanced reinforcement compared to the bias group - though small sample size (2 vs 8 subjects) limits confidence in this result. Overall, these results suggest that subjects only employed repeats as a strategy when other information was poor, and otherwise relied on other cues such as visual contrast. However, the response to prior reinforcement presents a puzzle, matching a win-shift/lose-stay strategy. While animals may be readily conditioned to win-shift/lose-stay [20] it is surprising to observe it as a spontaneous strategy in the absence of reinforcement, and runs counter to expectations under simple associative learning models.

### Positional / directional bias

I anecdotally observed in Chapter 2 that there appeared to be a population-wide preference towards stimuli right and down that transcended conditioned biases. To test this, I formally analysed the distribution of choices subjects made and compared this to

model expectations. Confirming this anecdotal observation, I found that there was a small but highly significant population bias to choose lower and right positions more frequently than predicted. This pattern held true when subjects' choices were decomposed into their coordinate elements (i.e., horizontal, and vertical deviation) and was positively correlated with task difficulty. Overall, subjects chose the lower-right position a mean-average of 6.5%±1.8SE more trials than expected. This bias was present at all stages of training and did not align with any bias in the training schedule. Given the lack of any consistent factors that might have accidentally biased subjects, I hypothesized that this was either a consequence of early life events or represented a consistent lateralization of behaviour and/or cognition.

## Chapter 4

### Protocol acquisition

Developing on the methods used in previous chapters, experiments were carried out in a separate tank to that subjects were housed in. This approach permitted greater flexibility in the design of housing tanks both in terms of swimming volume and the inclusion of enrichment objects while ensuring consistency of training across subjects. However, this modification did necessitate additional training steps to acclimatise subjects to being moved between aquaria. Subjects were subsequently conditioned to move between compartments in the experimental tank and bite at a dummy stimulus before being enrolled in the experiment proper. Once enrolled, subjects were trained to discriminate a compound stimulus from three distracts in a 4AFC task. The positive stimulus differed from the distractors in two features, which were presented at high quality, intermediate quality, or varying quality; with one always being presented at high quality.

A pilot was carried out with a small group of subjects (n=5), with the intention to increase the sample size and fully balance condition assignment in future. Of the initial population,

only 3 completed pre-training as described above. The remaining subjects were unable to consistently complete a full session (10 trials) of movements between compartments, losing motivation and exhibiting stress responses. They were consequently removed from the experiment.

Of the three subjects that completed pre-training, none reached the pre-assigned learning criterion in the allocated time. Two sequential modifications to the experiment were made, addressing possible factors of subject motivation, limits of visual resolution, and cognitive load. These modifications did not afford any change in subject behaviour or improvement in performance, and the experiment was concluded due to time constraints.

# Significance

## Background

The ray-finned fishes are an extraordinarily diverse clade whose sister-taxon status and radically differing ecology relative to most tetrapods makes them a fascinating research subject, both in their own right and in counterpoint to more derived taxa. However, the cognitive and behavioural repertoire of this extraordinary group continues to be underestimated in both popular and scientific discourse, with both practical and scientific consequences. Practically, fish constitute a large and growing fraction of the global diet, a demand which is increasingly being met by aquaculture [21]. Furthermore, species such as the zebrafish have become essential models in biomedical research. However, the question of fish welfare in both contexts has only been the subject of serious inquiry since the turn of the millennium, and continues to be highly controversial [22], [23]. The concept of environmental enrichment for fishes has only recently begun to gain traction, e.g. [24]. Apart from ethical considerations, an understanding of how fish interact with their environment can be essential in aquaculture [25], [26] (reviewed in [27]) and research [28].

Fish present an ecological contrast to more typical models, in that most species can move with relative freedom within a three-dimensional volume. Most terrestrial organisms can be described as moving on a horizontal plane with three degrees of freedom (two of translation, one of rotation; yaw). In contrast, swimming fishes move through a 3D volume with 6 degrees of freedom (three each of translation and rotation) [29], [30], resulting in significant differences in how they encode space [31]. Furthermore, in contrast to flying species which must constantly expend energy to remain airborne most bony fishes are neutrally buoyant in water, allowing them to maintain position in a water column without ongoing expenditure of energy, and some species are capable of swimming backwards and sideways with little effort [32]. These differences have radical effects not only for movement ecology, but all aspects of behaviour that interact with movement – from foraging to mate selection [33].

## My findings

### Triggerfish as models

In this thesis, I have been able to further develop the range of experimental tasks to which triggerfish can be conditioned, as well as uncovering some potential practical and cognitive limitations in this domain. I have been able to demonstrate than not only can triggerfish be effectively conditioned to higher-order discrimination tasks but can be effectively maintained in laboratory settings for prolonged periods without negative welfare impacts. Both the speed of learning and memory retention abilities of Picasso triggerfish compare very favourably to other piscine models, further underscoring the value of this taxon for research in comparative animal cognition. The ability of triggerfish to learn tasks and stimuli that are both outside their evolved sphere of experience and that lack ecological relevance further underscored the cognitive flexibility of this species, and capacity for further research in future. However, the anomalous results uncovered in Chapter 4 also

underscore the importance of considering all known features of an organism's perceptual and cognitive processes in the design of experiments.

## Sensory updating

Both humans and non-human animals can be framed as probabilistic as opposed to deterministic decision-makers. Indeed, this may reflect a fundamental truth of cognition [34]. Decisions are made with uncertainty, and behavioural strategies shift when animals are confronted with more information [5]. While this has been most exhaustively investigated in the context of optimal foraging theory [3], [35]–[37], it has been documented in other contexts such as mate selection [38], [39], dispersal [40], and even web construction [41]. In order to increase the robustness of their environmental estimates, animals may actively sample the environment before making decisions [40], [42]. The underlying mathematical or cognitive rules underlying these processes remains a matter of intense debate. Bayesian statistics remains a highly popular model [5], [43]–[45], proposed to represent a mathematically optimal approach to updating prior knowledge with incoming information [4], [43], [46]. As a predictive model, it has seen some success in both human and non-human animals, for example in both vertebrate [37], [47], [48] and invertebrate [49] foraging, and in collective decision making [50], [51]. However, the application of Bayesian statistics in not without criticism. These typically focus on two main points – the assumption that Bayesian updating is computationally demanding compared to alternatives, and that the relative utility of prior information is highly sensitive to the degree of environmental stochasticity. Regarding the former, simpler alternatives such as linear operators often predict animal behaviour comparably while losing little in optimality [36], [52], [53]. That said, some recent formulations of Bayesian computations can be as parsimonious as simple associative learning accounts [34]. The latter criticism corresponds to several studies that have demonstrated animals' use of prior information to be highly contingent on both scale and timing [39], [54].

So far, the bulk of work has focused on field studies and manipulated animals' information/sensorium for only a limited time. Consequently, it remains unclear both how flexible updating is, and the relative importance of evolved priors [55], [56].

Here, I have demonstrated not only that triggerfish are able to use and weight both prior and incoming sensory information according to reliability but are also able to apply this ability in the context of a highly abstracted task featuring stimuli lacking in ecological significance, in which the role of evolved priors can be discounted. This ability to apply updating flexibly across contexts is notable – for comparison, even in humans the ability to incorporate prior probability in abstract tasks is contingent on experiment design, and is particularly sensitive to incentives [57].

Scientific dismissal of piscine cognitive abilities has been at least partially driven by the comparative simplicity of fish brains, and especially the absence of a mammalian neocortex [58], [59]. Nevertheless, research with different fish species continues to uncover cognitive abilities previously thought to be entirely seated within the neocortex. For example, elephantnose fish *(Gnathonemus petersii)* exhibit spontaneous cross-modal recognition of objects using vision previously only perceived using their active electric sense [60], [61]. Research into neural correlates basis of decision-making continues to be biased towards mammalian species and structures, e.g. [62], [63]. Results presented here further justify expanding breadth of model species for work on decision making.

### Multiple strategy use

Sequence effects, side biases, alternation, and various other similar schedule-based phenomena have been observed in structured trial experiments of all kinds with both human and non-human animal subjects for many decades [64]. In most fields, researchers have treated them as mere nuisance behaviours, as they have the capacity to disrupt experiments and distort outcomes. Consequently, much effort has been spent findings

tools to minimise or work around them [65]–[68]; with comparatively little attention paid to their cause or what they might tell us about decision-making.

Within an associative learning paradigm, it is presumed that the incidence of these phenomena will decline as subjects learn to attend to the relevant aspects of the task, and training schedules are constructed to facilitate this process. The assumption underlying this approach is that output biases can be eliminated by the removal of input biases. However, this assumption has been proved to be incorrect - even when input biases eliminated, response biases often persist [69].

To some extent, such fixed biases may represent a failure of animals to adapt to the artificial conditions of laboratory experiments – for example, balanced factorial schedules may not favour lateralized responding, but nature has no such compunction; and may favour lateralization of responses [70]. However, the frequency with which the same patterns recur across taxa and experiments favours a different explanation – that these phenomena are different expression of a conserved toolkit of simple, flexible heuristics that may be applied in a range of contexts [55]. Surprisingly, this view is quite old. Nearly a century ago, Krechevsky (1932) framed what we would now describe as sequence effects as the subject exploring different “hypotheses” about the experiment [71]. However, such interpretations were largely abandoned with the rise of behaviourism, and only recently are these strategies being considered as potentially informative about learning and decision-making [72].

During the testing phase of experiments discussed in Chapters 2 and 3, triggerfish exhibited behaviour consistent with the use of multiple distinct hypotheses or strategies. Subjects conditioned with a distributional bias weighted this prior information against visual information, but also exhibited a contextual position repeating response. These two strategies appear to have operated in parallel and combined explain more variance in

subjects' responses than either alone. While this joint approach was flexible, this flexibility was also clearly limited, as no subject reached optimal behaviour. This could be due to incomplete conditioning, or a limitation on triggerfish's ability to adapt to the artificial statistical properties of the task, *c.f.* [7]. The latter interpretation is supported by the differences in task acquisition across experiments, which suggest that triggers are less able to learn positional cues than visual contrast cues. The apparent win-shift/lose-stay response was also maladaptive in the context of the task and may represent an adaptive constraint.

While the relative effectiveness of win-stay/lose-shift and Bayesian updating strategies have been compared in non-human animal models (for example in nesting ecology, [73]), to the best of my knowledge the use of two strategies in parallel has not been studied outside humans [74].

## Unanswered Questions

### Bayesian modelling

As originally planned, initial experiments intended to specifically test alternative models of cue weighting in triggerfish, with reference to Bayesian updating. However, it transpired that this was functionally beyond the scope of the thesis, both in terms of time allocation and available human expertise.

### Speed-accuracy trade-offs

A further aspect of decision-making and information use not addressed in this thesis is how triggerfish accumulate evidence in order to decide on the scale of an individual trial. The process of decision-making in the context of 2AFC tasks has seen extensive experimental and theoretical attention in both humans and non-human animals [75], [76]. Since assessing stimuli is a time-dependent process, there is an inherent tension for a reward-maximising agent between speed and accuracy in decision-making, termed the

speed-accuracy trade-off (SAT). Both humans and non-human animals exhibit SATs, and are observed to shift their strategies to maximise reward rate in discrimination tasks [75], [77]–[79]. At an analytic level, the typical formulation for optimal decision-making is the drift-diffusion model [80], in which evidence accumulates through a series of time steps for either alternative, with a decision being made once evidence for one or the other alternative reaches a certain threshold. Noise and uncertainty in evidence accumulation are accounted for by the inclusion of a noise term to evidence accumulation, rendering the process stochastic [81]. In this formulation, an agent adjusts its SAT by shifting the decision threshold up or down. Drift diffusion models are formulated for 2AFC tasks and incorporate assumptions such as uniform parameters across trials [77]. The experiments discussed in this thesis represent a departure from this model, but extensions can be incorporated to account for these factors, e.g. [82].

Analysing the relationship between accuracy and speed in discrimination tasks can open a window into numerous aspects of an animal's ecology and cognition, as well as informing experiment design. While SATs are common a range of species and contexts, extensive variation has been observed in the ability to adaptively adjust SATs, both between species and individuals. In the former case, organisms from highly static environments might be expected to be inflexible in their SAT trade-off. In the latter case, differences in SATs are often correlated with personality differences along the shyness-boldness axis [75]. For example, animals are expected to increase decision time when error costs are high. This may include predation threat - guppies from high-predation regimes take more time to make decisions [83]. Individual differences may even bypass such trade-offs – in both sticklebacks *(Gasterosteus aculeatus)* [84] and elephantnose fish *(Gnathonemus petersii)* [85], bold individuals make decisions more rapidly than shy conspecifics without sacrificing accuracy. The advantage of bold individuals in the context demands an explanation for the

prevalence of shy individuals in populations. The manner in which subjects adjust their SAT thresholds is also subject to variation. For example, in a discrimination task with variable signal quality, humans adopted a single SAT threshold that was nearly optimal across most signal qualities [77]. By contrast, in a similar task rats adjusted their SAT threshold dynamically according to signal quality [82].

## Anomalous results

The apparent inability of triggerfish to acquire the discrimination tasks set out in Chapter 4 represents a fascinating anomaly in the context of prior work with this species, including my own in Chapter 2. At present, the most promising hypothesis for this outcome relates to an apparent limitation in triggerfish sensory processing. Mitchell *et al.* [16] manipulated the quality of a stimulus using either achromatic or chromatic noise to test triggerfish vision. They demonstrated that triggerfish are less able to segregate objects based on lightness than chromaticity. They suggest that the preference for chromatic cues is an ecological adaptation to the visual environment of reefs, in which wave action and shadows create highly irregular illumination.

However, while triggerfish in Mitchell *et al.*'s paper exhibited a clear preference in their experiment, all subjects were eventually able to distinguish stimuli disguised by achromatic noise, so this factor cannot not fully account for my observations. Whether subjects' inability to distinguish the stimuli represents a failure of the experimental design or due to as-yet unknown factors of sensory processing remains unclear, and I hope further work will shed light on this result.

# Final observations and future directions

## Frustration

### Observations

In several experiments discussed in this thesis (Specifically, Chapter 2 experiment 3, and Chapter 4), subjects were slow to learn the discrimination task, exhibiting extremely gradual learning curves. One approach I used to ameliorate this was to repeat trials exhibiting incorrect responses until a correct response was elicited (i.e., trial iteration). While successful, one consequence of this is that subjects transiently experienced very low reinforcement ratios. This often directly followed from a higher reinforcement ratio during pre-training. This resulted in low subject motivation, with an increase in overall decision latencies and no-go responses. Additionally, several subjects exhibited a range of disruptive behaviours that I informally characterised as frustrative.

#### *Avoidance*

- Increased delay to decision
- Increased delay to enter and exit the choice chamber before/after trials
- Increase in no-go responses
- Refusal to enter the jug for transport to experimental tank

#### *Agitation*

- A shift from slow-efficient balistiform swimming to fast-costly caudal fin swimming, typical of pursuit and escape responses
- Repetitive movements: circling, up-and-down swimming along tank edges
- Hiding behind the presentation board

#### *Aggression*

- Agonistic behaviour: If a subject remained in the experimental compartment for too long at the end of a trial, a net was used to gently usher it back to the start box. Typically, subjects simply avoided the net. During “frustrative” periods, subjects instead vigorously directed attacks upon the net, including biting and tail-beating

### *Distraction*

- Water jetting: jetting is used by several species of triggerfish to expose prey buried in loose sediment [87], and captive triggerfish readily learn to direct this technique at the water surface in the presence of an experimenter in anticipation of feeding. This is curiously reminiscent of the behaviour of archerfish (Toxotidae). These jets can be remarkably powerful, being ejected tens of centimetres above the water surface. Whether by accident or design, triggerfish were remarkably adept at striking the experimenter with these jets.
- Substrate-moving activity – moving particles of tank coral substrate (diameter approximately 10mm) around the tank, either by carrying them in the mouth, or by wedging the body into a corner of the tank and using vigorous beating movements of the caudal fin to blow gravel away. While a common part of the behavioural repertoire outside of experiments, during experiments triggerfish only engaged with this activity following multiple unreinforced trials.

In some cases, these behaviours were so extensive that subjects had to be removed from the experiment, both due to the disruption and as they received insufficient nutrition in experiments to ensure their health due to their poor performance. There were significant individual differences in the behaviours exhibited – agonistic behaviour was only observed in two subjects, possibly corresponding to a bold-shy spectrum in the population.

## Theory and practice

"Frustrative" behaviours in conditioning experiments have been documented since the early days of the discipline, e.g. [88], [89]. Typically, these occur when a reward is omitted followed a stimulus presentation – such as during an extinction protocol. Frustration can take several forms, but common presentations include an increase in agonistic behaviour, vocalisations, anxiety-like and escape responses, and overall activity. The appearance of many of these is remarkably consistent across diverse bird and mammal species [90]. Huston and colleagues [91] have interpreted these behaviours as falling into two categories of response – (a) seeking alternative behaviours that elicit the same reinforcer, and (b) seeking alternative reinforcers.

Frustrative behaviours belong to – and often co-occur with – a larger category of phenomena known as reward-schedule effects (RSEs), which concern how subjects respond to the omission or devaluation of an expected reward [92]. Prominent examples of this are the partial reinforcement extinction effect (PREE) and successive negative contrast (SNC). PREE is the counter-intuitive finding that animals are more resistant to extinction following a partial reinforcement schedule where not all CS-US pairings are reinforced. SNC refers to the finding that when the magnitude or quality of a reinforcer is reduced, performance will transiently drop to a level below what would be found if the animal had been consistently conditioned with the lower amount – an overreach effect. Both phenomena are typically accompanied by frustrative behaviours.

Reward-schedule effects are typically not well accounted for in classical formulations of associative learning theory. For example, the popular Rescorla-Wagner model [93] can be used to calculate the expected change in associative strength between a conditioned and unconditioned stimulus using the following formula:

$$\Delta V = \alpha\beta(\lambda - V)$$

Where ΔV is the change in the strength of the association on a single trial, α is the salience of the CS, β is the associative value of the US, and λ is the intensity of the US (β and λ both 0-1). An unreinforced trial is represented by a λ of 0, which necessarily produces a negative change in associative strength – a weakening of the response. This account cannot explain RSEs [94][1]. The most popular model for accounting for these phenomena is Amsel's [95], [96] frustration theory. This model set out three essential postulates, as follows:

[1] Although the Rescorla-Wagner model can be extended to incorporate frustration theory [129].

1. Omission or reduction of an expected reward induces an immediate aversive affective state known as primary frustration. This produces the immediate frustrative responses known as aftereffects.
2. With ongoing conditioning, the subject learns to anticipate this aversive state, producing in an avoidant response to the stimulus. This aversion disrupts the conditioned response, resulting in SNC.
3. If reinforcement continues intermittently, secondary frustration becomes paired with reinforcement, and the subject becomes desensitised to secondary frustration (and reward omission) through the process of counterconditioning. This desensitization accounts for PREE.

When applied to my observations described above, the model readily explains the range of behaviour discussed. An increase in aggression and anxiety-like responses, as well as alternative reward-seeking behaviour are documented aftereffects of primary frustration. Avoidant behaviours – including increased delay between trials, and escape-like swimming – can be interpreted as secondary frustration disrupting the conditioned stimulus-approach response.

Frustration theory was developed largely based on research in mammals, especially rats. By comparison, results from other taxa have been extremely mixed, and evidence for both the behavioural phenomena (PREE, SNC, etc.), as well as the underlying mechanisms (i.e. primary and secondary frustration) remains limited in non-mammals [94]. Focusing on fish, experimental demonstrations – or even documentation – of aftereffects are few. Response delays following omission may indicate primary frustration in both goldfish *(Carassius auratus)* [97] and the blackchin tilapia *(Sarotherodon melanotheron)*[2] [98], [99]. However, this disparity could be explained by other factors, such as satiation resulting in supressed

[2] Under the binomen *Tilapia macrocephala*, a junior synonym

response following reinforcement. More recently, elevated aggression in response to food omission has been observed in Atlantic salmon *(Salmo salar)* [25], [26]. However, elevated corticosterone – characteristic of frustration in mammals – was not [25].

To my knowledge only a single paper has explicitly tested fish for negative successive contrast, with negative results [100]. Following a 10-fold reward reduction in reward magnitude goldfish reduced their performance in-line with a small-reward condition without overreach or disruption. This matches RW model predictions, corresponding to a decrease in associative strength following a reduction in both $\beta$ and $\lambda$. Evidence for the partial reinforcement extinction effect is mixed in fish. In the two species studied (goldfish and blackchin tilapia), results have either been negative [101], [102] or mixed [97], [99]. In most cases partial reinforcement reduced resistance to extinction rather than strengthening it, and the positive results appeared to be highly dependent on experimental contingencies such as inter-trial interval.

Alongside similar results in amphibians *(Rhinella arenarum)* [103] and reptiles *(Chrysemys picta)* [104], this has led to the inference that the mechanisms underlying frustration – especially the egocentric aspects (secondary frustration) – are derived features characteristic of the mammalian lineage (although results with birds complicate this picture) [94]. Considering this orthodoxy, my observations become more anomalous. While primary frustration can explain the "emotional" behaviour immediately following an omission (i.e., as aftereffects), it does not suffice to predict the avoidant behaviour that I observed in anticipation of a trial – which would require secondary frustration. Trace memory effects [105] may conflate aftereffects and anticipation between trials within a session, but not anticipatory behaviour separated by a day or more from the previous trial.

While my observations are informal, when considered in conjunction with the relative paucity of studies and narrow range of species covered, the subject of frustration in fish

seems ripe for re-evaluation. Such a re-evaluation would potentially have both practical and ethical implications. The PREE is frequently exploited by psychologists to minimise unwanted changes to performance in unreinforced probe trials. However, the approach is not typically used with fish due to the (assumed) reverse PREE response – indeed, partial reinforcement was not used in the experiments described in this thesis for that very reason. If new research were to contradict this, it would recommend a change in practice. In addition, animals' response to reward omission has in recent decades been considered as a tool for evaluating animal welfare [106]. Given the burgeoning concern with fish welfare in aquaculture, further research could potentially provide useful evaluation tools.

Pursuant of this re-evaluation, I would recommend that new basic research building on the early work of the 60s and 70s be conducted, either with triggerfish or other model species such as elephantnose fish *(Gnathonemus petersii).* Experiments for evaluating extinction have been highly refined in mammals [107], and useful design innovations drawn from these may be implemented. Instrumental consummatory negative contrast (cSNC) or spaced trial experiments would be especially suitable and require little modification of existing techniques in use with fish.

## The caudal apparatus

The caudal apparatus (moniker my own) is an anatomical structure located on the caudal peduncle of Picasso triggerfish. It consists of three or more rows of raised scale-like spines on either side of the caudal peduncle, resembling nothing so much as the elevated slots of a cheese grater (figure 1). I have observed triggerfish "swiping" at targets with this structure when appearing threatened, such as when an experimenter cleaned the triggerfish's tank. Based on photographic evidence, I believe similar (though less prominent) structures are present in the other members of *Rhinecanthus*, as well as some other balistids in the genera *Balistapus*, *Balistoides*, and *Sufflamen* – species widely-distributed

within the Balistidae according to recent phylogenetic models [108]. While the structure is both prominent and unusual, I have been unable to find any reference to it in the academic literature. An article on poisonous fishes in the October 1952 edition of the American Bureau of Naval Personnel's information bulletin *All Hands* makes reference to a "tail spine" in the "grotesque... Trigger Fish" (*sic.*) [109]. However, it is clear from context that this erroneous description refers to the first dorsal spine (Triggerfish are incidentally, neither poisonous nor venomous, ciguatera notwithstanding). The only actual reference to the structure I have found is a cursory mention in a hobbyist article on the website *The Spruce*, which describes "....in front of the tail... prickly, spike-like rows of spines" [110]. However, no further information is provided.



*Figure 1:* photos showing the left caudal apparatus in four different Picasso triggerfish. Kindly modelled by (clockwise from top left): Gog, Randir, Fëanor, Magog.

Given its observed position, structure, and use; the apparatus appears functionally analogous to the eponymous caudal spines of the tangs and surgeonfishes (Acanthuridae). Acanthurids use these spines (which are erectable, in contrast to the caudal apparatus) during agonistic interactions with conspecifics, raking targets with the erected spine producing lesions and skin abrasions [111], and may also use them to deter predators. Given the observed similarities in the structures, and the comparable familiarity of both taxa to divers and researchers both, it is surprising that the Acanthurid caudal spine has received so much more attention than the caudal apparatus. Research into the taxonomic distribution, ontogeny and behavioural function of the caudal apparatus is warranted.

## Ecology, evolutionary history, and cognition

Both the movement and feeding ecology of an organism can affect learning style and how space is encoded in the brain [112]. Consequently, a logical extension of the research presented here would be to examine differences in cognition between related species of triggerfish exhibiting different ecologies. While primarily demersal and associated with rocky reef structures, some species and even populations of the balistidae exhibit highly divergent lifestyles. For example, the ocean triggerfish (*Canthidermis sufflamen*) is at least partially epipelagic [113] and employs an oscillatory form of balistiform locomotion in contrast to the undulatory form exhibited by the Picasso triggerfish [114]. In contrast, the sargassum triggerfish *(Xanthichthys ringens)* exhibits an ontogenetic shift in habitat from pelagic Sargassum drifts to coral reefs through its life [115]. This species also exhibits extensive inter-population variation in feeding ecology and morphology, with some populations being entirely planktivorous with a partially superior mouth, while others predominantly consume a more typical triggerfish diet of benthic hard-shelled prey with a terminal mouth [116], [117]. While the interrelationships of the taxon are disputed, the monophyly of the balistidae is not in question, and the group is known to be a relatively recent radiation [118].

## Metacognition and uncertainty monitoring

My research demonstrates that triggerfish can assess the uncertainty in a cue and altering how they use that information accordingly. Given the observations that various fish species exhibit a range of sophisticated abilities in sensory processing and uncertainty assessment, a clear extension of this research would be to assess the possibility of uncertainty monitoring and metacognition in triggers.

In brief, metacognition can be defined as cognition about cognition [119] - the presence of executive function to assess and control both perception and memory. In humans, assessment of this function is fundamentally grounded in the capacity to self-report one's conscious state – a tool that as yet we cannot extend to animal models [120]. Without the ability to self-report, experimental demonstrations of uncertainty monitoring (and by extension, metacognition) are highly challenging, and remain controversial even in our closest relatives, with some investigators claiming that observed results could be achieved through simple associative learning. The most common approach to this problem is an implementation of a relational discrimination task such as sparse-dense with the inclusion of an "uncertainty response" (UR), which permits subjects to decline a trial [121]. Even this approach is controversial, with critics arguing that it is equivalent to an "intermediate" response that could be acquired through simple associative learning. More recent work has attempted to distinguish these alternatives however, e.g. [122].

The key to a robust demonstration of uncertainty monitoring is the capacity of subjects to learn relational rules. While I did not directly address relational learning in this thesis, there is some evidence of relational learning in fish, albeit mixed. Both Malawian cichlids *(Pseudotropheus sp.)* [123] and goldfish *(Carassius auratus)* [124] can learn match-to-sample tasks, while the former can also learn to discriminate abstract categories [125]. Meanwhile, guppies *(Poecilia reticulata)* can be conditioned to complete both absolute and

relative numerical discriminations in highly abstracted contexts [126]. By contrast, both archerfish *(Toxotes chatareus)* [127] and Malawian cichlids *(Pseudotropheus zebra)* [128] were unable to learn a same/different rule.

In short, there is reasonable evidence that the fundamental cognitive capabilities and tools required to assess metacognition are present in fish, and while unlikely given current knowledge of the distribution of metacognitive abilities, even a non-successful experimental investigation could provide fascinating insight into piscine learning.

## Final remarks

In this thesis, I have demonstrated that triggerfish can combine and weight incoming sensory and prior information to produce weighted sensory estimates and can use this facility to inform their decision-making in a novel context. To the best of my knowledge, this constitutes the first demonstration of flexible updating in the Teleostei. Further, I found that triggerfish appear to use multiple strategies in concert to solve discrimination tasks even when associative strength is high, suggesting that evidential weighting extends beyond information to behavioural strategies themselves. While the neuro-physical basis of these abilities remains unknown, the presence of such flexible and powerful decision-making in a non-cortical vertebrate demands a re-consideration of our understanding of the evolution of cognition in vertebrates.

While I was not able to fully address all of the questions I set out to, even the anomalous results produced informative data that raise interesting questions about the perceptual tuning and cognitive basis of object recognition in triggerfish and suggest several avenues of future research.

The 4AFC discrimination task advanced here a statistically powerful and flexible approach to interrogating perceptual and cognitive abilities both in triggers and potentially other species and is a very promising approach for future research.

All of the above was possible thanks to a synthesis of concepts and methods derived from many different fields from experimental psychology to ethology, psychophysics, and behavioural ecology. This development highlights the benefit and need for continued collaboration and crosstalk across and between disciplines

## References

[1] M. W. Self and S. Zeki, "The integration of colour and motion by the human visual brain," *Cereb. Cortex*, vol. 15, no. 8, pp. 1270–1279, Aug. 2005.

[2] N. E. Munoz and D. T. Blumstein, "Multisensory perception in uncertain environments," *Behav. Ecol.*, vol. 23, no. 3, pp. 457–462, May 2012.

[3] T. J. Valone, "Are Animals Capable of Bayesian Updating ? An Empirical Review," *Oikos*, vol. 112, no. 2, pp. 252–259, Feb. 2006.

[4] J. M. McNamara, R. F. Green, and O. Olsson, "Bayes' theorem and its applications in animal behaviour," *Oikos*, vol. 112, no. 2, pp. 243–251, Feb. 2006.

[5] J. M. McNamara and A. I. Houston, "The application of statistical decision theory to animal behaviour," *J. Theor. Biol.*, vol. 85, no. 4, pp. 673–690, Aug. 1980.

[6] R. F. Green, "A simpler, more general method of finding the optimal foraging strategy for Bayesian birds," *Oikos*, vol. 112, no. 2, pp. 274–284, Feb. 2006.

[7] A. S. Dunlap, D. R. Papaj, and A. Dornhaus, "Sampling and tracking a changing environment: persistence and reward in the foraging decisions of bumblebees," *Interface Focus*, vol. 7, no. 3, p. 20160149, Jun. 2017.

[8] S. P. Collin and N. S. Hart, "Vision and photoentrainment in fishes: The effects of natural and anthropogenic perturbation," *Integrative Zoology*, vol. 10, no. 1. Wiley-Blackwell, pp. 15–28, 01-Jan-2015.

[9] E. Ingraham, N. D. Anderson, P. L. Hurd, and T. J. Hamilton, "Twelve-Day Reinforcement-Based Memory Retention in African Cichlids (Labidochromis caeruleus)," *Front. Behav. Neurosci.*, vol. 10, p. 157, 2016.

[10] V. Pignatelli, C. Champ, N. J. Marshall, and M. Vorobyev, "Double cones are used for colour discrimination in the reef fish, Rhinecanthus aculeatus.," *Biol. Lett.*, vol. 6, no. 4, pp. 537–9, Aug. 2010.

[11] C. Champ, G. M. Wallis, M. Vorobyev, U. E. Siebeck, and N. J. Marshall, "Visual acuity in a species of coral reef fish: Rhinecanthus aculeatus.," *Brain. Behav. Evol.*, vol. 83, no. 1, pp. 31–42, 2014.

[12] C. M. Champ, M. Vorobyev, and N. J. Marshall, "Colour thresholds in a coral reef fish," *R. Soc. Open Sci.*, vol. 3, no. 9, Sep. 2016.

[13] V. Pignatelli and J. Marshall, "Morphological characterization of retinal bipolar cells in the marine teleost rhinecanthus aculeatus," *J. Comp. Neurol.*, vol. 518, no. 15, pp. 3117–3129, Mar. 2010.

[14] E. E. Simpson, N. J. Marshall, and K. L. Cheney, "Coral reef fish perceive lightness illusions," *Sci. Rep.*, vol. 6, Oct. 2016.

[15] C. Newport *et al.*, "Fish use colour to learn compound visual signals," *Anim. Behav.*, vol. 125, pp. 93–100, Mar. 2017.

[16] L. Mitchell, K. L. Cheney, F. Cortesi, N. J. Marshall, and M. Vorobyev, "Triggerfish uses chromaticity and lightness for object segregation," *R. Soc. Open Sci.*, vol. 4, no. 12, p. 171440, Dec. 2017.

[17] N. F. Green, H. H. Urquhart, C. P. Van Den Berg, N. J. Marshall, and K. L. Cheney, "Pattern edges improve predator learning of aposematic signals," *Behav. Ecol.*, vol. 29, no. 6, pp. 1481–1486, Jul. 2018.

[18] K. L. Cheney *et al.*, "An Ishihara-style test of animal colour vision.," *J. Exp. Biol.*, vol. 222, no. Pt 1, p. jeb189787, Jan. 2019.

[19] K. Vancleef, J. C. A. Read, W. Herbert, N. Goodship, M. Woodhouse, and I. Serrano-Pedraza, "Two choices good, four choices better: For measuring stereoacuity in children, a four-alternative forced-choice paradigm is more efficient than two," *PLoS One*, vol. 13, no. 7, p. e0201366, Jul. 2018.

[20] C. K. Randall and T. R. Zentall, "Win-stay/lose-shift and win-shift/lose-stay learning by pigeons in the absence of overt response mediation," *Behav. Processes*, vol. 41, no. 3, pp. 227–236, Dec. 1997.

[21] FAO, "In brief, The State of World Fisheries and Aquaculture, 2018," 2018.

[22] V. A. Braithwaite, *Do Fish Feel Pain?*, 1st ed. Oxford: Oxford University Press, 2010.

[23] F. Chatigny, "The Controversy on Fish Pain: A Veterinarian's Perspective," *Journal of Applied Animal Welfare Science*, vol. 22, no. 4. Routledge, pp. 400–410, 2019.

[24] C. DePasquale, S. Fettrow, J. Sturgill, and V. A. Braithwaite, "The impact of flow and physical enrichment on preferences in zebrafish," *Appl. Anim. Behav. Sci.*, vol. 215, pp. 77–81, 2019.

[25] M. A. Vindas *et al.*, "Omission of expected reward agitates Atlantic salmon (Salmo salar)," *Anim. Cogn.*, vol. 15, no. 5, pp. 903–911, Sep. 2012.

[26] M. A. Vindas *et al.*, "Frustrative reward omission increases aggressive behaviour of inferior fighters," *Proc. R. Soc. B Biol. Sci.*, vol. 281, no. 1784, Apr. 2014.

[27] A. Fernö, G. Huse, P. J. Jakobsen, T. S. Kristiansen, and J. Nilsson, "Fish Behaviour, Learning, Aquaculture and Fisheries," in *Fish Cognition and Behavior*, Oxford, UK: wiley, 2011, pp. 359–404.

[28] V. A. Braithwaite and I. Ahlbeck Bergendahl, "The Effects of Early Life Experience on Behavioural Development in Captive Fish Species," in *The Welfare of Fish*, 1st ed., T. S. Kristiansen, A. Fernö, M. A. Pavlidis, and H. van de Vis, Eds. Springer, Cham, 2020, pp. 111–127.

[29] R. I. Holbrook and T. Burt de Perera, "Three-dimensional spatial cognition: Freely swimming fish accurately learn and remember metric information in a volume," *Anim. Behav.*, vol. 86, no. 5, pp. 1077–1083, Nov. 2013.

[30] R. I. Holbrook and T. Burt de Perera, "Separate encoding of vertical and horizontal components of space during orientation in fish," *Anim. Behav.*, vol. 78, no. 2, pp. 241–245, Aug. 2009.

[31] V. A. Davis, R. I. Holbrook, and T. Burt de Perera, "The influence of locomotory style on three-dimensional spatial learning," *Animal Behaviour*, vol. 142. Academic Press, pp. 39–47, 01-Aug-2018.

[32] R. Ruiz-Torres, O. M. Curet, G. V. Lauder, and M. A. MacIver, "Kinematics of the ribbon fin in hovering and swimming of the electric ghost knifefish.," *J. Exp. Biol.*, vol. 216, no. Pt 5, pp. 823–34, Mar. 2013.

[33] S. D. Healy, I. E. Bacon, O. Haggis, A. P. Harris, and L. A. Kelley, "Explanations for variation in cognitive ability: Behavioural ecology meets comparative cognition," *Behavioural Processes*, vol. 80, no. 3. Elsevier, pp. 288–294, 01-Mar-2009.

[34] B. R. Foley and P. Marjoram, “Sure enough: efficient Bayesian learning and choice,” *Anim. Cogn.*, vol. 20, no. 5, pp. 867–880, Sep. 2017.

[35] M. L. Wildhaber, R. F. Green, and L. B. Crowder, “Bluegills continuously update patch giving-up times based on foraging experience,” *Anim. Behav.*, vol. 47, no. 3, pp. 501–513, Mar. 1994.

[36] L. D. Devenport and J. A. Devenport, “Time-dependent averaging of foraging information in least chipmunks and golden-mantled ground squirrels,” *Anim. Behav.*, vol. 47, no. 4, pp. 787–802, 1994.

[37] J. C. Alonso, J. A. Alonso, L. M. Baustita, and R. Munoz-Pulido, “Patch use in cranes: a field test of optimal foraging predictions,” *Anim. Behav.*, vol. 49, no. 5, pp. 1367–1379, May 1995.

[38] W. Hunte, R. A. Myers, and R. W. Doyle, “Bayesian mating decisions in an amphipod, Gammarus lawrencianus Bousfield,” *Anim. Behav.*, vol. 33, no. 2, pp. 366–372, 1985.

[39] B. Luttbeg and R. R. Warner, “Reproductive decision-making by female peacock wrasses: flexible versus fixed behavioral rules in variable environments,” *Behav. Ecol.*, vol. 10, no. 6, pp. 666–674, Nov. 1999.

[40] V. Selonen and I. K. Hanski, “Decision making in dispersing Siberian flying squirrels,” *Behav. Ecol.*, vol. 21, no. 2, pp. 219–225, Mar. 2010.

[41] W. Maciejewski, “Orb-web spiders as Bayesian learners,” *Sci. Nat.*, vol. 106, no. 5–6, p. 22, Jun. 2019.

[42] S. R. X. Dall and R. A. Johnstone, “Managing uncertainty: Information and insurance under the risk of starvation,” *Philos. Trans. R. Soc. B Biol. Sci.*, vol. 357, no. 1427, pp. 1519–1526, 2002.

[43] K. P. Körding, “Decision theory: What ‘should’ the nervous system do?,” *Science*, vol. 318, no. 5850. pp. 606–610, 26-Oct-2007.

[44] A. C. Courville, N. D. Daw, and D. S. Touretzky, “Bayesian theories of conditioning in a changing world,” *Trends Cogn. Sci.*, vol. 10, no. 7, pp. 294–300, 2006.

[45] S. R. X. Dall, L.-A. Giraldeau, O. Olsson, J. M. McNamara, and D. W. Stephens, “Information and its use by animals in evolutionary ecology,” *Trends Ecol. Evol.*, vol. 20, no. 4, pp. 187–193, 2005.

[46] S. Okasha, “The evolution of Bayesian updating,” *Philos. Sci.*, vol. 80, no. 5, pp. 745–757, Dec. 2013.

[47] B. Luttbeg and T. A. Langen, “Comparing alternative models to empirical data: cognitive models of Western Scrub-Jay foraging behaviour,” *Am. Nat.*, vol. 163, no. 2, pp. 263–276, Feb. 2004.

[48] J. A. van Gils, I. W. Schenk, O. Bos, and T. Piersma, “Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches.,” *Am. Nat.*, vol. 161, no. 5, pp. 777–93, May 2003.

[49] J. M. Biernaskie, S. C. Walker, and R. J. Gegear, “Bumblebees learn to forage like bayesians,” *Am. Nat.*, vol. 174, no. 3, pp. 413–423, Sep. 2009.

[50] S. Arganda, A. Pérez-Escudero, and G. G. de Polavieja, “A common rule for decision making in animal collectives across species.,” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 109, no. 50, pp. 20508–13, Dec. 2012.

[51] P. Nonacs and J. L. Soriano, “Patch sampling behaviour and future foraging expectations in Argentine ants, Linepithema humile,” *Anim. Behav.*, vol. 55, no. 3, pp. 519–527, Mar. 1998.

[52] A. Lange and R. Dukas, “Bayesian approximations and extensions: Optimal decisions for small brains and possibly big ones too,” *J. Theor. Biol.*, vol. 259, no. 3, pp. 503–516, Aug. 2009.

[53] A. D. Higginson, T. W. Fawcett, A. I. Houston, and J. M. McNamara, “Trust your gut: Using physiological states as a source of information is almost as effective as optimal bayesian learning,” *Proc. R. Soc. B Biol. Sci.*, vol. 285, no. 1871, p. 20172411, Jan. 2018.

[54] H. H. Marshall, A. J. Carter, A. Ashford, J. M. Rowcliffe, and G. Cowlishaw, "How do foragers decide when to leave a patch? A test of alternative models under natural and experimental conditions," *J. Anim. Ecol.*, vol. 82, no. 4, pp. 894–902, Jul. 2013.

[55] T. W. Fawcett, S. Hamblin, and L.-A. Giraldeau, "Exposing the behavioral gambit: the evolution of learning and decision rules," *Behav. Ecol.*, vol. 24, no. 1, pp. 2–11, Jan. 2013.

[56] J. C. Ramírez and J. A. R. Marshall, "Can natural selection encode Bayesian priors?," *J. Theor. Biol.*, vol. 426, pp. 57–66, Aug. 2017.

[57] C. A. Holt and A. M. Smith, "An update on Bayesian updating," *J. Econ. Behav. Organ.*, vol. 69, no. 2, pp. 125–134, Feb. 2009.

[58] T. Burne *et al.*, "Big ideas for small brains: What can psychiatry learn from worms, flies, bees and fish," *Molecular Psychiatry*, vol. 16, no. 1. Nature Publishing Group, pp. 7–16, 30-Jan-2011.

[59] J. D. Rose, "The Neurobehavioral Nature of Fishes and the Question of Awareness and Pain," *Reviews in Fisheries Science*, vol. 10, no. 1. CRC Press LLC, pp. 1–38, 2002.

[60] S. Schumacher, T. Burt de Perera, and G. von der Emde, "Electrosensory capture during multisensory discrimination of nearby objects in the weakly electric fish Gnathonemus petersii.," *Sci. Rep.*, vol. 7, p. 43665, Mar. 2017.

[61] S. Schumacher, T. Burt de Perera, J. Thenert, and G. von der Emde, "Cross-modal object recognition and dynamic weighting of sensory inputs in a fish.," *Proc. Natl. Acad. Sci. U. S. A.*, vol. 113, no. 27, pp. 7638–43, Jul. 2016.

[62] L. P. Sugrue, G. S. Corrado, and W. T. Newsome, "Matching behavior and the representation of value in the parietal cortex," *Science (80-. ).*, vol. 304, no. 5678, pp. 1782–1787, 2004.

[63] M. F. S. Rushworth and T. E. J. Behrens, "Choice, uncertainty and value in prefrontal and cingulate cortex," *Nature Neuroscience*, vol. 11, no. 4. Nature Publishing Group, pp. 389–397, 26-Apr-2008.

[64] B. J. Fellows, "Chance Stimulus Sequences for Discrimination Tasks," *Psychol. Bull.*, vol. 67, no. 2, pp. 87–92, 1967.

[65] L. W. Gellermann, "Chance Orders of Alternating Stimuli in Visual Discrimination Experiments," *Pedagog. Semin. J. Genet. Psychol.*, vol. 42, no. 1, pp. 206–208, Mar. 1933.

[66] D. Herrera and M. Treviño, "Undesirable choice biases with small differences in the spatial structure of chance stimulus sequences," *PLoS One*, vol. 10, no. 8, pp. 1–13, Aug. 2015.

[67] C. J. Gerard, H. A. Mackay, B. Thompson, and W. J. McIlvane, "Rapid generation of balanced trial distributions for discrimination learning procedures: a technical note.," *J. Exp. Anal. Behav.*, vol. 101, no. 1, pp. 171–8, Jan. 2014.

[68] C. A. Catania, "Reinforcement Schedules: The Role of Responses Preceding the One that Produces the Reinforcer," *J. Exp. Anal. Behav.*, vol. 15, no. 3, pp. 271–287, 1971.

[69] L. Lebovich, R. Darshan, Y. Lavi, D. Hansel, and Y. Loewenstein, "Idiosyncratic choice bias naturally emerges from intrinsic stochasticity in neuronal dynamics," *Nat. Hum. Behav.*, vol. 3, no. 11, pp. 1190–1202, Nov. 2019.

[70] G. Vallortigara and L. J. Rogers, "Survival with an asymmetrical brain: Advantages and disadvantages of cerebral lateralization," *Behav. Brain Sci.*, vol. 28, no. 4, pp. 575–633, 2005.

[71] I. Krechevsky, "'Hypotheses' in rats," *Psychol. Rev.*, vol. 39, no. 6, pp. 516–532, Nov. 1932.

[72] M. Jones, T. Curran, M. C. Mozer, and M. H. Wilder, "Sequential Effects in Response Time Reveal Learning Mechanisms and Event Representations," 2013.

[73] K. A. Schmidt and C. J. Whelan, "Nesting in an uncertain world: Information and sampling the future," *Oikos*, vol. 119, no. 2, pp. 245–253, Feb. 2010.

[74] G. Charness and D. Levin, "When optimal choices feel wrong: A laboratory study of bayesian updating, complexity, and affect," *American Economic Review*, vol. 95, no. 4. Kahneman and Tversky, pp. 1300–1309, 2005.

[75] L. Chittka, P. Skorupski, and N. E. Raine, *Speed-accuracy tradeoffs in animal decision making*, vol. 24, no. 7. 2009, pp. 400–407.

[76] R. Bogacz, P. T. Hu, P. J. Holmes, and J. D. Cohen, "Do humans produce the speed-accuracy trade-off that maximizes reward rate?," *Q. J. Exp. Psychol.*, vol. 63, no. 5, pp. 863–891, May 2010.

[77] F. Balci *et al.*, "Acquisition of decision making criteria: Reward rate ultimately beats accuracy," *Attention, Perception, Psychophys.*, vol. 73, no. 2, pp. 640–657, Feb. 2011.

[78] J. Drugowitsch, G. C. Deangelis, D. E. Angelaki, and A. Pouget, "Tuning the speed-accuracy trade-off to maximize reward rate in multisensory decision-making," *Elife*, vol. 4, no. JUNE2015, 2015.

[79] P. Simen, J. D. Cohen, and P. Holmes, "Rapid decision threshold modulation by reward rate in a neural network," *Neural Networks*, vol. 19, no. 8, pp. 1013–1026, Oct. 2006.

[80] R. Bogacz, E. T. Brown, J. Moehlis, P. J. Holmes, and J. D. Cohen, "The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks.," *Psychol. Rev.*, vol. 113, no. 4, pp. 700–765, 2006.

[81] R. Ratcliff and J. N. Rouder, "Modeling Response Times for Two-Choice Decisions," *Psychol. Sci.*, vol. 9, no. 5, pp. 347–356, Sep. 1998.

[82] P. Reinagel, "Speed and Accuracy of Visual Motion Discrimination by Rats," *PLoS One*, vol. 8, no. 6, p. 200, Dec. 2013.

[83] J. G. Burns and F. H. Rodd, "Hastiness, brain size and predation regime affect the performance of wild guppies in a spatial memory task," *Anim. Behav.*, vol. 76, no. 3, pp. 911–922, Sep. 2008.

[84] D. Mamuneas, A. J. Spence, A. Manica, and A. J. King, "Bolder stickleback fish make faster decisions, but they are not less accurate," *Behav. Ecol.*, vol. 26, no. 1, pp. 91–96, Jan. 2015.

[85] K. Kareklas, R. W. Elwood, and R. A. Holland, "Personality effects on spatial learning: Comparisons between visual conditions in a weakly electric fish," *Ethology*, vol. 123, no. 8, pp. 551–559, Aug. 2017.

[86] A. G. Dyer and L. Chittka, "Bumblebees (Bombus terrestris) sacrifice foraging speed to solve difficult colour discrimination tasks," *J. Comp. Physiol. A Neuroethol. Sensory, Neural, Behav. Physiol.*, vol. 190, no. 9, pp. 759–763, 2004.

[87] T. K. Frazer, W. J. Lindberg, and G. R. Stanton, "Predation on sand dollars by gray triggerfish, Balistes capriscus, in the northeastern Gulf of Mexico," *Bull. Mar. Sci.*, vol. 48, no. 1, pp. 159–164, 1991.

[88] N. E. Miller and S. S. Stevenson, "Agitated behavior of rats during experimental extinction and a curve of spontaneous recovery," *J. Comp. Psychol.*, vol. 21, no. 2, pp. 205–231, Apr. 1936.

[89] M. H. Elliott, "The effect of change of reward on the maze performance of rats," *Univ. Calif. Publ.*, vol. 4, no. 2, pp. 185–188, 1928.

[90] M. R. Papini and R. T. Dudley, "Consequences of surprising reward omissions.," *Rev. Gen. Psychol.*, vol. 1, no. 2, pp. 175–197, 1997.

[91] J. P. Huston, M. A. d. S. Silva, M. Komorowski, D. Schulz, and B. Topic, "Animal models of extinction-induced depression: Loss of reward and its consequences," *Neuroscience and Biobehavioral Reviews*, vol. 37, no. 9. Pergamon, pp. 2059–2070, 01-Nov-2013.

[92] M. R. Papini, "Diversity of Adjustments to Reward Downshifts in Vertebrates," *Int. J. Comp. Psychol.*, vol. 27, no. 3, pp. 420–445, 2014.

[93] R. A. Rescorla and A. R. Wagner, "A theory of classical conditioning: Variations in the effectiveness of reinforcement and nonreinforcement.," *Class. Cond. II Curr. Res. Theory*, vol. 21, no. 6, pp. 64–99, 1972.

[94] M. R. Papini, "Role of Surprising Nonreward in Associative Learning," *Japanese J. Anim. Psychol.*, vol. 56, no. 1, pp. 35–54, 2006.

[95] A. Amsel, "The role of frustrative nonreward in noncontinuous reward situations," *Psychol. Bull.*, vol. 55, no. 2, pp. 102–119, Mar. 1958.

[96] A. Amsel, "Frustration theory - Many years later," *Psychol. Bull.*, vol. 112, no. 3, pp. 396–399, 1992.

[97] S. R. Scobie, D. C. Gold, and D. Fallon, "Reward and reward omission: Time-dependent aftereffects in rats and fish," *Bull. Psychon. Soc.*, vol. 3, no. 6, pp. 452–454, 1974.

[98] R. C. Gonzalez, R. M. Eskin, and M. E. Bitterman, "Alternating and random partial reinforcement in the fish, with some observations on asymptotic resistance to extinction.," *Am. J. Psychol.*, vol. 74, pp. 561–568, 1961.

[99] R. C. Gonzales, R. M. Eskin, and M. E. Bitterman, "Extinction in the fish after partial and consistent reinforcement with number of reinforcements equated," *J. Comp. Physiol. Psychol.*, vol. 55, no. 3, pp. 381–386, 1962.

[100] G. Lowes and M. E. Bitterman, "Reward and learning in the goldfish," *Science (80-. ).*, vol. 157, no. 3787, pp. 455–457, Jul. 1967.

[101] J. Wodinsky and M. E. Bitterman, "Resistance to extinction in the fish after extensive training with partial reinforcement.," *Am. J. Psychol.*, vol. 73, pp. 429–434, 1960.

[102]R. Gonzalez, C., N. K. Holmes, and M. E. Bitterman, "Resistance to extinction in the goldfish as a function of frequency and amount of reward.," *Am. J. Psychol.*, vol. 80, no. 2, pp. 269–275, 1967.

[103]R. N. Muzio, E. Ruetti, and M. R. Papini, "Determinants of instrumental extinction in terrestrial toads (Bufo arenarum)," *Learn. Motiv.*, vol. 37, no. 4, pp. 346–356, Nov. 2006.

[104]A. Pert and M. E. Bitterman, "Reward and learning in the turtle," *Learn. Motiv.*, vol. 1, no. 1, pp. 121–128, Feb. 1970.

[105]M. R. Papini, "Comparative psychology of surprising nonreward," in *Brain, Behavior and Evolution*, 2003, vol. 62, no. 2, pp. 83–95.

[106]O. H. P. Burman, R. M. A. Parker, E. S. Paul, and M. Mendl, "Sensitivity to reward loss as an indicator of animal emotion and welfare," *Biol. Lett.*, vol. 4, no. 4, pp. 330–333, Aug. 2008.

[107]A. Jiménez-García *et al.*, "Consummatory Successive Negative Contrast in Rats," *BIO-PROTOCOL*, vol. 9, no. 7, 2019.

[108]C. L. McCord and M. W. Westneat, "Phylogenetic relationships and the evolution of BMP4 in triggerfishes and filefishes (Balistoidea)," *Mol. Phylogenet. Evol.*, vol. 94, pp. 397–409, Jan. 2016.

[109]Unattributed, "One Man's Fish is Another's Poison," *All Hands Bur. Nav. Pers. Inf. Mag.*, vol. 428, no. October 1952, pp. 14–17, 1952.

[110] S. Hauter and D. Hauter, "Triggerfish Family Balistidae Profile," *The Spruce*, 2019. [Online]. Available: https://www.thesprucepets.com/triggerfish-family-balistidae-profile-2925851. [Accessed: 25-Feb-2020].

[111] U. M. Schober and H. Ditrich, "Anatomy and use of the caudal spines in the aggressive behaviour of a surgeonfish (Osteichthyes: Acanthuridae)," *Mar. Behav. Physiol.*, vol. 21, no. 4, pp. 277–284, 1992.

[112] V. A. Davis, R. I. Holbrook, S. Schumacher, T. Guilford, and T. B. de Perera, "Three-dimensional spatial cognition in a benthic fish, corydoras aeneus," *Behav. Processes*, vol. 109, no. PB, pp. 151–156, Nov. 2014.

[113] N. Lezama-Ochoa *et al.*, "Present and future potential habitat distribution of carcharhinus falciformis and canthidermis maculata by-catch species in the tropical tuna purse-seine fishery under climate change," *Front. Mar. Sci.*, vol. 3, no. MAR, p. 34, Mar. 2016.

[114] H. Loofbourrow, "Hydrodynamics of balistiform swimming in the picasso triggerfish, Rhinecanthus aculeatus," University of British Columbia, 2006.

[115] E. Lieske and R. Myers, *Collins Pocket Guide – Coral Reef Fishes of the Indo-Pacific and Carribean: Indo-Pacific and Caribbean*. Collins, 1994.

[116] E. S. Hobson, "Feeding relationship of teleostean fishes on coral reefs in Kona, Hawaii," *Fish. Bull.*, vol. 72, no. 4, pp. 915–1031, 1974.

[117] R. G. Turingan, P. C. Wainwright, D. A. Hensley, R. G. Turingan, R. C. Wainwright, and D. A. Hensley, "Interpopulation variation in prey use and feeding biomechanics in Caribbean triggerfishes," *Oecologia*, vol. 102, no. 3, pp. 296–304, Jun. 1995.

[118] K. Matsuura, "Taxonomy and systematics of tetraodontiform fishes: a review focusing primarily on progress in the period from 1980 to 2014," *Ichthyological Research*, vol. 62, no. 1. Springer-Verlag Tokyo, pp. 72–113, 2014.

[119] J. D. Smith and D. A. Washburn, "Uncertainty monitoring and metacognition by animals," *Curr. Dir. Psychol. Sci.*, vol. 14, no. 1, pp. 19–24, Feb. 2005.

[120] J. D. Smith, "The study of animal metacognition," *Trends in Cognitive Sciences*, vol. 13, no. 9. Elsevier Current Trends, pp. 389–396, 01-Sep-2009.

[121] J. J. Couchman, M. V. C. Coutinho, M. J. Beran, and J. D. Smith, "Beyond stimulus cues and reinforcement signals: A new approach to animal metacognition," *J. Comp. Psychol.*, vol. 124, no. 4, pp. 356–368, Nov. 2010.

[122] J. David Smith, M. V. C. Coutinho, B. A. Church, and M. J. Beran, "Executive-attentional uncertainty responses by rhesus macaques(Macaca mulatta)," *J. Exp. Psychol. Gen.*, vol. 142, no. 2, pp. 458–475, 2013.

[123] S. Gierszewski, H. Bleckmann, and V. Schluessel, "Cognitive Abilities in Malawi Cichlids (Pseudotropheus sp.): Matching-to-Sample and Image/Mirror-Image Discriminations," *PLoS One*, vol. 8, no. 2, p. e57363, Feb. 2013.

[124] M. Goldman and S. Shapiro, "Matching-to-sample and oddity-from-sample in goldfish," *J. Exp. Anal. Behav.*, vol. 31, no. 2, pp. 259–266, Mar. 1979.

[125] V. Schluessel, G. Fricke, and H. Bleckmann, "Visual discrimination and object categorization in the cichlid Pseudotropheus sp.," *Anim. Cogn.*, vol. 15, no. 4, pp. 525–537, Jul. 2012.

[126] M. E. Miletto Petrazzini, C. Agrillo, V. Izard, and A. Bisazza, "Relative versus absolute numerical representation in fish: Can guppies represent 'fourness'?," *Anim. Cogn.*, vol. 18, no. 5, 2015.

[127] C. Newport, G. Wallis, and U. E. Siebeck, "Same/different abstract concept learning by archerfish (Toxotes chatareus)," *PLoS One*, vol. 10, no. 11, p. e0143401, Nov. 2015.

[128] T. Fuss, L. John, and V. Schluessel, "Same or different? Abstract relational concept use in juvenile bamboo sharks and Malawi cichlids," *Curr. Zool.*, Jul. 2018.

[129] H. B. Daly and J. T. Daly, "A mathematical model of reward and aversive nonreward: Its application in over 30 appetitive learning," *J. Exp. Psychol. Gen.*, vol. 111, no. 4, pp. 441–480, 1982.