

THE STARLING AS A RATIONAL DECISION-MAKER

CYNTHIA SCHUCK-PAIM

WOLFSON COLLEGE
&
DEPARTMENT OF ZOOLOGY
UNIVERSITY OF OXFORD



A THESIS
SUBMITTED FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
HILARY TERM 2003

The Starling as a Rational Decision-Maker

**Cynthia Schuck-Paim,
Wolfson College**

**Thesis submitted for the Degree of Doctor of Philosophy,
Hilary Term, 2003**

A central question in behavioural and evolutionary ecology is to understand how animals make decisions between, for instance, potential mates, nesting sites, foraging patches and territories. Normative models of choice usually predict preferences between alternatives by computing their value according to some criterion and then identifying the alternative with greatest value. An important consequence of this procedure is captured in the economic concept of rationality, defined through a number of principles that are necessary for the existence of a scale of value upon which organisms base their choices. Violations of rationality are nonetheless well documented in psychological and economic studies of human choice and consumer behaviour, and have forced a reinterpretation of much of the existing data and models.

Although largely unexplored in the study of animal decision-making, the systematic observation of irrationality would similarly pose serious challenges for functional approaches to behaviour. In this thesis I explore the possibility that violations of rational axioms may also be found in animal choices, using the European starling (*Sturnus vulgaris*) as a model species. My objectives were threefold. Firstly, I investigated the prevalence of rationality across distinct foraging paradigms, in situations involving multialternative choices, structured choice sets, choices between alternatives described by multiple attributes and risk-sensitive decisions. In a number of distinct experiments, the preferences of the starlings were consistent and stable across contexts, conforming to basic rational principles such as transitivity and regularity. A second objective was to explore possible factors underlying reported violations of rational axioms by animals. Amongst potential mechanisms, I review and examine the implications of the use of hierarchical and higher order choice rules, as well as the presence of constraints on the perception of rewards. Finally, I examine the likely effect of contextual changes on an organism's state, and consequently choice behaviour, and experimentally confirm the expectation that state-dependence in foraging preferences can underlie the observation of seemingly irrational behaviour. Altogether, my results suggest that, rather than being a common phenomenon, breaches of rationality in animals might be restricted to specific sets of parameters and conditions. They also emphasize the importance of considering the potential multitude of factors underlying violations of rationality in animal choices, and suggest that students of economic rationality in animal behaviour should also view preferences as a dynamic, state-dependent measure.

Acknowledgements

I would like to take this opportunity to thank all of the many people who have, in one way or another, helped me to develop this thesis. Alex Kacelnik gave me the opportunity to come to Oxford, and suggested that I worked on the topic of economic rationality in animal choices. From my first days in Oxford Alex has always been a friend and unfailing in his support and encouragement for my work, having contributed in many ways to what appears in the following lines.

The support of (past and present) colleagues from the Behavioural Ecology Group was also very important for me. Special thanks go to Jackie Chappell, Esteban Fernandez-Juricic, Barnaby Marsh, Lorena Pompilio, Martin Shapiro and Steven Siller who, apart from being great company, were always keen to discuss ideas and a source of practical advices. I feel particularly indebt to Lorena Pompilio, for the many hours of discussion and the inspiration for the ideas that appear in the thesis. Lorena has also helped me to run the experiments described in Chapter 6, and Barnaby worked collaboratively with me for the production of Chapter 7. I am also very grateful to Alan Grafen, Martin Rossi and Steven Siller for their so much appreciated statistical advice and help. Alan Grafen and David Raubenheimer have kindly agreed to be part of my advisory committee, and Itamar Simonson, Jerry Green, Duncan Luce and Michael Davidson were also very helpful in clarifying some thoughts. This thesis has also benefited from the comments of John Krebs and Innes Cuthill, who have also kindly agreed to be part of my examining committee.

I was also very fortunate in having wonderful friends who made life in Oxford a very rich and enjoyable experience. Although too many to list, a special thanks go to Sarah and Nigel, Ghinwa and the members of 'Chez Bop Bizz', Henrique and Ivy, Anita and Sinisa, Mallica and Matt, the friends in the Brazilian, Portuguese, Spanish and Argentinean communities and the neighbours in 'E-Block'. I would also like to express my enormous gratitude to my family and parents, Nelson and Vera, for their love, devotion and constant support throughout my entire life. Thank you for caring so much.

Finally, I would like to thank my husband and friend, Wladimir, who taught me to enjoy science, and have supported me in every possible way throughout these last eight years. Thank you for being so understanding about the many weekends and holidays I had to be in the lab running experiments and looking after the birds. Thank you for your love, warmth, fun and friendship. Life with you has been the best thing I could have ever hoped for.

My stay in Oxford and development of this thesis would not have been possible without the grant offered by the Brazilian government, a doctoral scholarship from CAPES (Coordenação de Aperfeiçoamento de Pessoal de Nível Superior). I have also benefited from grants and financial aid provided by The Association for the Study of Animal Behaviour (ASAB), Wolfson College and the Behavioural Ecology Group to participate in conferences and meetings.

For Vladimir

Table of Contents

CHAPTER 1. INTRODUCTION	01
CHAPTER 2. RATIONALITY IN RISK-SENSITIVE FORAGING CHOICES BY STARLINGS	03
2.1 Introduction	04
2.2 Methodology	07
2.2.1 Subjects	08
2.2.2 Apparatus	09
2.2.3 Training	09
2.2.4 Experimental Protocol	09
2.3. Data Analysis	11
2.3.1 Preference Scores	11
2.3.2 Context-dependence	12
2.3.3 Transitivity	13
2.4. Results	14
2.4.1 Knowledge about the alternatives	14
2.4.2 Choice trials	16
2.4.2.1 Preference for Variability	16
2.4.2.2 Context-Dependence	17
2.4.2.3 Transitivity	18
2.4.3. No-choice trials	19
2.4.3.1 Preference for Variability	19
2.4.3.2 Context-Dependence	20
2.4.3.3 Transitivity	20
2.5 Discussion	21
2.6 References	25
CHAPTER 3. LUCE'S CHOICE AXIOM AND CONTEXT-DEPENDENT PREFERENCES IN ANIMAL CHOICES	27
3.1 Introduction	28
3.2 Structured Choice Sets	31
3.3 Presence of Uncorrelated Choices	32
3.3.1 The Sampling Case	32
3.3.2 The "Noise" Case	33
3.4 Constraints on Perception	34
3.5 Conclusion	39
3.6 References	40
CHAPTER 4. STARLINGS SHOW INVARIANT PREFERENCE WITH STRUCTURED CHOICE SETS	42
4.1 Introduction	43
4.2 Methodology	46
4.2.1 Experimental Design	46
4.2.2 Subjects	47
4.2.3 Apparatus	48

4.2.4 Training	48
4.2.5 Experimental Protocol	48
4.3 Data Analysis	51
4.3.1 General Procedures	51
4.3.2 Preference Scores for the Fixed Options	51
4.3.3 Relative Preferences	51
4.3.4 Absolute Preferences	52
4.3.4.1 Regularity	53
4.3.4.2 Transitivity	53
4.4 Results	54
4.4.1 Knowledge about the Alternatives	54
4.4.2 Preference for Variability	57
4.4.3 Relative Choice Proportions	57
4.4.4 Absolute Choice Proportions	60
4.4.4.1 Regularity	60
4.4.4.2 Transitivity	62
4.5 Discussion	64
4.6 References	69
CHAPTER 5. STATE-DEPENDENCE AND ECONOMIC RATIONALITY IN ANIMAL CHOICES	71
5.1 Introduction	72
5.2 When Foraging is State-Dependent	73
5.2.1 Minimizing the Probability of Costly Errors	74
5.2.2 Minimizing the Probability of Starvation	77
5.3 Conclusion	80
5.3.1 Testing Rationality in Animal Choices	82
5.4 References	83
CHAPTER 6. TESTING RATIONALITY IN ANIMAL BEHAVIOUR: STATE-DEPENDENT VERSUS CONTEXT-DEPENDENT PREFERENCES	85
6.1 Introduction	86
6.2 General Methodology	88
6.2.1 Experimental Design	88
6.2.2 Subjects	91
6.2.3 Apparatus	91
6.2.4 Training	92
6.2.5 Experimental Protocol	92
6.3 Analysis of Data	93
6.3.1 Comparison of Treatments	93
6.3.2 Within-Treatment Analyses	94
6.4 Experiment 1: Discrimination of Amounts	94
6.4.1 Methods	94
6.4.2 Results and Discussion	95
6.5 Experiment 2: Decoy Effects on Choice	96
6.5.1 Methods	96
6.5.2 Results	97
6.5.2.1 Comparison of Treatments	97
6.5.2.2 Within-Treatment Analyses	97
6.5.3 Discussion	99

6.6 Experiment 3: Effects of Intake Rate on Choice	102
6.6.1 Methods	102
6.6.2 Results and Discussion	103
6.6.2.1 Preference Scores	103
6.6.2.2 Stability in Preferences	104
6.7 Experiment 4: Decoy Effects with Intake Controlled	104
6.7.1 Methods	104
6.7.2 Results And Discussion	106
6.7.2.1 Comparison of Treatments	106
6.7.2.2 Within-Treatment Comparison	107
6.7.2.3 Differences in Intake Rate	108
6.7.2.4 Stability in Preferences	110
6.8 General Discussion	112
6.8.1 State-Dependent versus Context-Dependent Preferences	114
6.8.1.1 Energetic State and Rationality in Jays	114
6.8.1.2 Energetic State and Rationality in Hummingbirds	115
6.8.2 Rate Currencies	117
6.8.3 Conclusion	119
6.9 References	121
CHAPTER 7. STATE-DEPENDENT VALUATION OF FORAGING OPTIONS BY STARLINGS	124
7.1 Introduction	125
7.2 Methods	127
7.2.1 Pre-Training	127
7.2.2 Training	128
7.2.3 Choice Trials	129
7.3 Results and Discussion	129
7.4 References	132
CHAPTER 8. CONCLUSION	133
8.1 Concluding Remarks	134
APPENDIX. OTHER ACADEMIC WORK PUBLISHED DURING THE DEVELOPMENT OF THE THESIS	138

Introduction

The research described in the following chapters falls in the overlapping fields of behavioural ecology, experimental economics and experimental psychology, and blends theoretical and experimental analyses to investigate animal decision-making. Specifically, I focus on the analysis of economic rationality in animal choices, namely the notion that by making use of effective available means to achieve value maximization individual choices should be consistent across different choice contexts.

Each of the chapters was written so as to stand alone, starting with a brief review of the topic. This approach is aimed at allowing the judgement of each chapter as a publication, but has the drawback of causing a certain degree of repetition between chapters.

The thesis is organized as follows. In Chapter 2 I investigate the presence of context-dependence in the risk-sensitive foraging preferences of the starlings. My primary aim was two-fold. Firstly, to determine the contingency of foraging behaviour upon the context of choice and secondly, to test the implicit assumption made by studies in the area that the simple and tractable scenarios upon which empirical research was based do reflect preferences in more realistic, multialternative cases.

In Chapter 3 I discuss a number of ways whereby a restrictive axiom of rational choice theory (the Choice Axiom of Luce) might be violated, providing alternative interpretations for the observation of such departures. Specifically, I consider instances where animals make use of higher order and hierarchical choice rules, as well as cases where a certain proportion of choices ('uncorrelated choices') result from causes other than those related to the properties of the rewards. I also explore the implications of mechanistic models of choice that assume the existence of constraints on perception, and more generally compare the predictions of the axiom with the optimal expectation in situations involving choice sets structured as a function of the degree of similarity of the options. Some of the possibilities discussed in Chapter 3 are tested experimentally in Chapter 4.

The remaining chapters are dedicated to discuss the potential influence of an organism's state (as defined by a number of correlate measures such as level of energy reserves, body temperature and gut contents) on its foraging preferences, together with its implication for the study of economic rationality in animals. I initially (Chapter 5) consider

possible means whereby apparent breaches of rationality may emerge if the likely effect of context on an organism's state is considered together with the observation that the fitness consequences of foraging choices are likely to be state-dependent. Next (Chapter 6), I demonstrate experimentally that irrationality in choices may emerge as a consequence of state-dependence in decision-making. Finally, in Chapter 7 I investigate the mechanisms whereby the state of an animal may affect the valuation of a payoff.

I conclude the thesis with a brief general discussion dominated by my overall opinion that, on the basis of my results, the assumption of rationality is still justified in most cases, and when violated, a deeper and more critical analysis than that usually conducted by studies in the area is generally required to reveal the causes underlying the observed inconsistencies in behaviour.

CHAPTER 2

2.1 INTRODUCTION	4
2.2 METHODOLOGY	7
2.2.1 SUBJECTS	8
2.2.2 APPARATUS	9
2.2.3 TRAINING	9
2.2.4 EXPERIMENTAL PROTOCOL	9
2.3. DATA ANALYSIS	11
2.3.1 PREFERENCE SCORES	11
2.3.2 CONTEXT-DEPENDENCE	12
2.3.3 TRANSITIVITY	13
2.4. RESULTS	14
2.4.1 KNOWLEDGE ABOUT THE ALTERNATIVES	14
2.4.2 CHOICE TRIALS	16
2.4.2.1 PREFERENCE FOR VARIABILITY	16
2.4.2.2 CONTEXT-DEPENDENCE	17
2.4.2.3 TRANSITIVITY	18
2.4.3. NO-CHOICE TRIALS	19
2.4.3.1 PREFERENCE FOR VARIABILITY	19
2.4.3.2 CONTEXT-DEPENDENCE	20
2.4.3.3 TRANSITIVITY	20
2.5 DISCUSSION	21
2.6 REFERENCES	25

RATIONALITY IN RISK-SENSITIVE FORAGING CHOICES BY STARLINGS*

*SCHUCK-PAIM, C. & KACELNIK, K. 2002. *Rationality in risk-sensitive foraging choices by starlings. Animal Behaviour, 64: 869-879.*

(as the version presented here is basically the same as that published, the use of the first person plural was maintained)

2.1 Introduction

Theories of decision-making developed by evolutionary biologists and by economists share many features and have interesting differences. One shared feature is the expectation of behavioural coherence, namely that a subject's preference system should behave "rationally". Here we use the term rational in the economists' sense, meaning only that preferences are consistent across contexts, regardless of what these preferences are or the process by which the subject takes decisions. Mas-Collel et al. (1995, pg.6) encapsulate the idea thus: "*The hypothesis of rationality is embodied in two basic assumptions about the preference relation: completeness and transitivity*". Completeness implies that there are well defined preferences between all possible pairs of options, and transitivity that if 'a' is preferred to 'b' and 'b' to 'c' then 'a' should be preferred to 'c'. For cases when subjects are faced with repeated choices and preference is not absolute but partial, transitivity is usually redefined in probabilistic terms as Stochastic Transitivity. Regularity is another property frequently invoked as a test of rationality (Luce 1977, Simonson & Tversky 1992, Tversky & Simonson 1993). We will treat regularity in more detail later, but its basic postulate is that the addition of alternatives to the choice set can never increase the choice probability of a given option.

Breaches of rationality present an interesting challenge because there are few normative frameworks in biology that are compatible with such violations (see Houston 1997 for an important exception), and hence it makes sense to investigate how prevalent these violations are. Recent reports of violations of rationality by hummingbirds, honeybees and gray jays (Shafir 1994, Hurly & Oseen 1999, Waite 2001a, Waite 2001b, Bateson et al. 2002, Shafir et al. 2002) open the possibility that irrational choice may be common. If this were the case the predictive value of normative (optimality) modeling should be questioned. To follow this issue we investigate transitivity, regularity and context-dependence in risky choice using European starlings (*Sturnus vulgaris*) choosing between distinct levels of variance in delay to food rewards.

First of all consider why is rationality commonly expected by normative theorists. If the decision maker assigns a score to each alternative according to an ordinal scale of "taste", "worth", "satisfaction", "utility" or, in the case of biology, expected fitness gain, then decisions based on maximizing these criteria should not breach rationality. The normative modeler ranks the options according to a conceptually orderly scale based on the

fitness consequences of each choice. Violations of rationality question the existence of any scale of preference, and this goes against the spirit of the game.

Consider for instance an intuitive explanation for the expectation of regularity. As described before, regularity is defined by the fact that adding further options into a choice set should not increase the proportion of choices for any pre-existing option. Say that a starling feeding in a field containing two prey types, worms and beetles, prefers worms 80% of the opportunities. Now imagine that berries in the same field become ripe, and hence join the set of available food items. If berries are not included in the diet, choices of the two previous prey types should remain unaltered, and if they were included, they would steal opportunities from either or both of them. Unless the value of worms or beetles is changed by the presence of berries, they can only lose by the presence of a further competitor. If regularity is shown empirically to be violated, one must conclude that the prey types do not have an absolute worth but one that is in some way relative to the available alternatives. As at least in humans these violations do occur (Huber et al. 1982, Simonson & Tversky 1992, Tversky & Simonson 1993), the issue is worth further thinking.

There are at least two kinds of possible accounts for violations of rationality. First, subjects are unlikely to be tuned to all the features of all the alternatives in any choice set, and hence we may expect them to take decisions under a different amount of knowledge in different occasions. This is in part akin to say that subjects will often make mistakes (which is theoretically trivial), but it is also a reminder that for a test of rationality to be informative, one must ensure that the attributes of an alternative are not only judged from the perspective of the omniscient experimenter, but from that of the choice maker given its state and knowledge. This could be important in many subtle ways. For instance, Houston (1997) demonstrates that, because optimal action in a series of choices is a function of future choices, if the present opportunity set is used by the decision maker to infer what her future opportunities will be, and hence what her future actions will be (including possible errors), then adding a suboptimal option to the choice set can in fact modify what is the best choice right now, and lead to intransitive behaviour. This issue implies that training conditions may be just as important as the conditions of the preference test itself simply because the training conditions allow the animals to predict what they may expect in future choices, hence having an effect on what is presently optimal. In the experimental investigation reported below we handle this problem explicitly by running experiments that avoid confounding the effects of training with those of the choice scenario.

Second, in contrast with our strategy as scientists for identifying optimal behaviour, organisms are unlikely to take decisions by assignments of worth in a one-dimensional scale. Instead, natural selection shapes motivational mechanisms of choice that tend to result in favourable outcomes under the most frequent ecological conditions of the species. The optimality models tell us how favourable different outcomes are under specific assumptions for ecological circumstances, but not how animals process information. It should not be surprising that mechanisms of choice produce irrational outcomes under some circumstances. This concept is familiar to students of optimal foraging, where the notion of “rule of thumb” has been present for a long time (Houston et al. 1982, Stephens & Krebs 1986), and of human decision making, where the advantage of “fast and frugal” heuristic rules that cope well (but not perfectly) with decisions in normal circumstances is widely recognized (Gigerenzer et al. 1999). Breaches of rationality of many different sorts are well documented in humans and explained in terms of the psychological mechanisms involved in generating action. Some authors have emphasized the adaptive value of simplified heuristics, pointing that rules usually referred to as ‘comparative’ (through which options are compared along their attributes instead of being assigned a single, absolute value) do perform well in many cases but can also lead to intransitive and context-dependent choices (Tversky 1969, Tversky 1972, Shafir et al. 1989, Wedell 1991, Gigerenzer et al. 1999).

In the context of risky choice in animals, one interesting result is that reported by Hurly & Oseen (1999). These authors studied the preferences of rufous hummingbirds (*Selasphorus rufus*) between three kinds of artificial flowers offering the same mean volume of sucrose, but three distinct variance levels. They found that the birds preferred intermediate variance to both no variance and high variance when all three options were available, but preferred no variance to intermediate variance when the high variance option was absent, in binary tests. In many studies of risk sensitivity using binary choices it has been found that no variance in food amount is preferred to variance, as reported for the hummingbirds in binary choices. In contrast, variance is always preferred to no variance in delays to food (Kacelnik & Bateson 1996). Hurly and Oseen’s results might imply that all these experiments, rather than reflecting a general attitude towards or against variance, may be artifactual results of using only binary choice designs.

Here we investigate the putative presence of context-dependence with starlings facing choices between options differing in variance in delay to food. We examine both trinary and binary choices and separate training from choice context effects. Our objective is

two-fold: firstly, to investigate the extent to which the inclusion of a third alternative into a binary choice set can affect the starlings' preferences between the original foraging options and secondly, to explore the generality of previous research on risk-sensitive foraging.

2.2 Methodology

The experiments were planned to allow for the distinction between two possible causes of inconsistent choice: background and choice set. The former refers to the effect of the set of options known to the subjects, whether they are or not present at the time of choice (background context) and the second to the effect of the number of options at the moment of choice (choice context). To achieve this, we explore preferences in three experiments, one in which the birds were trained and tested in binary scenarios, one in which they were trained and tested in a trinary scenario, and a third one in which subjects learnt about three options, but had only two of them available at the moment of each choice (the third option was what is usually referred to as a 'phantom decoy' in marketing research, that is, a known option which is unavailable at the time choice is presented; Pratkanis & Farquhar 1992). The complementary fourth design, training in binary and testing in trinary situations cannot be done because subjects would be ignorant of the properties of the third option at the time of choice.

Variance in rate of energetic gain can be generated by controlling variability either in the amount (or energy density) of food offered or in the delay to receive this food. We chose the second option for two reasons: better control of the knowledge of the decision makers and greater consistency in the existing results from previous binary choice studies. While there is currently no direct way of evaluating an animal's knowledge of amounts (we can only measure preferences between amounts), there are precise procedures to measure an animal's knowledge of temporal parameters. In terms of preferences in binary choices, experiments using amount variability have shown conflicting results: there is a weak tendency towards risk aversion but some studies have reported the opposite and yet others have found indifference. Variability in delay, in contrast, is preferred virtually always.

Overall, there were three variance levels: 'none', 'medium' and 'high'. Experiment 1 establishes the baseline preferences in three binary treatments (NM: none vs. medium; MH: medium vs. high; NH: none vs. high). All subjects experienced the three treatments. In experiment 2 subjects were trained and tested with all three alternatives. In experiment 3 subjects were trained with the three alternatives but tested in binary choices with all three

possible pairs, one pair at a time. Therefore, both in experiment 1 and 3 subjects chose between two alternatives, but in experiment 3 choice pairs were subsets of a larger training set comprising three alternatives. If changes in the strength of preference depend on the number of options at the time of choice, then preferences in experiment 1 and 3 should be similar, but different from those of experiment 2. If they depend on the subject's knowledge of the set of options in the environment, then experiments 2 and 3 should yield similar results, but experiment 1 should differ.

2.2.1 Subjects

The subjects were eighteen wild-caught European starlings, captured in Oxford, UK (licence 19990420, English Nature) as adults on July 1999. After capture the birds were kept in two outdoor aviaries (3 m high x 3.2 m wide x 4 m long) in groups of up to 14 birds each. While in the aviary the starlings were fed on ad libitum food, a mixture of turkey crumbs, Orlux© pellets and mealworms (*Tenebrio sp.*). Drinking and bathing water was always available and replaced daily. Ten days before the beginning of the training sessions the birds were transferred from the outdoor aviary to the laboratory and housed in individual cages (120 cm x 60 cm x 50 cm) that served both as home cages and experimental chambers. Lights in the experimental rooms were gradually switched on at 0600 and off at 1800 h, and temperature ranged from 11°C to 16°C. Subjects were visually but not acoustically isolated. During this period of adaptation all birds had free access to water and the same ad libitum food as received in the aviaries. During the experimental period, the starlings were given free access to turkey starter crumbs after the end of the daily sessions, from 1430 until 1700 h. At 1700 h they were supplemented with 5 mealworms and then food was removed. From previous experience this regime is known to allow the starlings' body weights to remain stable at approximately 90% of their free feeding value or above (Bateson 1993). Fresh drinking water was always available in the cages and bathing trays were provided twice a week. Subjects were assigned to three groups of six members (three males and three females per group), one group being used in each experiment. Each group spent a maximum of 6 weeks in the laboratory, and then was transferred back to the outdoor aviary. The experiments were initiated in November 1999 and completed in March 2000. There was no mortality. All subjects were released back into the wild (University Parks, Oxford, UK) during the spring. Although the experiments were not run simultaneously but over a total period of 15 weeks, we analyse them jointly under the assumption that putative differences between groups would be due to treatment. It is arguable that groups of birds ran at slightly different times

may show different preferences because of seasonal effects, but we consider this extremely unlikely given the kind of effect studied, and the fact that all subjects were tested in winter and after acclimatization to controlled laboratory conditions.

2.2.2 Apparatus

The experimental cages had a panel with a centrally mounted food hopper and two response keys 6 cm from either side of the hopper. The cages had two perches, 85 cm apart from each other. The panels used in experiment 2 had a third response key placed 3 cm above the hopper. An Acorn Risk PC 600 computer running Arachnid® software (Paul Fray Ltd.) controlled the contingencies and collected the data. During the experimental sessions food rewards were units of turkey starter crumbs sieved to a standardized size (mean \pm S.D.: 0.011 ± 0.003 g). The crumbs were delivered by a pellet dispenser (Campden Instruments®) at a rate of 1 unit/s.

2.2.3 Training

After training the subjects to retrieve food from the hopper, an auto-shaping procedure was introduced in which the illumination of a white pecking key (side determined randomly) for 8 s was followed by the delivery of 4 units of turkey crumbs and an intertrial interval (ITI) of 60 s. The birds received 2 daily sessions in this condition (100 trials per session), one at 0700 and another at 1300 h, after which they were given ad libitum food. Once all birds pecked reliably at the keys, food delivery was made conditional on one key pecking. In these sessions, if the birds did not respond during the 8 s interval the key light was switched off and the ITI started. The training program was interrupted when all birds were pecking in at least 80% of the trials.

2.2.4. Experimental Protocol

The levels of variability in delay to reward were: ‘None’: delay=20 s (Coefficient of Variation=0); ‘Medium’: equiprobable delays of either 12 s or 28 s (C.V.=40%) and ‘High’: equiprobable delays of either 4 s or 36 s (C.V.=80%). Variability levels were indicated by the colours transilluminating the pecking keys (pink, orange and blue). Colours were balanced across birds, but for a given bird one colour was always associated with one particular level of variability. Although subjects had participated in an earlier experiment, they had not experienced colours as discriminative stimuli before.

A discrete trials procedure with a fixed ITI of 40 s was employed. There were two types of trials: no-choice trials and choice trials. No-choice trials were intended to provide the birds with information about the mean and variance of each alternative. They started with one of the pecking keys blinking (0.7 s on, 0.3 s off). The first peck caused the light to stay on for the programmed delay, during which pecks on the key were recorded in 1 s bins. The first peck after the delay had elapsed triggered the delivery of 5 units of food, and a new ITI started. Choice trials began with all keys (2 or 3, depending on the experiment) blinking. The first peck on any of them caused the chosen key to turn continuously on for its respective delay and the others to turn off. The first peck after the delay had elapsed triggered the delivery of the standard reward and the beginning of a new ITI.

Subjects experienced two sessions per day, starting at 0700 and 1300 h. Each session comprised a series of blocks, each consisting of a number of no-choice trials followed by one choice trial (details for each experiment are given below). To ensure that all subjects had equal experience of all options each option appeared in two of the no-choice trials within each block, so that both of its outcomes could be experienced. The order (and sides) in which the options were presented in the no-choice trials was randomly determined. Each experiment ended after all subjects had stable preferences (as judged by eye) for at least six sessions, provided they had experienced at least 15 sessions. In experiments 2 and 3 we used the data from the last six sessions (total of 72 choice trials) for analysis. In experiment 1 we adopted the same stopping criterion, but used data from the last four sessions, so as to analyze the same number of choice trials ($n=72$) in each experiment.

In experiment 1 (Training: binary - Testing: binary) each session consisted of 18 blocks of 5 trials each (total of 90 trials per session): 4 no-choice trials (2 levels of variability \times 2 possible outcomes; option 'none', associated with only one outcome, was presented twice, so that subjects had equal experience of each alternative) followed by one choice trial. The order in which each treatment (none vs. medium; medium vs. high; none vs. high) was experienced was balanced across subjects, with each bird experiencing one of the six possible orders. Subjects were given three resting days with ad libitum food between treatments. In experiment 2 (Training: trinary - Testing: trinary) each session comprised 12 blocks of 7 trials each (total of 84 trials per session): 6 no-choice trials (3 levels of variability \times 2 possible outcomes) followed by one choice trial. Experiment 3 (Training: trinary; Testing: binary) used the same schedule as described for experiment 2 except for the

presentation of only two keys in the choice trials. Choice pairs were randomly determined in each block.

2.3 Data Analysis

2.3.1 Preference Scores

Preference for the options was assessed using the proportion of choices in choice trials and the latencies to peck in no-choice trials.

In a binary choice, the null hypothesis states that preferences should not differ from 50%. In a trinary choice, however, the random expectation is one-third for each alternative, preventing a direct quantitative comparison of preference levels between the two scenarios. To overcome this difficulty we modified the procedure used by Tversky and Simonson (1993) and transformed the choice proportions of experiment 2 (trinary choices) by considering only two of the options at a time, so that the percentage of choices for the pair concerned added up to 100%. For instance, preference for ‘none’ relative to ‘medium’ inferred from the choice set {None, Medium, High} and depicted as $P^C(N,M; \{N,M,H\})$ is

$$P^C(N,M;\{N,M,H\}) = \frac{C(N;\{N,M,H\})}{C(N;\{N,M,H\}) + C(M;\{N,M,H\})} \times 100 \quad [1]$$

where N, M and H stand for ‘none’, ‘medium’ and ‘high’, respectively. $C(N;\{N,M,H\})$ is the number of choices for ‘none’ when presented with ‘medium’ and ‘high’ and $C(M;\{N,M,H\})$ is the number of choices for ‘medium’ when presented with ‘none’ and ‘high’. Thus, preference for ‘none’ over ‘medium’ in a trinary context is defined as the number of choices for ‘none’ divided by the number of choices for ‘none’ and ‘medium’ added together. The superscript in P^C indicates that this measure is inferred from choices and not from latencies, as explained later on.

Preferences were also assessed using latencies to peck in no-choice trials, with shorter latencies corresponding to a higher level of preference. We additionally tested whether the relative length of latencies matched preferences as shown in choice trials. Relative preference between pairs of options in no-choice trials was calculated for each subject in the following way:

$$P^L(N, M) = \left(1 - \frac{L_N}{L_N + L_M}\right) \times 100 \quad [2]$$

where $P^L(N, M)$ is the relative preference between ‘none’ and ‘medium’ and L_N and L_M are the subject’s mean latencies to peck ‘none’ and ‘medium’ in no-choice trials. Relative preferences between ‘medium’ and ‘high’ and between ‘none’ and ‘high’ were calculated in the same way.

2.3.2 Context-dependence

One of the most general assumptions of probabilistic theories of choice is that choice probabilities are some function of the underlying utility of the alternatives being considered. As a consequence, preferences are expected to conform to some observable properties that are necessary conditions for the existence of such a utility scale. As previously mentioned, one such property is regularity. Formally, regularity is satisfied if the probability of taking an item x from a subset $A = \{x, y\}$ that is part of $B = \{x, y, z\}$ is such that, for all z ,

$$P^C(x; \{x, y\}) \geq P^C(x; \{x, y, z\}) \quad [3]$$

In words, the condition asserts that the proportion of choices for an option cannot be increased when a new option is included in the choice set. The possibility of this condition being violated is dependent on the proportion of choices allocated for the new alternative z . If, for instance, z is chosen in a large proportion of the trials, the total number of choices for x and y together will be very small, and an absolute increase in the proportion of choices for x (i.e. a violation of regularity) could be only achieved if accompanied by a large change in the relative attractiveness between x and y (so that z takes almost exclusively from y).

Another property that results from assuming the existence of a utility scale upon which preferences are based is known as ‘independence from irrelevant alternatives’ (IIA). IIA is best captured in the choice axiom of Luce (1959), also known as the constant-ratio rule (Clarke 1957), a strong probabilistic version of IIA (Luce & Suppes 1965) that states that the ratio of the probabilities of any two alternatives is entirely independent from the choice set. As a consequence of the axiom, the inclusion of a new option into a choice set is assumed to take from pre-existing options in proportion to their original shares (Huber et al. 1982). By considering the relative shares of only two options at a time, the use of preferences as calculated by equation 1 therefore allows us to compare relative preferences between experimental contexts and test IIA. Here we used as the critical test for the

detection of context effects the interaction between Experimental context (Experiments 1, 2 and 3) and Choice pair (none-medium, medium-high, none-high), as it tests the extent to which context influences the strength of preference (as given by equation 1) across the three possible choice pairs, and can detect preference reversals and relative changes in proportions between experimental contexts. The same test was used to verify whether preferences, as revealed by latencies, changed between contexts. Yet, some caveats need being raised when comparing relative preferences between contexts for testing context-dependence in the assignment of value to options.

Underlying the logic of the choice axiom is the assumption of ‘scalability’, or ‘proportionality’, that is, that the probability of choosing an option is directly proportional to its value (in fact, Tversky 1972 shows that the choice axiom is a special case of simple scalability). Similarly to the choice axiom, equation 1 predicts that the relative preference between two options in a binary context is unaltered when a third option is included in the choice set only if preference for the original options is proportional to their value. However, it is easy to imagine situations in which the relative preference for an option would differ between a binary and trinary context because proportionality is not satisfied, and yet the value assigned to the options is fixed and independent of the choice set (see Debreu 1960 for some examples). As Luce & Suppes (1965) highlight, the validity of simple scalability (and therefore of the choice axiom) can be questioned whenever alternatives are perceived to be similar. Similarly, if animals use a maximizing principle only “choosing” one option, with elections for other options representing a constant amount of ‘errors’ or noise between contexts, then relative preferences would be distinct in a trinary situation without implying that the value, or attractiveness, of an option is dependent on the context. Consequently, caution should be shown when interpreting the mechanisms causing significant differences in relative preferences between contexts.

2.3.3 Transitivity

We analysed both weak stochastic transitivity (i.e. presence of cycles, WST) and strong stochastic transitivity (SST) with the data from the choice trials and no-choice trials. SST states that if alternatives are ranked in an order such that ‘a’ is preferred to ‘b’, and ‘b’ to ‘c’, then preference for ‘a’ over ‘c’ should exceed or equal the strongest of the two other preferences, that is $p(a,c) \geq \max [p(a,b), p(b,c)]$, where $p(a,b)$ is the probability of choosing option ‘a’ from the choice set (a,b) (Fishburn 1991).

2.4. Results

2.4.1 Knowledge about the alternatives

To characterize the accuracy of the subjects' knowledge of the temporal properties of each option we analysed the pecking patterns during the delays experienced in the no-choice trials, by computing number of pecks in 1s time bins (Fig. 1). Only the long delays of the two variable options were included. For the option with no variability we used a random half of the data points, so that sample sizes were the same for all variability levels. For presentation economy only the mean peck rate is shown in the graph, but the figure represents well the subjects' individual functions.

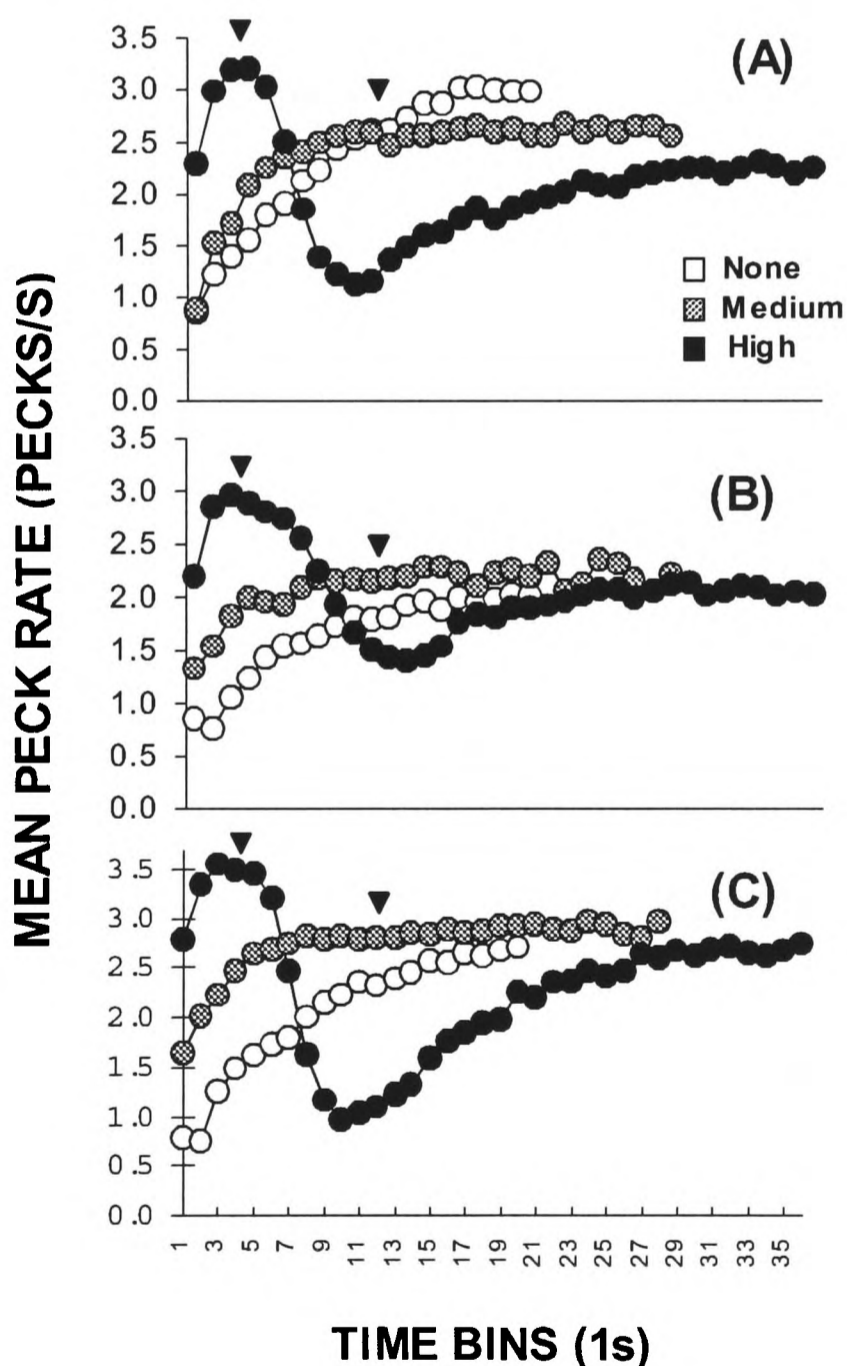


Figure 1. Mean peck rate (pecks/s) for each of the variability levels in experiments 1, 2 and 3 (Figs. 1A, 1B and 1C, respectively). The arrows point to the time bins where food was potentially due (4 s, 12 s).

In the three experiments, the function depicting the response rate for ‘high’ shows a clear peak slightly before time bin 4, when a reward could be potentially obtained, followed by a steep decline and then a gentler rise towards the end of the 36 s delay, when reward was due again. We tested whether the first peak (observed time bin corresponding to maximum peck rate) differed between contexts, but could not detect any significant difference (Kruskal-Wallis test: $\chi^2_2 = 3.05$, $P = \text{NS}$). We also tested whether the peck rate when food was potentially due (time bins 4 and 36) changed between contexts, but did not find significant differences in either case (Repeated-measures ANOVA for 4 and 36, respectively: $F_{2,10} = 0.45$, $P = \text{NS}$; $F_{2,10} = 0.90$, $P = \text{NS}$). From the peck pattern it is therefore possible to infer that starlings were capable of correctly timing the delays associated with ‘high’, and that the experimental context did not noticeably change such perception.

Response rate for ‘medium’ does not show the discernible peaks seen for ‘high’. This is not surprising: according to the scalar property of behavioural timing, the spread of the response distribution is proportional to the length of the interval being timed, and therefore the relative distance between two intervals determines the distinctiveness of peaks of responding at the corresponding reinforcement times (Gibbon 1977). In the present experiments, the ratio between the two delays corresponding to intermediate variability (12 s and 28 s) falls in the range where no discernible peaks of responding were observed by other authors (Catania & Reynolds 1968, Leak & Gibbon 1995). Peck rates at the time bins when food was potentially due (12 and 28) were not different between experimental contexts (Repeated-measures ANOVA for 12 and 28, respectively: $F_{2,10} = 0.457$, $P = \text{NS}$; $F_{2,10} = 0.945$, $P = \text{NS}$).

Finally, the response curve associated with ‘none’ gradually rises towards 20s, the end of the interval. The slope of the curve was not significantly different between the three contexts (Repeated-measures ANOVA, square-root transformed data: $F_{2,10} = 0.123$, $P = \text{NS}$), nor was the peck rate at the last time bin (20), when food was imminent (Repeated-measures ANOVA: $F_{2,10} = 1.44$, $P = \text{NS}$).

We additionally tested whether peck rates differed between treatments in experiment 1. As in this experiment we performed all three possible binary treatments, we obtained two response curves for each variability level (e.g. peck rate for ‘high’ was obtained from the trials experienced in treatments (none, high) and (medium, high)). Therefore, for each variability level we tested whether the peck rate differed between treatments in the same

way as that described to detect differences between experimental contexts. No significant difference between treatments was detected in the pecking pattern for ‘high’ (peck rates at time bins 4 and 36, respectively, paired-t tests: $t_5 = -2.12$, $P = \text{NS}$ and $t_5 = -0.34$, $P = \text{NS}$; time bin corresponding to max. peck rate: $t_5 = 0.67$, $P = \text{NS}$), ‘medium’ (peck rate at time bins 12 and 28, respectively: $t_5 = 1.58$, $P = \text{NS}$ and $t_5 = 1.92$, $P = \text{NS}$) or ‘none’ (homogeneity of slopes, square-root transformed data: $t_5 = -0.504$, $P = \text{NS}$; peck rate at time bin 20: $t_5 = -1.82$, $P = \text{NS}$).

The pecking patterns thus suggest that starlings were able to recall the time of occurrence of rewards, and that their perception of each option was not distorted by the number of foraging alternatives either experienced or available for choice. Response functions were also qualitatively consistent with the scalar property as they were progressively shallower as the delays increased.




Figure 2. Mean percentage of choices (\pm S.E.) for the most variable option of each choice pair in each of the three experiments. Relative proportions in experiment 2 (trinary-trinary) are given by equation (1).

2.4.2 Choice trials

2.4.2.1 Preference for variability

Figure 2 shows the mean percentage of choices for the most variable option in each choice pair for each experimental context. Absolute preferences for each option in the three experimental contexts are shown in figure 3. Subjects were strongly prone to risk in all contexts. In experiments 1 and 2 (hereafter binary-binary and trinary-trinary, respectively)

all subjects significantly preferred (binomial tests, $P < 0.05$) the most variable option from the choice pairs (medium, high) and (none, high), with two and four birds being indifferent between the options ‘none’ and ‘medium’ in the binary-binary and trinary-trinary contexts, respectively. In experiment 3 (hereafter trinary-binary) subjects were also risk-prone, significantly preferring the most variable option in all choice pairs, with the only exceptions of one bird that did not reach significance in the (medium, high) pair and another that did not reach significance in the (none, medium) pair. Overall, these results reveal a very strong preference for variability in delay to reward when the choice set includes the high variance option, and a more moderate preference when subjects face a decision between ‘none’ and ‘medium’.




Figure 3. Mean percentage of choices (\pm S.E.) for each option in each of the three experiments (choice sets: (N,M), (M,H), (N,H) and (N,M,H), where N stands for ‘none’, M for ‘medium’ and H for ‘high’).

2.4.2.2. Context-dependence

To test the interaction between experimental context and the level of preference expressed in each choice pair we conducted a repeated-measures ANOVA with ‘choice pair’ as the within-subject, random factor, ‘experimental context’ as the between-subject, fixed factor and proportion of choices (arcsine square root transformed) for the most variable option as the dependent variable. The interaction between experimental context and choice pair was not significant ($F_{2.8, 21.3} = 0.106$, $P = 0.95$, Greenhouse-Geiser correction for sphericity). To minimize the probability of incurring a type II error (i.e. accepting a false null hypothesis) we conducted a compromise power analysis (Faul & Erdfelder 1992, Erdfelder et al. 1996)

to determine the critical F-ratio at which the test would have a power of at least 70% of detecting a medium effect size (Cohen 1988). This was done by setting the ratio of the type I error probability (α) over the type II error probability (β) at 1 (i.e. same probability of making either error type). This yields a critical F-ratio of $F_{2.8, 21.3} = 1.02$, above which one should reject the null hypothesis. Although this procedure greatly increases the likelihood of rejecting the hypothesis of consistence among contexts by increasing α , the results were still far from the required values, indicating that preferences in each possible choice pair were not reliably different across experimental contexts. Although preferences were not significantly different, we calculated eta-square values – indicators of the effect-size or amount of variance explained by the interaction – to estimate the actual differences between experimental contexts. The effect size was extremely small (eta-squared=0.011, see Cohen 1988 for effect size conventions), indicating a high level of similarity in the subjects' preferences for variability across the three contexts.

Additionally, subjects did not significantly increase the absolute proportion of choices for any of the options in the trinary-trinary context. Out of the 12 possible comparisons (comparison of choice proportions for each option in the contexts binary-binary and trinary-binary with those in the trinary-trinary), there was a significant decrease in the absolute proportion of choices in 4 cases in the trinary-trinary context (Student t-test for independent samples; $P < 0.05$ following Bonferroni adjustments for multiple comparisons), but never an increase. These tests were again conducted taking into account the results of a compromise power analysis (power of at least 70% for detecting a medium effect size, α/β ratio = 1, critical t: $t_{10} = 2.12$), but none of the results allowed us to reject regularity.

2.4.2.3 Transitivity

Preferences as revealed in choice trials satisfied WST for all subjects, analysed individually. Starlings preferred 'high' to 'medium', 'medium' to 'none' and 'high' to 'none', therefore ranking options in a scale such that 'high' > 'medium' > 'none'. This ordering was consistent for all birds in the three experimental contexts. Quantitatively, individual preferences also conformed to SST. For none of the birds preference for 'high' in the pair (none, high) was significantly lower than in the other two choice pairs ($P < 0.05$, log-likelihood G-tests, power ranged from 61% to 99%).

2.4.3 No-choice trials

2.4.3.1 Preference for variability

As we performed all three possible binary treatments in experiment 1, we obtained two measures of latency for each variability level. For example: latency to peck ‘none’ was obtained from the trials experienced in the (1) ‘none’ vs. ‘medium’ and (2) ‘none’ vs. ‘high’ treatments. Therefore, for each variability level we tested whether latencies during no-choice trials were different between the two treatments in which they occurred, as the preference for a particular variability level could be affected by the nature of the option with which it was paired. No significant difference between treatments was detected (Student t-test for ‘none’, ‘medium’ and ‘high’, respectively, log-transformed data: $t_5 = -0.979$, $P = \text{NS}$; $t_5 = 1.123$, $P = \text{NS}$; $t_5 = 0.112$, $P = \text{NS}$). The two latency replicates obtained for each option were thus pooled in subsequent analyses.

Figure 4 shows average latencies to accept each option in the no-choice trials. For analysis, latencies were log-transformed and a repeated-measures ANOVA performed on each experimental context. Latencies were significantly different between variability levels in all experiments (Experiments 1, 2 and 3, respectively: $F_{2,10} = 13.06$, $P < 0.01$; $F_{2,10} = 28.70$, $P < 0.001$; $F_{2,10} = 12.26$, $P < 0.01$). Post-hoc comparisons (Bonferroni) show that the latency to peck ‘high’ was significantly shorter than that to peck ‘none’ and ‘medium’ ($P < 0.05$ in all cases). Latency to peck ‘medium’ was also shorter than that to peck ‘none’ in experiment 3, but not in experiments 1 and 2. Overall, these results show that in the no-choice trials subjects consistently responded earlier to ‘high’ than to ‘medium’ and ‘none’, and that latency difference for ‘medium’ in relation to ‘none’ was not so extreme. There was also a significant positive correlation between the relative preference as calculated by equation 2 and that shown in the choice trials (Relative preference for (none,medium), (medium,high) and (none,high), respectively, $n = 18$ subjects (contexts pooled) in all cases: $r^2 = 0.62$, $P < 0.001$; $r^2 = 0.21$, $P = 0.051$; $r^2 = 0.33$, $P < 0.01$). Starlings’ proneness to respond in the no-choice trials was thus consistent with preferences observed in the choice trials, confirming the reliability of latencies as a measure of preference.




Figure 4. Mean latency to peck (s) (\pm S.E.) each of the variability levels in each of the three experiments.

2.4.3.2 Context-dependence

We also tested whether preferences, as revealed by latencies, were consistent across contexts in the same way as previously described for choice proportions. The critical test for the detection of context effects – the interaction between Experimental Context and Variability level – was not significant ($F_{4,30} = 1.158$, $P = \text{NS}$). To minimize the possibility of a type II error, we conducted a compromise power analysis in the same way and adopting the same parameters as described for the analysis of context-dependence with the choice data (α/β ratio = 1). The analysis yielded a critical F-ratio of $F_{4,30} = 1.21$, still not enough to reject the null hypothesis although increasing the type I error probability (α). These results indicate that the effect variability had on latencies was not reliably different across contexts.

2.4.3.3 Transitivity

All subjects satisfied WST, ranking options (according to readiness to peck) in a scale such that response was earlier for ‘high’ than for ‘medium’, for ‘medium’ than for ‘none’ and for ‘high’ than for ‘none’. SST was checked for each subject using the relative preference as given by equation 2. This analysis again confirms SST, as for none of the birds was relative preference for ‘high’ as compared to ‘none’ lower than in the other two possible pairs.

2.5 Discussion

In three experiments starlings showed a regular and consistent pattern of preference for variability in delay to reward, which was unaltered by whether the training or the choice context had two or three alternatives. Preferences satisfied both weak and strong forms of stochastic transitivity. These results extend the generality of previous research on risk-sensitivity to choices with more than two alternatives, and suggest that violations of rationality in foraging contexts may be restricted to a specific set of parameters and conditions.

In terms of the direction of risk preference shown by the starlings, the results are consistent with previous research and theories of risk-sensitivity based on the cognitive processes underlying the assignment of value to each option. For instance, both the local rate maximization model (Caraco et al. 1992, Bateson & Kacelnik 1996) and scalar utility theory (Kacelnik & Brito e Abreu 1998) lead to risk proneness for delays, as observed. Preference levels were nonetheless more extreme than those predicted by the simplest form of SUT, which is based on each decision being based on a single random sample from memory (see also Bateson & Kacelnik 1995). This failure of the model to predict extreme risk proneness for delays (while being successful in predicting the sign of preference or certainty equivalent) is pervasive, and discussed elsewhere (Kacelnik & Brito e Abreu 1998).

In quantitative terms, the strength of preference between ‘none’ and ‘medium’ was the less extreme and between ‘none’ and ‘high’ the most extreme. According to local rate maximization or expectation of the ratios (EoR), an option’s computed attractiveness can be defined as the expected ratio of amount of food (A) over the time (t) taken to acquire it such that $EoR = E\left(\frac{A}{t}\right)$. EoR does not provide an efficient algorithm to rank options according to their rates of reward, but it has been shown to be fairly robust in accounting for observed choices (Mazur 1984, Mazur 1986, Bateson & Kacelnik 1996). The computed EoR for ‘none’, ‘medium’ and ‘high’ (excluding ITI and adding reward delivery time) were respectively 0.20, 0.22 and 0.34 units/s, hence one could have expected a similar pattern if preference were an increasing function of EoR.

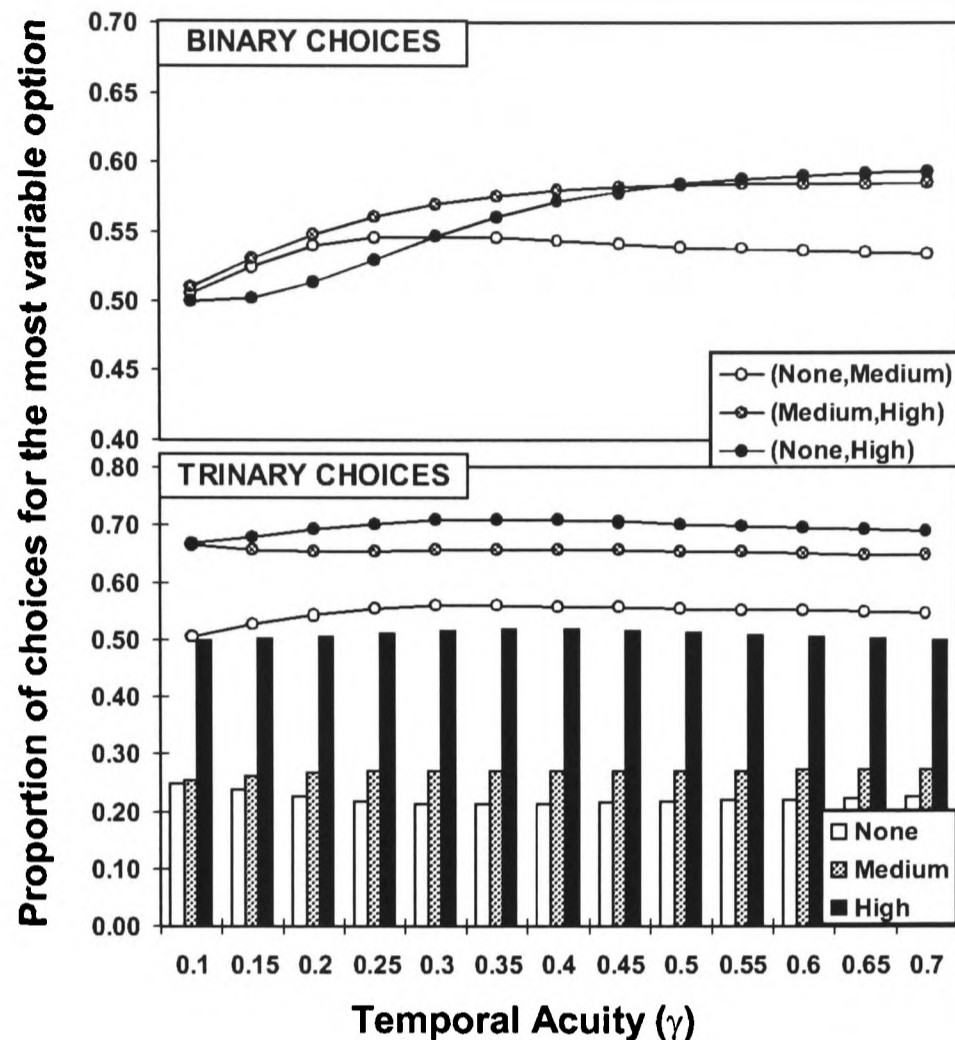


Figure 5. Predictions of scalar utility theory (SUT; see Kacelnik & Brito e Abreu 1998 for a detailed description) for choices between three levels of variability in delay to reward ('none': 20s, 'medium': equiprobable delays of 12 and 28 s and 'high': equiprobable delays of 4 and 36 s) in binary and trinary contexts. Lines represent the predicted proportion of choices for the most variable option in each possible choice pair (relative proportions in the trinary context are given by equation 1). Bars in the trinary context represent the absolute proportion of choices for each variability level.

The predictions of SUT are, on the other hand, dependent on the temporal acuity (γ) of the subjects, and violate rationality under some reasonable circumstances (*M. Bateson, personal communication*). Such predictions are in this sense interesting, as while many sensible decision mechanisms can produce irrationalities when options are defined in more than one dimension (see Wedell 1991 for a review), few do so for options that differ in only one dimension (e.g. Parducci 1965, Janiszewski & Lichtenstein 1999). Figure 5 illustrates SUT predictions for the variability levels used in this study. Preferences as shown in the figure conform to weak stochastic transitivity in both binary and trinary scenarios. Strong stochastic transitivity is however violated in binary contexts whenever γ is lower than 0.5 (since preference between the extreme options 'none' and 'high' is lower than in the other two choice pairs). Relative preferences (bars in the graph, as calculated by equation 1) are also higher in the trinary than in the binary situation, and an interaction between

experimental context (binary or trinary) and choice pair (as that used as a test of context-dependence in this study) should be expected for $\gamma < 0.5$, since in such cases the relative strength of preference between the choice pairs depends on whether the context is binary or trinary. Such predictions are nevertheless not confirmed but the results reported here. Estimates of γ for starlings have been previously obtained by using production tasks, in which the subjects reproduce a behavioural index of a previously experienced time interval. On the basis of available evidence (Brunner et al. 1992, Kacelnik & Brunner 2002) it can be assumed that for time estimates $0.1 < \gamma < 0.6$. Therefore, according to SUT, context effects and violations of strong stochastic transitivity are to be expected.

If multialternative situations per se do not bring about irregularities in behaviour it is possible that the apparent abnormalities reported by other authors derived from features of their experimental design that differed from ours. One possibility relates to the attribute of the alternatives that was variable in each study. Hurly & Oseen (1999) tested preference for variability in amount, whereas we tested preference for variability in delay to reward. The direction of risk preference is the reverse between these dimensions (Kacelnik & Bateson 1996). In addition, preference levels tend to be higher and more consistent for delays, but a direct comparison is difficult because levels of preference are related to the difference in the coefficient of variation between options (Shafir 2000), and few studies have equalized coefficient of variations in amount and delays. Hurly and Oseen used coefficients of variation for the no variance, intermediate variance and high variance options of 0%, 33.3% and 66.6% respectively, slightly lower but in the same broad range than ours (respectively 0%, 40% and 80%). The fact that preference levels tend to be less extreme when variability is in food amount and difference in variance is lower also means that it may have been easier to detect irrationalities under these conditions. Hurly and Oseen do not inform in their study what the percentage of choices for each alternative was, but it is possible to infer from their graphs that they were not as high as ours. As we worked with delays, we found a strong preference for variability, and this might have dwarfed potential weak context effects.

Another difficulty in comparing with previous work is that when dealing with amounts (Shafir 1994, Hurly & Oseen 1999, Waite 2001a, Waite 2001b, Bateson et al. 2002, Shafir et al. 2002) or other dimensions of reward such as flower corolla length (Shafir 1994, Shafir et al. 2002) or distance to a reward positioned inside a tube (Waite 2001a, Waite 2001b) it is difficult to precise whether subjects have accurate knowledge of the

programmed outcomes and their variability. While we showed that the starlings were able to correctly time the delay to food independently of the experimental context, it is difficult to know whether the information learnt by the hummingbirds (Hurly & Oseen 1999), bees (Shafir 1994, Shafir et al. 2002) or jays (Waite 2001a, Waite 2001b) in studies showing violations of a rationality principle differed between choice contexts.

Ensuring that experimental subjects acquire the same amount of information about alternatives before preferences are tested can be important because predicted choices may depend on the subject's knowledge. As Houston (1997) points out, the expected fitness of each choice in a sequence depends on the environment to be faced later. If the subject takes the alternatives present in each choice as a cue to what is to be found in the future, preference between options presented in pairs may be non-transitive because each pair of alternatives may be seen as a predictor of a different future. The argument, yet, does not apply when the future to be faced is the same, as its impact on present choices is the same whatever the choice pair concerned. In the trinary-binary experiment, the future scenario that is inferable by the birds is identical (i.e. the future contains the three options) regardless of the type of binary choice presented, hence having the same impact on present choice. Thus, by ensuring that the shadow of different futures did not influence the results we may have additionally prevented a violation of rationality as predicted in Houston's terms.

Students from areas as diverse as psychology, marketing and political sciences see violations of normative axioms as useful tools for the investigation of the cognitive processes underlying decision-making, and as devices to aid marketing and political campaigning by incorporating strategic factors in the design and presentation of new products and candidates. In the case of animal decision-making, the widespread occurrence of irrational choice in the wild could pose serious difficulties for the integration of functional and mechanistic analysis, and hence it is important not only to document their occurrence, but also their prevalence across distinct paradigms and species. Our results do not pose major challenges to previous findings and their prevalent theoretical explanations, but they are negative with respect to irrationality. This of course does not mean that irrationalities are rare, only that they are not an obvious problem for the investigation of sensitivity for time-induced risk on the basis of binary choices.

2.6 References

- Bateson, M. 1993. Currencies for decision making: the foraging starling as a model animal. Oxford: University of Oxford. Unpublished DPhil. Thesis.
- Bateson, M., Healy, S.D. & Hurly, A. 2002. Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, 63: 587-596.
- Bateson, M. & Kacelnik, A. 1995. Preferences for fixed and variable food sources: variability in amount and delay. *Journal of Experimental Analysis of Behavior*, 63: 313-329.
- Bateson, M. & Kacelnik, A. 1996. Rate currencies and the foraging starling: the fallacy of the averages revisited. *Behavioral Ecology*, 7: 341-352.
- Brunner, D., Kacelnik, A. & Gibbon, J. 1992. Optimal foraging and timing processes in the starling, *Stumus vulgaris*: effect of inter-capture interval. *Animal Behaviour*, 44: 597-613.
- Caraco, T., Kacelnik, A., Mesnik, N. & Smulewitz, M. 1992. Short-term rate maximization when rewards and delays covary. *Animal Behaviour*, 44: 441-447.
- Catania, C.A. & Reynolds, G.S. 1968. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of Experimental Analysis of Behavior*, 11: 327-385.
- Clarke, F.R. 1957. Constant-ratio rule for confusion matrices in speech communication. *Journal of the Acoustical Society of America*, 29: 715-720.
- Cohen, J. 1988. *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Debreu, G. 1960. Review of R.D.Luce: Individual choice behaviour: A theoretical analysis. *American Economic Review*, 50: 186-188.
- Erdfelder, E., Faul, F. & Buchner, A. 1996. GPower: a general power analysis program. *Behavior Research Methods, Instruments and Computers*, 28: 1-11.
- Faul, F. & Erdfelder, E. 1992. GPOWER: A priori, post-hoc, and compromise power analyses for MS-DOS [Computer program]. Bonn, FRG: Bonn University.
- Fishburn, P.C. 1991. Nontransitive preferences in decision-theory. *Journal of Risk and Uncertainty*, 4: 113-134.
- Gibbon, J. 1977. Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, 84: 279-325.
- Gigerenzer, G., Todd, P.M. & the ABC Research Group. 1999. *Simple heuristics that make us smart*. New York: Oxford University Press.
- Houston, A.I. 1997. Natural selection and context-dependent values. *Proceedings of the Royal Society of London Series B-Biological Sciences*, 264: 1539-1541.
- Houston, A.I., Kacelnik, A. & McNamara, J. 1982. Some learning rules for acquiring information. In: *Functional Ontogeny* (Ed. by McFarland, D.J.), pp. 140-191. Boston: Pitman.
- Huber, J., Payne, J.W. & Puto, C. 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9: 90-98.
- Hurly, T.A. & Oseen, M.D. 1999. Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal Behaviour*, 58: 59-66.
- Janiszewski, C. & Lichtenstein, D.R. 1999. A range theory account of price perception. *Journal of Consumer Research*, 25: 353-368.
- Kacelnik, A. & Bateson, M. 1996. Risky theories - the effects of variance on foraging decisions. *American Zoologist*, 36: 402-434.
- Kacelnik, A. & Brito e Abreu, F. 1998. Risky choice and Weber's Law. *Journal of Theoretical Biology*, 194: 289-298.
- Kacelnik, A. & Brunner, D. 2002. Timing and foraging: Gibbon's scalar expectancy theory and optimal patch exploitation. *Learning and Motivation*, 33: 175-195.

- Leak, T.M. & Gibbon, J. 1995. Simultaneous timing of multiple intervals: implications of the scalar property. *Journal of Experimental Psychology - Animal Behavior Processes*, 21: 3-19.
- Luce, R.D. 1959. *Individual Choice Behaviour: A Theoretical Analysis*. New York: Wiley.
- Luce, R.D. 1977. The choice axiom after twenty years. *Journal of Mathematical Psychology*, 15: 215-233.
- Luce, R.D. & Suppes, P. 1965. Preference, utility, and subjective probability. In: *Handbook of Mathematical Psychology* (Ed. by Luce, R.D., Bush, R. & Galanter, E.). New York: John Wiley.
- Mas-Collel, A., Whinston, M.D. & Green, J.R. 1995. *Microeconomic Theory*. New York: Oxford University Press.
- Mazur, J.E. 1984. Tests of an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology - Animal Behavior*, 10: 426-436.
- Mazur, J.E. 1986. Fixed and variable ratios and delays: further tests of an equivalence rule. *Journal of Experimental Analysis of Behavior*, 12: 116-124.
- Parducci, A. 1965. Category judgement: a range-frequency model. *Psychological Review*, 72: 407-418.
- Pratkanis, A.R. & Farquhar, P.H. 1992. A brief-history of research on phantom alternatives - evidence for 7 empirical generalizations about phantoms. *Basic and Applied Social Psychology*, 13: 103-122.
- Shafir, E.B., Osherson, D. & Smith, E.O. 1989. An advantage model of choice. *Journal of Behavioral Decision Making*, 2: 1-23.
- Shafir, S. 1994. Intransitivity of preferences in honey-bees - support for comparative-evaluation of foraging options. *Animal Behaviour*, 48: 55-67.
- Shafir, S. 2000. Risk-sensitive foraging: the effect of relative variability. *Oikos*, 88: 663-669.
- Shafir, S., Waite, T.A. & Smith, B.H. 2002. Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 51: 180-187.
- Simonson, I. & Tversky, A. 1992. Choice in context: trade-off contrast and extremeness aversion. *Journal of Marketing Research*, 29: 281-295.
- Stephens, D.W. & Krebs, J.R. 1986. *Foraging Theory*. Princeton: Princeton University Press.
- Tversky, A. 1969. The intransitivity of preferences. *Psychological Review*, 76: 31-48.
- Tversky, A. 1972. Elimination by aspects: A theory of choice. *Psychological Review*, 79: 281-299.
- Tversky, A. & Simonson, I. 1993. Context-dependent preferences. *Management Science*, 39: 1179-1189.
- Waite, T.A. 2001a. Background context and decision making in hoarding gray jays. *Behavioral Ecology*, 12: 318-324.

CHAPTER 3

3.1 INTRODUCTION	28
3.2 STRUCTURED CHOICE SETS	31
3.3 PRESENCE OF UNCORRELATED CHOICES	32
3.3.1 THE SAMPLING CASE	32
3.3.2 THE "NOISE" CASE	33
3.4 CONSTRAINTS ON PERCEPTION	34
3.5 CONCLUSION	39
3.6 REFERENCES	40

LUCE'S CHOICE AXIOM AND CONTEXT-DEPENDENCE IN ANIMAL CHOICES

3.1 Introduction

Nowadays the study of decision-making is an inevitable part of a broad range of disciplines. While in economics it is essential to understand and model decisions to predict the behaviour of markets, investors and organizations, in psychology and the social sciences the cognitive aspects underlying an individual's choice behaviour are of major interest. Similarly, evolutionary and behavioural biology have progressed, on the one hand, by focusing on the role of natural selection on the evolution of the organisms' preferences and use of choice rules in a wide range of situations, and on the other, by analysing the effects of animals' preferences on the evolution of the selected traits.

A general assumption shared by choice theories in all fields is that individuals should behave so as to maximize the expected value of a given utility function under the constraints they face¹. This means that preferences should reflect the positioning of the options on a scale determined by the consequences of each choice for the decision-maker, for example in terms of 'monetary gains', 'satisfaction' or, in the case of biology, expected fitness gain.

In spite of its intuitive appeal, the assumption that individuals are perfect maximizers in the sense previously defined is usually challenged by the observation that they often fail to exhibit all-or-none choices on repeated choices, selecting different alternatives under conditions that are supposedly identical. If on the one hand it is possible that such inconsistencies result from individuals responding to different aspects of the choice problem, on the other in most cases these aspects cannot be identified, and choice behaviour is viewed as a probabilistic process whereby the choice probability of each option is expressed as some function of its underlying scale value, so the idea of a numerical scale of value governing choices is maintained.

Within this framework, a major concern of empiricists has been the investigation of a number of properties that would be necessary conditions for the existence of such a scale. One such property underlies most theoretical work on probabilistic choice, and is known as

¹It is important to notice, however, that the notion of 'utility' might be used in conceptually different ways depending on the field of research. Thus, while in economics utility is generally used to describe the function an agent is maximizing, based on the observation of her past choices, in biology it is commonly employed normatively, to predict the preferences individuals ought to have given some functional criteria (Krebs & McCleery 1984).

'independence from irrelevant alternatives' (IIA), namely the expectation that the ratio of the choice probabilities of any two alternatives is entirely unaffected by the utilities of other options (Ben-Akiva & Lerman 1985). This is equivalent to saying that the preference between alternatives should be independent of the context (as defined by the opportunity set) in which they occur. IIA is best captured in the choice axiom of Luce (1959), also known as the 'constant ratio rule' (Clarke 1957), a central model to most probabilistic theories of choice behaviour. The axiom predicts that the inclusion of a new option into a choice set should take from pre-existing options in proportion to their original shares, and thus relative choice proportions between any pair of options (i.e., the ratio of choice probabilities of any two alternatives) should be the same in all contexts. Formally, this means that the choice axiom is satisfied if the probability of taking an item A from a set $B_i = \{A, B\}$ is the same as the relative probability of taking A as opposed to B from a set $Tri = \{A, B, C\}$, such that for all C

$$\frac{p_{B_i}(A)}{p_{B_i}(A) + p_{B_i}(B)} = \frac{p_{Tri}(A)}{p_{Tri}(A) + p_{Tri}(B)} \quad [1]$$

where $p_{B_i}(A)$ and $p_{Tri}(A)$ represent the probability of taking A from the binary and trinary choice sets, respectively.

Breaches of the choice axiom have however been frequently reported in the psychological and economic literature (Huber et al. 1982, Wedell 1991, Payne et al. 1992, Simonson & Tversky 1992, Tversky & Simonson 1993, Slaughter et al. 1999), and more recently in studies of animal behaviour. For instance, studying the foraging preferences of honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*), Shafir et al. (2002) showed that the relative preference of the bees and jays between two original foraging options changed with the introduction of an additional, unattractive alternative. Similarly, Bateson et al. (2002) analysed the preferences of wild rufous hummingbirds (*Selasphorus rufus*) for artificial flowers, showing that the relative preferences of the birds for one of the two target options was higher in a trinary context (where a third option, or 'decoy', was available) than in binary choices (only two options on offer). The same was true in a study designed to test the risk-sensitive foraging preferences of European starlings (*Sturnus vulgaris*) for variability in amount of food (Bateson 2002): the relative preferences of the

starlings for the less variable options were always higher in a trinary than in a binary context.

Violations of the choice axiom, such as the above mentioned, are usually taken to imply the use of 'comparative' mechanisms of choice, whereby the value of an option is not only determined by its intrinsic properties and consequences as assumed by the normative modelling in biology, but instead is constructed at the moment of choice according to the presence and attributes of other options² (Shafir 1994, Hurly & Oseen 1999, Waite 2001, Bateson 2002, Bateson et al. 2002, Shafir et al. 2002; see also Shafir et al. 1989, Wedell 1991, Payne et al. 1992 for some possible accounts of context-dependence based on the use of comparative mechanisms of choice). For example, discussing the changes in the relative level of preference of the bees and jays for the available foraging options, as mentioned above, Shafir et al. (2002) conclude: "*Our findings support a 'comparative' method of evaluating options. According to this method, options are not evaluated separately and assigned fixed utilities, but rather the relative value of an option depends on the options with which it is compared*". Similarly, Bateson et al. (2002) interpret their results as being "*hard to reconcile with the use of absolute [independent] evaluation mechanisms, and suggest that animals may well evaluate alternative options using comparative mechanisms, whereby the value assigned to a given option is dependent on the other options available at the time the evaluation is made*".

My objective here is however to argue that, in the case of violations of the choice axiom, the previous interpretation is not unique, as a number of models compatible with the assumption of independent choice currencies can also lead to such violations in the absence of context-dependence as defined before. I start by briefly reviewing some of the arguments put forth in the psychological literature on choice behaviour and by considering instances where choices are uncorrelated to the properties of the rewards being offered. I then concentrate the discussion on mechanistic models of choice that incorporate constraints on perception.

² The conclusion follows from the assumption that the ratio of the choice probabilities (say, 'p') of any two alternatives is some function of the ratio ('v') of their scale values, and thus any observed change in 'p' would imply that the presence of additional alternatives enhanced – or weakened – the relative attractiveness between the existing options ('v').

3.2 Structured Choice Sets

The validity of the choice axiom has some well-recognized difficulties whenever the choice set can be structured as a function of the distinct degrees of similarity (or substitutability) between its component options, or when a hierarchy of decisions is involved in the choice process (Luce & Suppes 1965, Tversky & Russo 1969, Tversky 1972). Consider for example the well-known possibility illustrated in the 'red bus–blue bus' paradox (Debreu 1960). Suppose a decision-maker (DM) has to decide on whether to take a bus or a car, and that she is indifferent between these options. Now let her choice possibilities include an additional alternative, so that at this time she faces a choice between a car, a blue bus and a red bus. A natural expectation if DM is also indifferent between the blue and red buses, is that the probability of choosing a car will be $P(\text{car}) = \frac{1}{2}$, and that $P(\text{blue bus}) = P(\text{red bus}) = \frac{1}{4}$. However, since DM is indifferent between the options in pairwise choices, the choice axiom requires the probability of choosing any option to be the same when the three alternatives are presented simultaneously, that is to say, $P(\text{car}) = P(\text{red bus}) = P(\text{blue bus}) = \frac{1}{3}$. The paradox thus shows that preferences in choice situations with more than two alternatives may violate the axiom owing to the use of choice rules that group similar alternatives together, a possibility later developed by Tversky & Russo (1969) and known as the 'similarity hypothesis'. According to it, a third alternative introduced into a binary choice set would take its share of choices mainly from the option most similar to it, and therefore the relative preference for the latter would be reduced.

Tests of the similarity hypothesis confirmed it in a number of situations (see e.g. Tversky & Gati 1978), but were mostly restricted to experiments dealing with human choices where a natural grouping of the options could be unambiguously established. In the case of animal choices, however, a priori knowledge of the degree of equivalence between the options for the subjects may be harder to ensure, especially in those cases where options are composed of more than one attribute. Yet, if on the one hand unequivocal tests of the hypothesis may be more difficult to design, on the other the possibility that animals use choices rules that are distinct from that assumed by the choice axiom cannot be ignored. This is especially important if we consider that, in fact, the use of choice rules strictly in line with Luce's theory can lead to suboptimal choices and preference reversals between categories of options whenever the opportunity set is to some extent structured. To see that, suppose that in the paradox previously described elections for car represents the optimal (or simply the 'preferred') solution, with DM showing a higher probabilistic bias towards car

than bus in binary choices. Now let the choice set include more buses of different colours, all of which are suboptimal (or less 'preferred') relative to car. According to the choice axiom, each new bus included in the choice set should take its share of choices from all previously existing options, including car. As a consequence, the proportion of choices for car should gradually drop with the inclusion of buses, to the point that the overall probability of choosing car becomes lower than that of choosing a bus (of any colour) – a preference reversal, or violation of the optimal expectation if the choice DM ought to make is predicted a priori. In such cases, the choice behaviour predicted by the similarity hypothesis (which states that the probability of choosing the categories 'car' versus 'bus' – regardless of colour – should remain relatively constant) would represent a better solution to the problem.

3.3 Presence of uncorrelated choices

Violations of the choice axiom may also be obtained, without implying context-dependence in the assignment of value to options, whenever a certain proportion of choices (hereafter 'uncorrelated choices') result from causes other than those related to the properties of the rewards. In this section I illustrate this possibility by considering the following two cases.

3.3.1 The sampling case

If a decision maker faces two options of different absolute value the optimal allocation is for exclusive use of the better alternative. This model is descriptively challenged by all observations that indicate that allocation of behaviour is rarely exclusive, and many models exist to give account for this failure. One normative approach is that a fraction of the responses are allocated to alternatives other than the best one as a means to track environmental changing opportunities (Krebs et al. 1978, Dow & Lea 1987). If this allocation is distributed by any possible pattern between the less-than-best alternatives, it follows that the relative allocation of behaviour to the best one against any of the others will increase as the number of sub-optimal alternatives increases, in violation of the axiom. To see this, say that over a total of N choices, s are "sampling" responses distributed among the sub-maximal options. When only alternatives A and B are present, with A always representing the highest payoff, $N_A = N - s$, and $N_B = s$. Hence the relative proportion of choices for A is $P_B(A) = (N-s)/s$. If a third sub-maximal alternative is added, now the number of choices for A is unaltered but that for B will be reduced so that now $P_{AC}(B) = s - P_{AB}(C)$. Thus, the relative proportion of choices for A over B becomes $P_{BC}(A) = (N - s) / (s$

– $P_{AB}(C)$), that is necessarily bigger than that observed in the binary situation. Of course, under different assumptions (e.g. each extra patch is sampled with frequency s – namely, the total sampling frequency is increasingly higher with the enlargement of the choice set) we could be lead to the opposite conclusion – that relative preference for the most profitable patch would be lower in the trinary than in the binary context.

3.3.2 The “noise” case

Say now that the reason for non-exclusive choice of the best alternative is that a fraction p of the choices are assigned randomly to all existing options (Schuck-Paim & Kacelnik 2002, see also Bateson 2002). I call these “uncorrelated” choices because a proportion of choices is not correlated to the payoffs. Now, a fraction u of choices is assigned in an uncorrelated fashion. Each option will receive u/n choices from this process, where n is the number of options. Choices for the best option are thus

$$P_B(A) = N - \left((n-1) * \frac{u}{n} \right) \quad [2]$$

where N is the total number of choices, while choices for each sub-maximal option is u/n . In the case where $n=2$, $P_B(A)$ reduces to $N - (u/2)$. In binary choices the relative preference for the best option over its alternative is thus

$$\frac{P_B(A)}{P_A(B)} = \frac{N - \left(\frac{u}{2} \right)}{\frac{u}{2}} \quad [3]$$

Now $P_X(A)$ represents the fraction of choices for A in a set of three or more. Since

$$\frac{N - \left((n-1) * \frac{u}{n} \right)}{\frac{u}{n}} > \frac{N - \left(\frac{u}{2} \right)}{\frac{u}{2}} \quad [4]$$

for any $n > 2$, $P_X(A)/P_X(B) > P_B(A)/P_A(B)$. Notice here that although this process makes the relative preference for the best option over any sub-optimal action to be higher the larger the choice set, the absolute proportion of choices for A should decrease with the inclusion of more options into a set, since $N - (u/2) > N - ((n-1)*(u/n))$ for any $n > 2$.

3.4 Constraints on Perception

Here I argue that violations of the choice axiom can also be explained by considering the accuracy with which animals process information about available alternatives. I start by considering existing models and their underlying assumptions. I then proceed by examining their predictions in relation to the choice axiom.

The idea of imperfect discrimination as applied to choice problems dates back to Thurstone's law of comparative judgement (Thurstone 1927). In the most widely used case of his theory (Thurstone case V), Thurstone assumed that choice objects could be represented by what he termed 'discriminal processes' (U), that is, variables $U_i = u_i + X_i$, with u representing the parameter value of the choice objects and X representing independent random variables with a normal distribution function and equal variance. The idea was that, when faced with a choice, subjects would choose the objects with the largest discriminial process, or in other words, the object with the largest perceived value. Within the foraging literature, a number of models of information processing were also developed to predict the effects of errors in memory or perception on animals' decisions. This is the approach that underlies, for instance, Scalar Utility Theory (SUT), a process-based model developed under the framework of risk-sensitive foraging (Reboreda & Kacelnik 1991, Kacelnik & Brito e Abreu 1998). SUT assumes that memory of reward attributes is inaccurate, and that the variability in the animal's internal representation of rewards is best represented by Weber's Law. This means that errors are distributed following a normal probability distribution with a coefficient of variation (standard deviation/mean) that is constant across distinct magnitudes (i.e. the higher the magnitude of a stimulus, the greater a subject's absolute error). When making a choice between two alternatives, a subject would draw a sample from its internal representation of each reward (but see Kacelnik & Brito e Abreu 1998 for a discussion of other possible choice mechanisms), choosing the one yielding the most favourable outcome. Under a similar perspective, Yoccoz et al. (1993) also developed a model in which the presence of constraints on perception is taken into account. The model makes a series of predictions concerning optimal prey choice by considering the energetic gain (G) and handling time (T) associated with each prey type. Gains and times are however perceived with errors, accounted for by the addition of a normally distributed random variable with mean of zero to G and T . Finally, a similar approach is also taken by McNamara & Houston (1987). These authors developed a model to explain why individuals frequently show partial levels of preference for foraging options

instead of an absolute preference for the option yielding the highest payoff. In their model, animals compare the payoff (H) associated to each option, but their perception of H is subject to errors, accounted for by the addition of a random variable (X , distributed following a double exponential probability density function) to H , such that $H'=H+X$. The main postulate underlying the model is that the probability of choosing an option is a function of the fitness costs associated with the suboptimal action, which is in turn given as a function of the comparison of 'inaccurate' H 's.

Although all models are structurally similar in assuming that when faced with a decision individuals choose the option with the best perceived value – where the perception of value is subject to error –, they differ in some relevant aspects. Firstly, Thurstone's law of judgement, SUT and Yoccoz's model assume that it is the perception of primary attributes describing the available alternatives – such as time, amount, distance, volume, etc – that is subject to error, and that subjects choose the options associated with the best primary percepts. Conversely, the cost model of McNamara & Houston assumes that secondary combinations of attributes, namely the final payoff associated to each alternative, are perceived with error. Although the difference might be trivial if one considers that the presence of constraints in the perception of primary dimensions necessarily produces error in the perception of their secondary combinations, the distinction can be important to the extent that it can lead to qualitatively distinct predictions about the choice behaviour of individuals.

More importantly however for the present discussion, a crucial difference between these models seems to lie on their assumptions about the function describing the distribution of perceptual errors. While, for instance, SUT assumes that errors are distributed in accordance with Weber's Law, Yoccoz et al. (1993) use a normal distribution with a variance term that is independent of the parameters' magnitude. More uniquely perhaps, McNamara and Houston (1987) assume that the random error variable X follows a double exponential distribution. The reason for this choice was not arbitrary, but rather based on Yellot's (1977) mathematical proof that this distribution would be the only one for which choices between more than two alternatives would conform to the choice axiom. Yellot's arguments were based on the comparison of Thurstone's Case V with Luce's axiom, and showed that while for comparisons of pairs of options a number of distributions of perceptual errors would produce choice probabilities compatible with the requirements of

the axiom, the double exponential would be the only one able to do so in experiments involving the comparison of binary and trinary choices.

The assumption that errors are distributed following a double exponential distribution is however unrealistic. In fact, many years of empirical research in psychophysics have shown Weber's Law (or Weber-Fechner Law) to hold almost universally in discrimination tasks involving a large number of different species and protocols (see e.g. Church & Gibbon 1982, Cheng & Roberts 1991, Wearden 1991 for examples with rats, pigeons and humans, respectively). The use of the double exponential distribution is therefore inappropriate, the main implication being – as Yellot elegantly demonstrated – that violations of the axiom should be expected for any model relying on a distinct assumption.

As an example, here I develop the predictions of SUT for choices between delayed rewards in binary and trinary contexts. Specifically, I compare the predicted relative choice proportions (calculated as shown in equation 1) for a target choice pair consisting of options T_1 and T_2 (where T_1 and T_2 differ only in the delay subjects have to wait to receive a reward) in binary and in trinary choice situations, where a third option (or 'decoy') is included in the set of available opportunities. Figure 1 illustrates the results. The predictions were calculated for different levels of timing accuracy (parameter γ in the figure, represented by the differently coloured lines), with $\gamma=0.1$ and $\gamma=1$ representing relatively accurate and inaccurate levels of time perception, respectively. The figure shows the difference in the relative choice probabilities (T_1 relative to T_2) between the two situations, that is

$$D = \frac{p_{Tri}(T_1)}{p_{Tri}(T_1) + p_{Tri}(T_2)} - \frac{p_{Bi}(T_1)}{p_{Bi}(T_1) + p_{Bi}(T_2)} \quad [2]$$

where D is the difference in relative choice proportions between the binary and trinary situations. Positive and negative values indicate a higher level of relative preference in the trinary and binary contexts, respectively, whereas no difference (values equal to zero) represents the predictions of the choice axiom. In a majority of cases, the inclusion of a decoy causes relative choice proportions for T_1 to drop. This effect is especially strong when the decoy is better than either of the targets (section I of the graphs), that is, when it offers the shortest delay (decoys with delays that are intermediate and longer than those offered by

the target options, sections II and III, have a weaker effect on preferences; in fact, it can be seen that the longer the decoy's delay the smaller the difference in relative preferences between contexts). SUT predictions are however dependent on the timing accuracy assumed, with very low values of gamma ($\gamma=0.1$ to 0.2) leading to an increase in the level of preference for the preferred target in the trinary context depending on the delay associated with the decoy, especially so when the difference between the target options is smaller (fig.1a). Comparison of figures 1a, 1b and 1c also shows that the higher the difference in the delay associated with each target option, the more pronounced the decrease in relative preferences in the trinary context. The example thus shows, under SUT assumptions, the direction in which the inclusion of a decoy into a choice set can affect preferences, and illustrate what Yellot has proved more than 25 years ago: that violations of the choice axiom are to be expected whenever the perception of rewards are subject to perceptual constraints – probably a realistic assumption for most choice problems encountered by animals.

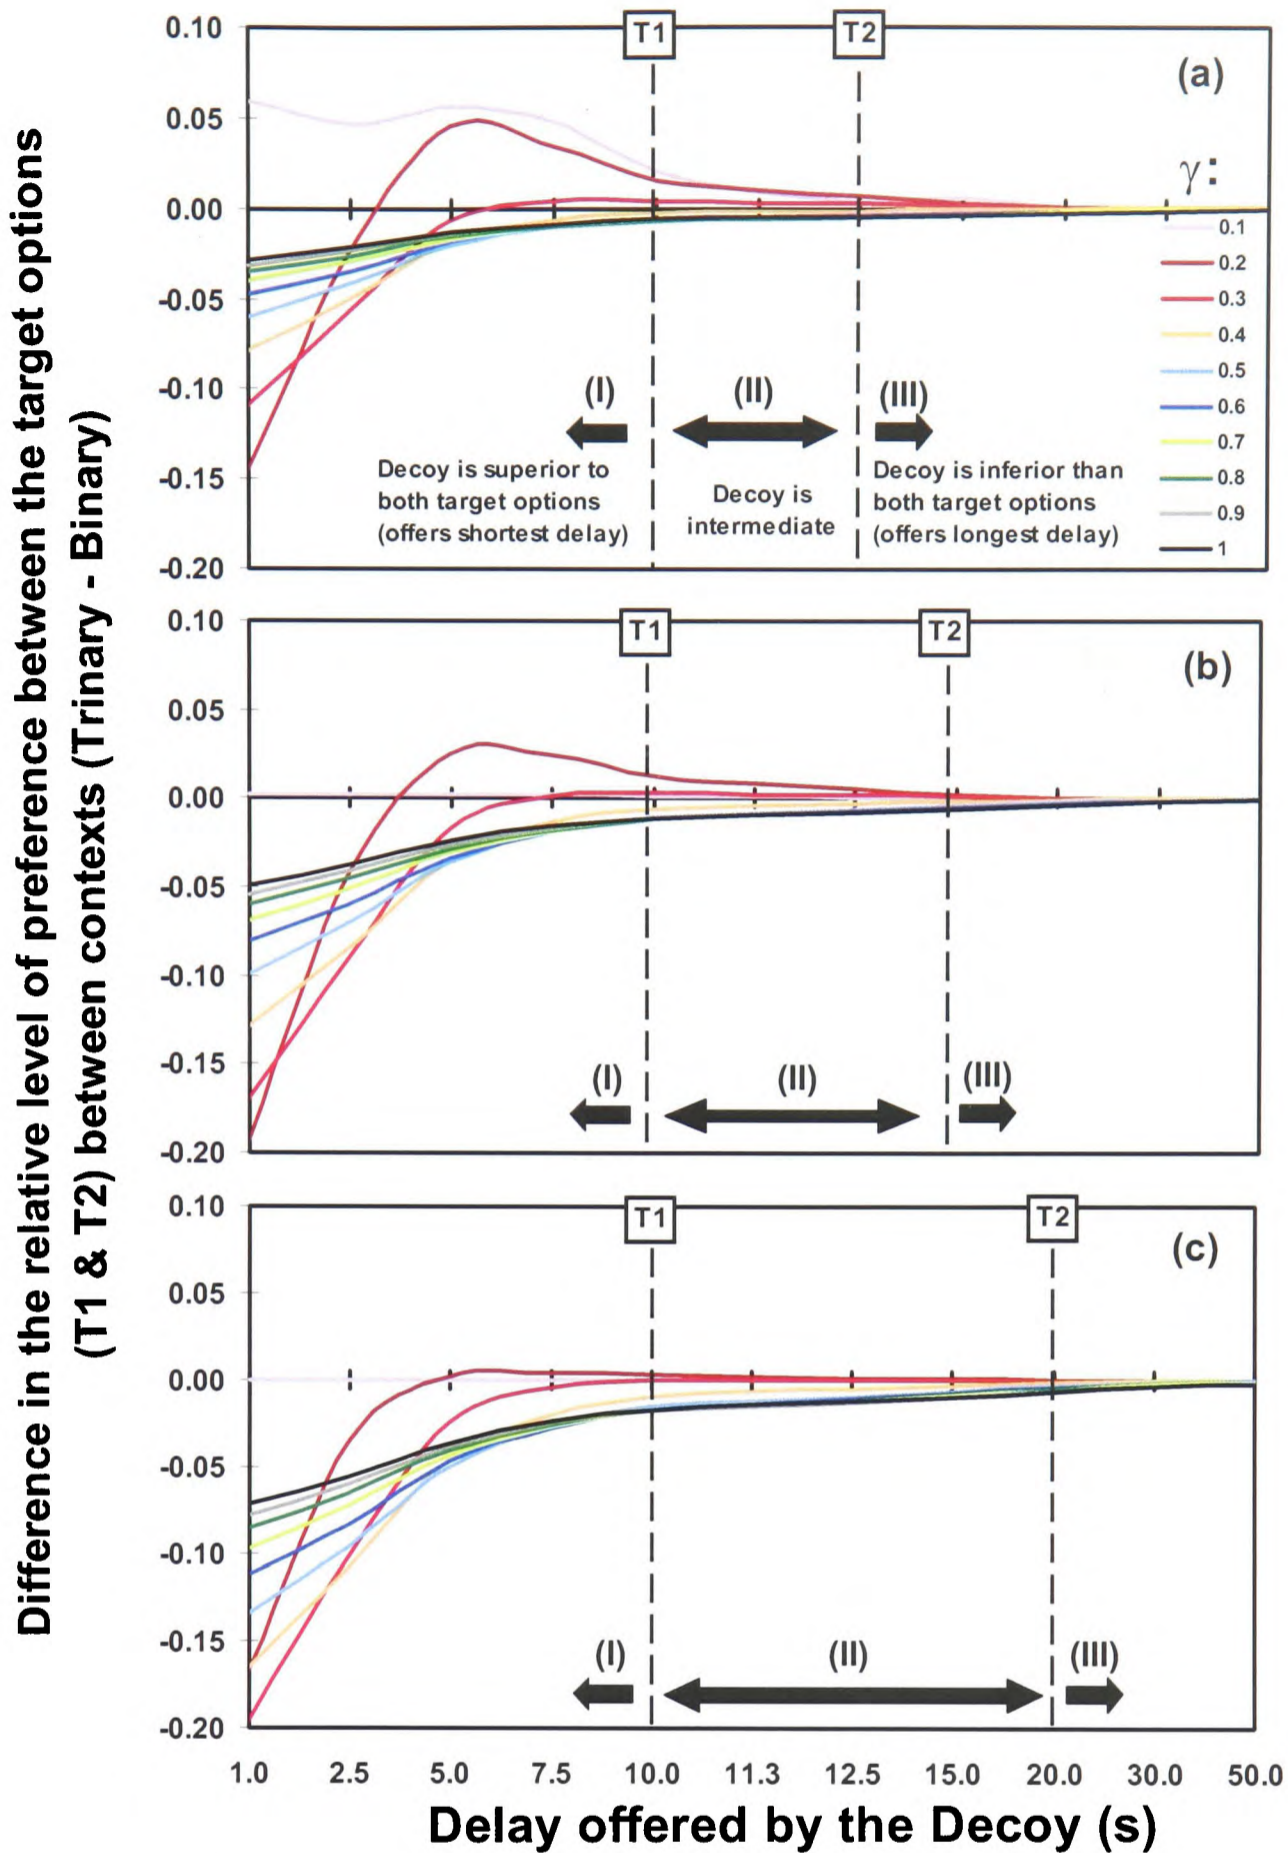


Figure 1. Predictions of Scalar Utility Theory (SUT, Reboreda & Kacelnik 1991, Kacelnik & Brito e Abreu 1998) regarding the relative level of preference between two target options (T1 and T2) in binary and trinary contexts, when a decoy (third option) is included in the choice set. The target options differ in the delay that subjects have to wait to receive a reward after they make a choice. Subjects are assumed to choose in each trial the option with the shortest perceived delay. The graphs show the difference in the relative level of preference for T1 over T2 between the two contexts (Trinary - Binary) as a function of the delay offered by the decoy presented in the trinary set. Positive and negative values in the ordinate indicate a higher relative preference in the trinary and binary situations, respectively. The coloured lines show the predictions assuming different timing accuracies (γ), ranging from 0.1 (relatively accurate) to 1.0 (relatively inaccurate). In the three graphs T1 offers a delay of 10s. In (a), (b) and (c) T2 offers a delay of 12.5s ($T1=0.8*T2$), 15s ($T1=0.66*T2$) and 20s ($T1=0.5*T2$), respectively.

3.5 Conclusion

Most probabilistic theories of choice rely on the assumption that choice objects are assigned a value on a single scale, with choice probabilities being represented by some monotonic function of their scale values (Krantz 1967). Luce's axiom require the additional, and stronger, assumption that the magnitude of preferences between two objects can be directly determined by the ratio of their scale values (Tversky & Russo 1969). Contrary to previous claims in the literature on animal choices, here I have argued that violations of this assumption do not necessarily imply the use of comparative mechanisms of choice, whereby the value of an option is constructed as a function of the presence and value of other alternatives. Of course, I do not dispute that violations are compatible with the use of such mechanisms, nor do I claim that available data would be better fit if analysed under the alternative frameworks suggested. However, I emphasize that a number of hypotheses in line with the assumption of independent choice currencies can also lead to departures of the axiom. This is the case for instance of models that are based on the presence of perceptual errors in the evaluation of choice objects. Alternatively, the presence of a few uncorrelated choices – characterized by the use of higher order rules (such as 'sampling') or as the outcome of the effect of uncontrolled variability in the choice process – can also lead to changes in relative choice proportions between contexts. Finally, I emphasize that strict compliance with Luce's theory might lead to suboptimal choices whenever the alternatives of a set are substitutable or similar to each other. The distinction between these possibilities and true context-dependence is an important one and I believe should be acknowledged in future research. Firstly, they tell us different stories about the computational rules and information processing modes employed by animals when making decisions. Additionally, the former do not pose major challenges to the normative modelling of animal behaviour.

3.6 References

- Bateson, M. 2002. Context-dependent foraging choices in risk-sensitive starlings. *Animal Behaviour*, 64: 251-260.
- Bateson, M., Healy, S.D. & Hurly, A. 2002. Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, 63: 587-596.
- Ben-Akiva, M.E. & Lerman, S.R. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand*. Cambridge: MIT Press.
- Cheng, K. & Roberts, W.A. 1991. Three psychophysical principles of timing in pigeons. *Learning and Motivation*, 22: 112-128.
- Church, R.M. & Gibbon, J. 1982. Temporal generalization. *Journal of Experimental Psychology - Animal Behavior Processes*, 8: 165-186.
- Clarke, F.R. 1957. Constant-ratio rule for confusion matrices in speech communication. *Journal of the Acoustical Society of America*, 29: 715-720.
- Debreu, G. 1960. Review of R.D.Luce: Individual choice behaviour: A theoretical analysis. *American Economic Review*, 50: 186-188.
- Dow, S.M. & Lea, S.E.G. 1987. Sampling of scheduled parameters by pigeons: tests of optimization theory. *Animal Behaviour*, 35: 102-114.
- Huber, J., Payne, J.W. & Puto, C. 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9: 90-98.
- Hurly, T.A. & Oseen, M.D. 1999. Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal Behaviour*, 58: 59-66.
- Kacelnik, A. & Brito e Abreu, F. 1998. Risky choice and Weber's Law. *Journal of Theoretical Biology*, 194: 289-298.
- Krantz, D.H. 1967. Rational distance functions for multidimensional scaling. *Journal of Mathematical Psychology*, 4: 226-245.
- Krebs, J.R., Kacelnik, A. & Taylor, A. 1978. Tests of optimal sampling by foraging great tits. *Nature*, 275: 27-31.
- Krebs, J.R. & McCleery, R.H. 1984. Optimization in behavioural ecology. In: *Behavioural Ecology: and evolutionary approach* (Ed. by Krebs, J.R. & Davies, N.B.), pp. 91-121. Massachusetts: Sinauer.
- Luce, R.D. 1959. *Individual choice behaviour: A theoretical analysis*. New York: Wiley.
- Luce, R.D. & Suppes, P. 1965. Preference, utility, and subjective probability. In: *Handbook of Mathematical Psychology* (Ed. by Luce, R.D., Bush, R. & Galanter, E.). New York: John Wiley.
- McNamara, J.M. & Houston, A.I. 1987. Partial preferences and foraging. *Animal Behaviour*, 35: 1084-1099.
- Payne, J.W., Bettman, J.R. & Johnson, E.J. 1992. Behavioral decision research - a constructive processing perspective. *Annual Review of Psychology*, 43: 87-131.
- Reboreda, J.C. & Kacelnik, A. 1991. Risk sensitivity in starlings: variability in food amount and food delay. *Behavioral Ecology*, 2: 301-308.
- Schuck-Paim, C. & Kacelnik, A. 2002. Rationality in risk-sensitive foraging choices by starlings. *Animal Behaviour*, 64: 869-879.
- Shafir, E.B., Osherson, D. & Smith, E.O. 1989. An advantage model of choice. *Journal of Behavioral Decision Making*, 2: 1-23.
- Shafir, S. 1994. Intransitivity of preferences in honey-bees - support for comparative-evaluation of foraging options. *Animal Behaviour*, 48: 55-67.

- Shafir, S., Waite, T.A. & Smith, B.H. 2002. Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 51: 180-187.
- Simonson, I. & Tversky, A. 1992. Choice in context: trade-off contrast and extremeness aversion. *Journal of Marketing Research*, 29: 281-295.
- Slaughter, J.E., Sinar, E.F. & Highhouse, S. 1999. Decoy effects and attribute-level inferences. *Journal of Applied Psychology*, 84: 823-828.
- Thurstone, L.L. 1927. A law of comparative judgement. *Psychological Review*, 34: 273-286.
- Tversky, A. 1972. Elimination by aspects: A theory of choice. *Psychological Review*, 79: 281-299.
- Tversky, A. & Gati, I. 1978. Studies of similarity. In: *Cognition and Categorization* (Ed. by E. Rosch & B. Lloyd), pp. 81-98. Hillsdale, NJ: Erlbaum.
- Tversky, A. & Russo, J.E. 1969. Substitutability and similarity in binary choices. *Journal of Mathematical Psychology*, 6: 1-12.
- Tversky, A. & Simonson, I. 1993. Context-dependent preferences. *Management Science*, 39: 1179-1189.
- Waite, T.A. 2001. Background context and decision making in hoarding gray jays. *Behavioral Ecology*, 12: 318-324.
- Wearden, J.H. 1991. Do humans possess an internal clock with scalar properties? *Learning and Motivation*, 22: 59-83.
- Wedell, D.H. 1991. Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology - Learning Memory and Cognition*, 17: 767-778.
- Yellott, J.I. 1977. The relationship between Luce's choice axiom, Thurstone's theory of comparative judgment, and the double exponential distribution. *Journal of Mathematical Psychology*, 15: 109-144.
- Yoccoz, N.G., Engen, S. & Stenseth, N.C. 1993. Optimal foraging: the importance of environmental stochasticity and accuracy in parameter estimation. *American Naturalist*, 141: 139-157.

CHAPTER 4

4.1 INTRODUCTION	43
4.2 METHODOLOGY	46
4.2.1 EXPERIMENTAL DESIGN	46
4.2.2 SUBJECTS	47
4.2.3 APPARATUS	48
4.2.4 TRAINING	48
4.2.5 EXPERIMENTAL PROTOCOL	48
4.3 DATA ANALYSIS	51
4.3.1 GENERAL PROCEDURES	51
4.3.2 PREFERENCE SCORES FOR THE FIXED OPTIONS	51
4.3.3 RELATIVE PREFERENCES	51
4.3.4 ABSOLUTE PREFERENCES	52
4.3.4.1 <i>REGULARITY</i>	53
4.3.4.2 <i>TRANSITIVITY</i>	53
4.4 RESULTS	54
4.4.1 KNOWLEDGE ABOUT THE ALTERNATIVES	54
4.4.2 PREFERENCE FOR VARIABILITY	57
4.4.3 RELATIVE CHOICE PROPORTIONS	57
4.4.4 ABSOLUTE CHOICE PROPORTIONS	60
4.4.4.1 <i>REGULARITY</i>	60
4.4.4.2 <i>TRANSITIVITY</i>	62
4.5 DISCUSSION	64
4.6 REFERENCES	69

STARLINGS SHOW INVARIANT PREFERENCES WHEN CHOICE SETS ARE STRUCTURED

4.1 Introduction

Imagine that a migrating bird has just arrived at its breeding island, and now needs to look for a suitable site where to build a nest. The area is restricted and spatially structured, all vacant sites will be occupied in a short period by the arriving birds, and the decision on whether to settle on the central region or near its boundaries will determine the exposure to predators and prey levels experienced. In the centre nests are safer from the danger posed by foxes but far from the feeding grounds, and in the edge the opposite occurs. On balance, both kinds of location offer equal net benefit. While overflying the island the bird spots three opportunities, one on the central region and another two in the periphery. We can envisage the process of choice as a hierarchical one (centre vs. periphery, then choices within area) or as a simultaneous one, in which all opportunities are considered at once. The former would lead to a probability of choice of one half for the centre and a further halving to $\frac{1}{4}$ for each peripheral site. The latter would yield a probability of $\frac{1}{3}$ for each location. Notice that the simultaneous handling of choices implies that the probability of choosing the central option drops by increasing the number of peripheral sites (e.g., had the bird spotted four sites on the edge and one in the centre, the overall probability of choosing centre would be $\frac{1}{5}$ as opposed to $\frac{4}{5}$ of choosing the periphery), whereas the hierarchical choice process implies that the probability of choosing centre will always be the same, regardless of the number of sites in either region.

Knowing which outcome is the most likely requires, then, some knowledge of the mechanisms and choice rules adopted by the bird. Here I investigate experimentally the two possibilities illustrated in the example. The first one (also known as ‘similarity hypothesis’ in the context of human choices; Tversky & Russo 1969) assumes the ability to categorize or group the available options in terms of their degree of similarity or substitutability, and that the choice process is hierarchical. For instance, in the previous example the bird would firstly decide between the central and marginal regions and, if choosing the latter, decide in a second step between the two options contained in this category. As a result of the adoption of such groupings, the inclusion of any additional option into a choice set should hurt more (in terms of choice probabilities) the original alternative most similar to it (notice that under the hierarchical process the probability of choosing a periphery site over the centre is $\frac{1}{2}$, but drops to $\frac{1}{4}$ if there are two sites available in the edge). The second possibility, conversely, derives directly from the notion that the preference between any two alternatives should be entirely independent of the context (as defined by the presence and number of other alternatives) in

which they occur. This principle, also known as ‘independence from irrelevant alternatives’ (IIA), underlies most theoretical work on choice behaviour. In its most restrictive form – captured by Luce’s choice axiom (Luce 1959) – it states that the ratio of the choice probabilities of any two alternatives should be entirely unaffected by the presence and properties of other options. Hence, the inclusion of a new option into a choice set should take from pre-existing options in proportion to their original shares, and therefore the relative preference between two options should always be the same. Accordingly, the second possible outcome of our bird’s decision problem shows that the relative probability of choosing each one of the peripheral sites over the central one would be the same in binary (two options available) and trinary (three options available) choices.

The analysis of the extent to which relative preferences differ amongst contexts provides an interesting tool for the understanding of the processes underlying animal choices. Many normative models in biology assume that options are assigned a value based on the potential fitness gain (or surrogates of it) associated to each option, and that choices should be such as to maximize this value. However, the adoption one or another of the choice mechanisms described above necessarily leads to different solutions to the same problem. To make this even clearer, suppose that in the problem faced by the nesting bird, the central site offers a net benefit slightly higher than that gained by nesting in the periphery, so the bird should be expected to either choose the centre exclusively or, assuming choice is not absolute but probabilistic, to show a higher probabilistic bias of settling in the centre than in the edge. Under the former assumption (exclusive choices) centre would be always chosen over edge, regardless of the mechanism employed. However, animals often show partial instead of exclusive levels of preference for options on repeated choice opportunities (McNamara & Houston 1987), and in this latter (probabilistic) context the distinction between the two possibilities matters. If adopting the hierarchical choice process, the probability of nesting in the centre should always be higher than that of nesting in the border, as discussed. Conversely, the use of choice rules in line with the choice axiom should make the probability of settling in the centre to increasingly drop with increases in the availability of sites in the border; in other words, the deviation from the optimal expectation (choice of ‘centre’) should be increasingly larger.

Also relevant is the fact that the observation of violations of Luce’s axiom (i.e., changes in relative choice proportions between contexts) is usually taken as support for the

idea that the value of an option is not only determined by its intrinsic properties and consequences as assumed by normative modelling in biology, but instead is constructed depending on the presence and attributes of other options (preferences are said to be ‘context-dependent’; Real 1991, Tversky & Simonson 1993, Shafir 1994, Hurly & Oseen 1999, Bateson 2002, Bateson et al. 2002, Shafir et al. 2002, but see also Houston 1997). The similarity hypothesis previously described, in turn, offers a compelling alternative to this interpretation. By taking into account the structure of the choice set in terms of the outcomes that are substitutable for one another, it offers a possible explanation for changes in choice proportions between contexts independently of the potential effect other options might have on preferences. In the experimental investigation reported below I handle this problem explicitly by using a design that enables the distinction of these two interpretations (context-dependence and similarity hypothesis).

Animals are often faced with the need to make choices between a number of alternatives, such as foraging patches, prey types, nest sites, territories and sexual partners. It is likely that in many of these circumstances the choice sets are to some extent structured, with the available options sharing different degrees of similarity with each other. Yet, the effect that such a natural structuring would have on animal choices is still poorly understood. Amongst the reasons, a priori knowledge of the degree of similarity/substitutability between the options for an animal may be hard to ensure in complex choice tasks, especially when options are composed by more than one attribute. In these cases, the relative importance of similar and distinctive attributes (as well as their relative weighting) may be different from that assumed by the experimenter, preventing an unambiguous analysis of the effect of similarity relations on choice.

In this study I examine the foraging preferences of European starlings (*Sturnus vulgaris*) in situations where the choice set was characterized by obvious differences and similarities between its elements. By analysing how preferences were expressed according to the structure of the choice set, I aimed to understand the mechanisms and rules used by the starlings in multialternative choice situations, together with their implications for the functional analysis of behaviour.

4.2 Methodology

4.2.1 Experimental Design

To examine whether the starlings were more likely to employ a hierarchical process of choice or a simultaneous one, as described in the introduction, I compared their risk-sensitive foraging preferences between binary and trinary choice contexts, so as to test whether the inclusion of new options into the choice set changed their relative level of preference for the options. Choice sets were designed so as to impose clear differences and resemblances between the options. To that end, each set comprised three alternatives, two of which shared exactly the same parameter values, being perfectly substitutable by each other (a decision problem essentially similar to that faced by the nesting bird illustrated in the introduction). The three alternatives were characterized by differences in the delay the animals had to wait to receive food. The two equivalent options offered the same fixed delay to food, whereas the third offered a variable delay with two equiprobable outcomes.

The design was planned so as to also enable the interpretation of potential changes in relative levels of preference between contexts as due to the structure of the choice set (as postulated by the similarity hypothesis) or to context-dependence in the assignment of value to the available alternatives. Under the latter framework, a widely used explanation for such changes (also referred to as the ‘comparative approach’, see e.g. Shafir et al. 1989, 1993) is based on the idea that options are compared along each of their relevant attributes before a decision is reached, instead of being assigned an independent value on a single dimension. As different attributes may have different weights, perception of differences of each attribute may depend upon the whole range of values being considered and the overall dominance rank of an alternative may be built up according to its dominance position in each of the attributes considered (Wedell 1991), the overall attractiveness of an option can be determined by the specific value of competing alternatives in each of the attributes. The approach thus offers a possible explanation for changes in preference between contexts mostly in those cases where the alternatives differ in at least two attributes. Hence, to minimize the possibility that potential changes in relative preferences were due to these mechanisms (instead of the employment of a hierarchical choice process), in the present investigation I offered the subjects with options that differed only in one characteristic. Additionally, trinary choices were presented interspersed with the three types of binary choices (each corresponding to one of the three possible choice pairs). In this way, I

intended to equalize the background context (as defined by the set of options known to the subjects) in the binary and trinary scenarios, thus controlling for the possibility of context-dependent effects due to the construction of preferences under different backgrounds (Pratkanis & Farquhar 1992, Tversky & Simonson 1993, Waite 2001).

The choice to manipulate variability in delay to reward was based on the observation that partial, rather than all-or-none, preferences are usually obtained for repeated risky choices (see reviews in Kacelnik & Bateson 1996 and Shafir 2000), therefore minimizing the possibility of ceiling effects and in this way facilitating the detection of differences in choice proportions between contexts. Subjects were presented with five treatments characterized by differences in the delay offered by the two identical fixed options. The fixed delay to food varied from a relatively short (as compared to the mean delay offered by the variable option) to a long one, thus enabling the analysis of choice probabilities and the models' predictions along a wide range of partial levels of preference (i.e., from when the fixed option was preferred over the variable to when the variable was strongly preferred over the fixed).

4.2.2 Subjects

The subjects were ten wild-caught European starlings, captured in Oxford, UK (under an English Nature licence) as adults on March 2001. Prior to the experiments the starlings were kept in two outdoor aviaries (3 m high x 3.2 m wide x 4 m long). While in the aviary the starlings were fed on ad libitum food, a mixture of turkey crumbs, Orlux© pellets and mealworms (*Tenebrio* sp.). Drinking and bathing water was always available and replaced daily. Five days before the beginning of the training sessions the birds were transferred from the outdoor aviary to the laboratory and housed in individual cages (120 cm x 60 cm x 50 cm) that served both as home cages and experimental chambers. During this period of adaptation all birds had free access to water and the same ad libitum food as received in the aviaries. During the experimental period, the starlings were given free access to food (turkey starter crumbs) after the end of the daily sessions, from 1700 until 1900 h. At 1900 h they were supplemented with 10 mealworms and then food was removed. From previous experience this regime is known to allow the starlings' body weights to remain stable at approximately 90% of their free feeding value or above (Bateson 1993). Fresh

drinking water was always available in the cages and bathing tubs were provided daily. There was no mortality. The experiments were initiated in November 2001 and completed in March 2002. All subjects were released back into the wild (University Parks, Oxford, UK) during the spring.

4.2.3 Apparatus

The experimental cages had a panel with a centrally mounted food hopper and three response keys (at either side of the hopper and above it). The cages were also supplied with two perches, 85 cm apart from each other. An Acorn Risk PC 600 computer running Arachnid® software (Paul Fray Ltd.) controlled the contingencies and collected the data. During the experimental sessions food rewards were units of Noyes precision pellets (0.02g). The pellets were delivered by a pellet dispenser (Campden Instruments®) at a rate of 1 unit/s. Lights in the experimental rooms were gradually switched on at 0530 and off at 1930 h, and temperature ranged from 11°C to 15°C. Subjects were visually but not acoustically isolated.

4.2.4 Training

Subjects were first trained to retrieve food from the hopper. Birds' responses were then induced by a standard auto-shaping procedure, in which the illumination of one of the pecking keys (side determined randomly) for 10s was followed by the delivery of two precision pellets and an intertrial interval (ITI) of 60s. The birds received two daily sessions in this condition (100 trials per session), one at 0800 h and another at 1200 h, after which they were fed with the ad libitum food. After all birds were pecking the keys reliably food delivery was made conditional at one key pecking. In these sessions, if the birds did not respond during the 10s interval the light went off and the ITI started. The training program was interrupted when all birds were pecking in at least 80% of the trials. This was achieved after six sessions. All birds had already been used in an operant setup before and were accustomed to peck keys.

4.2.5 Experimental Protocol

All subjects experienced five treatments. In each treatment the subjects were presented with three options differing in the level of variability in delay to reward: a variable option, offering equiprobable delays of either 12 s or 28 s (Coefficient of Variation=40%) and two fixed

options ($CV = 0$) which offered the same delay to reward. In treatments 1, 2, 3, 4 and 5, the delay offered by the fixed options was 12s, 15s, 18s, 22s and 28s, respectively. The variable option was the same in all treatments. The order of treatments was balanced across subjects following a balanced Latin-square design. As I had an odd number of treatments ($n=5$) I used two separate arrangements comprising 5 orders each (Namboodiri 1972), as shown in table 1. Each of the options was associated with one colour trans-illuminating the pecking keys (green, red and yellow). The association of colours with options was balanced across birds, but for a given bird one colour was always associated with one particular option. Although the subjects were already familiar with the experimental apparatus, they had not experienced colours as discriminative stimuli before.

Table 1. Order in which subjects experienced the five treatments (the treatments are described in terms of the delay of the fixed options)

Bird	1st	2nd	3rd	4th	5th
1	12	15	28	18	22
2	15	18	12	22	28
3	18	22	15	28	12
4	22	28	18	12	15
5	28	12	22	15	18
6	22	18	28	15	12
7	28	22	12	18	15
8	12	28	15	22	18
9	15	12	18	28	22
10	18	15	22	12	28

A discrete trials procedure with a fixed ITI of 60s was employed. There were three types of trials: no-choice trials, choice trials and probe trials. No-choice trials were intended to provide the birds with information about each alternative. They started with one of the pecking keys (randomly determined) blinking (0.7s on, 0.3s off). The first peck on the key caused the light to be on for the delay corresponding to the option presented, during which pecks on the key were counted. The first peck after the delay had timed out triggered the delivery of the standard food reward (two precision pellets) in the hopper, and a new ITI started. Choice trials began with all pecking keys (two or three, depending on the type of choice trial) blinking. The first peck on any of the keys caused the chosen key to be continuously on for its respective delay and the remaining key lights to turn off. The first peck after the delay had elapsed triggered the delivery of the standard reward and the beginning of a new ITI. Probe trials were used to measure the subjects' knowledge of the temporal properties of each option. Similarly

to no-choice trials, they also started with one of the pecking keys flashing. The first peck on the key caused the light to be on for a delay three times longer than the delay corresponding to the option presented, after which no reward was given and a new ITI started.

Subjects experienced four sessions (maximum of 2 hours) per day, starting at 0600, 0900, 1200 and 1500 h. Each session comprised 5 blocks of 15 trials each. Blocks were structured as follows: each block started with one probe trial, where one of the three options was presented. The order in which the options were presented in probe trials was randomly determined, with the only restriction that each option had to be presented twice per session. Probe trials were followed by three no-choice trials, one corresponding to each option (variable, fixed₁ and fixed₂), presented in a random order. After completion of the no-choice trials subjects experienced four choice trials in a random order: three binary choice trials, where each of the three possible binary choice pairs (variable vs. fixed₁, variable vs. fixed₂ and fixed₁ vs. fixed₂) was presented once, and one trinary choice trial (all options simultaneously available). Choice trials were followed by another series of three no-choice trials and four choice trials, with the only difference that in the second series of no-choice trials I programmed the variable option to offer the outcome that had not been experienced in the previous series (for instance, if a subject had experienced the short delay of the variable option in the first series of no-choice trials, the long delay would be presented in the second series). This procedure was employed to ensure that all subjects experienced the two possible outcomes of the variable option equally often during the no-choice trials. The sides where the options appeared were randomly determined in all trial types.

The two morning sessions and the two afternoon sessions were grouped for the purpose of analysis. For each bird, the experiment ended when the regression coefficient of the choice proportions of six consecutive sessions (against session number) was not significantly different from zero and the standard deviation of the choice percentages of these six sessions did not exceed 10. This criterion had to be reached in at least two out of the four types of choice trials for the treatment to end. The birds were given a minimum of eight and maximum of twenty sessions. Data from the last six sessions were used for analysis. Subjects were given three resting days with ad libitum food between treatments.

4.3 Data Analysis

4.3.1 General Procedures

I initially analysed the subjects' knowledge of the temporal properties of the rewards and their preferences for variability across the treatments. For most analyses of choice probabilities I plotted the individual choice proportions for each treatment and choice pair as a function of the choice proportion predicted by the various principles tested (described below). The use of such a procedure enabled the analysis of the general trend for the group of subjects without losing information on individual idiosyncrasies.

4.3.2 Preference Scores for the Fixed Options

The two fixed options offered the same delay to reward in all treatments, differing only in the colour of the discriminative label associated to each of them. However, it is usually the case that the influence of external sources of variability on preferences is maximal when the difference between the options is minimal. Accordingly, perfect indifference between the fixed options was not observed in most cases. Although only two subjects showed a slight colour bias (not significant) in the beginning of the experiment (measured through the difference in preference for the fixed options in the first 50 binary trials of this type) and averaged preference levels between the two fixed options were close to indifference (50%), most individual subjects showed small deviations from indifference in more advanced stages of the experiment that, although not significant in most cases, were correlated across treatments with the colour of the discriminative label associated to each option. Therefore, for the purpose of analysis, the fixed alternatives $fixed_1$ and $fixed_2$ were classified for each subject and treatment in two new categories, as Preferred fixed (P) and Non-Preferred fixed (N).

4.3.3 Relative Preferences

To compare the subjects' relative preferences for each choice pair between contexts I transformed choice proportions in the trinary situation by considering only two options at a time, so that the percentage of choices for the pair added up to 100%. For instance, preference for the preferred fixed option as compared to the variable in the trinary situation – $p_T(P, V)$ – is:

$$p_T(P, V) = \frac{p_T(P)}{p_T(P) + p_T(V)} \times 100 \quad [1]$$

where P and V stand for ‘preferred’ fixed and ‘variable’, respectively, and $p_T(P)$ and $p_T(V)$ represent the number of choices for ‘preferred’ and ‘variable’ options, respectively, in the trinary context. Thus, preference for P over V in a trinary context is defined as the number of choices for P divided by the number of choices for P and V added together. The same procedure was adopted to calculate the relative preference between the options (P,N) and (N,V) in the trinary situation.

If Luce’s choice axiom holds and relative preferences are not different between contexts, I should expect the slope of the regression of the relative proportion of choices for any given pair of options in a trinary situation (as calculated by equation 1) against the relative proportion of choices for this same pair in a binary situation to be equal to one and its intercept not different from zero. To test this prediction I analysed the relationship between relative preferences in the two contexts. In all analyses I considered the relative proportion of choices in the binary and trinary situations to be the independent and dependent variables, respectively (choice proportions were square root arcsine transformed). To take into account the fact that each bird contributed with five datapoints to the regression (one corresponding to each treatment), I calculated the slope and intercept of the regression for each individual bird separately, subsequently testing whether the groups of ten intercepts and ten slopes (one intercept and one slope corresponding to each subject) were significantly different from zero and one, respectively.

4.3.4 Absolute Preferences

The economic concept of rationality is usually defined through a number of principles (such as transitivity and regularity – explained below – and the notion of IIA previously mentioned) that are viewed as necessary consequences of the existence of a utility scale upon which organisms base their choices. Although violations of rationality are normally viewed as a challenge for the existence of such a scale (preferences would be context-dependent), the similarity hypothesis, along with the hypotheses described in Chapters 3, predict differences in relative preferences between the binary and trinary contexts in the absence of context-

dependence, providing some possible accounts for violations of IIA. Yet, they cannot explain violations of principles based on less restrictive conditions than those imposed by the axiom, such as weak and moderate stochastic transitivity, or regularity. Hence, finding such violations would constitute strong evidence for the use of a context-dependent, or ‘comparative’ mechanism of choice. I therefore tested whether the starlings’ preferences were both transitive and regular.

4.3.4.1 Regularity

The principle of regularity is one of the weakest forms of the notion of independence from irrelevant alternatives (Luce & Suppes 1965), and asserts that the absolute proportion of choices for an option cannot be increased when a new option is included in the choice set. Formally, regularity is satisfied if the probability of taking an item A from a subset $B_i = \{A, B\}$ that is part of $T = \{A, B, C\}$ is such that, for all C,

$$p_{B_i}(A) \geq p_T(A) \quad [2]$$

I therefore plotted, for each choice pair and option, each subject’s absolute choice proportion in the binary situation as a function of the absolute proportion of choices for the same option in the trinary context and analysed the observed deviations.

4.3.4.2 Transitivity

I analysed the degree to which choices were transitive by testing whether the subjects’ preferences conformed to weak, moderate, strong and perfect stochastic transitivity with the data from the choice trials. Weak stochastic transitivity (WST) is one of the less restrictive rational principles, guaranteeing the existence of at least an ordinal scale of values upon which preferences are based. It states that if $p_B(A) \geq 0.5$ and $p_C(B) \geq 0.5$, then $p_C(A) \geq 0.5$, where $p_B(A)$, $p_C(B)$ and $p_C(A)$ are the probabilities of choosing A when paired with B, B when paired with C and A when paired with C, respectively. In words, the condition asserts that if you prefer ‘A’ over ‘B’, and ‘B’ over ‘C’, you should also prefer ‘A’ over ‘C’. Moderate stochastic transitivity (MST) also guarantees the existence of an ordinal scale of value, and requires that if alternatives are ranked in an order such that ‘A’ is preferred to ‘B’, and ‘B’ to ‘C’, then $p_C(A) \geq \min [p_B(A), p_C(B)]$. In a numerical example, if $p_B(A) = 0.6$ and $p_C(B) = 0.8$, then by MST $p_C(A)$ must be at least 0.6. Strong stochastic transitivity (SST)

imposes the more restrictive condition that preference for ‘A’ over ‘C’ should exceed or equal the strongest of the two other preferences, that is $p_C(A) \geq \max [p_B(A), p_C(B)]$. As stressed by Sumpter et al. (1999), confirmation of SST not only implies the existence of an ordinal scale, but also that the strength of preference between pairs of options can be predicted from their relative scale values. Finally, perfect stochastic transitivity (PST), also known as the ‘product rule’ (Luce & Suppes 1965), additionally ensures that the strength of preference between pairs of options can be directly computed from the difference of their respective scale values. According to it,

$$\frac{p_C(A)}{p_A(C)} = \frac{p_B(A)}{p_A(B)} \times \frac{p_C(B)}{p_B(C)} \quad [3]$$

It is worth noting that, as Luce & Suppes (1965) highlight, PST necessarily implies SST, which in turn implies MST that implies WST. Here I plotted the predictions of each of form of transitivity (but those concerning WST) for each bird in every possible choice situation as a function of their observed preferences, subsequently testing whether the observed individual deviations were statistically significant.

4.4 Results

4.4.1 Knowledge about the alternatives

To characterize the subjects’ knowledge of the temporal properties of each option I analysed their pecking patterns during the delays experienced in the probe trials, by computing the number of pecks in 1s time bins. Figure 1 shows the mean peck rate (pecks/s) for the variable option in the five possible treatments. Similarly to the results reported by Schuck-Paim & Kacelnik (2002), there were no discernible peaks at the time bins when food was potentially due. As highlighted in this previous study, this pattern of responding follows from the scalar property of behavioural timing: since it has been shown that the spread of the response distribution is proportional to the length of the interval being timed, the presence of distinct peaks of responding should depend on the ratio between these intervals. In the present experiments, the ratio between the two delays corresponding to the variable option (12 s and 28 s) was in the same range where no discernible peaks were observed by other authors (Catania & Reynolds 1968, Leak & Gibbon 1995).

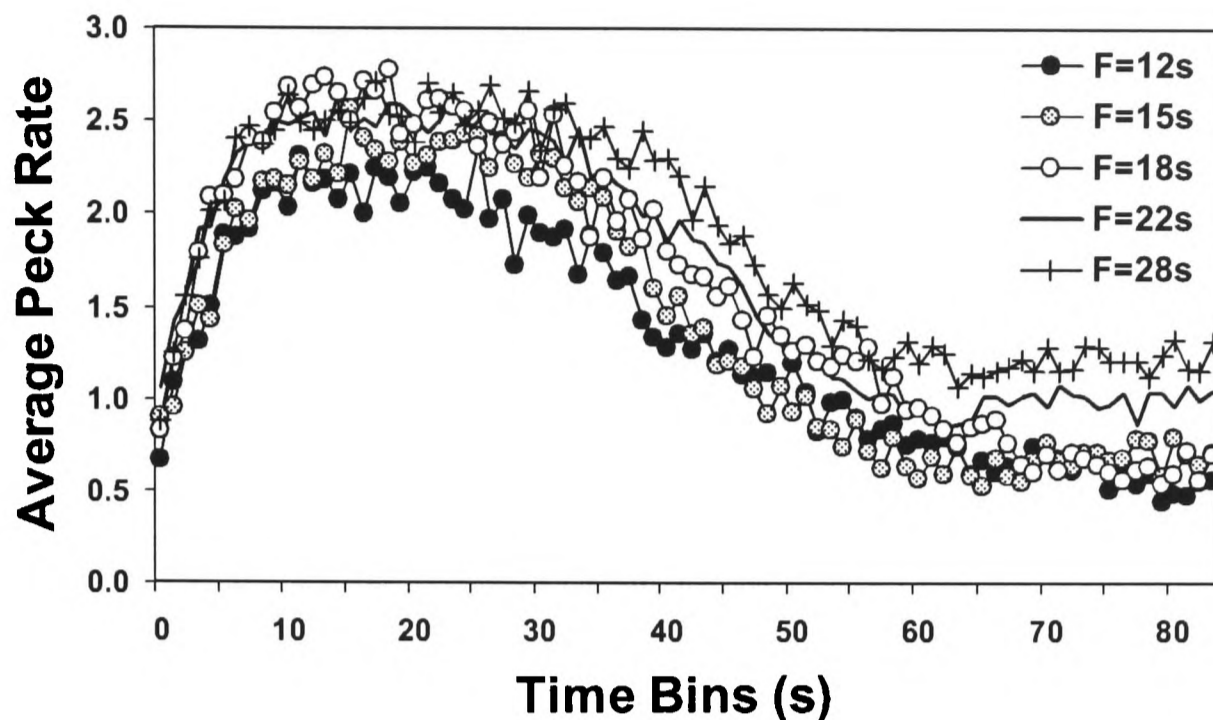


Figure 1. Mean peck rate (pecks/s) for the variable option during probe trials (total delay = 84s) in the five treatments employed. Food was potentially due at the time bins 12 and 28.

Figure 1 also shows that although the shape of the function is maintained across treatments, the magnitude of the peck rates for the variable option appears to progressively increase as the delay of the fixed options increases (i.e. as preference for the variable option increases). To test this effect we conducted a repeated-measures ANOVA (on square-root transformed rates) having treatment and time bins as within-subjects, fixed factors. As expected, time bins had a significant effect on peck rates ($F_{83,747} = 53.2, P < 0.001$), but the effect of treatment, although in the right direction, was not significant ($F_{4,36} = 2.05, P = 0.13$).

I also analysed the pecking rates of the fixed options (fig.2). It is clear from the figure that for both options the response functions were progressively shallower as the delays increased, in accordance with the scalar property of timing (Gibbon 1977). More importantly perhaps, the graphs on the right side of the figure – where time bins are expressed as a proportion of the bin when reward was due, and pecking rates as a proportion of the rate exhibited at the peak time – show that the pecking curves superpose nicely when the axes are standardized, as also predicted by the scalar property. Taken together, these results strongly indicate that subjects were able to correctly time the delays associated with all options, and that their timing accuracy varied in accordance with the scalar property.

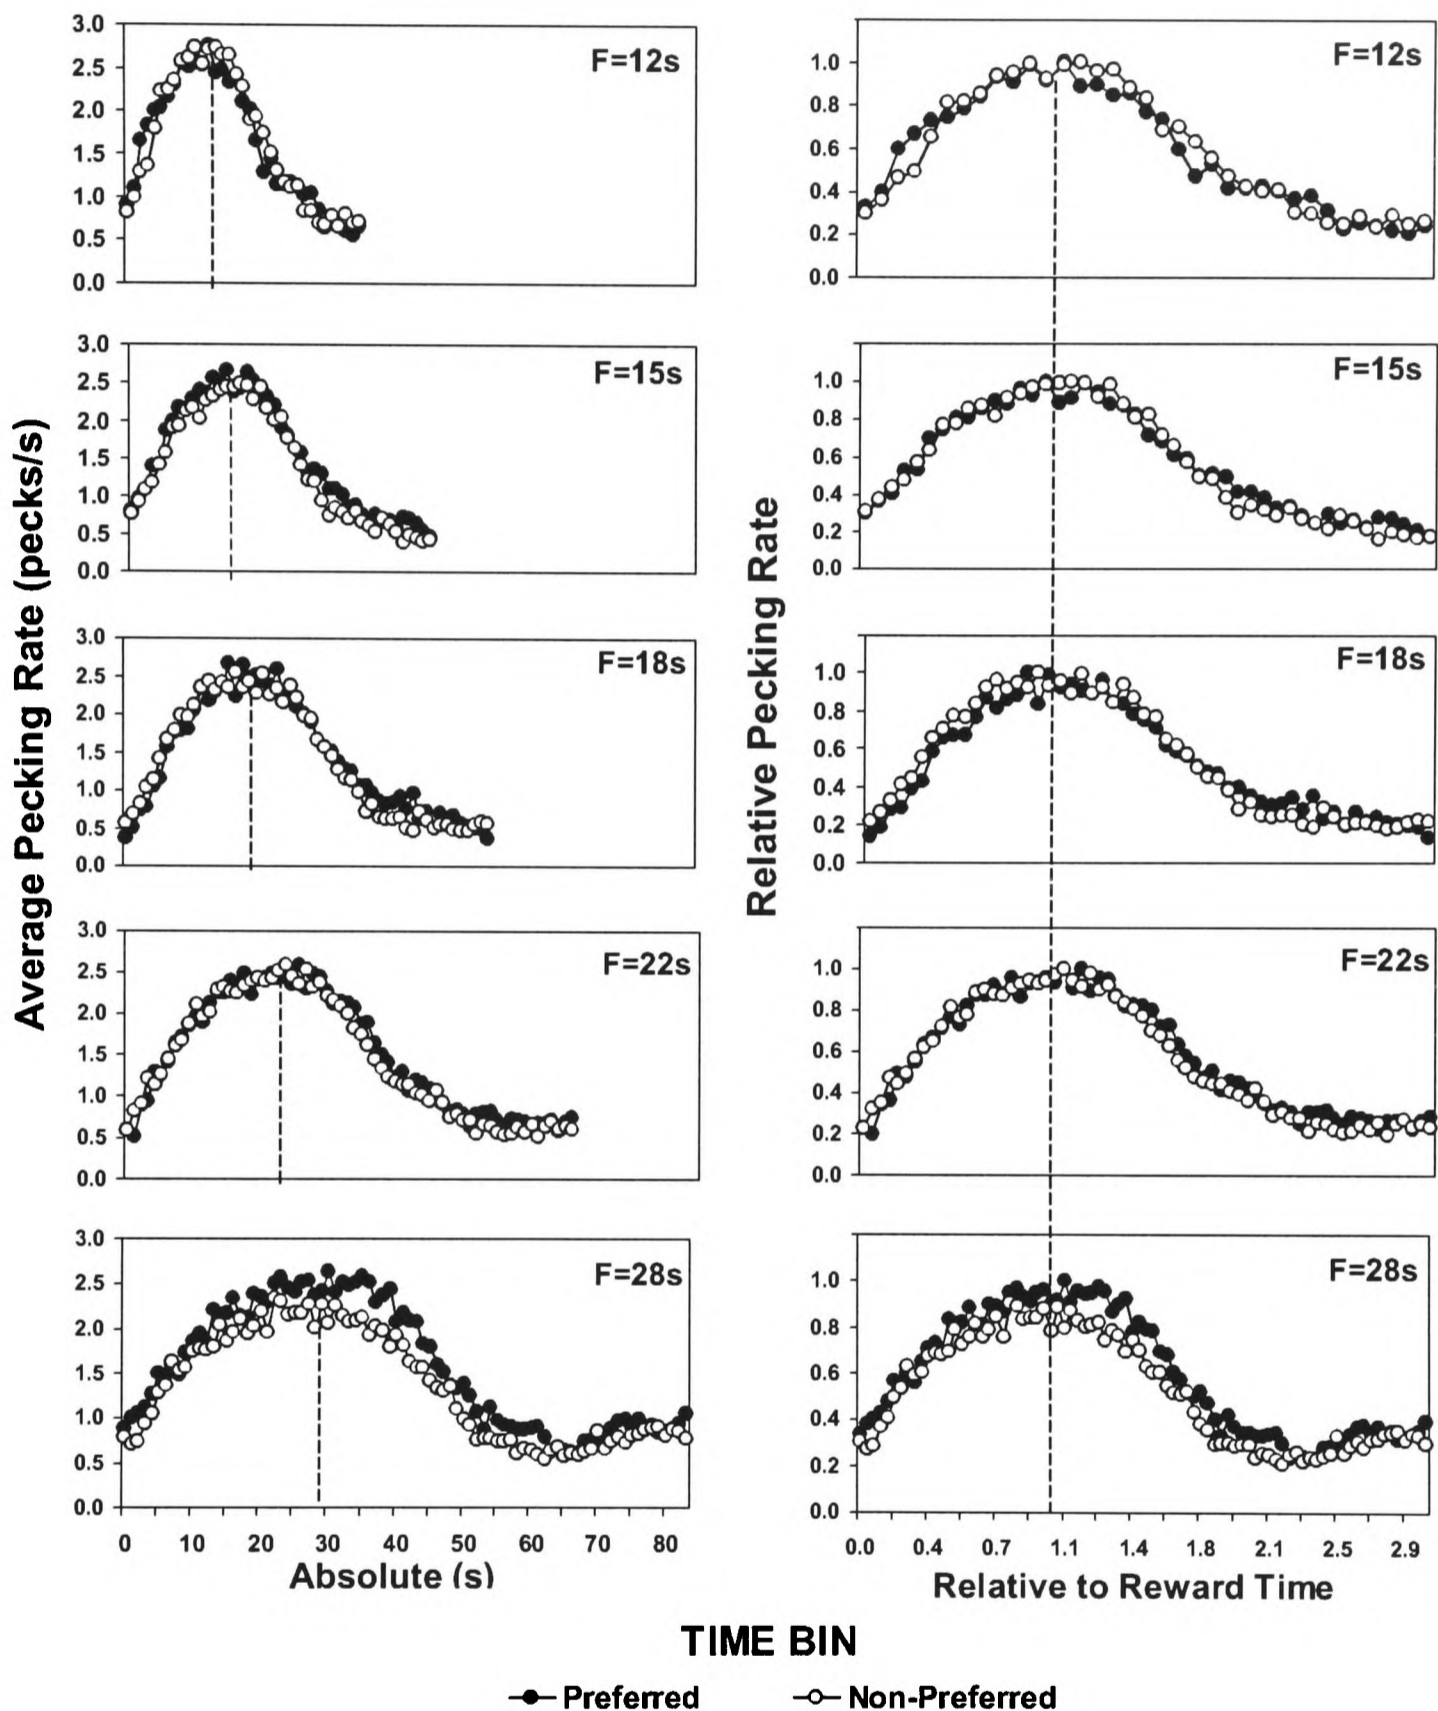


Figure 2. Mean peck rate (pecks/s, graphs on the left) and standardized peck rate (graphs on the right) for the fixed options during probe trials (total delay = 3 x F) in the five treatments employed (dark circles = preferred fixed option, white circles = non-preferred fixed option). The delay corresponding to the treatment (F) indicates the time bin when food was potentially due. The graphs on the right side were obtained by calculating relative peck rates as a proportion of the rate at the peak time, and relative time bins as a proportion of the time bin corresponding to the delay when food was potentially due.

4.4.2 Preference for Variability

Preference for the fixed options decreased as their delay increased. Figure 3 shows the mean proportion of choices for both fixed options when each was paired with the variable in binary choices. Although the use of an averaging procedure can result in partial levels of preference that are not representative of the values exhibited by each bird, it is possible to see in the figure that the degree of variability associated with the average value was relatively low for all treatments and thus in this case also illustrative of individual preferences. Preference levels were marginally higher when the variable option was paired with P than when it was paired with N (Repeated-measures ANOVA on arcsine square-root transformed data, having option and treatment as within-subjects factors: $F_{1,9} = 4.12$, $P = 0.07$). Accordingly, indifference in the former case was obtained when the delay of the fixed option was approximately 16s, as compared to a value of approximately 15s in the latter case.

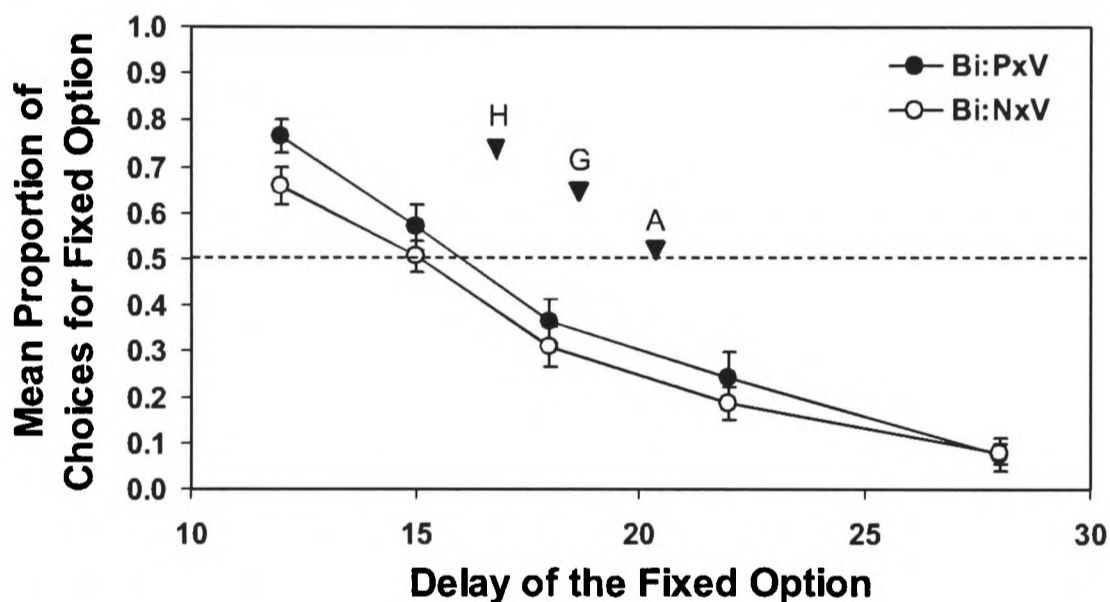


Figure 3. Mean proportion of choices (\pm S.E.) for both fixed options (P and N) when paired with the variable (V) in binary choices. The point indicated by 'A' (the arithmetic mean of the variable delay) indicates where a risk-neutral subject would cross the indifference line. H and G denote the harmonic (16.8s) and geometric (18.3s) means of the variable delay, indicating where a subject would cross the indifference line according respectively to the process-based models of risk-sensitive choice EoR ('expectation of the ratios; Bateson & Kacelnik 1996) and SUT (Scalar Utility Theory; Kacelnik & Brito e Abreu 1998).

4.4.3 Relative Choice Proportions

The relationship between the relative choice proportions in the binary and trinary contexts for the three possible choice pairs is illustrated in figure 4. For the choice pair (P,N) the group of individual slopes ('s', see 'Data Analysis') and intercepts ('c') was not significantly different

from one and zero, respectively ($s=0.85 \pm 0.3$; $t_9 = -0.49$, $P = \text{NS}$ and $c=0.12 \pm 0.26$; $t_9 = 0.47$, $P = \text{NS}$). The same analysis was used for the choice pairs (P,V) and (N,V). In the first case, both the intercept of the regression tended to be significantly different from zero ($c=0.17 \pm 0.08$; $t_9 = 2.1$, $P = 0.06$) and the slope was significantly different from one ($s=0.75 \pm 0.09$; $t_9 = -2.5$, $P < 0.05$). Since the test of both parameters yielded significant results, for illustrative purposes I drew the best-fitting line in the respective section of the figure. It indicates that when the variable option is preferred over the fixed (mainly on treatments 22 and 28, left corner of the graph), the relative preference for P tends to be higher in the trinary than in the binary context. Conversely, when the fixed options are preferred over the variable (mainly on treatments 12 and 15), the relative preference for P over V is higher in the binary than in the trinary context. These results thus suggest that when the fixed options are chosen on a high proportion of the trials, they compete more with each other than with the variable option. However, although the trend in this direction is significant, the magnitude of the deviations is small, and not confirmed with the choice pair (N,V). In this latter case, I cannot reject either the hypothesis that the intercept of the regression is different from zero ($c=-0.001 \pm 0.06$; $t_9 = -0.02$, $P = \text{NS}$) or the hypothesis the slope is different from one ($s=0.96 \pm 0.09$; $t_9 = -0.35$, $P = \text{NS}$). The figure also shows the presence of some preference reversals in all treatments, where the relative preference for one option over another is reversed between contexts. These reversals are however concentrated around the region of indifference between the options, and should be expected unless the relationship was devoid of variability.

To confirm the observed results, I conducted a Repeated-measures ANOVA for each of the choice pairs having ‘treatment’ and ‘context’ and within-subject factors. For the choice pair (P,N), there was no effect of either treatment or context ($F_{4,28} = 0.98$, $P = \text{NS}$ and $F_{1,7} = 2.58$, $P = \text{NS}$, respectively) on preferences. For the choice pairs (P,V) and (N,V) there was, as expected, a strong effect of treatment ($F_{4,28} = 42.1$, $P < 0.001$ and $F_{4,28} = 47.6$, $P < 0.001$, respectively), but no effect of context ($F_{1,7} = 0.90$, $P = \text{NS}$ and $F_{1,7} = 0.6$, $P = \text{NS}$, respectively). For none of the choice pairs was the interaction between ‘treatment’ and ‘context’ significant ($F_{4,28} = 1.44$, $P = \text{NS}$ and $F_{4,28} = 1.18$, $P = \text{NS}$, respectively). Therefore, in light of these results I do not have strong evidence to conclude that, overall, relative preferences deviate significantly from those predicted by the choice axiom. Although there was one significant trend towards similar options taking more from each other when they were valued highly, this trend was not confirmed by subsequent results and analyses.

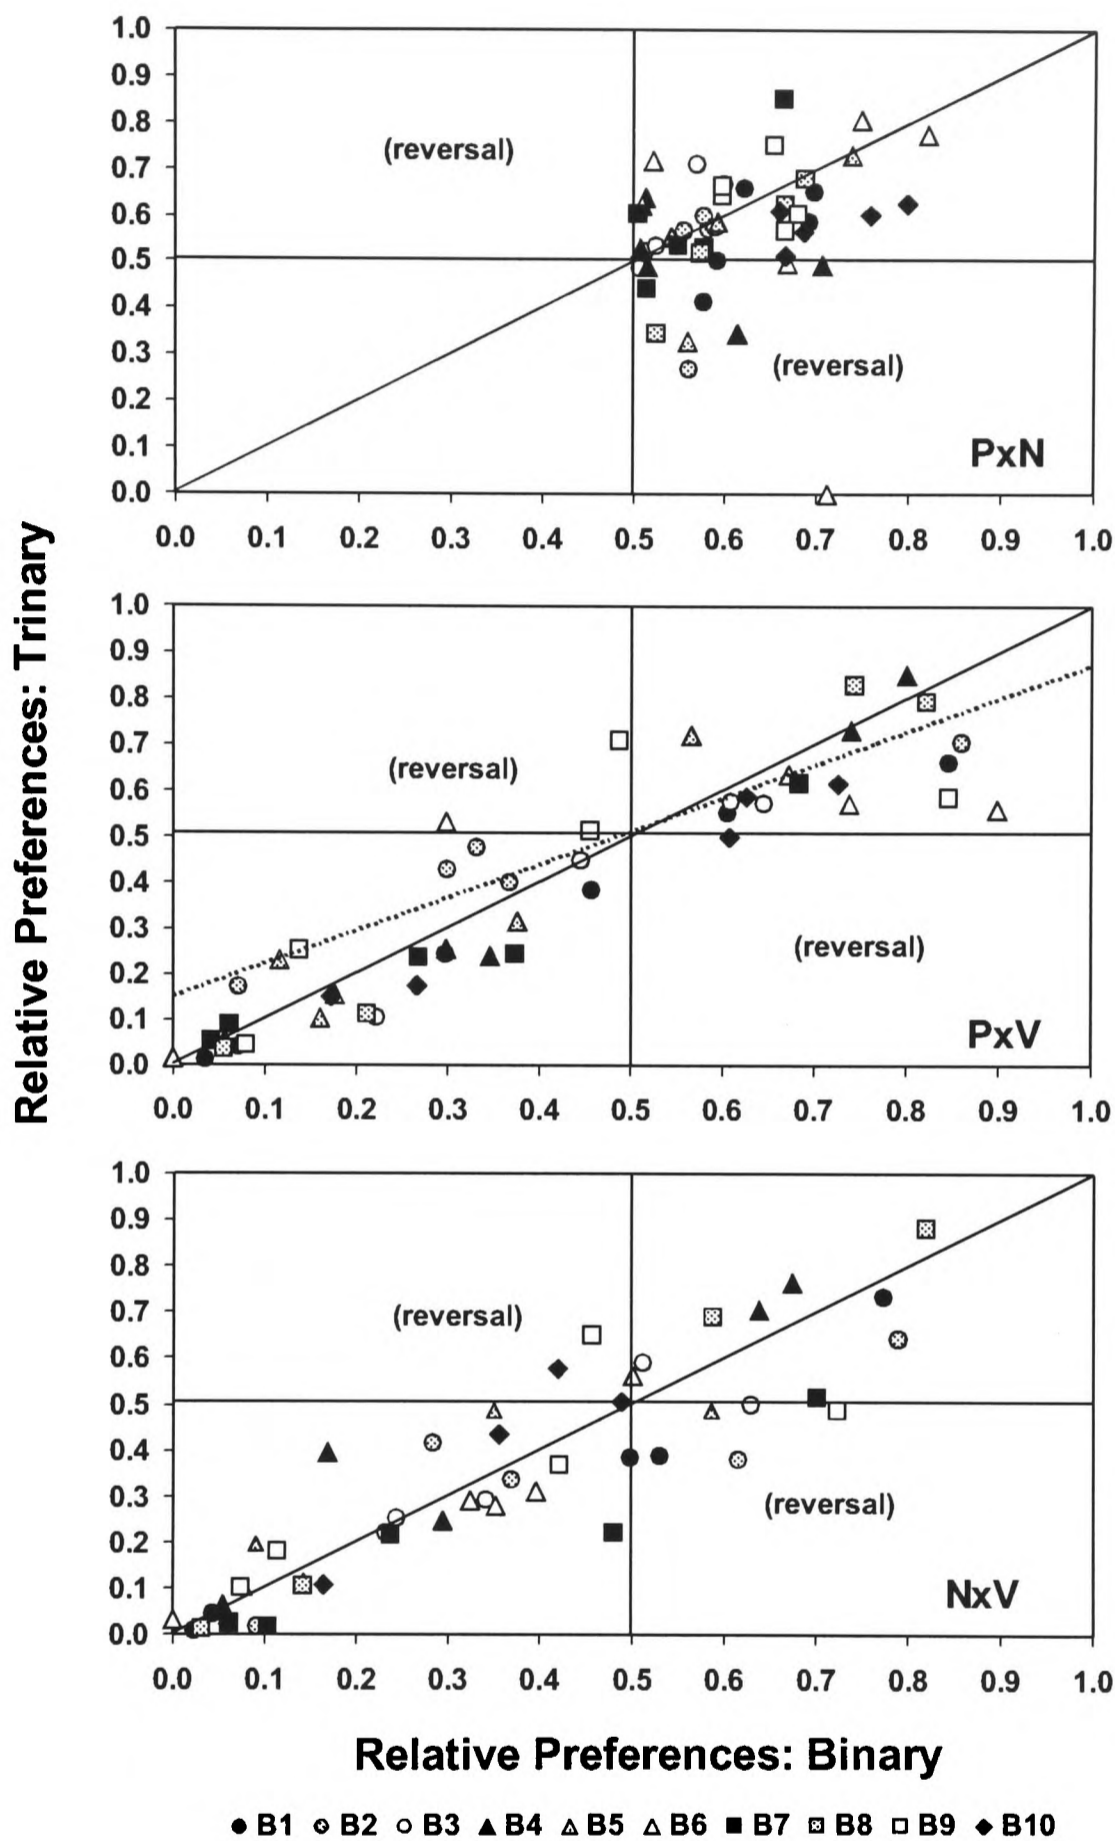


Figure 4. Relative preferences of subjects in the binary context for each of the three choice pairs plotted as a function of relative preferences in the trinary contexts for the same options. The diagonal lines have a slope of 1.0 and intercept of 0. The datapoints falling on the line represent cases where relative preferences were the same in both contexts. Datapoints falling above and below the line represent cases where relative preferences were higher and lower, respectively, in the trinary than in the binary context. The sections in the upper left and lower right corners of the graphs correspond to the areas of where preference reversals between contexts occurred. The dotted line in the PxV graph represent the slope of the regression of relative preferences in the binary context against those observed in the trinary. The different symbols represent each of the ten birds employed (B1...B10).

4.4.4 Absolute Choice Proportions

4.4.4.1 Regularity

The relationship between the absolute proportion of choices for each option in the binary and trinary contexts is depicted in the graphs of figure 5. The group of starlings could potentially violate regularity in up to six possible ways since, for each of the three options, the absolute proportion of choices could increase in relation to either of the two binary contexts in which it was presented. To test whether regularity was violated I conducted a series of paired-t tests on absolute choice proportions (arcsine square root transformed) for each of the 6 possible situations depicted in the figure. I conducted a different test for each of the five possible treatments, since I did not want to take the risk of accepting the null hypothesis of regularity for the grouped data (i.e. all treatments considered together) when there could be the possibility of rejecting it for an individual treatment. Overall, subjects did not significantly increase the absolute proportion of choices for any of the options in the trinary context in any of the five treatments. Out of the thirty possible tests (five treatments x six possible comparisons; paired-t tests), there were twenty-three instances where absolute choice proportions were significantly lower in the trinary context (or twelve instances if the alpha level to reject the null hypothesis is adjusted – Bonferroni procedure – for multiple comparisons), but not a single instance where it was higher. Therefore, the hypothesis of regularity cannot be rejected when subjects are considered together.

The diagonal line drawn in each of the graphs represents the points where absolute choice proportions are the same in both contexts. Data points above the line represent an increase in the absolute proportion of choices in the trinary context by an individual subject and, thus a potential violation of regularity. Although it is not possible to establish the significance of these deviations for each bird individually (there are no replicate measures), one should expect to find, out of the three hundred possible violations (six comparisons x five treatments x ten subjects), fifteen significant violations of regularity by chance alone. There were a total of thirty-one data points for which absolute preferences were higher (by any degree) in the trinary context. A closer inspection of this data reveals that most of the deviant points corresponded to choice proportions lower than 0.5. This was to be expected, since in these cases the number of choices for the focal option (the one for which regularity was tested) was very low, and differences of one or two choices between contexts could cause changes in proportion. For instance, in treatment 28 one bird chose option P only twice when

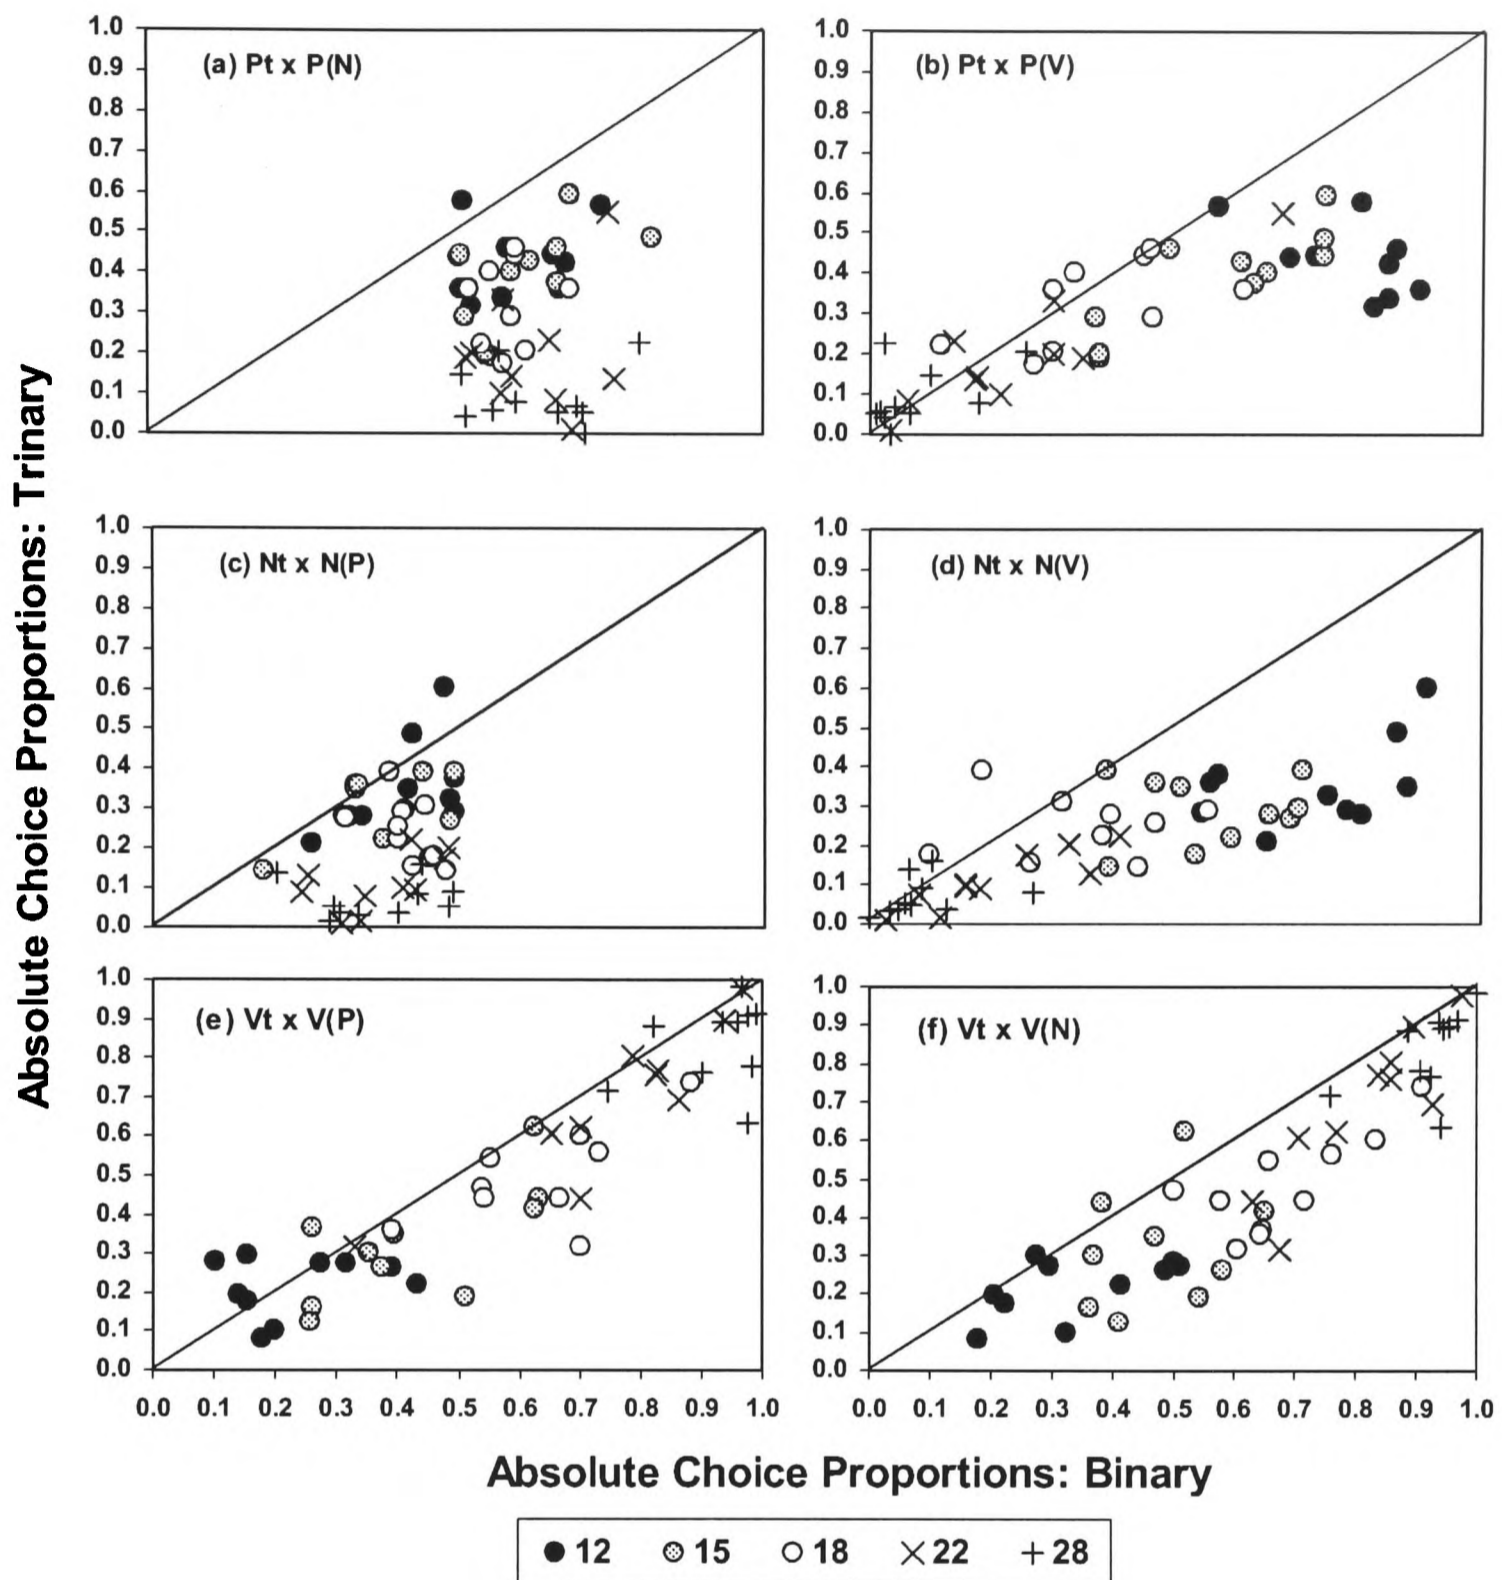


Figure 5. Absolute choice proportions of subjects for each of the options (in each of the two possible binary choice pairs where the option was offered) plotted as a function of the absolute proportion of choices for the same option in the trinary contexts (e.g. $P_T \times P(N)$ denotes the comparison of absolute choice proportions for P in the trinary context with that for P when paired with N in the binary). The diagonal lines have a slope of 1.0 and intercept of 0. The datapoints positioned above the line represent instances where choice proportions were higher in the trinary context, i.e. potential violations of regularity. The different symbols represent the different treatments employed.

paired with V (1.8% of the bird's choices), but three times in the trinary context (2.7%), causing preferences to be higher by 50% in this latter situation. Since for many data points we find that both the magnitude of the deviation and sample sizes are small, the rejection of the null hypothesis would be unlikely in most of these instances. I therefore do not have

evidence to conclude that regularity was violated at the individual level more frequently than what would be expected by chance.

4.4.4.2 Transitivity

None of the birds violated WST. The order in which the birds ranked the options depended, as expected, on the treatment employed, with preference for the variable option progressively increasing as the delay of the fixed option became longer. The fit between the observed proportions and those predicted by PST, SST and MST is illustrated in figure 6. In the case of PST, the null hypothesis that the slope of the regression is not significantly different from one ($s=0.99 \pm 0.1$; $t_9 = -0.06$, $P = \text{NS}$) or the intercept different from zero ($c=0.04 \pm 0.11$; $t_9 = 0.38$, $P = \text{NS}$) cannot be rejected. A repeated-measures ANOVA having treatment and 'nature of proportion' (i.e., whether predicted or observed) as within-subject factors showed that, although the result was not significant at $\alpha=0.05$, observed choice proportions tended to be slightly lower than those predicted according to the product rule ($F_{1,9} = 3.73$, $P = 0.08$; Greenhouse-Geisser correction for sphericity applied). As expected, there was a strong effect of treatment ($F_{2,28,20.1} = 6.52$, $P < 0.01$; Greenhouse-Geisser corrected). The interaction between the two factors was not significant ($F_{4,36} = 1.08$, $P = \text{NS}$).

SST and MST are potentially violated in those instances where the observed proportion of choices for the extreme options is lower than the predicted by either principle, that is, for those datapoints positioned below the diagonal lines drawn in the respective graphs. In the case of SST, it can be seen in the graph that most of the deviations correspond to those cases where choice proportions in the extreme pair (namely, the pair for which preference should be the most extreme) were either very high or close to indifference. If we were to find deviations of strong stochastic transitivity by chance, such pattern should be expected: when preference levels are already high in choices between adjacent options³, any variability in the subjects' choices for the extreme options can bring choice proportions down and lead to departures of SST. In the same way, since at the indifference point differences in choice proportions between the options are, by definition, small, even a small degree of variability or noise in the subject's choices for the extreme options could cause SST to be violated. I tested the significance of individual departures as shown in the graph. Out of fourteen deviations, only one was statistically significant ($\chi^2_1 = 5.64$, $P < 0.05$), less than what would be expected by chance (2.5 deviations). Therefore, I cannot reject the hypothesis that

³ Consider for instance the ranking $P > N > V$. In this case, P and N are adjacent to each other, as well as N and V. P and V, in turn, represent the extreme options.

the birds' preferences in binary choices satisfied SST. Similarly, none of the deviations of MST were significant (Chi-square tests corrected for continuity: $P = \text{NS}$).




Figure 6. Observed proportion of choices of individual birds in the extreme binary choice pair plotted as a function of the proportion of choices predicted if perfect (PST), strong (SST) and moderate stochastic transitivity (MST) were satisfied. The diagonal lines have a slope of 1.0 and intercept of 0. In the case of PST, datapoints falling on the line represent cases where the observed choice proportion was the same as the predicted. In the case of SST and MST, the datapoints falling on the line represent those cases where the choice proportions for the extreme choice pair was the same as the maximum and minimum proportion observed for the adjacent choice pairs, respectively. Thus, all points above the line represent those cases for which SST and MST were satisfied. The different symbols correspond to the distinct treatments employed (represented by the delay offered by the fixed option).

4.5 Discussion

The similarity hypothesis (Tversky & Russo 1969) predicts that a new option that becomes available into a choice set should take its share of choices mainly from the existing alternative most similar to it. My results showed, however, that the strength of preference exhibited by the starlings for the options did not depend on the structure or context of choice. Although the choice sets were characterized by the presence of unambiguous relations of similarity between its elements (two out of three options were always identical in terms of their parameter values, and thus perfectly substitutable by each other), preferences did not reflect the employment of a hierarchical sequence of choices, or the grouping of the available alternatives within distinct categories for the purpose of choice. One possibility is that, although the fixed options shared the same parameter values, the starlings were not able to perceive them as belonging to the same category, preventing the observation of hierarchical choices. However, a vast literature on categorisation (strongly based on experimental paradigms such as ‘matching-to-sample’; Herrnstein 1990, Pearce 1997) have shown animals to be generally able to identify options as belonging to separate groups according to their common and distinctive features. Hence, if the starlings possess such ability, they do not appear to employ it when choosing between different degrees of time-induced risk.

As discussed in Chapter 3, the presence of constraints in the perception of rewards could also lead to the observation of results qualitatively different from those predicted by the choice axiom. One of these models, Scalar Utility theory (SUT; Reboreda & Kacelnik 1991, Kacelnik & Brito e Abreu 1998), predicts for the present case levels of preference that are intermediate to those predicted by the similarity hypothesis and choice axiom, being closer to the former or latter model depending on the subject’s accuracy in the perception of time. Figure 8 illustrates the predictions of SUT and the similarity hypothesis – as compared to Luce’s choice axiom – for the relationship between the relative preferences in the binary and trinary contexts in this study. A higher level of preference for the fixed options in the binary context than in the trinary is predicted both by the similarity hypothesis and SUT. The deviations predicted by the latter model are however of smaller magnitude, especially in those cases where the accuracy of a subject’s time discrimination (γ) is high. In such cases, they are closer to the predictions of choice axiom than those yielded by the similarity hypothesis. Therefore, it would have been possible that the experimental tests conducted

here were not powerful enough to detect such a small effect size. Nevertheless, the direction in which observed preferences tended to deviate from the choice axiom was not consistent with this possibility (see observed slope in fig.4b): while SUT predicts that the relative preference for the variable option should be always higher in the trinary context, the observed levels of preference for the variable alternative were higher in the trinary context only when the fixed options were preferred, but lower otherwise.

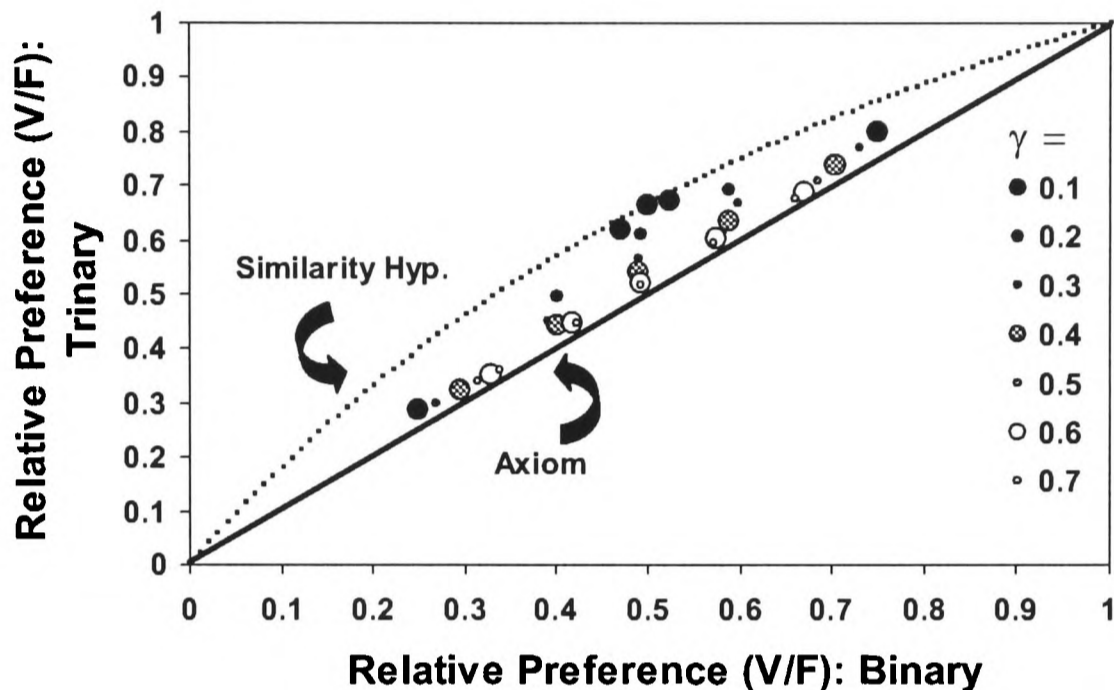


Figure 8. Predictions of SUT (see Chapter 3 for a detailed description) for choices between two levels of variability in delay to reward (variable, V: equiprobable delays of 12 and 28 s and fixed, F: constant delay of either 12 or 15 or 18 or 22 or 28 s depending on the treatment employed). The predictions vary according to the timing accuracy (γ) assumed in the model. The dashed and solid lines indicate the predictions of the similarity hypothesis and choice axiom, respectively.

Overall, the starlings' preferences were also transitive and regular across contexts. Although I did detect some departures from transitivity at the individual level, they were not more frequent than those expected by chance alone. The significance of individual breaches of regularity is however more difficult to assess. Here, the birds showed a potential breach of regularity (namely, a higher absolute preference for one of the alternatives in the trinary context as compared to the binary) on a total of thirty-one occasions. Although the absence of replicate measures per bird prevents one establishing the significance of these increases in the absolute level of preference, significant departures from regularity would be unlikely in most of these cases due to the extremely reduced magnitude and sample sizes associated with the observed effects (see 'Results'). Since there should be a total of fifteen significant breaches of regularity by chance alone, more than fifteen of these thirty-one cases would

have to be significant to reject the hypothesis of regularity, an unlikely result given the previous arguments. Previous research seems to have however overlooked the importance of analysing the likelihood of significant departures from regularity. For instance, in an experimental investigation of risk-sensitivity, Bateson (2002) compared the preferences of European starlings for variability in amount of food in binary and trinary choice situations, having observed a higher absolute preference for one option in the latter context on three occasions (out of the 48 possible comparisons). But although this number was close to the random expectation (2.5 significant deviations by chance alone) and the magnitude of the reported differences was small (differences of 1% and 2% in choice proportions between contexts), the results were interpreted as valid violations. Another study compared the foraging preferences of rufous hummingbirds (*Selasphorus rufus*) for artificial flowers differing in the volume and concentration of sucrose offered in trinary and binary contexts (Bateson et al. 2002), having also found increases in absolute preferences for one flower type in the trinary situation. Yet, the statistic tests employed were inadequate for testing regularity. The authors conducted contingency table analysis on the number of choices made by individual birds to each option offered in the binary treatment against the number of choices made to each of these same two options in the trinary. By ignoring choices for the third option included in the latter context, the authors however tested whether there were differences in the relative – and not absolute – proportion of choices between the treatments. Finally, Shafir et al. (2002) and Hurly & Oseen (1999) also reported an overall increase in the absolute preferences of gray jays (*Perisoreus canadensis*) and rufous hummingbirds, respectively, for food rewards when these were presented in a trinary choice context. Although the observed effects were robust in both cases, no supporting statistics were provided. In general, a thorough analysis of changes in preferences is paramount if one considers that observed choice proportions are only estimates of the real probabilities of selecting an option and thus, while I do not dispute the possibility that the reported results were valid, I feel that further analyses and studies are required to establish the generality and extent of the phenomenon.

The consistency in the starlings' behaviour observed in the present study does not pose major structural challenges to previous accounts of risk-sensitive foraging based on simpler, more tractable scenarios (reviewed in Kacelnik & Bateson 1996), and generalizes previous empirical research to multialternative choice situations, replicating the results of Schuck-Paim & Kacelnik (2002). The strength of preference for variability was increasingly

high as the delay of the fixed option increased, with the starlings showing a clear risk-prone strategy around the values in which both options offered the same mean rate of reward (see figure 3).

If on the one hand the observed stability in the starlings' behaviour is reassuring, on the other the fact that the third options took their share of choices from existing alternatives in proportion to their original shares – in accordance with Luce's theory – is more intriguing from a functional perspective. For instance, a number of models of risk-sensitive foraging (see review in McNamara & Houston 1992) predict what an animal's strategies towards risk should be if it was to maximize some surrogate of fitness. While it is not my aim here to discuss such models or their fit to my data, if we assume for the moment that the choice strategies exhibited by the starlings in one of the contexts (binary or trinary) approached the optimal behaviour, then we face a dilemma when comparing such strategies with those employed in the other context. Figure 9 makes this point clearer. It shows for each of the treatments the average preference for the fixed options in the binary situation and the sum of the preferences for the fixed option in the trinary, enabling the comparison of the total proportion of choices taken by the constant option in either context. We can see that the overall response to risk, as expressed by the overall level of preference for the fixed alternative, was different between the two choice situations – and in accordance with the predictions of the choice axiom. In one of the treatments (where the fixed option offered 18s of delay to food), preferences were even reversed: although the starlings were prone to risk in the binary situation they were risk-averse in the trinary, as in this latter context most of their choices were taken by the fixed alternatives (despite the fact that the birds showed a higher preference for the variable option as compared to each fixed alternative individually). Thus, although relative preferences between any pair of options were invariant and stable when considered separately, the overall allocation of choices between the two types of options (fixed and variable) changed with the size of the choice set. This effect should be even stronger should the predictions of the choice axiom hold for multialternative situations with larger choice sets, to the point of producing higher overall levels of preference for a type of option present in large numbers even when this option is clearly inferior to the alternative type. One possibility is that, while the observed behaviour is not functionally meaningful under the current framework, process-based models of choice – adaptive under a broader perspective – yielding response probabilities directly proportional to the properties of the rewards, independently of what else is available, could account for the phenomenon. I

am, however, not aware of any such model that could be presently applied in the context of risk-sensitive foraging.

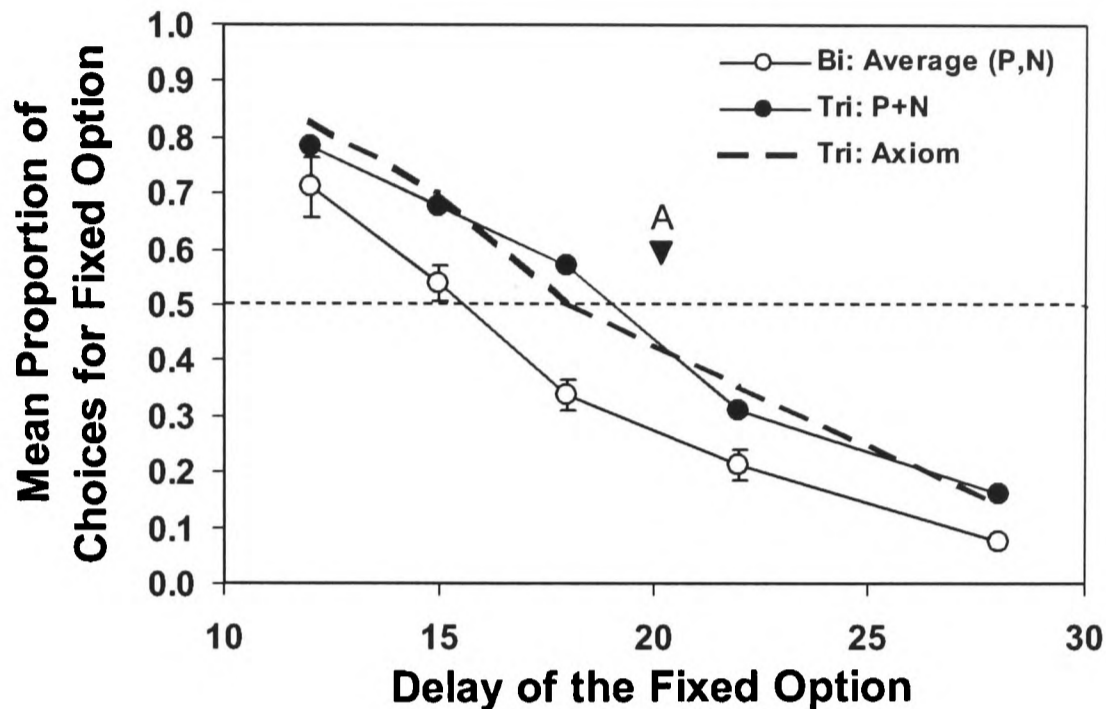


Figure 9. Mean proportion of choices (\pm S.E.) for the fixed options (P, N) when paired with the variable (V) in binary choices and the sum of the proportion of choices for the fixed options in the trinary context in the five treatments employed. The predictions of Luce's axiom for the trinary context were calculated based on the average proportion of choices for the fixed options in the binary situation. The similarity hypothesis predicts that preferences in the trinary context should be the same as in the binary. The point indicated by 'A' (the arithmetic mean of the variable delay) indicates where a risk-neutral subject would cross the indifference line.

Going back to the problem faced by the nesting bird discussed earlier, the present results indicate that the increase in the availability of potential sites in the margins of the field would have caused the probability of choosing the central region to increasingly drop, even if the latter offered a higher net benefit. Although satisfying a strong axiom of probabilistic theories of rational choice, this outcome clearly deviates from the optimal expectation.

4.6 References

- Bateson, M. 1993. Currencies for decision making: the foraging starling as a model animal. Oxford: University of Oxford. Unpublished DPhil. Thesis.
- Bateson, M. 2002. Context-dependent foraging choices in risk-sensitive starlings. *Animal Behaviour*, 64: 251-260.
- Bateson, M., Healy, S.D. & Hurly, A. 2002. Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, 63: 587-596.
- Bateson, M. & Kacelnik, A. 1996. Rate currencies and the foraging starling: the fallacy of the averages revisited. *Behavioral Ecology*, 7: 341-352.
- Catania, C.A. & Reynolds, G.S. 1968. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *Journal of Experimental Analysis of Behavior*, 11: 327-385.
- Gibbon, J. 1977. Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, 84: 279-325.
- Hermstein, R.J. 1990. Levels of stimulus control. *Cognition*, 37: 133-166.
- Houston, A.I. 1997. Natural selection and context-dependent values. *Proceedings of the Royal Society of London Series B - Biological Sciences*, 264: 1539-1541.
- Hurly, T.A. & Oseen, M.D. 1999. Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal Behaviour*, 58: 59-66.
- Kacelnik, A. & Bateson, M. 1996. Risky theories - the effects of variance on foraging decisions. *American Zoologist*, 36: 402-434.
- Kacelnik, A. & Brito e Abreu, F. 1998. Risky choice and Weber's Law. *Journal of Theoretical Biology*, 194: 289-298.
- Leak, T.M. & Gibbon, J. 1995. Simultaneous timing of multiple intervals: implications of the scalar property. *Journal of Experimental Psychology - Animal Behavior Processes*, 21: 3-19.
- Luce, R.D. 1959. *Individual Choice Behaviour: A Theoretical Analysis*. New York: Wiley.
- Luce, R.D. & Suppes, P. 1965. Preference, utility and subjective probability. In: *Handbook of Mathematical Psychology* (Ed. by Luce, R.D., Bush, R. & Galanter, E.). New York: John Wiley.
- McNamara, J.M. & Houston, A.I. 1987. Partial preferences and foraging. *Animal Behaviour*, 35: 1084-1099.
- McNamara, J.M. & Houston, A.I. 1992. Risk-sensitive foraging: a review of the theory. *Bulletin of Mathematical Biology*, 54: 355-378.
- Namoodiri, N.K. 1972. Experimental designs in which each subject is used repeatedly. *Psychological Bulletin*, 77: 54-64.
- Pearce, J.M. 1997. *Animal Learning and Conditioning: an introduction*. East Sussex: Psychological Press.
- Pratkanis, A.R. & Farquhar, P.H. 1992. A brief history of research on phantom alternatives: evidence for seven empirical generalizations about phantoms. *Basic and Applied Social Psychology*, 13: 103-122.
- Real, L.A. 1991. Animal choice behavior and the evolution of cognitive architecture. *Science*, 253: 980-986.
- Reboreda, J.C. & Kacelnik, A. 1991. Risk sensitivity in starlings: variability in food amount and food delay. *Behavioral Ecology*, 2: 301-308.
- Schuck-Paim, C. & Kacelnik, A. 2002. Rationality in risk-sensitive foraging choices by starlings. *Animal Behaviour*, 64: 869-879.
- Shafir, E.B., Osherson, D. & Smith, E.O. 1989. An advantage model of choice. *Journal of Behavioral Decision Making*, 2: 1-23.

- Shafir, E.B., Osherson, D.N. & Smith, E.E. 1993. The advantage model - a comparative theory of evaluation and choice under risk. *Organizational Behavior and Human Decision Processes*, 55: 325-378.
- Shafir, S. 1994. Intransitivity of preferences in honey-bees: support for 'comparative' evaluation of foraging options. *Animal Behaviour*, 48: 55-67.
- Shafir, S. 2000. Risk-sensitive foraging: the effect of relative variability. *Oikos*, 88: 663-669.
- Shafir, S., Waite, T.A. & Smith, B.H. 2002. Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 51: 180-187.
- Sumpter, C.E., Temple, W. & Foster, T.M. 1999. The transitivity of choices between different response requirements. *Journal of Experimental Analysis of Behavior*, 72: 235-249.
- Tversky, A. & Russo, J.E. 1969. Substitutability and similarity in binary choices. *Journal of Mathematical Psychology*, 6: 1-12.
- Tversky, A. & Simonson, I. 1993. Context-dependent preferences. *Management Science*, 39: 1179-1189.
- Waite, T.A. 2001. Background context and decision making in hoarding gray jays. *Behavioral Ecology*, 12: 318-324.
- Wedell, D.H. 1991. Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology - Learning Memory and Cognition*, 17: 767-778.

CHAPTER 5

5.1 INTRODUCTION	72
5.2 WHEN FORAGING IS STATE-DEPENDENT	73
5.2.1 MINIMIZING THE PROBABILITY OF COSTLY ERRORS	74
5.2.2 MINIMIZING THE PROBABILITY OF STARVATION	77
5.3 CONCLUSION	80
5.3.1 TESTING RATIONALITY IN ANIMAL CHOICES	82
5.4 REFERENCES	83

STATE-DEPENDENCE AND ECONOMIC RATIONALITY IN ANIMAL CHOICES

5.1 Introduction

The study of animal decision-making has undeniably progressed through the use of insights borrowed from normative thinking in economy. A particularly good example of this exchange of ideas is the application of optimality theory in the evolutionary analysis of behaviour. By assuming that, as a result of natural selection, animals should behave so as to maximize their lifetime fitness (or surrogates of it under existing constraints), the approach provided a common framework for the analysis of animal choice behaviour in terms of the benefits and costs associated with different courses of action.

One recent threat, however, to both economic and biological theories of decision-making is based on the finding that individuals often fail to behave 'rationally' (Tversky & Simonson 1993, McFadden 1999, Waite 2001a, Waite 2001b, Shafir et al. 2002), as required under the assumption of value maximization. Economic rationality embraces a number of principles that are necessary consequences of the existence of a scale of value upon which humans, and non-humans alike, base their choices. Transitivity, for instance, is a hallmark of rational models. It requires, as one can intuitively see, that if 'a' is preferred to 'b', and 'b' to 'c', that 'a' should also be preferred over 'c'. Another rational principle underlying most theoretical work on choice is that of independence from irrelevant alternatives (or simply IIA), namely the expectation that options should be evaluated independently of the presence or absence of other options.

As I highlight in previous chapters, violations of rationality are usually taken to imply context-dependence, or the idea that the value of an option would not be only determined by its intrinsic properties and consequences as assumed by most normative models (but see Houston 1997), but instead would be constructed at the moment of choice as a function of the value of other options available. In fact, a number of psychological mechanisms have been proposed to explain how such dependence could be achieved. According to them, the observed anomalies, or failure to exhibit consistent preferences (independently of what these preferences are or their fit to optimal expectations), would be attributable to the specific ways whereby individuals process information, or to the specific heuristics used for making choices (Parducci 1965, Tversky 1969, Shafir et al. 1989, Wedell 1991, Payne et al. 1992, Simonson & Tversky 1992). For instance, one such class of models

referred to as 'comparative' (Shafir et al. 1989) proposes that, instead of being assigned a single and independent value summarizing all its properties, an option would be compared along each of its relevant dimensions (or attributes) with another before a decision is reached. The framework thus enables the value of an option to depend on the option to which it is compared, providing a possible mechanism for departures of rationality whenever options are defined by more than one attribute.

The central concept underlying the previously described and similar accounts of irrational behaviour is that preferences would be constructed, rather than revealed, during the choice process (Payne et al. 1992). In this Chapter I however argue that breaches of rationality may also emerge under the assumption of independent choice currencies if the likely effect of context on an organism's state is considered concomitantly with the observation that the fitness consequences of choices are likely to be state-dependent (McFarland & Houston 1981, Houston & McNamara 1999). Unless an organism's environment was devoid of any variability between generations (which would allow estimates of its conditions to be genetically determined), animals often need to experience the available alternatives before acquiring any degree of knowledge about their respective outcomes. If we acknowledge that experience with different contexts, and thus contingencies, should differentially affect correlate measures of state (e.g., physiological parameters such as levels of energy reserves, body temperature and gut contents, as well as the degree of information about the environment, McNamara & Houston 1986), and that the fitness consequences of decisions are not fixed, but contingent upon state, then we should expect preferences to strongly depend on the context of choice in natural situations. Here I illustrate this argument by considering some possible examples of foraging situations in which individuals behaving so as to maximize their probability of survival and reproduction violate rational principles. It is my aim to show that the observation of seemingly irrational behaviour is also compatible with state-dependence in decision-making.

5.2 When Foraging is State-Dependent

In this section I look into some possible effects of changes in context on an organism's energetic state and, consequently, choice behaviour. I assume a scenario in which an animal searching for food has to repeatedly choose between foraging patches during a given period.

For simplicity, I compare two situations: one in which the animal has to choose between two patches of food (A and B) with another in which a third patch (C) becomes available. In all cases, I suppose that the average profitability of the patches (μ) is such that $\mu_A > \mu_B > \mu_C$.

5.2.1 Minimizing the probability of costly errors

I start by considering a similar situation as that presented in Chapter 3, in which choices for less profitable patches may result from the presence of errors in the perception of rewards. Following McNamara & Houston (1986, 1987), I assume that the probability of choosing the suboptimal action is a function of the difference between the expected terminal fitness values of each alternative. The main idea underlying their approach is that deviations from the optimal policy should be more common whenever the costs of such deviations are small, whereas costly errors should be rare. Here I build on this premise to analyse some possible effects of context on foraging preferences. Figure 1 illustrates the rationale. In the binary situation, the mean energetic state of the animal is given by X_{ab} , and determined as a function of the proportion of choices for both options A and B (that is, the profitability of the patches – weighted by the relative frequency with which they are chosen – affects some correlate measure of the forager's energetic state). In the trinary situation, where C is also available, the mean energetic state of the animal is given by X_{abc} , and determined as a function of the proportion of choices for options A, B and C. We can see (fig.1b) that, if the relative proportion of choices between A and B (computed both in the binary and trinary situations as the proportion of choices for A over the sum of choices for A and B) is unaltered between the two contexts, the mean energetic state of an individual will be lower in the trinary situation ($X_{abc} < X_{ab}$) simply because option C is energetically poorer. This should hold whichever the proportion of choices for C (assuming $p(C) > 0$ and that in both situations the individual has the same total number of foraging opportunities), but the decrease in the subject's energetic state should be more pronounced the higher the number of elections for C. In fact, in Chapter 3 I have shown that the inclusion of a third option into a choice set should, overall, decrease the relative proportion of choices for the most profitable option (the proportion of choices for A relative to B, in the presence case), leading to an even higher decrease in the forager's energetic state. However, I shall conservatively assume here that the relative proportion of choices for the two original patches A and B remains unaltered when C becomes available, so that the effect of state changes on

preferences can be analysed independently of the perceptual effect described in that Chapter. I next consider the potential improvement in state produced by pursuing either A or B in the two scenarios (to make the comparison clearer, I have omitted changes in state produced by choosing C in the trinary context). As shown in the graph, the difference between the improvement in state caused by the election of either A or B is the same in the binary (ΔX_{ab}) and trinary (ΔX_{abc}) situations, and so the cost of a suboptimal election (B) should also be the same in these two contexts. Yet, this should not be the case if changes in fitness (or some other positively valued measure) and not purely state, affected choices (some evidence towards this is given by Kacelnik & Marsh 2002, and in Chapters 6 and 7). This is illustrated in figure 1a. Following Kacelnik & Marsh (2002), I plotted fitness as a function of state with the abscissa drawn vertically, to emphasize the link between the two graphs. It is likely that potential improvements in fitness associated to each option's gains are a diminishing function of an organism's initial state, being concave in shape (for instance, if state represents an organism's level of reserves, then the contribution of a food item to his future reproductive success should decrease with increasing reserves in the presence of physiological limits to store energy – see also McNamara & Houston 1982 for a thorough discussion of the foundations of the use of such a function). Under this assumption, the figure shows that the canonical cost (δ , defined as fitness cost paid by choosing the suboptimal action; McNamara & Houston 1986) of pursuing option B is higher in the trinary context, which should lead to a higher level of preference for A relative to B if the probability of costly errors is lower than that of less costly ones. It is worth noting that this outcome could be likewise expected under a more proximal (mechanistic) perspective: in agreement with the observation that the function characterizing an animal's learning rate across time tends to be steeper under conditions of higher deprivation (Capaldi & Havancik 1973, Tarpy & Mayer 1978), under poorer states an animal could learn to discriminate the properties associated with each option more rapidly, and so costly errors would be eliminated before less costly ones (see also Houston 1987).

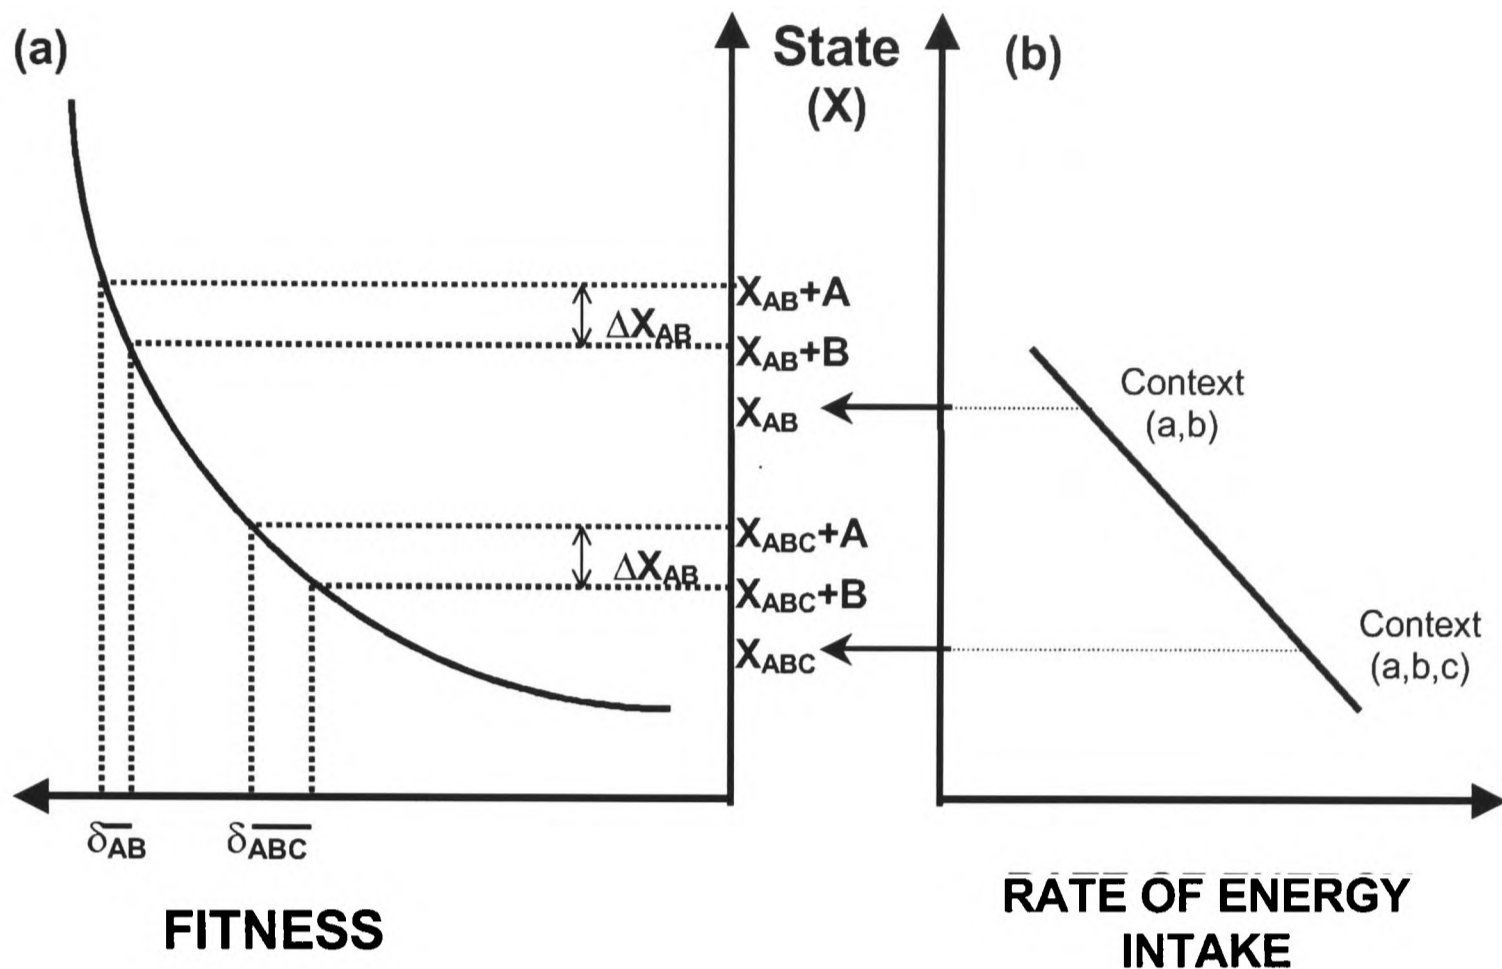


Figure 1. A functional model of how changes in state can lead to differences in partial levels of preferences (or ‘context-dependent’ preferences). In (b) a correlate measure of state is assumed to decrease as mean rate of energy intake is reduced (due to the inclusion of less profitable options in the choice set). Two contexts are considered: a binary context, where only options A and B are available (a,b) and a trinary context where a third option C is also present (a,b,c). In all cases, the mean profitability of the patches (μ) is such that $\mu_A > \mu_B > \mu_C$. Exposure to the trinary context therefore leads to a poorer state (X_{abc}) than that reached after exposure to the binary (X_{ab}) if $p(C) > 0$. In (a) the figure shows fitness as a concave function of an organism’s initial state. ‘ $X_{ab}+A$ ’ and ‘ $X_{ab}+B$ ’ denote the state reached by a forager in the binary scenario after pursuing options A and B, respectively. Similarly, ‘ $X_{abc}+A$ ’ and ‘ $X_{abc}+B$ ’ denote the state reached by a forager in the trinary scenario after pursuing options A and B, respectively. The fitness cost of choosing option B is higher in the trinary (δ_{abc}) than in the binary situation (δ_{ab}).

The hypothetical situation described above thus shows that, in violation of the principle of independence from irrelevant alternatives, the inclusion of additional options into a set can affect the strength of preference for existing options whenever it concomitantly affects the state of individuals. Although the example concerns the comparison of binary and trinary choices, it can be certainly generalized to other situations in which changes in state are likely to occur, such as those caused by previous exposure to distinct backgrounds (as defined by the number and nature of options experienced in the past).

As a final remark it is important to highlight that the above process should be however viewed as a dynamic one. For instance, by electing a poorer option (C in the

example) individuals might experience a poorer energetic state, making the likelihood of pursuing the most profitable option to increase (and that of choosing option C to decrease) under the framework previously discussed. This in turn could lead to an improvement of the animal's state and return to the homogeneity of states between contexts after a certain number of elections. If on the one hand such homogeneity could equalize preferences between contexts by lowering back the proportion of choices for the better option (A), on the other reaching a richer state could enhance the choice probability of the poorest option once again, reducing the individual's level of reserves, and reinitiating the cycle.

5.2.2 Minimizing the probability of starvation

I now consider another scenario in which a non-breeding animal needs to reach a certain energetic requirement by the end of a finite foraging period. Failing to do so leads to starvation. The animal is allowed to switch between options and is able to continuously assess his state and to adjust his behaviour accordingly. Possible analogies are those of a small bird in winter that needs to store enough energy before dusk to be able to survive during the night (see Houston & McNamara 1999 and references therein), or a food-hoarding animal that has to cache for when food is unavailable. As in the previous example, I compare two situations: one in which the animal has to choose between two food patches (A and B) with another in which a third patch (C) is also available, with the average profitability of the patches (μ) within this period being such that $\mu_A > \mu_B > \mu_C$.

To illustrate some possible context effects when foraging is state-dependent I consider a simplified situation, as diagrammatically sketched in figure 2, where an animal has a total of five foraging opportunities to reach a survival threshold (R) before foraging is interrupted. Figures 2a and 2b represent the binary and trinary contexts, respectively. In both situations choices for A lead to a positive net gain of energy, whereas the gain associated with choosing B equalizes the costs incurred (e.g., metabolic expenditure and handling costs), so the resulting net gain is null. In the trinary context the costs of electing C are higher than the benefits gained with the choice, resulting in a negative net gain of energy. Such parameter values were arbitrarily chosen for graphical simplicity, but the rationale described in this section equally applies to other choice situations where $\mu_A > \mu_B > \mu_C$. Here I set X_0 (the initial state of the animal) $< R$.

In both contexts the shaded area represents the space containing possible sequences of action, that is, those sequences enabling the forager to overcome the lethal boundary before foraging is interrupted. Choices for the less profitable patch result from the tendency of animals to continuously gather information about the world that might be useful in unforeseen circumstances (Inglis et al. 1997, Bean et al. 1999, Inglis et al. 2001). Following Inglis et al. (2001), I consider that possible sources of uncertainty relate to the presence of novel as well as ambiguous food patches (those for which the animal is uncertain either about its profitability or about whether it has visited it previously) and the possibility that patch profitability varies with time (even though it may not, as assumed in this simplified example).

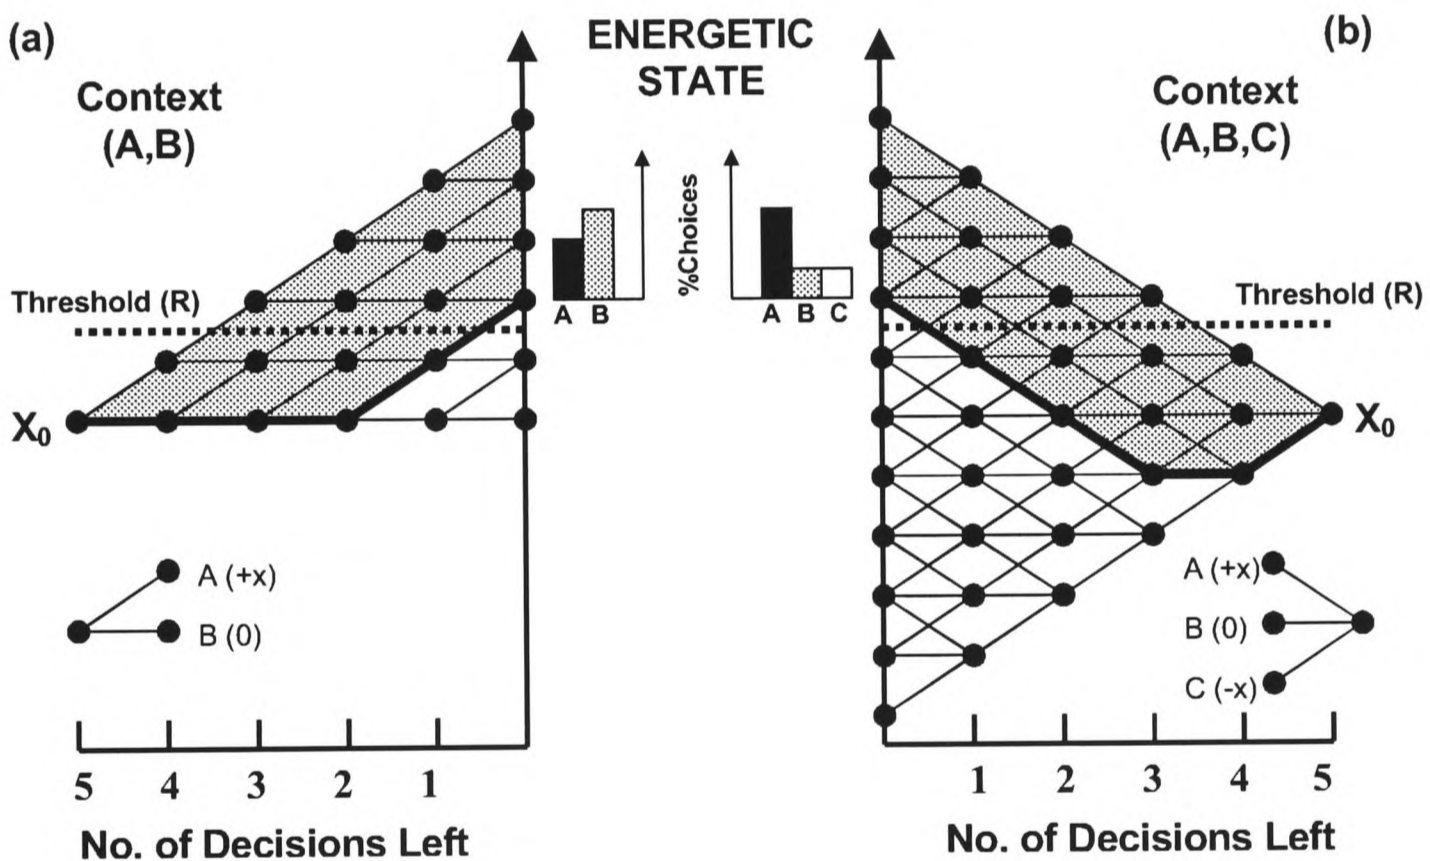


Figure 2. Diagrammatic representation of the choice problem faced by a forager that has to reach an energetic threshold (R) by the end of a foraging bout consisting of five choice opportunities. X_0 denotes the state of the animal at the beginning of the bout. Choices for A, B and C lead to a positive, null and negative net gain of energy, respectively. Each node in the graphs represents a possible action, and the shaded areas represent the space containing possible trajectories of choice under the requirement of reaching R . While in the binary context (a) there are always two possible available actions (A and B), in the trinary (b) foragers can choose between three alternatives at each decision node (A, B and C). In (a) the thick lines illustrates a possible sequence of decisions involving the maximum possible number of suboptimal elections (i.e., elections for option B) that enables the energetic boundary to be crossed. In (b) the thick lines also illustrates one trajectory involving the maximum possible number of suboptimal elections (elections for option B) if option C is pursued once when it becomes available. The bar graphs in the centre show, for each context, the resulting absolute proportion of choices for each option in the end of the foraging period under the previously described trajectories.

We can see in figure 2a that the minimum number of elections for A that enable the forager to reach the threshold is two, or 40% of the opportunities. Conversely, the minimum number of elections for A in the presence of option C (fig.2b) is three, or 60% of the foraging opportunities. Thus, if the new option C is pursued once when it becomes available, the minimum frequency with which the most profitable option (A) is chosen has to be increased to compensate for the energetic loss and enable the forager to be above the lethal-boundary. This possible increase in the absolute proportion of choices for A, needed to counterbalance an ‘unlucky’ decision, would represent a violation of regularity. Regularity is one of the most general and less restrictive rational principles. It is a weak form of the notion of independence from irrelevant alternatives (Luce & Suppes 1965), and asserts that the absolute proportion of choices for an option cannot be increased when a new option is included in the choice set. Yet, the example illustrates that it is hypothetically feasible to violate regularity whenever the context of choice leads to a deterioration of the state of a subject that has to be subsequently compensated.

The possibility previously described is however only one out of the many possible courses of action available to the forager, and not necessarily a likely one. It is based on the notion that animals would be able to select foraging patches other than the most profitable one with the aim of gaining information (potentially useful in the future) as long as being able to reach the required threshold. Such strategy could be achieved if, for instance, animals could monitor their level of deprivation (e.g. through their hunger state) and the time left for foraging, restricting sub-optimal elections to states corresponding to high level of reserves, and early in the foraging bout. Some evidence that this could be the case comes for instance from a study from Dall & Johnstone (2002), who have recently developed a state-dependent sampling model to examine choices between two foraging options differing in the level of variability associated with their payoffs. In their model a forager is allowed to reduce the uncertainty associated with the options by either collecting information or insuring against it by building up reserves. They show that sampling the more variable, or ‘risky’, option should decline with decreasing energy reserves and late in the foraging period; that is, it should be observed in those instances where the forager is able to compensate for an increase in the risk of an energetic shortfall in case the option sampled is energetically poor. Perhaps more related to the present description, evidence in this direction also comes from a number of studies in the literature on ‘contrafreeloading’, or the observation that animals will forage persistently for ‘earned food’, possibly with the aim of

gaining information (as previously described), even when identical free food is also available. Similarly to the results of Dall & Johnstone (2002), these studies show that information-gathering behaviour increasingly predominates over maximization of food intake as deprivation decreases (Inglis & Ferguson 1986, Bean et al. 1999).

Of course, under different assumptions one should observe different trajectories, which might or might not lead to a breach of IIA. For instance, a key assumption of the previously described example is that trajectories that imply violations of regularity would be more likely whenever fitness is a step function of energy intake, instead of increasing continuously to a maximum attainable value⁴. In this regard, two points are worth considering. Firstly, it might be the case that attempts to optimise for the present may limit an organism's future possibilities. Thus, although the assumption of a step function might be unlikely to hold in the short-term given the likelihood that extra gains should enhance an organism's probability of survival and reproduction (see Krebs & McCleery 1984 for a discussion), it is possible that in the long run animals benefit by allocating a proportion of time to explore further possibilities whenever it is safe (in energetic terms) to do so. Within this perspective, conditions such as the time left for foraging, level of energy reserves and degree of stochasticity in gains should determine the optimal balance between maximizing food intake and gaining information, and whether animals should sacrifice at all short-term gain for a potentially better long-term outcome. Secondly, the diagrammatic representation in figure 2 clearly oversimplifies the choice problem faced by a forager. Nevertheless, rather than being aimed at determining the specific circumstances whereby rationality could be violated, or advocating that such violations should be likely, my objective is to illustrate the idea that it would be possible to observe formal violations of rational axioms whenever energetic losses caused by contextual differences must be compensated by reliable gains (or, conversely, whenever extra profits enable the use of strategies other than exploitation).

5.3 Conclusion

The idea that the value, or 'utility' of an item should be dependent on previous achievements has been long recognized. As Bernoulli (1738) emphasized, the degree of

⁴ Notice that under this assumption the observation of trajectories implying violations of regularity could also be expected within a satisficing problem (Simon 1955), whereby the decision-maker is satisfied after meeting some minimal (energetic in the present case) requirement.

satisfaction obtained with a given improvement in one's assets should be inversely proportional to the amount previously owned. Analogously, evolutionary approaches to behaviour have often considered the dependency of an organism's state on the fitness consequences of its choices (McFarland & Houston 1981, Mangel & Clark 1988, Cuthill & Houston 1997, Houston & McNamara 1999). Building on this assumption, here I have argued that state-dependency in decision-making may lead to apparent violations of rational axioms if the likely effects of context on an organism's state are taken into account. The hypothetical situations described in this chapter represent only some illustrative outcomes within a wider range of possibilities that may emerge when context influences state measures. For instance, one of the most frequently cited models in the literature on risk-sensitive foraging, the 'energy budget rule' (Stephens 1981, Stephens & Krebs 1986), predicts the occurrence of preference reversals (from risk-proneness to risk-aversion and vice-versa) depending on the level of reserves of subjects relative to a critical threshold that enables survival. Although empirical evidence supporting the rule is extremely controversial (reviewed in Kacelnik & Bateson 1996), it would be possible, in theory, to find switches in risk-sensitive preference between contexts depending on the extent to which these contexts can differently affect overall budgets. Additionally, the arguments presented here may be particularly important in experimental paradigms that present subjects with forced (or 'no-choice') trials, in which individuals must experience all possible options when presented in isolation to be able to learn about them. In these designs, compulsory exposure to energetically poorer options introduced into the choice set may affect the state of the individual, leading to a decrease in the probability of sub-optimal elections (as defined in section 5.2.1) or to the employment of compensatory strategies (as defined in section 5.2.2) in choice trials, hence inducing breaches of rationality.

A somewhat different way of analysing the problem is to redefine the available choice objects not only as a function of their physical and temporal properties, but also as a function of the 'state' in which they are experienced. Under this framework – developed mainly for the study of choices under uncertainty and usually referred to as 'state-preference approach' in economics (Debreu 1959) – it is possible to specify each 'good' or 'reward' experienced in each state as a different 'commodity', and then use standard economic theories to analyse preferences. For instance, as discussed along the previous sections, acquisition of a particular reward in one state may provide a different marginal fitness gain

than does the same reward at a different state. In this case, we could define reward in state 1 as x_1 and reward in state 2 as x_2 , and a utility function, $U(x_1, x_2)$, over the two ‘commodities’ x_1 and x_2 . By redefining the utility function to be maximized over sets of state-dependent commodities the apparent problem associated with the observation of breaches of rationality (resulting from state-dependent preferences) becomes trivial, and the violations cease to exist.

5.3.1 Testing rationality in animal choices

Although relevant in both scenarios, the arguments put forth establish a practical distinction between human and non-human research, since while humans can realize the possible outcomes of choices without having to experience them, animals need to initially come into contact with the consequences of available actions in order to learn about them. If distinct contexts translate into exposure to different contingencies, which in turn can affect state, then individuals behaving so as to optimise their probability of survival and reproduction might appear to breach rationality.

While state-dependency may lead to the observation of apparent irrational behaviour independently of context-dependence in the assignment of value to options, I obviously do not claim that previous findings of irrationality were necessarily due to the mechanisms hypothesised here, nor do I dispute the potential influence of context-dependence on choice behaviour. However, the possibilities discussed in this Chapter emphasize that a thorough and deeper analysis of choice behaviour is paramount if we are to understand the mechanisms underlying animal decision-making and the means whereby animals process information. Here I have shown that differences in context may affect correlated measures of state, and unwittingly lead to the observation of economic irrationality. The distinction can be extremely important: state-dependency in choices is fundamentally different from context-dependent behaviour.

5.4 References

- Bean, D., Mason, G.J. & Bateson, M. 1999. Contrafreeloading in starlings: testing the information hypothesis. *Behaviour*, 136: 1267-1282.
- Bernoulli, D. 1738. Specimen Theoriae novae de Mensura Sortis (Exposition of a New Theory on the Measurement of Risk). Translated by L. Sommer. *Econometrica* (1954), 22: 23-36.
- Capaldi, E.D. & Havancik, J.R. 1973. Effects of previous weight level on rats' straight-alley performance. *Journal of Experimental Analysis of Behavior*, 97: 93-97.
- Cuthill, I.C. & Houston, A.I. 1997. Managing time and energy. In: *Behavioural Ecology* (Ed. by Krebs, J.R. & Davies, N.B.), pp. 97-120. Oxford: Blackwell.
- Dall, S.R.X. & Johnstone, R.A. 2002. Managing uncertainty: information and insurance under the risk of starvation. *Philosophical Transactions of the Royal Society of London Series B - Biological Sciences*, 357: 1519-1526.
- Debreu, G. 1959. *Theory of Value, an Axiomatic Analysis of Economic Equilibrium*. New Haven: Yale University Press.
- Houston, A.I. 1987. The control of foraging decisions. In: *Quantitative Analysis of Behavior (Foraging)* (Ed. by Commons, M.L., Kacelnik, A. & Shettleworth, S.J.), pp. 41-62. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Houston, A.I. 1997. Natural selection and context-dependent values. *Proceedings of the Royal Society of London Series B - Biological Sciences*, 264: 1539-1541.
- Houston, A.I. & McNamara, J.M. 1999. *Models of Adaptive Behaviour*. Cambridge: Cambridge University Press.
- Inglis, I.R. & Ferguson, N.J.K. 1986. Starlings search for food rather than eat readily available food. *Animal Behaviour*, 34: 614-616.
- Inglis, I.R., Forkman, B. & Lazarus, J. 1997. Free food or earned food? A review and fuzzy model of contrafreeloading. *Animal Behaviour*, 53: 1171-1191.
- Inglis, I.R., Langman, S., Forkman, B. & Lazarus, J. 2001. An information primary model of exploratory and foraging behaviour. *Animal Behaviour*, 62: 543-557.
- Kacelnik, A. & Bateson, M. 1996. Risky theories - the effects of variance on foraging decisions. *American Zoologist*, 36: 402-434.
- Kacelnik, A. & Marsh, B. 2002. Cost can increase preference in starlings. *Animal Behaviour*, 63: 245-250.
- Krebs, J.R. & McCleery, R.H. 1984. Optimization in behavioural ecology. In: *Behavioural Ecology: and evolutionary approach* (Ed. by Krebs, J.R. & Davies, N.B.), pp. 91-121. Massachusetts: Sinauer.
- Luce, R.D. & Suppes, P. 1965. Preference, utility, and subjective probability. In: *Handbook of Mathematical Psychology* (Ed. by Luce, R.D., Bush, R. & Galanter, E.). New York: John Wiley.
- Mangel, M. & Clark, C.W. 1988. *Dynamic Modelling in Behavioral Ecology*. Princeton: Princeton University Press.
- McFadden, D. 1999. Rationality for economists? *Journal of Risk and Uncertainty*, 19: 73-105.
- McFarland, D.J. & Houston, A.I. 1981. *Quantitative Ethology: the State Space Approach*. Boston: Pitman Advanced Publishing Program.
- McNamara, J.M. & Houston, A. 1982. Short-term behaviour and lifetime fitness. In: *Functional Ontogeny* (Ed. by McFarland, D.J.), pp. 60-87. London: Pitman.
- McNamara, J.M. & Houston, A.I. 1986. The common currency for behavioural decisions. *American Naturalist*, 127: 358-378.
- McNamara, J.M. & Houston, A.I. 1987. Partial preferences and foraging. *Animal Behaviour*, 35: 1084-1099.

- Parducci, A. 1965. Category judgement: a range-frequency model. *Psychological Review*, 72: 407-418.
- Payne, J.W., Bettman, J.R. & Johnson, E.J. 1992. Behavioral decision research - a constructive processing perspective. *Annual Review of Psychology*, 43: 87-131.
- Shafir, E.B., Osherson, D. & Smith, E.O. 1989. An advantage model of choice. *Journal of Behavioral Decision Making*, 2: 1-23.
- Shafir, S., Waite, T.A. & Smith, B.H. 2002. Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 51: 180-187.
- Simon, H.A. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69: 174-183.
- Simonson, I. & Tversky, A. 1992. Choice in context: trade-off contrast and extremeness aversion. *Journal of Marketing Research*, 29: 281-295.
- Stephens, D.W. 1981. The logic of risk-sensitive foraging preferences. *Animal Behaviour*, 29: 628-629.
- Stephens, D.W. & Krebs, J.R. 1986. *Foraging Theory*. Princeton: Princeton University Press.
- Tarpy, R.M. & Mayer, R.E. 1978. *Foundations of Learning and Memory*. Illinois: Scott, Foresman and Company.
- Tversky, A. 1969. The intransitivity of preferences. *Psychological Review*, 76: 31-48.
- Tversky, A. & Simonson, I. 1993. Context-dependent preferences. *Management Science*, 39: 1179-1189.
- Waite, T.A. 2001a. Background context and decision making in hoarding gray jays. *Behavioral Ecology*, 12: 318-324.
- Waite, T.A. 2001b. Intransitive preferences in hoarding gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 50: 116-121.
- Wedell, D.H. 1991. Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology - Learning Memory and Cognition*, 17: 767-778.

CHAPTER 6

6.1 INTRODUCTION	86
6.2 GENERAL METHODOLOGY	88
6.2.1 EXPERIMENTAL DESIGN	88
6.2.2 SUBJECTS	91
6.2.3 APPARATUS	91
6.2.4 TRAINING	92
6.2.5 EXPERIMENTAL PROTOCOL	92
6.3 ANALYSIS OF DATA	93
6.3.1 COMPARISON OF TREATMENTS	93
6.3.2 WITHIN-TREATMENT ANALYSES	93
6.4 EXPERIMENT 1: DISCRIMINATION OF AMOUNTS	95
6.4.1 METHODS	95
6.4.2 RESULTS AND DISCUSSION	95
6.5 EXPERIMENT 2: DECOY EFFECTS ON CHOICE	96
6.5.1 METHODS	96
6.5.2 RESULTS	97
6.5.2.1 <i>COMPARISON OF TREATMENTS</i>	97
6.5.2.2 <i>WITHIN-TREATMENT ANALYSES</i>	98
6.5.3 DISCUSSION	99
6.6 EXPERIMENT 3: EFFECTS OF INTAKE RATE ON CHOICE	102
6.6.1 METHODS	102
6.6.2 RESULTS AND DISCUSSION	103
6.6.2.1 <i>PREFERENCE SCORES</i>	103
6.6.2.2 <i>STABILITY IN PREFERENCES</i>	104
6.7 EXPERIMENT 4: DECOY EFFECTS WITH INTAKE CONTROLLED	104
6.7.1 METHODS	104
6.7.2 RESULTS AND DISCUSSION	106
6.7.2.1 <i>COMPARISON OF TREATMENTS</i>	106
6.7.2.2 <i>WITHIN-TREATMENT COMPARISON</i>	107
6.7.2.3 <i>DIFFERENCES IN INTAKE RATE</i>	108
6.7.2.4 <i>STABILITY IN PREFERENCES</i>	110
6.8 GENERAL DISCUSSION	112
6.8.1 STATE-DEPENDENT VERSUS CONTEXT-DEPENDENT PREFERENCES	114
6.8.1.1 <i>ENERGETIC STATE AND RATIONALITY IN JAYS</i>	114
6.8.1.2 <i>ENERGETIC STATE AND RATIONALITY IN HUMMINGBIRDS</i>	115
6.8.2 RATE CURRENCIES	117
6.8.3 CONCLUSION	119
6.9 REFERENCES	121

TESTING RATIONALITY IN ANIMAL BEHAVIOR: CONTEXT-DEPENDENT *VERSUS* STATE- DEPENDENT PREFERENCES

6.1 Introduction

During the past decades the study of animal behaviour has greatly benefited from the use of concepts borrowed from economic theory. This was the case, for instance, of the introduction of game theory ideas in the study of animal conflict (Maynard-Smith & Price 1973). Similarly, optimality modelling emerged based on the view of animals as utility maximizers – with utility being replaced by the notion of fitness, and Darwinian evolution playing the role of the omnipresent maximizing agent.

One common consequence of assuming that individuals should behave so as to maximize the expected value of a particular utility function is captured in the concept of rationality. Rationality encapsulates a number of principles that are necessary conditions for the existence of an absolute scale of value upon which humans, and non-humans alike, base their choices. Transitivity, for instance, is a hallmark of rational choice theories. It states that if ‘a’ is preferred to ‘b’, and ‘b’ to ‘c’, then – as the reader can intuitively surmise – ‘a’ should also be preferred over ‘c’. Another frequently invoked principle is that of independence from irrelevant alternatives (or simply IIA), namely the expectation that options should be evaluated independently of the presence or absence of other alternatives.

Breaches of rationality of different sorts are, however, well documented in the literature on human decision-making (Payne et al. 1992, Tversky & Simonson 1993, Gigerenzer et al. 1999). For instance, in violation of the principle of IIA, some authors (Huber et al. 1982, Simonson & Tversky 1992, Heath & Chatterjee 1995) have demonstrated that the addition of a new option into a binary choice set can change the attractiveness of pre-existing alternatives – an unexpected result under the assumption that the value of an option should be based uniquely on its intrinsic properties. To account for these and similar results, alternative models were developed (see e.g. Shafir et al. 1989, Wedell 1991) in which the subjective value of an option can be contingent upon the context of choice, as defined by the configuration of the opportunity set in terms of the number and nature of its options. Some of these models for instance (normally referred to as ‘comparative’), consider that options are evaluated and compared along each of their relevant attributes instead of being assigned an independent value on a single dimension, offering a possible mechanism for the occurrence of context effects whenever the elements of a set differ in more than one attribute. The central concept underlying the approach is that

preferences would be constructed, rather than revealed, during the choice process (Payne et al. 1992).

Violations of rationality by animals have also been recently reported in the literature (Shafir 1994, Hurly & Oseen 1999, Waite 2001b, Waite 2001a, Bateson et al. 2002, Shafir et al. 2002) and have, potentially, major implications for our current understanding of animal behaviour. On the one hand such paradoxical results can be used as a guiding tool for the investigation of the cognitive mechanisms underlying animal choices; on the other, they challenge theories of value maximisation often used by behavioural ecologists, in which the utility a particular resource, as determined by its properties, is often assumed to be a key determinant of its choice probability (McFarland & Houston 1981, Stephens & Krebs 1986). And indeed, most observations of irrational behaviour by animals are taken to imply a major challenge to the optimality approach (Shafir 1994, Hurly & Oseen 1999, Bateson 2002, Bateson et al. 2002, Shafir et al. 2002).

I believe however that this should not be necessarily the case if – before seeking for an alternative framework under which to develop our theoretical understanding of behaviour – one considers the specificities of the circumstances under which animals make decisions. Although the frameworks underlying violations of rationality by both animal and human subjects are to a large extent similar, raising the possibility that the same fundamental mechanisms of choice are responsible for their overt behaviour, direct comparisons of results in both fields is however impaired by critical procedural differences. Perhaps one of the most crucial difference derives from the fact that, while human subjects can be given one-shot decisions and linguistically instructed about the properties of the alternatives, animals must be trained and repeatedly exposed to the contingencies to learn about them. If on the one hand this procedure is necessary to ensure that the properties of the alternatives are correctly judged by animals, on the other repeated contact with different contexts, and thus contingencies, is likely to affect a number of correlate measures of an organism's state (e.g., physiological measures of energetic state, as well as the very ability to learn about the choice problem). This might be a key difference in light of the expectation that the consequences, and thus choice probability, of available courses of action are likely to be contingent upon an organism state, as long acknowledged by normative modelling of animal behaviour (McFarland & Houston 1981, Cuthill & Houston 1997, Houston & McNamara 1999). Some of the apparent inconsistencies in choice behaviour could be thus simply a

natural consequence of state differences resulting from the side effects of past experience with different outcomes.

Here I conducted a series of experiments to investigate the putative presence of context-dependence in animal choices, and the possibility that economically irrational behaviour can emerge as a consequence of state-dependent decision-making. Specifically, I compare the foraging preferences of European starlings (*Sturnus vulgaris*) when exposed to distinct contexts, designed so as to maximize the potential influence of decoys in the manner foreseen by comparative models of choice. The first two experiments provide the necessary background for the critical experiments that are conducted later, directly testing (experiment 2) the influence of context on preferences in the absence of controls for energetic state. The last two experiments investigate preferences for two focal options in the absence of contextual differences but presence of state differences (experiment 3), and in the presence of contextual differences but absence of state differences (experiment 4). The results are then used to reinterpret some of the findings of irrational behaviour previously reported in the literature.

6.2 General Methodology

6.2.1 Experimental Design

The basic design employed in the series of experiments reported in this study can be described as follows. Two treatments were employed where European starlings (*Sturnus vulgaris*) were presented with foraging options described by two attributes: amount of food (in units of pellets) and the interval between the action chosen by the subject and its outcome (hereafter 'delay'). In each treatment the starlings experienced three options: two target options (the focal choice pair) and a decoy, namely, an inferior alternative designed to affect preference between the two targets. While the focal choice pair was always the same, the decoy was specific for each treatment. Figure 1 details the value of the options in each attribute. T_A and T_D represent the options Target Amount and Target Delay, respectively. Target Amount was programmed to offer an amount of food larger than that offered by Target Delay, whereas the latter offered a shorter delay to reward. D_A and D_D represent the decoys 'Amount' and 'Delay', respectively, offered to the birds experiencing one treatment or another. While D_A offered a larger amount of food (the same as T_A), D_D offered the shortest delay (the same as T_D).




Figure 1. Amount and delay to food (right side of the graph) corresponding to the options used in the experiments. Each group of subjects experienced three options: T_A and T_D (the focal choice pair) and a decoy (either D_A or D_D). The target alternatives T_A and T_D offer the same short-term rate of food intake (solid lines) of 0.5 units/s, whereas the decoys D_A and D_D offer the same short-term rate of 0.25 units/s. The dashed lines (interrupted at 30s for space economy) indicate the long-term rate of food intake, considering the time interval between consecutive feeding opportunities (assuming an interval of 60s).

Figure 1 shows that the direction of preference between pairs of options should depend on the time interval considered by the birds. Classical optimal foraging models assume that animals should maximize the long-term rate of energy intake, defined as the amount of energy gained over the total time taken to acquire it, including intervals between feeding events and the delay between the decision to pursue an option and its final consumption. In fact, from an evolutionary perspective the expectation is that subjects should consider a time horizon longer than that associated with a single foraging bout, since the energy gained is likely to be invested only after a large number of opportunities is accomplished (Stephens & Krebs 1986, Stephens & Anderson 2002). Accordingly, maximization of long-term rate of intake should lead to a higher level of preference for the Target Amount (T_A) over the remaining alternatives, and similarly a higher preference for D_A over both T_D and D_D , and for T_D over D_D (slope of the dashed lines in fig.1). Although normatively compelling, a large body of empirical evidence is however at odds with this expectation. Many studies have shown that animals, including humans, tend to use more shortsighted rules that fail to account for all the time intervals involved in the acquisition and consumption of a prey, or reward (Green et al. 1981, Mazur 1987, Mazur 1988, Bateson

& Kacelnik 1996, Stephens & Anderson 2002). In general, the delay between a decision being made and the food being consumed is crucial in determining preferences, whereas the intervals between feeding opportunities tend to have little effect on behaviour. Considering this body of evidence, I therefore programmed the options forming the focal choice pair to offer the same short-term rate of food gain per unit of time (solid lines in fig.1). The idea was to make any potential influence of the decoy more apparent by presenting the subjects with a situation where the target options were equivalent, so that even a slight effect of the decoy could lead to preference for one of the targets. As an additional procedure to facilitate the detection of possible context effects, the decoy was planned to have opposite effects in each treatment (see below), thus doubling any potential effect size.

Under the framework proposed by comparative models of choice, there would be two possible ways whereby D_A could increase preference for T_A . Firstly, it has been previously shown (Parducci 1965, Janiszewski & Lichtenstein 1999) that a same difference in a physical attribute can have a greater effect in a narrow range than in a wide range of values. Thus, by increasing the range of values in the dimension where T_A is worse than T_D (delay to food), D_A could make this unfavourable difference subjectively less important. Secondly, D_A could favour T_A over T_D through its asymmetric relationship of dominance with these target options (where dominance is used as a synonym for superiority in terms of attribute values). In the design, D_A is completely dominated by T_A (as the latter offers the same amount of food after a shorter delay), but not by T_D (although the latter offers a shorter delay, D_A offers a larger amount of food). Thus, in the presence of D_A , the target T_A could be more attractive for being the only option that dominates it completely (see Huber et al. 1982, Tversky & Simonson 1993 and references therein for various examples of the influence of ‘asymmetrically dominated decoys’). The same mechanisms would be expected to make D_D to enhance the attractiveness of T_D over T_A .

An interesting alternative to the above model is however given by considering another approach, known as the ‘similarity hypothesis’ (Tversky & Russo 1969). The hypothesis states that the strength of preference for one option over its alternative may differ between two choice sets owing to the use of choice rules that group similar alternatives together. For instance, a decoy introduced into a binary choice set would take its share of choices mainly from the option most similar to it. As a consequence, the relative preference for the latter would be weakened. Here, the similarity hypothesis predicts results that are opposite in direction to those foreseen by comparative models. By stating that a new option

included in the choice set should hurt more similar options, it anticipates that when included in the choice set D_A should take its share of choices mainly from T_A , as these options share the same parameter values in the attribute 'Amount', and the delay offered by D_A is closer to that offered by T_A than to that offered by T_D . Accordingly, the presence of D_D should make the subjects' preference for T_D over T_A to disproportionately decrease.

6.2.2 Subjects

The subjects were forty wild-caught European starlings (*Sturnus vulgaris*), captured in walk-in traps in Sandwich and Oxford, UK, in February 2001 and July 2002 respectively under English Nature licences (numbers 20010082 and 2020068, respectively). After capture the birds were kept in two outdoor aviaries (3 m high \times 3.2 m wide \times 4 m long). While in the aviary the starlings were fed on ad libitum food, a mixture of turkey crumbs, Orlux© pellets and mealworms (*Tenebrio* sp.). Drinking and bathing water was always available and replaced daily. Ten days before the beginning of the training sessions the birds were transferred from the outdoor aviary to the laboratory and housed in individual cages (120 cm \times 60 cm \times 50 cm) that served both as home cages and experimental chambers. Lights in the experimental rooms were gradually switched on at 0500 and off at 1900 hours, and temperature ranged from 12°C to 16°C. Subjects were visually but not acoustically isolated. During this period of adaptation all birds had free access to water and the same ad libitum food as received in the aviaries. Fresh drinking water was always available in the cages and bathing trays were provided every other day. Each group spent a maximum of five weeks in the laboratory, and then was transferred back to the outdoor aviary. All subjects were experimentally naïve. The experiments were initiated in April 2002 and completed in October 2002. There was no mortality. Twenty-eight subjects were released back into the wild (Sandwich, UK) in July 2002. The twelve birds captured in Oxford were kept for future research.

6.2.3 Apparatus

All experimental cages had a panel with a centrally mounted food hopper and two response keys positioned 6 cm from either side of the hopper. The cages were also supplied with two perches, 85 cm apart from each other. The panels used in experiment 1 had a third response key on its centre, 3 cm above the hopper, which was used as an attention key. An Acorn Risk PC 600 computer running Arachnid experimental language (Paul Fray Ltd.) controlled the stimulus events and contingencies, and collected the data. During the experimental

sessions food rewards were units of Orlux© pellets, crushed and sieved to an even size ($0.025 \pm 0.005\text{g}$). Automatic pellet dispensers (Campden Instruments®) situated above the cages delivered the pellets at a rate of 1 unit/s.

6.2.4 Training

Subjects were first trained to retrieve food from the hopper. Responses were then induced by a standard auto-shaping procedure in which the illumination of one of the pecking keys (side determined randomly) for 8s was followed by the delivery of 2 pellets and an intertrial interval (ITI) of 60s. The birds received 2 daily sessions in this condition (100 trials per session), one at 0700 and another at 1300 hours, after which they were given ad libitum food. Once all birds pecked reliably the keys food delivery was made conditional at one key pecking. In these sessions, if the birds did not respond during the 8-s interval the key light went off and the ITI started. The training program was interrupted when all birds were pecking in at least 80% of the trials.

6.2.5 Experimental Protocol

In experiment 1 I used six birds. In experiments 2, 3 and 4 I used twelve, ten and twelve birds, respectively. Each of the last three experiments was conducted in two consecutive phases, with half of the birds being used in each phase. All treatments were balanced across phases. Colours and symbols were used as discriminative stimuli, and the pairing of colours and symbols with options was balanced across birds.

In all experiments I employed a discrete trials procedure with two types of trials: no-choice trials and choice trials. No-choice trials were intended to provide the birds with information about each alternative. They started with one of the pecking keys blinking (0.7 s on, 0.3 s off). The first peck caused the light to stay on for the programmed delay. The first peck after the delay had elapsed triggered the delivery of the amount of food corresponding to the option experienced, and an intertrial interval (ITI) started. Subjects also experienced choice trials, in which they were given a choice between two of three options. Choice trials began with all keys (two or three, depending on whether the choice trial was binary or trinary) simultaneously blinking. The first peck on any of them caused the chosen key to turn continuously on for its respective delay and the others to turn off. The first peck after the delay had elapsed triggered the delivery of the corresponding reward and the beginning of the ITI. The order (and sides) in which the options were presented was always randomly

determined. Subjects were given a maximum number of twenty experimental sessions. However, a treatment could be interrupted earlier if the subject reached the following stopping criteria (provided it had had a minimum of ten sessions): the regression coefficient of the choice proportions (for the focal choice pair in binary choices) of six consecutive sessions (against session number) was not significantly different from zero and the standard deviation of the choice proportions of these six sessions did not exceed 0.10. Data from the last six sessions was used for analysis. Details of each experiment are given in their corresponding sections.

6.3 Analysis Of Data

6.3.1 Comparison of Treatments

Differences in the subjects' preferences for Target Amount over Target Delay between treatments were analysed by directly comparing choice proportions for each target in binary choices between the two treatments. Additionally, I analysed whether preferences for the target options differed between treatments by analysing their relative preferences for the focal choice pair as expressed in trinary choices (where both target options and the decoy specific to the treatment were simultaneously available). To that end, I transformed choice proportions in the trinary situation by considering only two options at a time, so that the percentage of choices for the pair added up to 100%. Thus, preference for T_A over T_D in this context was calculated as:

$$p_T(T_A, T_D) = \frac{p_T(T_A)}{p_T(T_A) + p_T(T_D)} \times 100 \quad [1]$$

where $p_T(T_A)$ and $p_T(T_D)$ represent the proportion of choices for Target Amount and Target Delay, respectively, in the trinary context.

6.3.2 Within-treatment analyses

I also investigated whether the starlings' preferences for the focal choice pair differed, within the same treatment, between the binary and trinary choice contexts. I firstly tested the starlings' relative preferences for the pair, as given by equation 1. The expectation that relative preferences should be the same in both contexts derives directly from the notion that the preference between any two alternatives should be entirely independent of the context in which they occur. This notion is best captured in the choice axiom of Luce (Luce 1959),

which states that the ratio of the probabilities of any two alternatives should be entirely independent from the choice set. Hence, the inclusion of a new option into a choice set should take from pre-existing options in proportion to their original shares.

Finally, I compared the starlings' absolute preferences for each of the targets between the binary and trinary contexts to test for violations of regularity. The principle of regularity is one of the weakest forms of the notion of independence from irrelevant alternatives (Luce & Suppes 1965), and asserts that the absolute proportion of choices for an option cannot be increased when a new option is included in the choice set. Formally, regularity is satisfied if the probability of taking an item A from a subset $B=\{A,B\}$ that is part of $T=\{A,B,C\}$ is such that, for all C,

$$P_B(A) \geq p_T(A) \quad [2]$$

where $p_B(A)$ and $p_T(A)$ represent the probabilities of choosing A from S and T, respectively.

6.4 Experiment 1: Discrimination of Amounts

6.4.1 Methods

To ensure the birds were able to distinguish between the amounts planned in the experimental design previously described (there is strong evidence that they are able to discriminate between the delays used; Brunner et al. 1996) I initially conducted a discrimination experiment using six starlings, split into two groups of three birds each. Group 1 was presented with choices between options offering 1 and 2 units of food respectively, and Group 2 was presented with a choice between options offering 2 and 5 units of food, respectively. A fixed ITI of 60s was employed. Subjects experienced two sessions per day, one at 0800 and another at 1300 hours. Each session consisted of 84 trials, divided into twenty-one blocks of four trials each: two no-choice trials (each option presented once) followed by two choice trials. Subjects were allowed to feed ad libitum with turkey crumbs for 3 hours after the experimental sessions, after which they were supplemented with 10 mealworms and then food was removed.

6.4.2 Results and Discussion

All birds were able to discriminate between the programmed amounts. Figure 2 shows the proportion of choices of each bird to the option offering the largest amount of food. For all birds, preference levels were significantly higher than those expected by chance (binomial tests: $P < 0.01$), thus validating the use of the parameter values planned in the experimental design. It is however interesting to note that even in this simple choice problem, preference is not exclusive. Specifically, the variability in choice proportions is significantly higher ($F_{1,4} = 10.3$, $P < 0.05$) for those subjects faced with a choice between 2 and 5 units of food, even though the difference between the options in this latter situation is higher than that associated with a choice between 1 and 2 units of food (white bars in the graph). One possibility is that, by consuming a much larger amount of food during the experimental sessions due to the larger size of the rewards offered (at least 2.4g more per session), subjects in Group 2 were in a lower state of need than those in Group 1. Erroneous choices (i.e., choices for the option yielding the smallest amount of food) may have been thus more likely in the former group due to the fact that the acquisition of food was less of an essential necessity for these individuals. In fact, similar results on the foraging behaviour of birds support this expectation. For instance, gray jays (*Perisoreus canadensis*) were presented with choices between one and two units of food after different experimentally imposed

delays to access to these rewards (Waite & Field 2000). Consistently with the idea that errors should be more common the lower their cost (McNamara & Houston 1987), the jays failed to take the best option more frequently when subject to higher rates of energetic return (i.e., shorter delays). Although the authors discuss their results in terms of the cognitive constraints (such as limited attention) on the information-processing abilities of the jays, it is possible that differences in the subjects' state of hunger resulting from the manipulation of the intervals to food also contributed to the observed effects.




Figure 2. Proportion of choices (\pm S.E.) for the option offering the largest amount of food. Choice proportions are significantly different from random for all birds (binomial test, $P < 0.01$) Birds 1, 2 and 3 (white bars) were presented with choices between one and two units of food, and birds 4, 5 and 6 (black bars) with choices between two and five units of food.

6.5 Experiment 2: Decoy Effects on Choice

6.5.1 Methods

Twelve birds were randomly assigned to two treatment groups. Both groups were presented with target options T_A and T_D . However, while Group 1 experienced the decoy D_A , Group 2 experienced D_D . Each option was associated with one of three colours illuminating the pecking keys: yellow, red or green. A fixed ITI of 80s was employed.

Subjects experienced two sessions per day, at 0700 and at 1300 hours. Each session could last up to 3.5 hours and consisted of 63 trials, grouped into seven blocks of nine trials

each. To ensure experience of all options each block started with three no-choice trials (each option experienced once in a random order), followed by six choice trials (random order of presentation): two trinary choice trials (all options simultaneously available), two binary choice trials with the focal choice pair and two binary choice trials where each of the targets was presented once with the decoy. After the first session (at 1030 hours) subjects received 2.5 g of food (turkey crumbs). After the end of the second session subjects were allowed to feed ad libitum with turkey crumbs for 150 minutes, after which they were supplemented with 10 mealworms and then food was removed.

6.5.2 Results

6.5.2.1 Comparison of Treatments

Figure 3 shows how the relative preference for the target options was affected by experience with the decoys in the two experimental groups. In the binary choice trials, where only the focal choice pair was available, preference for T_A was significantly higher in the group experiencing the decoy D_D ($F_{1, 10} = 11.6$, $P < 0.01$). The same pattern was observed in the trinary trials, where the relative preference for T_A relative to T_D (as calculated by equation 1) was higher for the group experiencing D_D , although in this latter case differences were not significant ($F_{1, 10} = 3.57$, $P = 0.09$). Both target options were significantly preferred over the decoys (when presented with the latter in binary choices) in both experimental groups (one sample t-tests against a random expectation of 0.5: $P < 0.05$).




Figure 3. Mean proportion of choices (\pm S.E.) for the Target Amount (T_A) relative to Target Delay (T_D) in the two experimental groups. The bars for the trinary choice trials represent the relative preference for Target Amount over Target Delay in the presence of the decoy (as calculated by equation 1). † $P < 0.1$, ** $P < 0.01$.

6.5.2.2 Within-Treatment Analyses

Relative preferences for T_A over T_D were lower in the presence of the decoy D_D (in trinary choice trials) than when only the two targets were available in binary choices (paired-t tests: $t_5 = 3.4$, $P < 0.05$; comparison of white bars in fig.3). No differences in relative preferences between the binary and trinary contexts were, however, observed for the group where D_A was presented (paired-t test: $t_5=0.15$, $P = \text{NS}$; comparison of grey bars in fig.3).

I additionally compared the birds' absolute preferences for each option between the two contexts (binary and trinary). Figure 4 illustrates the results. Overall, there was no significant increase in the absolute proportion of choices for any of the options in the trinary context and, therefore, no violation of regularity. In the group experiencing the decoy D_D (chosen with an average proportion of 0.08 ± 0.03 , black bar in fig.4a), absolute choice proportions for T_D were not different between the binary and trinary contexts (paired-t test: $t_5 = -1.17$, $P = \text{NS}$; comparison of white bars in fig.4a), whereas choice proportions for T_A were significantly lower in the latter situation ($t_5 = 5.19$, $P < 0.01$; comparison of grey bars in fig.4a). In the group where D_A was presented, conversely, the absolute preference for T_D was marginally lower in the trinary context ($t_5 = 2.30$, $P = 0.07$; white bars in fig.4b), but preferences were not different between contexts in the case of T_A ($t_5 = 1.70$, $P = \text{NS}$; comparison of grey bars in fig.4b). In this group, the decoy took an average proportion of 0.22 ± 0.03 of the starlings' choices in the trinary context.

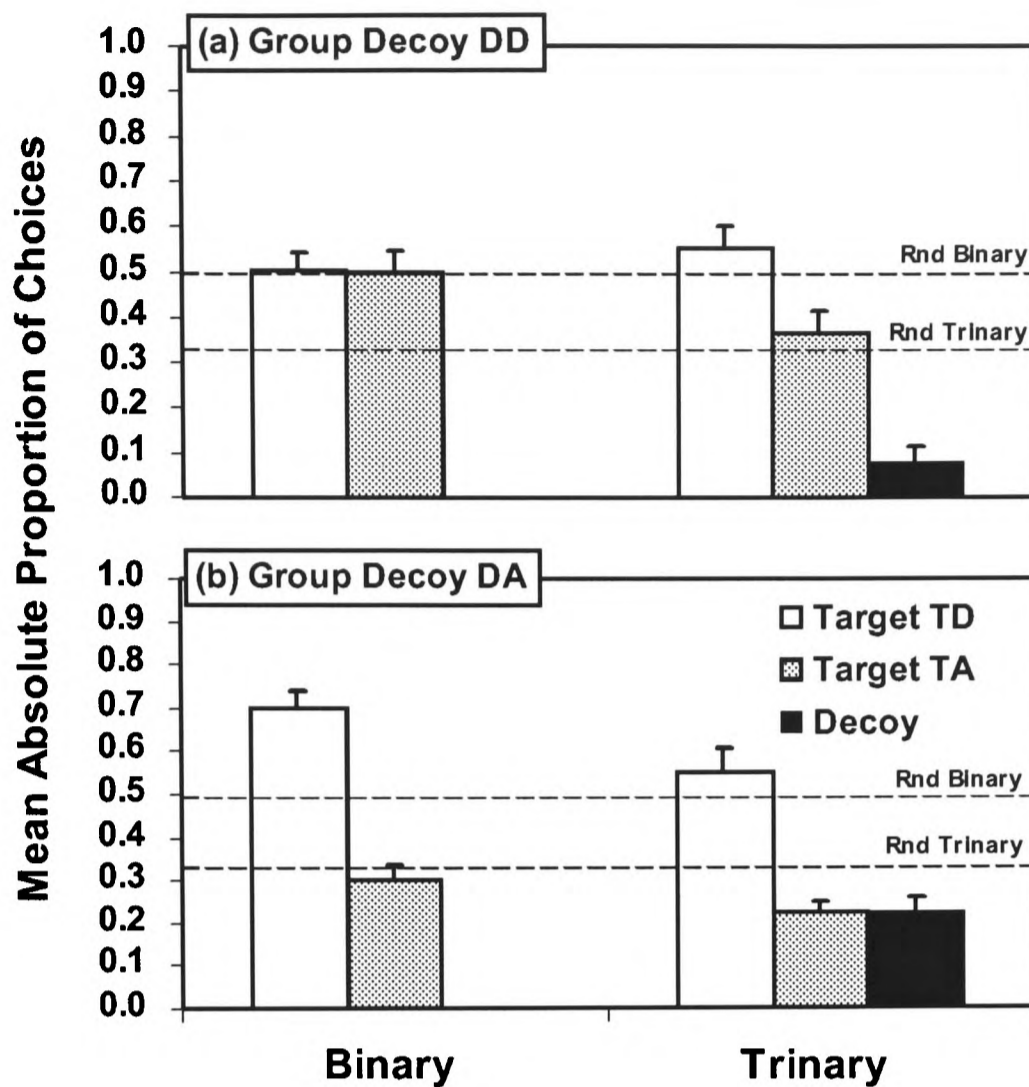


Figure 4. Mean absolute proportion of choices (\pm S.E.) for each option in the binary and trinary contexts in both experimental groups. The dashed lines indicate the preference levels corresponding to the random expectation in the binary (0.5) and trinary (0.33) contexts.

6.5.3 Discussion

Preferences for the target options were markedly different between the two experimental groups. While the group that experienced decoy D_D was on average indifferent between the two targets, those starlings presented with D_A were biased towards T_D . Although the differences were in the same direction as that predicted by the similarity hypothesis the results, however, do not provide unequivocal support for it. The hypothesis states that the inclusion of a new option into a choice set should hurt mainly the existing alternative most similar to it due to the assortment of some elements into different categories. As a consequence, preferences for a particular choice pair should vary between differently sized choice sets, or between sets with the same size but differing in at least one element. The similarity hypothesis could therefore explain changes in preference between the two groups as expressed in the trinary context, but not the marked differences obtained in the comparison of binary choices for the focal pair between the groups. Additionally, for one of the groups (the one experiencing the decoy D_D) the within-treatment analysis showed that

relative preferences were different between the binary and trinary contexts, but the deviations were in the direction predicted by the comparative model.

One possible account for the differences in preference between the two treatment groups is given by considering the variation in the real intake rate during a session caused by the repeated exposure of the subjects to distinct decoys. Both decoys offered the same short-term intake rate. However, the long-term rate of gain offered by D_A (including intervals between feeding events) was higher than that offered by D_D . Thus, in the end of a session where the decoy was, for instance, chosen on about 10% of the opportunities it was available (and the subject was indifferent between T_A and T_D), the difference between the total intake of a subject in Group D_A during a session was at least 6% higher than that of a subject in Group D_D . The higher the proportion of choices for the decoy, the more pronounced this difference was. Figure 5 shows the average intake rate of subjects in each group during the last six experimental sessions. Total intake rate (units of food per minute) is significantly different between the groups ($F_{1,10} = 15.524$, $P < 0.01$), approximately 13% higher in Group D_A . Notice however that the hypothesised effect is fundamentally different from the classical expectation of optimal foraging theory that long-term rate of gain should be the currency maximized in choices. Had this been the case, the birds should have ranked the options in an order such that $T_A > D_A > T_D > D_D$ (see fig.1), which was not observed. Yet, the indirect effect that repeated contact with energetically distinct decoys had on deprivation level may have affected preferences.




Figure 5. Total intake rate (total units of food received over total delay plus intertrial intervals experienced) in the last six experimental sessions in each group. ** $P < 0.01$.

Given the likely differences in deprivation level experienced by subjects in each group, coupled to the usual finding that learning rates tend to be slower under richer schedules (Tarpy & Mayer 1978), I decided to analyse whether both groups of birds had achieved a stable level of preferences for the target pair, or whether the level of stability reached was different between the groups. To that end, I verified the number of subjects for which the slope of the regression of choice proportions against session number was significantly different from zero (i.e., those subjects who did not reach the asymptotic behaviour), subsequently testing whether the slopes in one group were significantly different from those in the other. For all subjects, we could not reject the null hypothesis that the slope was not significantly different from zero. Moreover, these slopes were not significantly different between the two groups ($F_{1,10}=0.19$, $P = \text{NS}$). I then tested whether the variability in choice proportions (as measured by the standard deviation) across the last six experimental sessions (those used for analysis) was different between the groups, but could not detect any difference ($F_{1,10}=0.14$, $P = \text{NS}$). These results thus indicate that the subjects in both groups had already reached the asymptotic level of performance by the end of the experiments, and that the variability in their responses did not differ between groups.

To test the hypothesis that the differences between groups were induced by the observed variation in intake rate (and consequently, energetic state) experienced by the subjects I designed and conducted two further experiments. Both designs were based on the same framework as that described in figure 1. In one of them (experiment 3) I tested the foraging preferences of starlings for the focal choice pair in the absence of experience with decoys, with subjects being kept under different food consumption regimes. The objective of this experiment was to determine the starlings' choices for the target option under distinct energetic conditions, in the absence of potential contextual effects. In the second additional experimental (experiment 4) I tested the influence of the decoys D_A and D_D on the preference for the target options in two situations: when intake rate was and was not controlled. This last experiment was aimed at enabling the distinction of differences in preferences caused by the cognitive influence of the decoys from those caused by variation in intake rate.

6.6 Experiment 3: Effects of Intake Rate on Choice

6.6.1 Methods

I used a within-subjects design in which I measured the preference of ten birds for the focal choice pair under the following conditions: (1) intake rate similar to that obtained in the presence of the decoy D_A and (2) intake rate similar to that obtained in the presence of D_D . No decoys were however presented to the subjects. To achieve these two conditions I programmed the feeder dispensers to automatically deliver the reward corresponding to the appropriate decoy once per experimental block (details of the timings and spacing between deliveries along the sessions are given below). No action was needed from the part of the subjects to receive the reward, nor was any discriminative stimulus associated with it. The order in which the conditions were experienced was balanced across birds, with half of the birds experiencing condition \check{D}_A first and half experiencing condition \check{D}_D first. The within-subjects design was preferred owing to the usual high level of variability in the subjects' energetic requirements, which could dwarf the potential effects of the manipulation of intake rate. Each of the target options was associated with one of two colours illuminating the pecking keys: red or green. Colours were changed between treatments (to blue and yellow) so as to force the birds to learn again the association between the colours and their respective contingencies. The balancing of colours and treatments was however conducted in a way so that if generalizations of colours occurred (from red to yellow and green to blue, and vice-versa for the birds that experienced the reverse order), it should produce changes in preference between treatments in the opposite direction as that observed in experiment 2. All subjects were given one resting day with ad libitum food between treatments.

Subjects experienced three sessions per day, at 0600, 1000 and 1400 hours. Each session could last up to two and a half hours and consisted of 36 trials, grouped into twelve blocks of three trials each. Each block started with two no-choice trials where T_A and T_D were presented in a random order. The no-choice trials were then followed by a fixed ITI of 60s and the delay of the decoy corresponding to the condition experienced, after which the appropriate reward was delivered in the hopper, followed by the ITI. The last trial of the block was a choice trial where the focal choice pair was presented. With this protocol subjects experienced the reward (and delay) corresponding to the appropriate decoy on 25% of the block. In terms of comparison with the experimental blocks designed in experiment 2, this frequency would correspond to the decoy being experienced once in the forced trials

and once in one of the four choice trials where it was available in each block of experiment 2. After the daily sessions subjects were allowed to feed ad libitum with turkey crumbs for 150 minutes, after which they were supplemented with 10 mealworms and then food was removed.

6.6.2 Results and Discussion

6.6.2.1 Preference Scores

Preferences differed in the same direction as in Experiment 2 (fig. 6). The proportion of choices for T_A was significantly higher under the treatment simulating the presence of the decoy D_D (Repeated-measures ANOVA with treatment and order as within and between-subjects factors: $F_{1, 8}=13.5$, 1-tailed $P < 0.01$). The effect of order and the interaction between order and treatment were not statistically significant ($F_{1, 9}=0.03$, $P = \text{NS}$ and $F_{1,8}=3.88$, $P = \text{NS}$, respectively).

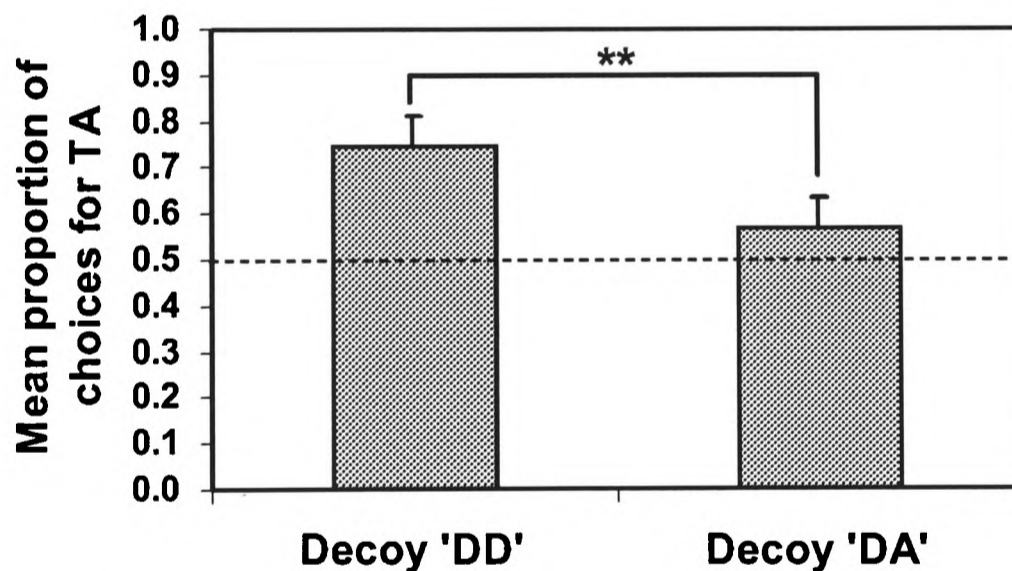


Figure 6. Mean proportion of choices (\pm S.E.) for the Target Amount (T_A) in the treatments simulating the presence of D_D and D_A , respectively. ****** $P < 0.01$.

The average intake rate experienced by the starlings during the last six sessions of the experiment also differed between treatments in the same direction as that observed in experiment 2 (fig.7). As programmed, intake rate was higher in the treatment simulating the presence of D_A (Repeated-measures ANOVA: $F_{1, 9}=17.5$, $P < 0.001$). Together, these results support the hypothesis that the differences in preference level between groups observed in Experiment 2 were also possibly mediated by state differences, with hungrier animals

showing a stronger bias towards options with a higher total intake rate (T_A) than more satiated animals.

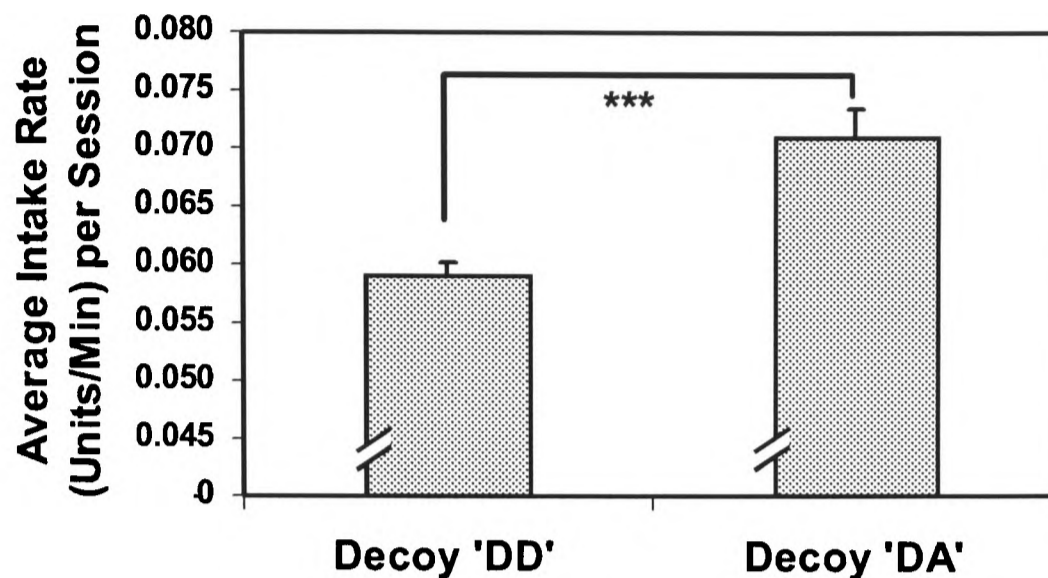


Figure 7. Total intake rate (total units of food received over total delay plus intertrial intervals experienced) in the last six experimental sessions in each treatment. *** $P < 0.001$.

6.6.2.2 Stability in Preferences

Following the same procedure as in experiment 2, I also investigated whether the level of stability in preferences for the focal choice pair was different between the two intake regimes. For only one bird (treatment D_A) was the slope of the regression marginally different from zero ($P = 0.07$). The group of slopes, however, was not different between treatments (Repeated-measures ANOVA, Greenhouse-Geisser correction for sphericity applied: $F_{1,9}=0.98$, $P = \text{NS}$). Similarly, the variability in choice proportions (as measured by the standard deviation) across the final sessions tended to be slightly higher for those subjects experiencing treatment D_A , but the differences were not significant ($F_{1,9}=2.04$, $P = \text{NS}$).

6.7 Experiment 4: Decoy Effects on Choice with Intake Controlled

6.7.1 Methods

Twelve birds were randomly assigned to two groups ('Fed' and 'Unfed', F and U, respectively, see below) of six birds each. Each group experienced two treatments (table 1), where preference for the focal choice pair was tested in the presence of the decoy D_A

(treatment D_A) and of the decoy D_D (treatment D_D). The order with which treatments were experienced was balanced across subjects.

Table 1. Distribution of subjects across groups and treatments. Each group consisted of six individuals. Each individual experienced both treatments (D_A and D_D), with the order in which treatments were experienced (between brackets) being balanced across subjects.

		Group	
		Fed	Unfed
Treatment	Decoy D_A	$3(1^{st}) + 3(2^{nd})$	$3(1^{st}) + 3(2^{nd})$
	Decoy D_D	$3(2^{nd}) + 3(1^{st})$	$3(2^{nd}) + 3(1^{st})$

In Group F intake rate was however controlled in a way such that subjects in this group experienced the same intake rate per session in both treatments (D_A and D_D). In Group U the differences in intake rate caused by the presence of different decoys in the two treatments were not controlled (this group thus reproduces the design of experiment 2). The within-subjects comparison of treatments was preferred owing to the potential variability in the energetic requirements of subjects, which would prevent the comparison of preferences for group F with intake rate (and more importantly the energetic state of subjects) being the same in both treatments. In the first treatment each option was associated with one of three colours illuminating the pecking keys. In the second treatment the colours were changed to reduce the carry over effect that could result from the development of preferences for a particular colour. I used green, white, orange, blue, red and pink. The pairing of colours with options was planned in a way that, if generalizations of colour occurred between treatments, they would produce preferences in the opposite direction as those predicted by the energetic effect. A fixed ITI of 60s was employed.

Subjects experienced three sessions per day, starting at 0500, 0900 and 1400 hours. Each session could last up to three and a half hours, and consisted of 63 trials, grouped into seven blocks of nine trials each. The blocks had the same structure as those in experiment 2: three no-choice trials followed by six choice trials presented in a random order: two trinary choice trials, two binary choice trials offering a choice between the target options and two binary choice trials where each of the targets was presented once with the decoy. To control intake rate between treatments, accounting for differences generated by the potentially distinct elections made by subjects, in one of the groups I adopted the following procedure: I calculated the maximum amount of food and delay per block for treatment D_A (the treatment

offering the potentially highest cumulative delay and amount per block) and in both treatments gave the subjects this amount and delay in each block (given the possibility that subjects could experience such maximum values as the outcome of their own choices, I would not be able to equalize intake per block if an amount or delay smaller than the maximum had been chosen). This was achieved by continuously computing the maximum amount and delay obtained by each subject during no-choice and choice trials and then giving subjects the difference between what they obtained and this maximum value. The feeder dispensers were programmed to automatically deliver the compensating amount twice per experimental block: in the middle of the block (after four trials) and in the end of it. These rewards were delivered after the ITI, added to the compensating delay. No action was needed from the part of the subjects to receive these rewards, nor any discriminative stimulus was associated with it. The deliveries in the middle of the block were always followed by an interval of 5 min to prevent satiation. Blocks were separated by 10 min intervals. After the daily sessions subjects were allowed to feed ad libitum with turkey crumbs for 60 min after which they were supplemented with 10 mealworms and then deprived.

6.7.2 Results and Discussion

6.7.2.1 Comparison of Treatments

Figure 8 shows, for each experimental group, whether and how preference for the target options was affected by experience with each decoy. In the group where intake rate was not controlled (Group 'U', fig.8a), preference in binary choices (where only the focal choice pair was available) for the Target Amount was higher in the treatment where the decoy D_D was offered (Repeated-measures ANOVA with treatment and order as within and between-subjects factors: $F_{1,4} = 5.1$, $P = 0.07$; neither order nor the interaction between order and treatment were significant: $F_{1,4} = 0.27$, $P = \text{NS}$ and $F_{1,4} = 1.03$, $P = \text{NS}$, respectively). Although the results are significant at $\alpha=0.07$, notice however that the direction of the results was predicted a priori, so employing a one-tailed test would double the obtained level of significance. The same pattern was observed in the trinary trials, where the relative preference for the Target Amount was higher for the group experiencing D_D , although in this latter case differences were not significant (effects of treatment, order and of the interaction between them, respectively: $F_{1,4} = 2.7$, $P = \text{NS}$, $F_{1,4} = 0.35$, $P = \text{NS}$ and $F_{1,10} = 0.17$, $P = \text{NS}$). Preference levels for the Target Amount therefore differed between treatments in the same direction and to the same extent as that seen in experiments 2 and 3.

For the subjects that had intake rate controlled during sessions (fig.8b), however, no differences between treatments were detected, either in the binary (Repeated-measures ANOVA with treatment and order as within and between-subjects factors: all $P > 0.5$) or trinary contexts (all $P > 0.5$). These results therefore strongly suggest that the differences in preference observed were the by-product of differences in intake rate caused by the presence of distinct decoys in each treatment.




Figure 8. Mean proportion of choices (\pm S.E.) for the Target Amount (T_A) relative to Target Delay (T_D) for the two experimental groups in the two treatment situations (presence of Decoys D_D and D_A). The bars for the trinary choice trials represent the relative preference for Target Amount over Target Delay in the presence of the decoy, as calculated by equation 1. $\dagger P < 0.1$.

6.7.2.2 Within-Treatment Comparison

In both groups (F and U) there were no differences in relative choice proportions between the trinary and binary contexts either in treatment D_A or treatment D_D (paired-t tests: $-0.07 < t_5 < 2.1$, $P = \text{NS}$ in all cases; see fig.8).

To test for violations of regularity, I also compared the birds' absolute preferences for each option between the two contexts (binary and trinary) for each treatment and group. Figure 9 illustrates the results. Overall, there was no significant increase in the absolute proportion of choices for any of the options in the trinary context and, therefore, no breach of regularity (paired-t tests: $-2.1 < t_5 < 0.65$, $P = \text{NS}$ in all cases). In both groups (F and U) each of the two target options was significantly preferred over the decoy (when presented with the latter in binary choices) in both treatments (one sample t-tests against a random expectation of 0.5: $P < 0.05$ in all cases).

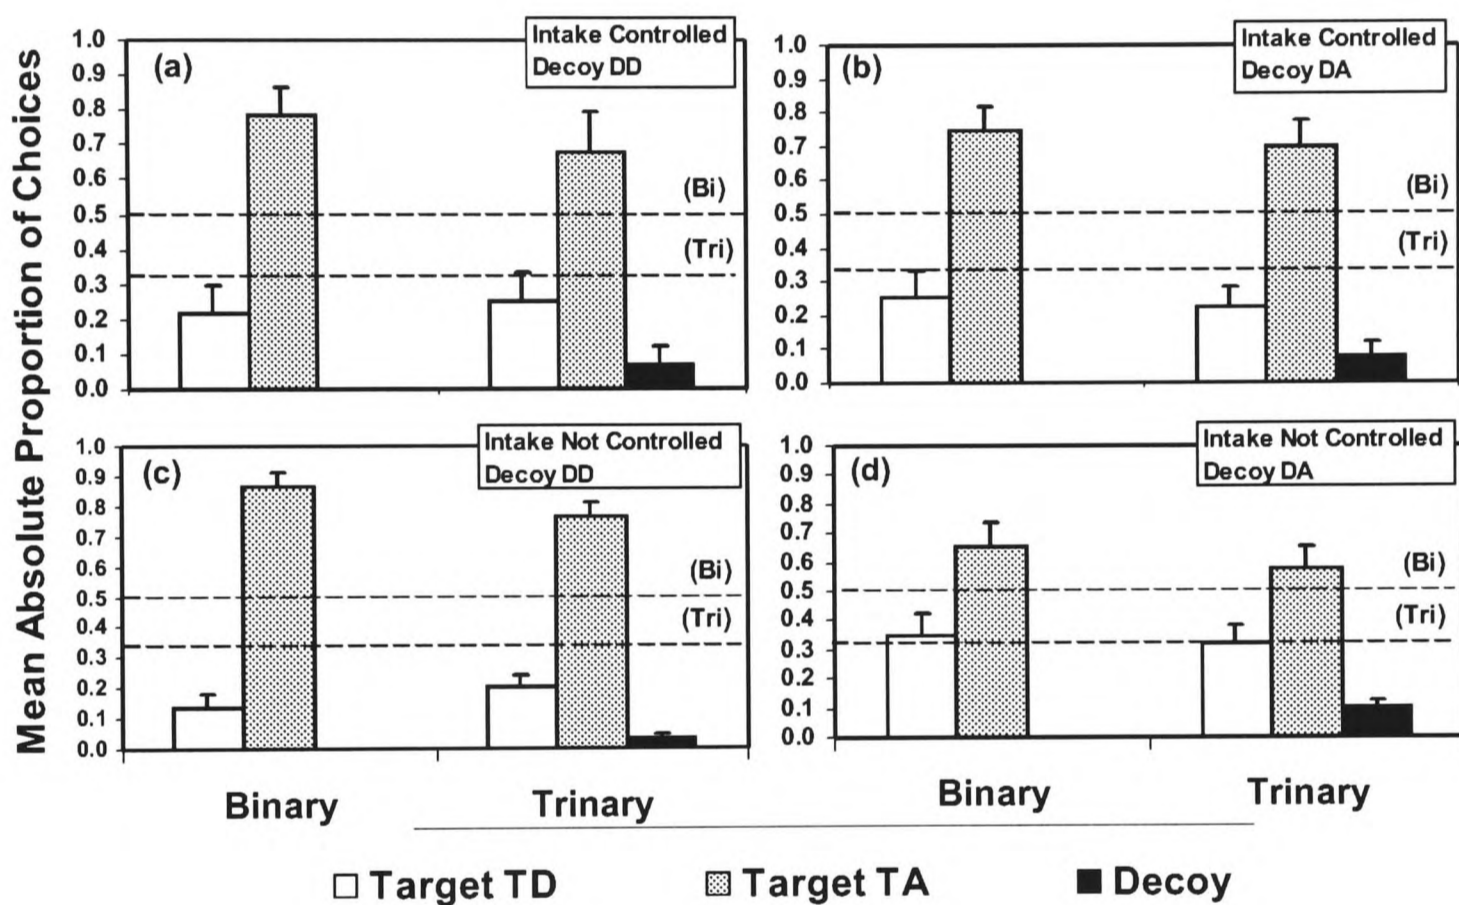


Figure 9. Mean absolute proportion of choices (\pm S.E.) shown by each group for each option in the binary and trinary contexts in both treatment situations. The decoy took an average proportion of (a) 0.07 ± 0.05 , (b) 0.07 ± 0.04 , (c) 0.03 ± 0.01 and (d) 0.10 ± 0.02 choices in the groups where intake rate was controlled (treatments D_A and D_D) and not controlled (treatments D_A and D_D), respectively. The dashed lines indicate the random expectation in the binary (Bi) and Trinary (Tri) contexts.

6.7.2.3 Differences in Intake Rate

I also measured the subjects' average intake rate during the last six experimental sessions for the group that did not have it controlled (fig.10). Similarly to the results reported for experiments 2 and 3, intake was higher in treatment D_A than in treatment D_D (Repeated-measures ANOVA with treatment and order as within and between-subjects factors:

$F_{1,4}=17.3$, $P < 0.01$; the effect of order and the interaction between the factors were not significant: $F_{1,4}=0.25$, $P = \text{NS}$ and $F_{1,4}=0.40$, $P = \text{NS}$, respectively).




Figure 10. Total intake rate experienced by subjects who did not have intake rate controlled during the last six experimental sessions in each treatment. ** $P < 0.01$.

The starlings were therefore more biased towards the option offering a higher long-term rate of gain (Target Amount) when hungrier than when more satiated. One aspect of my results would, however, seem to contrast with this observation and warrants further scrutiny. As a consequence of receiving the maximum possible amount of food in each block, those subjects assigned to the group where intake rate was controlled (Group F) experienced higher intake rates than those assigned to either treatment in Group U, where intake was not counterbalanced. In accordance with the previous results, one should thus expect the overall level of preference for T_A to be lower in the former group than in the latter if a subject's overall energetic state was to solely determine the strength of preference for the options. Nevertheless, no significant differences between groups were observed (Repeated-measures ANOVA with treatment and group as within and between-subjects effects, respectively: $F_{1,10}=0.05$, $P = \text{NS}$). Preference levels in the controlled group (fig.8b) were in fact similar to those observed for the birds under treatment D_A in the non-controlled group (grey bars in fig.8a). One possible account of this result is that overall intake in the latter treatment was already close to the subjects' physiological requirements and possible limits to energy budgets (see Bautista et al. 1998 for a thorough discussion), so that the higher rate available to the controlled group did not affect preferences. To explore this possibility I plotted the starlings' average food consumption (in grams of food per session) along the entire experiment for each one of the groups and treatments (fig.11). Subjects

under treatment D_D in Group U (grey circles) always showed a lower level of food consumption than those under treatment D_A (black circles), with differences between treatments being progressively smaller as – due to the increasingly frequency with which D_D was rejected, and T_A chosen, in choice trials – preferences in treatment D_D approached the maximum possible intake. Of more interest to the present analysis is the fact that, in treatment D_A , total consumption remained relatively stable at approximately 6.3 ± 0.2 g of food per session (or 18.9g per day) – clearly below the maximum potential consumption in that treatment (solid line), but within the range of total daily requirement reported for free-feeding adult starlings (Brito e Abreu & Kacelnik 1999). Confirmation of the hypothesis would have however required the inspection of total intake relative to each subject's individual requirements, as well as the continuous monitoring of variations in body mass. Since I did not conduct any of these measurements, this interpretation remains only as a possibility.

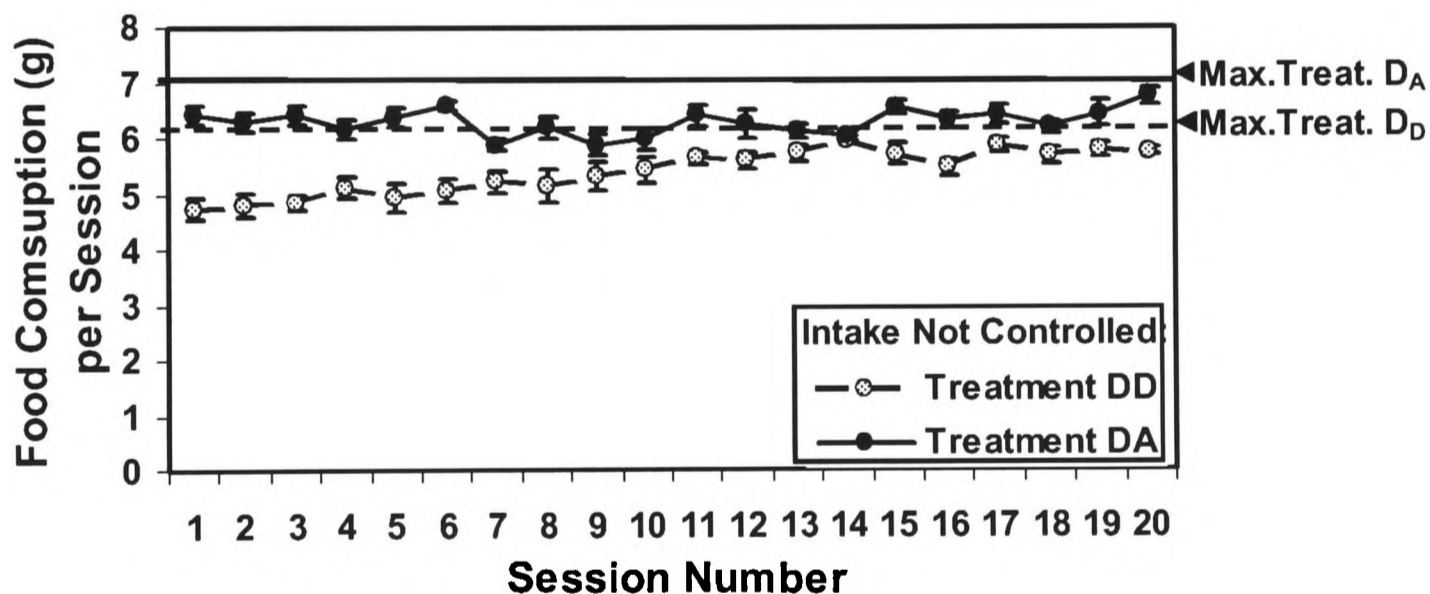


Figure 11. Average total food consumption (grams of Orlux© \pm S.E.) per session during the experiment. The datapoints represent total consumption of those subjects assigned to the group where intake rate was not controlled (Group U). The solid and dashed lines represent the possible maximum food consumption under treatments D_A and D_D , respectively. The upper solid line (maximum food consumption under treatment D_A) also represents the total amount of food per sessions offered to subjects in the controlled group (under both treatments).

6.7.2.4 Stability in Preferences

Similarly to experiment 2 and 3, I investigated whether the slopes of the regression of choice proportions for the focal choice pair against session number, as well as the variability in

choice proportions between sessions, were different when the starlings experienced different intake rates. For only one bird (Group F, treatment D_D) was the slope of the regression significantly different from zero. The slopes were, however, not different between treatments either in Group F (Repeated-measures ANOVA with Greenhouse-Geisser correction for sphericity: $F_{1, 5} = 0.09$, $P = \text{NS}$) or Group U ($F_{1, 5} = 0.39$, $P = \text{NS}$; Greenhouse-Geisser corrected). The variability in choice proportions (as measured by the standard deviation) across the final sessions, in turn, tended to be higher for those subjects experiencing treatment D_A (0.128 ± 0.023) as opposed to D_D (0.085 ± 0.017) in Group U, but the differences were not significant ($F_{1,5} = 1.9$, $P = \text{NS}$; Greenhouse-Geisser correction applied). There were no differences of variability between treatments in Group F ($F_{1, 5} = 0.18$, $P = \text{NS}$; Greenhouse-Geisser correction applied).

6.8 General Discussion

Theories of rational choice usually assume that the assignment of value to an option should be independent of the number and value of competing alternatives. As a consequence, preferences should be stable and consistent under different contexts. Contrary to this expectation, the starlings' preference for the focal choice pair was strongly affected by the presence of additional options that were irrelevant to the choice context. In a number of experiments, preferences for the target option offering the largest amount of food were reliably higher when subjects had experienced the decoy offering the shortest delay. Yet, the results also showed that, rather than being the outcome of the use of comparative or 'context-dependent' choice currencies, the observed deviations were the by-product of differences in intake rate caused by the birds' repeated exposure to energetically distinct decoys.

Why should preferences be modulated by differences in energetic budgets? One possible account is given by considering the costs (in terms of fitness loss) of choosing the option yielding the lower rate of gain under distinct states. As suggested by McNamara & Houston (1987), the probability of choosing a suboptimal action (Target Delay in the present case, as it offers a poorer long-term rate of gain) could be modelled as a function of the difference between the expected terminal fitness values of each of the alternatives. The main idea underlying their approach is that deviations from the optimal policy should be expected whenever the costs of such deviations are small. Here I build on this assumption and extend it to consider the simultaneous effects of variations in state. Figure 12 illustrates the idea. I firstly consider that repeated exposure to treatments offering lower and higher intake rates (corresponding to decoys D_D and D_A , respectively) causes some correlate measure of state (such as level of energy reserves) to be higher in the latter case (fig.12b). That is, state is assumed to be some positive function of rate of energy intake. I then follow the approach taken by Kacelnik & Marsh (2002) and consider the potential improvement in state produced by choosing each of the target options (T_A and T_D). As shown in the graph, the difference between the improvement in state caused by T_A and T_D is the same under both treatments (ΔS 's), and so the cost of taking the suboptimal option should also be the same in both contexts. Yet, this should not be the case if changes in fitness (or some other positively valued measure of utility), and not purely state, affected choices (some evidence for this is also provided in Chapter 7). Here I assume that potential improvements in fitness associated to each option's gains are likely to be a diminishing function of an organism's initial state,

being concave in shape (e.g., if state represents the level of reserves of an animal, then the contribution of a food item to his future reproductive success decreases with increasing reserves – see also McNamara & Houston 1982 for a thorough discussion of the foundations of the use of such a function). Figure 12a pictures this relationship. It shows that the canonical cost (defined as the difference between the expected fitness value of each option, McNamara & Houston 1986) of choosing the suboptimal option (T_D) is more severe under treatment D_D (δD_D) than under treatment D_A (δD_A). As a consequence, preference for T_A should be higher under treatment D_D if costly errors were less likely to occur. This possibility is further suggested by the observation that, in three out of the four experiments, preferences tended to be more variable under richer than poorer schedules.

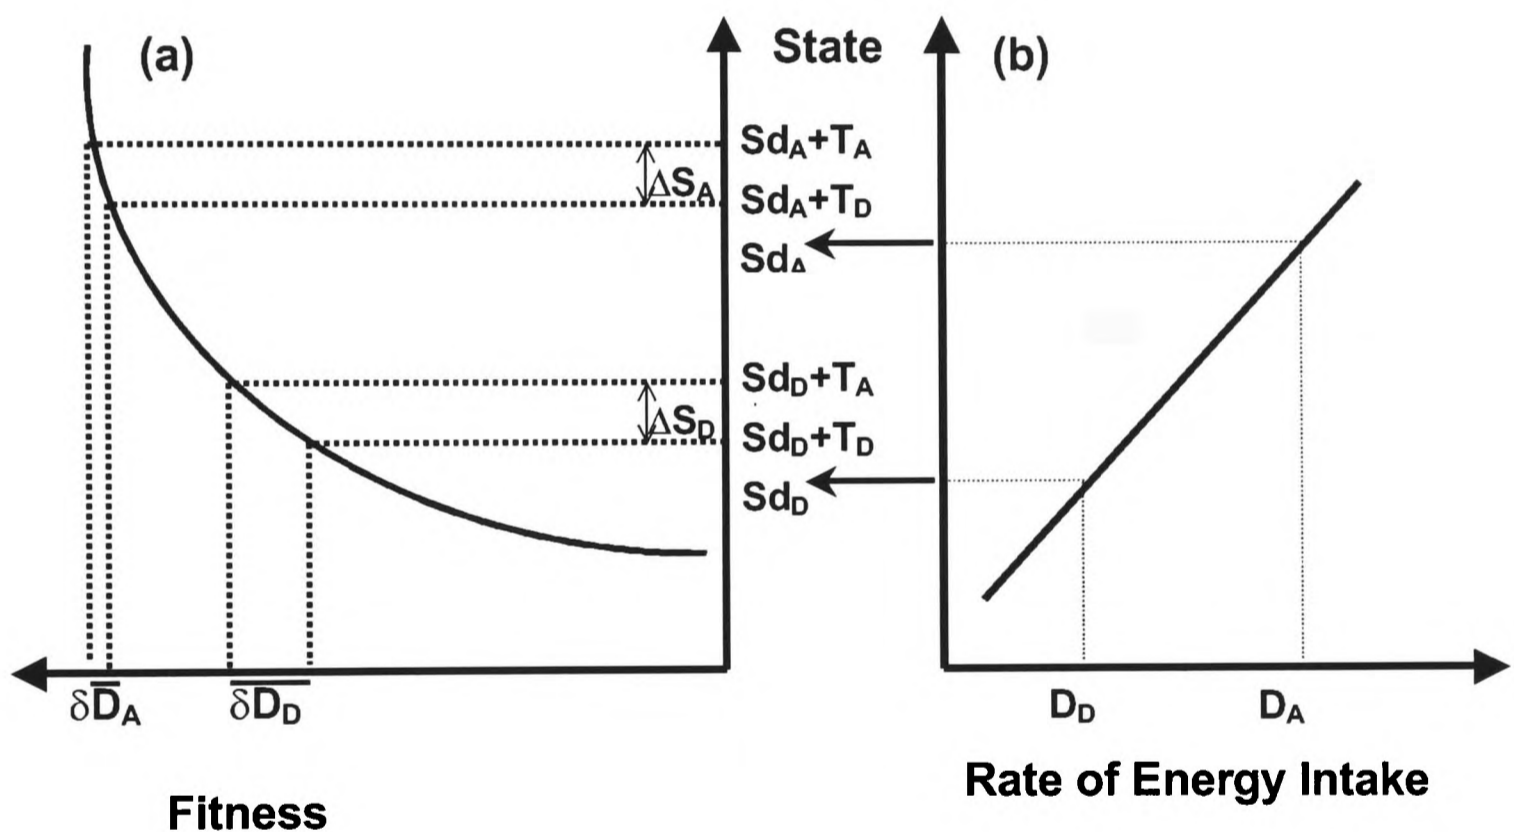


Figure 12. A functional model of how changes in state can lead to differences in partial levels of preferences. In (b) state is assumed to be some positive function of rate of energy intake. D_D and D_A represent the average intake rate experienced by subjects presented with the decoys D_D and D_A , respectively. Exposure to D_D leads to a poorer state (S_{d_D}) than that reached after exposure to D_A (S_{d_A}). In (a) fitness is assumed to be a concave function of an organism's initial state. $S_{d_D} + T_D$ and $S_{d_D} + T_A$ denote the state reached by subjects under treatment D_D after choosing target options T_D and T_A , respectively. Similarly, $S_{d_A} + T_D$ and $S_{d_A} + T_A$ represent the state reached by subjects under treatment D_A after choosing target options T_D and T_A , respectively. Although the difference in state between the options (T_A and T_D) is the same under either condition (S_{d_D} and S_{d_A}), the fitness cost of choosing T_D is higher under treatment D_D (δD_D) than under treatment D_A (δD_A). This should lead to a higher level of preference for T_A in the former treatment if costly errors were to be minimized.

Another possible explanation takes into account the potential differences in the way subjects learned the contingencies in each context. Based on the same paradigm as that described in figure 12, Houston (1987) argues that learning could be rapid when differences

in payoff between the options are larger, so that costly errors would be eliminated before less costly ones. From a mechanistic standpoint, we could likewise consider that learning rates may have been slower when the starlings were in a richer schedule. In fact, it has been long recognized that the function characterizing an animal's learning rate across time is steeper under conditions of higher deprivation, or 'drive' (Capaldi & Havancik 1973, Tarpay & Mayer 1978; see also Balleine 1992). It is thus possible that, along the experimental sessions, subjects were gradually learning the properties of the rewards and developing a stronger preferences for Target Amount, but those animals presented with treatment D_A took longer to do so, failing to reach the asymptotic performance with a higher frequency than those assigned to the alternative treatment. Yet, the analyses showed that there were no differences between groups in terms of the slopes of the regressions of choice proportions against session number. Hence, there is no evidence to presume that under richer schedules preferences would still evolve towards different levels had more time been given to the subjects. The argument is, however, extremely relevant whenever behaviour is compared between conditions inducing differences in state, especially in those designs employing a previously determined number of trials as opposed to the use of some stability criteria.

6.8.1 State-Dependent versus Context-Dependent Preferences

Breaches of rationality are well documented in the economic and psychological literature, and usually interpreted in terms of the psychological algorithms used to generate choices. One issue, however, sharply distinguishes human from non-human research, in that while humans can be presented with one-off choices and instructed of the alternatives verbally, non-humans must be repeatedly exposed to the contingencies in order to learn about them. Exposure to distinct contexts directly translates into exposure to distinct contingencies, which in turn can affect a suite of variables such as reserve levels and other physiological correlates of state. As a consequence, context-dependent behaviour may emerge simply as a by-product of state-dependency in decision-making.

6.8.1.1 Energetic State and Rationality in Jays

Consider for instance one recent study designed to test the effect of background context (as given by the set of options previously experienced) on the foraging preferences of food-hoarding grey jays (Waite 2001a). The jays were initially split into two groups and each group was given a series of twenty-five binary choices in a specific background context. One of the groups experienced context (A), being offered choices between one or three

raisins placed 0.5m inside a tube (where distance into the tube should correlate with predation risk). The second group was presented with context (B), being given binary choices between two identical options both offering one raisin 0.5m inside the tube. Both groups were subsequently (10 minutes later) presented with choices between either one raisin 0.3m into the tube or three raisins 0.7m into the tube. In violation of the principle of irrelevant alternatives, the strength of preference for the option offering three raisins was higher in the group that had experienced context B – an apparently paradoxical result interpreted as being consistent with the existence of cognitive biases leading to departures of value maximization (Waite 2001a). According to the author, an extended version of the model proposed by Tversky & Simonson (1993) on the context of human choices could account for his results by providing an explanation on how comparative evaluation of options varying in multiple attributes could lead to anomalous behaviour such as the one reported. Alternatively, however, experience with the distinct contexts and the consequent differences in the amount of food eaten and hoarded by the two groups may have led to the observed deviations. Those jays that had been previously given context A collected and transported an average of 62 raisins, approximately two and a half times more than those birds presented with context B (which collected and transported an average of 25 raisins). It is possible that, as observed in the present study, those jays previously presented with lower food supplies had more to gain by choosing the option yielding the larger amount of food (three raisins placed 0.7m into the tube). In fact, it has been shown that even if motivation to cache is not directly dependent upon hunger levels, previous caching can reduce subsequent caching of food, specially when the items available for storing are the same (Clayton & Dickinson 1999) – which could have produced the observed results.

6.8.1.2 Energetic State and Rationality in Hummingbirds

The comparison of binary and trinary choices usually employed in studies designed to test rationality in animals can also lead to altered states whenever these contexts are presented as treatments that are separated in time. For example, Bateson et al. (2002) compared the preferences of rufous hummingbirds (*Selasphorus rufus*) for three artificial flowers differing in the amount and concentration of sucrose (Target: 15µl, 40% sucrose; Competitor: 45µl, 30%; Decoy: 10µl, 35%) in binary (Target x Competitor) and trinary (Target x Competitor x Decoy) contexts. The birds did both treatments (binary and trinary) consecutively, with half of the birds experiencing the binary context first and half of the birds initially experiencing the trinary context. In both treatments, the birds had to make repeated choices between the

available flower types (placed on an artificial plate), until a minimum number of 150 choices for the Target and Competitor had been reached. The authors found that the strength of preferences for the Competitor over the Target significantly increased in the presence of the Decoy in the trinary context, having interpreted the results as being inconsistent with the use of absolute evaluation mechanisms as normally postulated by normative accounts of behaviour. Yet, the resulting differences could have been, similarly, caused by the exposure to energetically distinct contexts. Net rate of energy intake of the Target, Competitor and Decoy were respectively 81.9, 92.0 and 59.5 J/s. Considering for instance indifference between the first two options and an average proportion of choices for the Decoy of about 20% (as observed in the cited study), subjects would experience in the binary context an average intake rate approximately 7% higher than that experienced in the trinary. Cumulative gain along the 150 choices experienced was consequently likely to have been markedly lower in the trinary context, which could have led to an increase in preference for the option offering the higher net rate of energy intake in this situation – as the authors report.

The examples illustrate that controlling for the side effects that dissimilarities between choice sets may produce on an organism's state and motivation is paramount when investigating the influence of context on choice behaviour. This is even more important in light of the difficulty of predicting a priori how changes in state will affect preferences. For instance, the present results showed a higher level of preference for the larger, but more delayed reward when the starlings were under a poorer schedule – a result also previously reported by some authors (Christensen-Szalanski et al. 1980, Rechten et al. 1983, Bradshaw & Szabadi 1992). Conversely, the opposite pattern has been also observed in the literature (Snyderman 1983, Lucas et al. 1993), where increases in the level of deprivation had the effect of increasing preference for the more immediate, but poorer reward. While differences may be the artefact of the employment of distinct procedures, schedules and species, the fact that choice behaviour is found to be state-dependent seems, however, to be the rule.

If on the one hand changes in preference can be brought about by differences in state, on the other I do not dispute the possibility that comparative mechanisms of choice play a role in animal choices under some specific conditions. In the present study, both the binary and trinary choices were given under the same treatment in all experiments, where the state and intake rate experienced by the starlings was arguably the same. Still, in experiment 2 I observed a significant decrease in the subjects' relative preferences for one

target over another in the trinary situation (as compared to the binary). Although this result was not replicated for either of the two groups tested in experiment 4 – making it difficult to conclude about the generality of the phenomenon – it is interesting to observe that this context-dependent effect was produced exactly when subjects were indifferent between the targets in the binary context. It is thus possible that decoys are able to exert a more effective influence on behaviour whenever the specific choice of one or another option of a given choice pair is not relevant to the organism.

6.8.2 Rate Currencies

One final aspect of the starlings' behaviour is worth examining. The experiments were designed so that if the starlings were maximizing short-term rate of food intake they should have been indifferent between the target options. Nevertheless, only one group of birds (in experiment 2) showed, on average, indifference. In experiments 3 and 4, twenty-one (out of a total of twenty-four) subjects were biased towards Target Amount. One possibility is that the interval represented by the ITI had some effect on preferences. In fact, the addition of the ITI to the delay the birds had to wait to receive food would initially explain the overall higher level of preference for T_A over T_D (see dashed lines in figure 1), in accordance with the idea that, since this interval constrains the possibility of making new decisions, a perfect rate-maximizer should take it into account when assigning value to an option. However, had the starlings maximized long-term rate of gain considering the entire interval represented by the ITI, the strength of preference for D_A should have been considerably higher than that for T_D – contrary to the observation that, overall, we observed that $T_A > T_D > D_A > D_D$. It is thus possible that the effect of the ITI was best characterized by a hyperbolic function of the form

$$\hat{Rate} = \frac{Amount}{\alpha * ITI + Delay} \quad [3]$$

where \hat{Rate} represents the rate of food intake in units of amount per units of time, and α a parameter defining how sharply \hat{Rate} decreases as the ITI increases. The relative value of α thus determines the extent to which the predictions of classical foraging models were supported by the results, with $\alpha = 1$ representing the perfect fit. The above formulation would offer a possible account for the observed results whenever $\alpha \leq 0.11$ (i.e., relatively small effect of the ITI, calculated using the parameter values employed in the experimental

design). This is indicated by the slope of the dotted line in figure 13, which shows the extent to which the ITI that could have had an effect on preferences in our experiments.

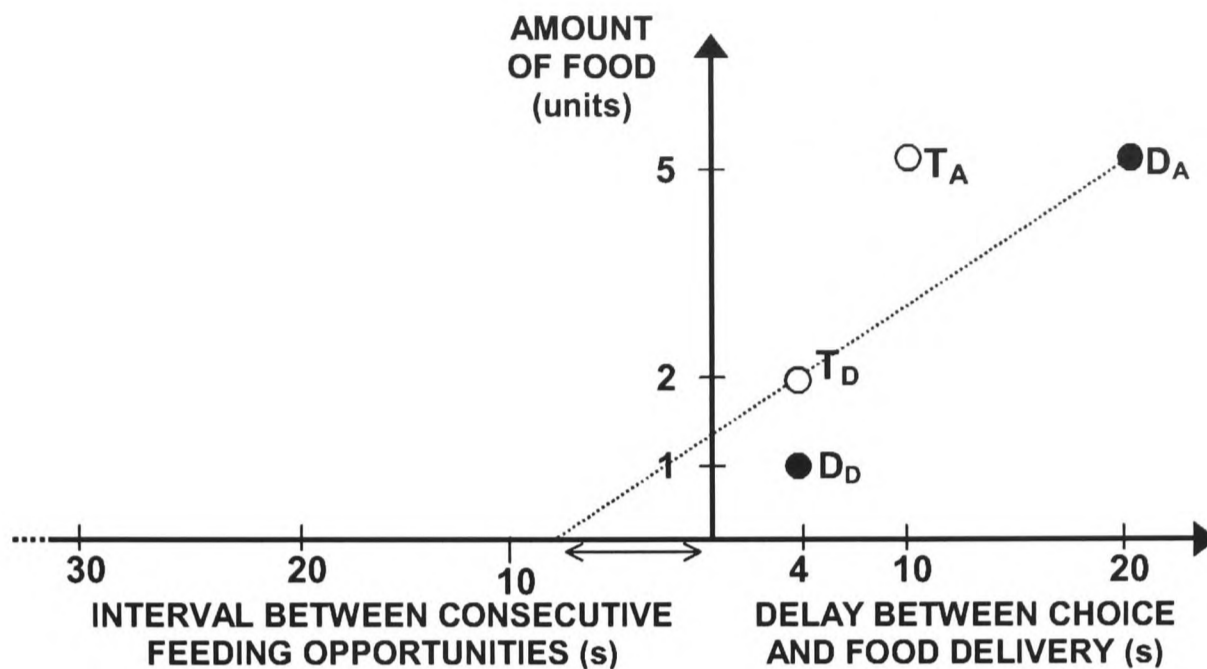


Figure 13. Each of the options used in the experiments, characterized as a function of the amount and delay to food (right side of the graph). The point where the dotted line crosses the abscissa indicates the ITI value for which subjects should be indifferent between T_D and D_A . Any ITI values (non-zero) shorter than the above should result in preferences between pairs of options such that $T_A > T_D > D_A > D_D$, as observed in experiments 3 and 4.

The results thus seem to agree with previous findings in the literature on operant behaviour (Mazur 1987, Mazur 1988, Mazur 2000), where it is usually found that the best relation describing the trade off between amount and delay to reinforcement has also the form of a hyperbolic function such that

$$V = \frac{A}{1 + KD} \quad [4]$$

where V represents the value assigned to the option, A and D the amount and delay (in seconds) to receive the reward, respectively, and K a dimensional constant (s^{-1}) with value close to 1 (notice that the above algorithm fits the observed pattern of preferences for almost any value of K larger than approximately $0.2s^{-1}$, but has the caveat of showing value in units of amount).

One group of subjects (Group D_A , experiment 2) showed however an overall bias towards target T_D , as opposed to T_A , a result seemingly at odds with the overall pattern of preference observed in the later experiments. It is possible, however, that the observed differences between experiment 2 and experiments 3 and 4 were an artefact of procedural

differences between them. While subjects in experiment 2 were given only two experimental sessions per day, being offered ad libitum food in between as well as after the sessions, those individuals tested in experiments 3 and 4 were presented with 3 sessions, receiving ad libitum food only at the end of the day. Hence, the latter experiments may have resembled more a closed economy schedule, and this difference could have affected the value attributed to the food earned during the sessions (Houston & McNamara 1989). In fact, since in experiment 2 subjects were freely provided with a fraction of their intake requirement twice per day, the value of obtaining the scheduled reinforcement may have been lower for these individuals, leading to a higher level of preference for the more immediate reward.

Finally, it is important to notice that if the state of an individual can affect the way it trades off amount and delay to food, as observed, then it is possible that variations in the very parameters defining the schedule of reinforcement given to subjects will affect the trade off as well. For instance, manipulations of the length of feeding interruptions (or ITI) could have an indirect effect on preferences whenever it causes differences in the deprivation level experienced by subjects. This putative state-driven effect of time intervals on preferences is fundamentally different from the expectation that long-term rates should be the currency maximized in choices (in fact, the present results showed that the intertrial intervals had little –experiments 3 and 4– or no effect –experiment 2– on the computation of the rate currency maximized by the birds). This prediction might be however important for the analysis of choice behaviour whenever procedures employing different intervals are compared, and although it is difficult to predict the direction in which changes in state will affect preferences, being aware of the possibility might help to understand empirical findings.

6.8.3 Conclusion

Following the growing body of evidence for irrational behaviour in the literature on human choices, recent reports on breaches of rationality in animals have also questioned the validity of the optimality approach, favouring the view that the reported inconsistencies would be the result of the use of choice currencies dependent on the context of choice. Although I am sympathetic to the possibility, this study suggests that a deeper and more critical scrutiny of the causes underlying observations of irrationality in animal choices is generally required. Amongst possible measures, crucial differences between human and non-human research

must be acknowledged. Here I emphasize the possibility that, due to need of repeatedly exposing animals to the contingencies of the choice problem, contextual changes may lead to variations in the overall state of individuals, which in turn can affect not only the amount of knowledge acquired by subjects but also the relevance of the decision and currency variables in terms of their consequences for the reproductive success of individuals. The present results confirm this expectation, emphasizing that in attempting to understand context-dependence in animal choices it is essential to control for such potential side effects. They also question the analysis of preference as a static variable, suggesting that research on economic rationality in animals should also view decision-making as the result of a dynamic, state-dependent process.

6.9 References

- Balleine, B. 1992. Instrumental performance following a shift in primary motivation depends on incentive learning. *Journal of Experimental Psychology - Animal Behavior Processes*, 18: 236-250.
- Bateson, M. 2002. Context-dependent foraging choices in risk-sensitive starlings. *Animal Behaviour*, 64: 251-260.
- Bateson, M., Healy, S.D. & Hurly, A. 2002. Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, 63: 587-596.
- Bateson, M. & Kacelnik, A. 1996. Rate currencies and the foraging starling: the fallacy of the averages revisited. *Behavioral Ecology*, 7: 341-352.
- Bautista, L.M., Tinbergen, J., Wiersma, P. & Kacelnik, A. 1998. Optimal foraging and beyond: How starlings cope with changes in food availability. *American Naturalist*, 152: 543-561.
- Bradshaw, C.M. & Szabadi, E. 1992. Choice between delayed reinforcers in a discrete-trials schedule - the effect of deprivation level. *Quarterly Journal of Experimental Psychology Section B- Comparative and Physiological Psychology*, 44B: 1-16.
- Brito e Abreu, F. & Kacelnik, A. 1999. Energy budgets and risk-sensitive foraging in starlings. *Behavioral Ecology*, 10: 338-345.
- Brunner, D., Fairhurst, S., Stolovitzky, G. & Gibbon, J. 1996. Mnemonics for variability: remembering food delay. *Journal of Experimental Psychology - Animal Behavior Processes*, 23: 68-83.
- Capaldi, E.D. & Havancik, J.R. 1973. Effects of previous weight level on rats' straight-alley performance. *Journal of Experimental Analysis of Behavior*, 97: 93-97.
- Christensen-Szalanski, J.J., Goldberg, A.D., Anderson, M.E. & Mitchell, T.R. 1980. Deprivation, delay of reinforcement and the selection of behavioural strategies. *Animal Behaviour*, 28: 341-346.
- Clayton, N.S. & Dickinson, A. 1999. Motivational control of caching behaviour in the scrub jay, *Aphelocoma coerulescens*. *Animal Behaviour*, 57: 435-444.
- Cuthill, I.C. & Houston, A.I. 1997. Managing time and energy. In: *Behavioural Ecology* (Ed. by Krebs, J.R. & Davies, N.B.), pp. 97-120. Oxford: Blackwell.
- Gigerenzer, G., Todd, P.M. & Gerd Gigerenzer, t.A.R. 1999. *Simple Heuristics that Make us Smart*. New York: Oxford University Press.
- Green, L., Fisher, E.B., Perlow, S. & Sherman, L. 1981. Preference reversal and self control: choice as a function of reward amount and delay. *Behaviour Analysis Letters*, 1: 43-51.
- Heath, T.B. & Chatterjee, S. 1995. Asymmetric decoy effects on lower-quality versus higher-quality brands: Meta-analytic and experimental evidence. *Journal of Consumer Research*, 22: 268-284.
- Houston, A.I. & McNamara, J.M. 1989. The value of food - effects of open and closed economies. *Animal Behaviour*, 37: 546-562.
- Houston, A.I. & McNamara, J.M. 1999. *Models of Adaptive Behaviour*. Cambridge: Cambridge University Press.
- Huber, J., Payne, J.W. & Puto, C. 1982. Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 9: 90-98.
- Hurly, T.A. & Oseen, M.D. 1999. Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal Behaviour*, 58: 59-66.
- Janiszewski, C. & Lichtenstein, D.R. 1999. A range theory account of price perception. *Journal of Consumer Research*, 25: 353-368.
- Kacelnik, A. & Marsh, B. 2002. Cost can increase preference in starlings. *Animal Behaviour*, 63: 245-250.

- Lucas, J.R., Peterson, L.J. & Boudinier, L.R. 1993. The effect of time constraints and changes in body mass and satiation on the simultaneous expression of caching and diet choice decisions. *Animal Behaviour*, 45: 639-658.
- Luce, R.D. 1959. *Individual Choice Behaviour: A Theoretical Analysis*. New York: Wiley.
- Luce, R.D. & Suppes, P. 1965. Preference, utility, and subjective probability. In: *Handbook of Mathematical Psychology* (Ed. by Luce, R.D., Bush, R. & Galanter, E.). New York: John Wiley.
- Maynard-Smith, J. & Price, G.R. 1973. The logic of animal conflict. *Nature*, 146: 15-18.
- Mazur, J.E. 1987. An adjusting procedure for studying delayed reinforcement. In: *Quantitative Analyses of Behaviour: The effect of delay and of intervening events on reinforcement value* (Ed. by Commons, M.L., Mazur, J.E., Nevin, J.A. & Rachlin, H.), pp. 55-73. Hillsdale, NJ: Erlbaum.
- Mazur, J.E. 1988. Estimation of indifference points with an adjusting-delay procedure. *Journal of Experimental Analysis of Behavior*, 49: 37-47.
- Mazur, J.E. 2000. Tradeoffs among delay, rate and amount of reinforcement. *Behavioural Processes*, 49: 1-10.
- McFarland, D.J. & Houston, A.I. 1981. *Quantitative Ethology: the State Space Approach*. Boston: Pitman Advanced Publishing Program.
- McNamara, J.M. & Houston, A. 1982. Short-term behaviour and lifetime fitness. In: *Functional Ontogeny* (Ed. by McFarland, D.J.), pp. 60-87. London: Pitman.
- McNamara, J.M. & Houston, A.I. 1986. The common currency for behavioural decisions. *American Naturalist*, 127: 358-378.
- McNamara, J.M. & Houston, A.I. 1987. Partial preferences and foraging. *Animal Behaviour*, 35: 1084-1099.
- Parducci, A. 1965. Category judgement: a range-frequency model. *Psychological Review*, 72: 407-418.
- Payne, J.W., Bettman, J.R. & Johnson, E.J. 1992. Behavioral decision research - a constructive processing perspective. *Annual Review of Psychology*, 43: 87-131.
- Rechten, C., Avery, M. & Stevens, A. 1983. Optimal prey selection - Why do great tits show partial preferences? *Animal Behaviour*, 31: 576-584.
- Shafir, E.B., Osherson, D. & Smith, E.O. 1989. An advantage model of choice. *Journal of Behavioral Decision Making*, 2: 1-23.
- Shafir, S. 1994. Intransitivity of preferences in honey-bees: support for comparative-evaluation of foraging options. *Animal Behaviour*, 48: 55-67.
- Shafir, S., Waite, T.A. & Smith, B.H. 2002. Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 51: 180-187.
- Simonson, I. & Tversky, A. 1992. Choice in context: trade-off contrast and extremeness aversion. *Journal of Marketing Research*, 29: 281-295.
- Snyderman, M. 1983. Optimal prey selection: the effects of food-deprivation. *Behaviour Analysis Letters*, 3: 359-369.
- Stephens, D.W. & Anderson, D. 2002. The adaptive value of preference for immediacy: when short-sighted rules have far-sighted consequences. *Behavioral Ecology*: in press.
- Stephens, D.W. & Krebs, J.R. 1986. *Foraging Theory*. Princeton: Princeton University Press.
- Tarpy, R.M. & Mayer, R.E. 1978. *Foundations of Learning and Memory*. Illinois: Scott, Foresman and Company.
- Tversky, A. & Russo, J.E. 1969. Substitutability and similarity in binary choices. *Journal of Mathematical Psychology*, 6: 1-12.
- Tversky, A. & Simonson, I. 1993. Context-dependent preferences. *Management Science*, 39: 1179-1189.

- Waite, T.A. 2001a. Background context and decision making in hoarding gray jays. *Behavioral Ecology*, 12: 318-324.
- Waite, T.A. 2001b. Intransitive preferences in hoarding gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 50: 116-121.
- Waite, T.A. & Field, K.L. 2000. Erroneous choice and foregone gains in hoarding gray jays (*Perisoreus canadensis*). *Animal Cognition*, 3: 127-134.
- Wedell, D.H. 1991. Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology - Learning Memory and Cognition*, 17: 767-778.

CHAPTER 7

7.1 INTRODUCTION	125
7.2 METHODS	127
7.2.1 PRE-TRAINING	127
7.2.2 TRAINING	128
7.2.3 CHOICE TRIALS	129
7.3 RESULTS AND DISCUSSION	129
7.4 REFERENCES	132

STATE-DEPENDENT VALUATION OF FORAGING OPTIONS BY STARLINGS*

**MARSH, B., SCHUCK-PAIM, C. & KACELNIK, K. State-dependent valuation of foraging options by starlings. Submitted to Behavioural Ecology.*
(as the version presented here is basically the same as that submitted, the use of the first person plural was maintained)

7.1 Introduction

Functional models of behaviour assume that animals are capable of making decisions according to the fitness value of different options, such as mates, feeding patches or prey types. Value, in turn, is often a function of the individual's state (see Mangel & Clark 1988 and Houston & McNamara 1999 for examples of state-dependent decision models). The corollary is that strictly speaking ideal subjects (those achieving the optimality solution at all times) would choose among options by a complex process involving both the physical attributes of each alternative and knowledge about its consequences as a function of state. Since this is unrealistic, biological analysis of behaviour assumes that alternatives are chosen on the basis of some simpler metric, accessible to the subject by normal psychological mechanisms, and showing a good correlation with fitness under most ecologically significant circumstances (Houston et al. 1982). Once this simplification is accepted, it becomes relevant to ask for the metric on which consequences of actions may be measured and compared.

One candidate metric, not necessarily optimal in all circumstances but probably quite a good compromise, may be some measure of the changes in the subject's state experienced as a consequence of taking an option, for example the increase in reserves consequent on deciding to forage in a given patch type. This is equivalent to remembering the absolute size of a reward. Another possibility, computationally more demanding, is for the subject to be sensitive to its own state at the time the payoffs were received, so that value depends on past marginal gains in fitness (or need reduction) consequent on choices of a given option. This would normally be correlated to the former but not necessarily linearly. Under the second possibility, two payoffs of equal size would be valued different if one occurred when reserves were lower, because the function of fitness versus reserves is often assumed to be concave. Both state changes and past marginal gains are intuitively appealing candidates for behavioural rules that could approximate optimal behaviour under many circumstances. A third possibility is that options are assigned value because they evoke the memory of a preferred state. Under this hypothesis, if an alternative is systematically experienced under a state of higher well-being (lower need) it should be preferred to one evoking a less desirable state.

In a recent study, Kacelnik and Marsh (2002) presented an explicit model to account for non-human similes of the so-called sunk-cost fallacy, a phenomenon related to

the Concorde Fallacy (Dawkins & Carlisle 1976). In their study, starlings obtained equal food amounts after either light or heavy work, and were later allowed to choose between the stimuli cueing for these food deliveries. They preferred the food source typically resulting from heavier work, and Kacelnik and Marsh pointed out that this may be understood if heavier work creates a state of greater need, so that the same food reward typically yields a greater need reduction when it takes hard work to achieve. Similar results had also been reported by Clement et al. (2000) working with pigeons. Both sets of results speak against a metric based exclusively on state change, because in both studies absolute gain in reserves (the reward size) was identical between options, thus favouring the possibility that some form of marginal gain valuation may be involved. The main limitation of Kacelnik and Marsh's test of their model is that in neither of the two studies was it shown that energy reserves decreased sufficiently with the amounts of work involved to justify a potential change in the slope of the fitness gains function.

In this study we continue exploring this issue by testing directly how an individual's state at the time of learning the properties of a novel food source influences preference for that food source in future foraging opportunities. The central hypothesis is that, options may be valued as a function of their past consequences in terms of changing the subject's need, and not as a function of the present change in needs, the past absolute change in the state of reserves or the desirability of the past state that the option evokes.

In spite of the importance of this issue for a broad class of functional models, we have been unable to find empirical or theoretical work that explicitly addresses it. In this communication we contrast the different hypotheses using the European Starling (*Sturnus vulgaris*) as the study animal. We trained starlings with identical food rewards from two artificial food sources that were encountered in two different states of food deprivation. Later, we gave them choices between the options while being in either of the two deprivation states. We consider three possible outcomes, related to the three hypotheses presented above. If choice is based on the physical properties of each option or on the absolute energetic state improvement associated with each option, the birds should be indifferent, because the rewards are identical. If need reduction at the time of exposure determines subjective value, we expect to see a preference for the option experienced in a greater state of need. Finally, if options acquire subjective value by association with the desirability of the state in which they are met, and which presumably they may evoke, they should prefer the option experienced in the higher state of reserves.

7.2 Methods

Subjects were 12 wild-caught adult European Starlings (licence 19990420/20010253 by English Nature). Before the beginning of the experiments (October 2000) the birds were kept in two outdoor aviaries, where they were fed ad libitum on a mixture of turkey crumbs, Orlux© pellets, and mealworms (*Tenebrio sp.*). Drinking and bathing water were always available and replaced daily. Five days before the beginning of the training sessions the birds were transferred from the outdoor aviary to the laboratory and housed in individual cages (120 cm x 60 cm x 50 cm) that served both as home cages and experimental chambers. Rooms were maintained at 18°C ($\pm 3^\circ\text{C}$), with lighting on a 12L:12D cycle with gradual transition periods at 0730 and 1930 h. Each cage had two perches (85 cm apart) and an operant panel with two circular response keys (3 cm diameter) and a central food hopper. The keys could be illuminated with white, red, or green. Food rewards were fixed for all treatments at 4 units of semi-crushed and sieved Orlux© pellets (ca 0.02 g per unit) and delivered at a rate of 1 unit/s by automatic pellet dispensers (Campden Instruments©). Birds were also permitted to feed ad libitum on turkey starter crumbs and supplementary mealworms at the beginning and end of each day, and for 15 minutes after the end of each experimental session. Orlux© (a preferred food item) was only available during the experimental sessions. Drinking water was always available and bathing trays were provided twice a week. There was no mortality or any sign of adverse state by any subject throughout the experiment. All subjects were released back into the wild (University Parks, Oxford, UK) during the following spring.

The birds were tested in two groups of six members each. All subjects had participated in an earlier experiment in which they pecked at black and white symbols to procure food rewards. For the experiment presented here (which used colours rather than symbols), the birds were transferred to new cages with different pecking panels, as described above.

7.2.1 Pre-training

Subjects were given one day to adapt to the new cages, followed by a day when they had to peck at keys lit randomly with respect to side (each using a white lamp) to release food. On this day all birds were observed to be pecking regularly to obtain food.

7.2.2 Training

This stage of the experiment was designed to generate in the subjects the subjective valuation of the food sources, in the absence of choice. From the morning of the third day onwards, the birds experienced a “cycle” of two sessions per day. The within-subjects design provided two session types: ‘hungry’ and ‘pre-fed’. To facilitate notation in the figures we refer to the first state as “H” and the second as “S” (standing for “satiated”). Each session lasted four hours and consisted of the following parts, in sequential order: (1) a food deprivation period (2 h 50 min); (2) a “manipulation” period (10 min); (3) a “key pecking” period, in which the birds needed to peck a coloured key to release food (this could last up to 45 minutes, depending on how long it took the birds to complete 10 trials); and (4) an *ad libitum* food period (turkey crumbs, for the final 15 min of each session) designed to reduce carry-over influences across sessions and equalise state for the following session.

The “manipulation” determined the kind of session: in ‘pre-fed’ sessions, the experimenter entered the room and provided *ad libitum* food (turkey crumbs) for nine minutes. In the ‘hungry’ sessions, the experimenter entered the room at the same times as in ‘pre-fed’ sessions, but no food was provided. In the “key pecking” period that immediately followed, a coloured lamp (red or green, balanced across subjects) corresponding with the session type was illuminated behind the pecking keys on either side of the panel. The side of presentation was random. The pecking key remained lit until a peck was registered. When pecked, the lamp was switched off and the food reward released, followed by an inter-trial-interval (ITI) of 80 seconds. The “key pecking” period terminated after the subject completed ten trials, or after 45 minutes, whichever came first.

The two daily sessions lasted four hours each, beginning at 0800 and 1200h. For each subject, one session of each type was given each day, with order of presentation balanced across subjects and reversing between days. The first group (n=6) had four days of training sessions. The second group (n=6) had five days, as on the third day the data were lost and so an additional day was added, as described above. This means that the number of observations available for both groups is the same but the amount of training received by the birds was slightly (c.20%) different. Three of the birds in the second group did not complete all of the “key pecking” opportunities in the four sessions of recorded data (they completed a total of 36, 37 and 37 trials out of the 40 opportunities in the ‘pre-

fed' sessions). Since however birds in this group had an extra training session, the total number of reinforcements during training slightly exceeded that of the first group.

7.2.3 Choice trials

After training, the birds experienced one day of “choice” sessions. In these sessions, in all trials both coloured lamps were simultaneously illuminated, so that subjects faced a simultaneous choice between the options previously associated to either treatment (‘hungry’ or ‘pre-fed’). The first peck to either key extinguished both lamps, released the food reward, and initiated the ITI. There was also a total of ten trials per session. The state of the subjects at the time of testing was balanced across sessions: three of the subjects of each group had the first session as in the ‘hungry’ treatment and the remaining three as in the ‘pre-fed’ treatment. This order was reversed in the afternoon session.

7.3 Results and Discussion

The purpose of the experiment was to test whether state at the time of training would affect preference for an option in the future and whether this effect might interact with state at the time of expressing the preference. Because of the difference in the amount of pre-training between the first and second group tested (4 and 5 days respectively, see Methods), when appropriate we consider the results of each group separately. We discuss two measures of choice bias: the proportion of subjects that favoured the “hungry” treatment option and the mean proportion of choices for this option across subjects. Out of the twenty choice opportunities each bird had, all twelve birds in both groups preferred the option associated with the “hungry” treatment more often (if a binomial test is applied to each group on its own $P = 0.031$, while for the 12 subjects together $P < 0.001$). Figure 1 shows the mean proportion of choices for the option associated with hunger during pre-training for at each state during testing). To test the influence of state during testing and the existence of group differences in choice proportions (arcsine square root transformed) we conducted a repeated-measures ANOVA, having state during testing (pre-fed or hungry) as a within-subjects variable and group (first or second) as a between-subjects variable. The state of the subjects at the time of testing did not significantly affect preferences ($F_{1, 10} = 0.25$, $P > 0.5$). Preference for the option associated with higher deprivation was stronger in the second group ($F_{1, 10} = 9.91$, $P = 0.01$). The only known difference between the two groups was the additional day of training used to compensate for one day of loss of data during training. If

one accepts as a reasonable hypothesis that this extra training may have contributed to this group's more extreme preference for the option associated with higher deprivation the implication would be that the closer the birds were to the asymptotic level of training the more pronounced the difference in preference between the options. The interaction between state at time of testing and group was not significant ($F_{1,10}=0.01$, $p>0.5$).

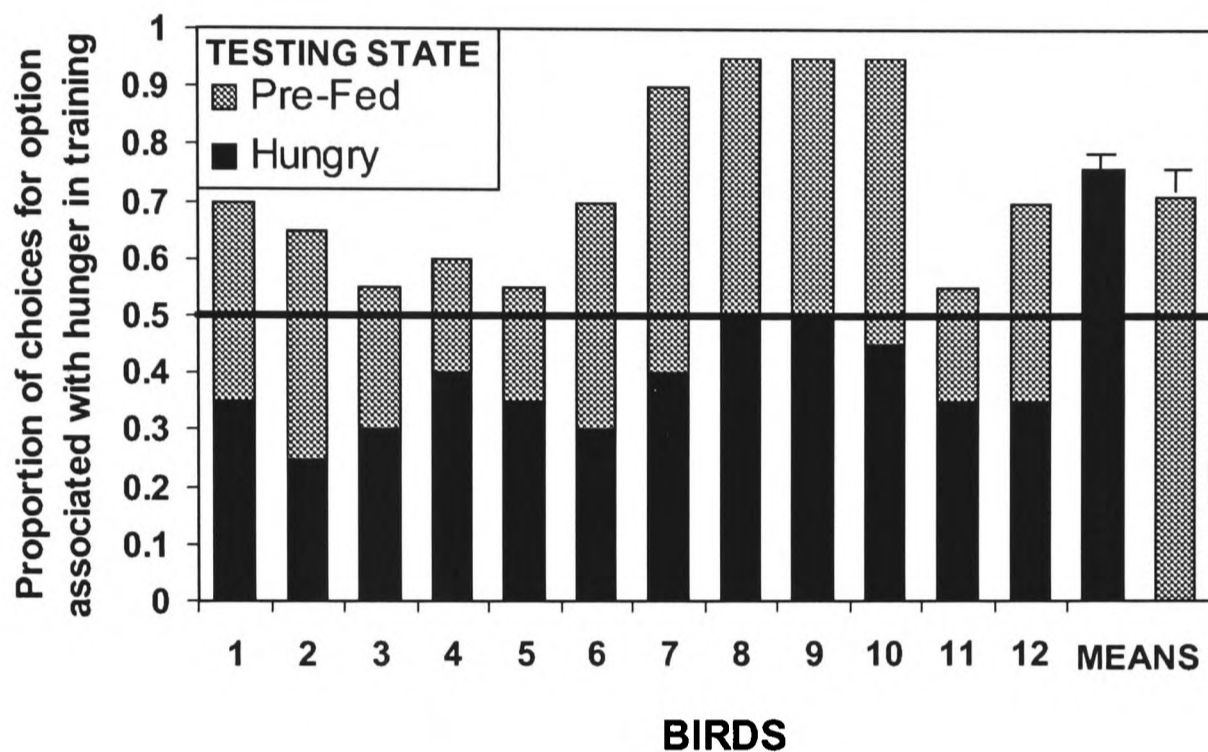


Figure 1. Proportion of choices for the option associated with higher deprivation during training. The twelve leftmost bars show proportion of choices for each individual. Since there was no significant effect of state at the time of testing, I show the result of adding choices over the two testing states by placing the bars on top of each other. The random expectation is 25% for each component and hence 50% for their sum. All 12 individuals chose the option associated with hunger more often than its alternative. Mean (\pm SE) proportions for both states during testing are shown on the two rightmost columns. For these bars the random expectation is 50%.

From a conditioning standpoint, rewards received under higher states of need should lead to faster acquisition of responding to their associated stimuli, but it is not clear whether this should translate into a higher level of choices when the options are encountered simultaneously, specially after learning has reached an asymptote. Indeed, recent developments in learning theory such as those fostered by Gallistel & Gibbon (2000) formalise the process of acquisition, but are mute respect to choice between stimuli either during acquisition or after asymptotic learning.

Work by Dickinson & Balleine (1994, 1995 and references therein) in the context of incentive learning also comes close to the issue investigated here. Although they did not address the problem of choice, in a series of experiments using rats, these authors show that the incentive value of a stimulus as expressed in rate of responding is modulated by the motivational state (e.g. hunger, thirst) of the animal while learning about it. Like the present results, theirs did not provide evidence for any effect of state at time of testing. Balleine and Dickinson propose that subjects learn the value of a resource in a given state through direct contact with it in that state. The results in this sense generalize and extend their findings to choice situations where options are identical in terms of their physical properties. The fact that we found no indication that state at time of testing influences choice speaks against the interpretation that internal state acts simply as a label for the specific conditioned stimuli (“red is better if state is good, green is better if state is bad”).

It is possible that the mechanism of assignment of value to options envisaged here, which can be summarised as remembering value as a function of the degree of “need reduction” or marginal benefit typically associated with each option, could sometimes lead animals to make seemingly ‘paradoxical’ choices, preferring food sources that deliver smaller rewards (lower absolute value) on account of their previous association with greater need deprivation. These data are not yet available, but if the prediction is confirmed it might give a direct measure of the potential deviation from the unconstrained optimum resulting from the use of a generally useful learning rule.

7.4 References

- Clement, T.S., Feltus, J.R., Kaiser, D.H. & Zentall, T.R. 2000. "Work ethic" in pigeons: Reward value is directly related to the effort or time required to obtain the reward. *Psychonomic Bulletin and Review*, 7: 100-106.
- Dawkins, R. & Carlisle, T.R. 1976. Parental investment, mate desertion and a fallacy. *Nature*, 262: 131-133.
- Dickinson, A. & Balleine, B. 1994. Motivational control of goal-directed action. *Animal Learning & Behavior*, 22: 1-18.
- Dickinson, A. & Balleine, B. 1995. Motivational control of instrumental action. *Current Directions in Psychological Science*, 4: 162-167.
- Gallistel, C.R. & Gibbon, J. 2000. Time, rate and conditioning. *Psychological Review*, 107: 289-344.
- Houston, A.I., Kacelnik, A. & McNamara, J. 1982. Some learning rules for acquiring information. In: *Functional Ontogeny* (Ed. by McFarland, D.J.), pp. 140-191. Boston: Pitman.
- Houston, A.I. & McNamara, J.M. 1999. *Models of Adaptive Behaviour*. Cambridge: Cambridge University Press.
- Kacelnik, A. & Marsh, B. 2002. Cost can increase preference in starlings. *Animal Behaviour*, 63: 245-250.
- Mangel, M. & Clark, C.W. 1988. *Dynamic Modelling in Behavioral Ecology*. Princeton: Princeton University Press.

CHAPTER 8

8.1 CONCLUDING REMARKS

134

CONCLUSION

8.1 Concluding Remarks

In this thesis I investigated the prevalence of economic rationality in animal choices using the European starling as a model species, and explored the potential implications of breaches of it for our understanding of the way animals make decisions. Here I briefly summarize the main results and conclusions presented in previous chapters, together with their collective significance.

In various experiments involving distinct foraging paradigms and levels of complexity in the design of the choice problem, the starlings showed a consistent and stable level of preference for the options made available to them. Overall, their preferences conformed to basic rational axioms, such as stochastic transitivity in its different levels, and the many forms of the principle of irrelevance from independent alternatives. Altogether, the recurrent consistency and stability in the starlings' behaviour of course does not imply that animals are strictly rational as assumed by economic theory, but speak against the expectation that '*context-dependent preferences may prove to be quite common, and could conceivably have pronounced ecological consequences*' (Shafir et al. 2002)⁵.

Although at first much less exciting than the observation of inconsistent and irregular behaviour, my initial results encouraged me to rethink possible causes underlying violations of rationality by animals as reported elsewhere, and their likely implications. Within this perspective, I firstly explored a number of instances where breaches of Luce's choice axiom, a restrictive form of the notion of independence of irrelevant alternatives, could be observed (Chapter 3). Luce's theory specifies how the utility of a particular alternative relates to its response probability, and predicts that the ratio of the choice probabilities of any two alternatives should be unaffected by the presence or value of new options included in the choice set. In this chapter, I argue however that if some proportion of choices is uncorrelated to the properties of the rewards – perhaps functionally explained by the need to track environmental fluctuations and reduce information uncertainty, or mechanistically due to the influence of uncontrolled sources of randomness – then changes in relative choice proportions between contexts might occur. Likewise, the existence of constraints on the perception of payoffs is also shown to be compatible with these

⁵ Shafir, S., Waite, T.A. & Smith, B.H. 2002. Context-dependent violations of rational choice in honeybees (*Apis mellifera*) and gray jays (*Perisoreus canadensis*). *Behavioural Ecology and Sociobiology*, 51: 180-187.

departures. I finally review the hypothesis that animals use higher order and hierarchical choices rules when making decisions, further contrasting the predictions of the choice axiom with the optimal expectation in the latter situation.

Chapter 4 experimentally tests some of the hypotheses postulated in Chapter 3, focusing on the possibility that the starlings use hierarchical choices rules, whereby choices are initially made between groups of alternatives (joined as a function of their degree of similarity or substitutability), and in a second step between alternatives within the group elected. The observed preferences were contrary to this expectation, but favourable to Luce's theory. Although compatible with this very restrictive rational axiom, these results were however intriguing from a functional perspective. They indicate that in choices between for instance, two different alternatives, the inclusion of new options into the choice set that are substitutable by the least preferred of the existing alternatives should gradually reduce the likelihood of pursuing the originally preferred one, to the point of generating a preference reversal between the two types of options. The counterintuitive conclusion (illustrated at the beginning of chapter 4 with the colonial nesting bird example) is therefore that, in this case, the birds' inability to use a more far-sighted rule (whereby choice probabilities result from considering the choice problem from a broader perspective) ironically leads to the observation of rational behaviour. I believe that these results open the possibility for future research to determine not only the generality of the phenomenon, but also the extent to which it can lead to deviations from the unconstrained optimum.

In the second half of the thesis the possibility that contextual changes in foraging scenarios affect an organism's overall energetic state, leading to a concomitant change in its choice behaviour, is explored. In Chapter 5 I argue that apparent breaches of rationality can be brought about in a number of ways as a result of state-dependency in decision-making whenever repeated contact with different contexts, and thus contingencies, affects the state of an animal. In one example, I show that it is possible to observe violations of regularity in those instances where energetic losses caused by experience with poorer options in a different context must be compensated by reliable gains. Indeed, in Chapter 6 I confirm some of these expectations, showing that context-dependence preferences are more likely to be observed when the energetic state of the subjects is not controlled. The economic definition of a rational decision-maker is predominantly a static one, and generally does not explicitly consider mechanisms of preference change. I suggest, however, that the *ceteris*

paribus assumption normally made in experimental tests of rationality may not hold in natural situations – the main implicating being that foraging preferences should be viewed as a dynamic, state-dependent measure also by students of economic rationality. Equally important perhaps, these results emphasise the distinction between human and non-human research, in that while the former can be instructed verbally about the specificities of the choice problem, the latter need to be pre-exposed to the contingencies to learn about them. A resulting drawback is that changes in the context of choice may lead to variations in the overall state of individuals, which in turn can affect both the amount of knowledge they can acquire and the importance of the different courses of action in terms of their consequences for the organism.

As emphasised in all the chapters, violations of rationality have strong implications for the understanding of the processes whereby animals make decisions, serving as useful tools for the analysis of the mechanisms used to process information. In general, they are taken to imply that the value of an option would not be only determined by its intrinsic properties and consequences as assumed by most normative models, but instead would be constructed at the moment of choice as a function of the value of other options available. The arguments put forth in this thesis show, however, that a number of alternative accounts other than the use of context-dependent mechanisms of choice (as usually suggested by students of economic rationality) can also underlie these deviations. As a whole, they indicate that a broader, and more critical, analysis of the choice problem faced by an animal will be usually required before unequivocal conclusions about the operation of particular cognitive mechanisms can be drawn.

As a final note, the following observation might be important. Throughout the thesis I have investigated and discussed choice situations the single alternatives offered to the subjects could be clearly ranked in at least an ordinal scale of value, as a function of their individual consequences for the animals according to some normative criteria. Nevertheless, in many cases it is possible that the optimal decision comprises a balanced combination of the available items, as opposed to a single element. For instance, one such case in which a set of items, or ‘consumption bundle’, is itself a commodity can result from the need of a balanced nutrient intake. In this situation, individual items might be pursued with different frequencies or choice probabilities, but they are only attributes or components of the total bundle being purchased. The relevant decisions in this context are those between the

possible combinations of individual food items, and not between single elements. Failure to acknowledge this might lead to the observation of (apparent) breaches of rationality. For example, it seems trivial otherwise that while bread may be preferred to salad when alone, salad may be preferred to bread if rice is also consumed, as rice provides the individual with the necessary carbohydrates and makes bread less of an essential intake. In this case, if choice behaviour were shaped so as to guarantee a proper intake of both vitamins (as provided by 'salad') and carbohydrates, the observed reversal of preference would be in line with the optimal expectation. Implicit in the rationale is therefore the idea that if the choice problem can be restated properly at a different level, then the seemingly violation of an axiom might cease to exist.

In conclusion, I have tried to emphasize along the thesis that, while the application of the economic methodology in the study of animal behaviour may be extremely useful for the understanding of animal decision-making, it requires adjustments to allow for the specificities of the biological and ecological scenarios under which animals operate. The necessity of this fine-tuning process might be useful to guide future research, and help us to fill the gap between normative predictions and observed behaviour.

**Other academic work published during
the development of the thesis (enclosed)**

ALONSO, W.J. & SCHUCK-PAIM, C. 2002. Sex-ratio conflicts, kin selection, and the evolution of altruism. *Proceedings of the National Academy of Sciences (PNAS)*, 99 (10): 6843-7.

FERNÁNDEZ-JURICIC, E., W.J. ALONSO & C. SCHUCK-PAIM. 2002. Peer-review to select academic job applicants. *Nature* 420: 16 (correspondence section).

SCHUCK-PAIM, C. & ALONSO, W.J. 2001. Deciding where to settle: conspecific attraction and web-site selection in the orb-web spider *Nephilengys cruentata*. *Animal Behaviour*, 62:1007-1012.

SCHUCK-PAIM, C. 2000. Orb-webs as extended phenotypes: web design and size assessment in agonistic conflicts between *Nephilengys cruentata* females (Araneae, Tetragnathidae). *Behaviour*, 37: 1331-1347.

SCHUCK-PAIM, C. 2003. Territoriality in spiders (in portuguese). *In: Ecology and Behaviour of Spiders* (ed. Gonzaga, M. & Japyassu, H.F.). FAPESP, Projeto Biota, São Paulo. *In review* (Book Chapter)⁶.

⁶ This book chapter is not enclosed as it is written in Portuguese.

The extracted article "Sex-ratio conflicts, kin selection, and the evolution of altruism" can be found at <http://dx.doi.org/10.1073/pnas.092584299>

impression. The original fingerprint was not damaged by this process.

These findings made sensational news. It was feared that astute money lenders would implant borrowers' thumbprints on fresh proforma and fill in larger amounts and/or higher rates of interest. When the case was referred to F. Brewster, the then Government Examiner of Questioned Documents, he submitted himself to a test prepared by Howrah Bar. The test consisted of 12 fingerprints — a random mix of genuine and transferred ones — on a single sheet of paper. He was asked to segregate them. In one attempt, Brewster identified the forged fingerprints. He noticed three differences: the transferred prints had diffuse lines while the originals had sharp patterns; the transferred prints were impregnated with gum; and the fibres on the part of the paper with transferred prints were disoriented.

Brewster therefore concluded that an observant fingerprint expert can easily differentiate an original mark from a transplanted one.

G. S. Sodhi*, **Jasjeet Kaur†**

*S. G. T. B. Khalsa College, University of Delhi, Delhi-110007, India

†Rajguru College of Applied Sciences for Women, University of Delhi, Delhi-110095, India

Science and government share anti-terrorist goals

Sir— Your news article headlined “National academies slam Bush proposal for data security” (*Nature* **419**, 769; 2002) mischaracterizes the debate here in Washington by its very title. It also leads readers to believe, erroneously, that I said that “without written guidelines, scientists can't accept [President Bush's science adviser, John] Marburger's assurances”.

The statement “Science and Security in an Age of Terrorism”, published on 18 October by the presidents of the National Academy of Sciences, the National Academy of Engineering and the Institute of Medicine, recognizes the need “to achieve an appropriate balance between scientific openness and restrictions on public information” when strategic secrets are at stake. But it also asks our government to maintain the current clear distinctions between classified and unclassified research, and recommends against poorly defined categories of “sensitive but unclassified” information that do not provide “precise guidance on what information should be restricted from public access”.

The statement also asks the Bush administration to reaffirm a national security directive signed by President

Ronald Reagan in 1985 which held that “no restrictions may be placed upon the conduct or reporting of federally funded fundamental research that has not received national security classification”.

In the wake of 11 September 2001, all of us in science and government have been forced to soberly reassess our roles concerning research touching on possible terrorist threats. No US scientist wants to publish research in a form that could be helpful to terrorists. Similarly, no US government official should want to hinder (or worse, stop) scientific research that might lead to effective tools against terrorism.

Scientists and government are listening to each others' legitimate concerns, and the government wants to enlist scientists in its anti-terrorism policies. In January, the National Academies plan to host a town meeting in which scientists, scientific publishers, national-security experts and government officials can talk face to face.

Marburger, director of the White House Office of Science and Technology Policy, is a key partner in these discussions. He informs us of the administration's concerns as well as communicating scientists' and scientific organizations' concerns to the White House. This is why I was disturbed at your implication that I said scientists couldn't rely on Marburger's assurances. That is simply wrong. The Bush administration has not yet even formulated its ultimate policy on this issue.

At the National Academies, we have never doubted Marburger's intentions or his goodwill. We share the same goal: to harness and focus the considerable energies of the science, engineering and health communities in the complex, rapidly changing challenge of counter-terrorism.

E. William Colglazier

National Academy of Sciences and National Research Council, 2101 Constitution Avenue Northwest, Washington DC 20418, USA

Peer review to select academic job applicants

Sir— Assessing the quality of candidates for academic positions has been the subject of controversy. Problems include nepotism and inbreeding, lack of impartiality among committee members — as M. Soler noted in Correspondence (*Nature* **411**, 132; 2001) — and the increasing use of impact factors as indicators of research performance (see, as one of many examples, *Nature* **415**, 726–729; 2002).

We suggest, as an alternative, that peer review could be a fairer method of research evaluation when scientists are being

assessed for new jobs. To ensure the expertise and impartiality of peers, appointment committees should include internationally recognized scientists who are not affiliated with the institution where the position is offered.

Operational costs could be kept low by use of e-mail, the Internet, videoconferencing and so on. Such cost-effective implementation would be of particular benefit to academic institutions in developing countries, which could call on a broader choice of specialists to comprise virtual appointment committees.

Rigorous criteria, relevant to the position, could be laid down before the job is advertised. These criteria might include originality of ideas, diversity of approaches, appropriateness of methods, statistical design and analyses, and so on. Candidates could then be assessed on the kind of quantitative scale that many journals use.

To that end, a comprehensive account of the candidate's research contributions, together with relevant published papers, can be presented without the journals' or the candidate's identities being revealed, so research committees can focus on scientific quality irrespective of journals' impact factors and potential personal biases. Moreover, the identity of the committee should be made public, to safeguard the impartiality of the selection process and the confidence of candidates.

Peer review has undeniably contributed to the advancement of science by providing a reliable system of quality control validated by experts. Some of the criticisms of its use (such as plagiarism and competition) are unlikely to apply to the selection of new academic staff, as it is usually a candidate's past achievements — rather than ideas, methods or data — that are under assessment.

No system is foolproof, and other criticisms of peer-review may apply to individual assessment — for example, sexism, opposing interests and lack of support for controversial ideas. However, these can be minimized by a policy of blind review, by having candidates specify potential conflicts of interest beforehand, and by including referees with different perspectives.

In summary, we believe that electronic information technologies can enhance the quality of recruitment systems in academic institutions by enabling the widespread use of peer review. This would be an improvement on current methods such as the impact factor, which is not a reliable measure of scientific excellence.

Esteban Fernández-Juricic, Wladimir J. Alonso, Cynthia Schuck-Paim

Department of Zoology, University of Oxford, Oxford OX1 3PS, UK

The extracted article "Deciding where to settle: conspecific attraction and web site selection in the orb-web spider *Nephilengys cruentata*" can be found at <http://dx.doi.org/10.1006/anbe.2001.1841>

The extracted article "Orb-webs as extended-phenotypes: Web design and size assessment in contests between *Nephilengys cruentata* females (Araneae, Tetragnathidae)" can be found at <http://dx.doi.org/10.1163/156853900501953>