

The Differential Influence of Strategy Selection and Implementation on Probability
Matching Behaviour in a Binary Prediction Task

by

Greta James

A thesis
presented to the University Of Waterloo
in fulfilment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Psychology

Waterloo, Ontario, Canada, 2017

© Greta James 2017

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner

Dr. John Minda
Associate Professor

Supervisor(s)

Dr. Derek Koehler
Professor

Dr. Daniel Smilek
Professor

Internal Member(s)

Dr. Mike Dixon
Professor

Dr. Jonathan Fugelsang
Professor

Internal-external Member

Dr. Michael Barnett-Cowan
Assistant Professor

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Probability Matching is a common and suboptimal strategy often used by participants in a Binary Prediction Task. A plethora of research has examined why participants engage in Probability Matching and also what factors improve the optimality of their choices. However, understanding of Probability Matching has been limited by a failure to differentiate between Strategy Choice (the conscious Strategy that a participant chooses to implement) and Strategy Implementation (the efficacy with which participants implement their chosen strategy). As a result, manipulations that primarily or partially cause effects by impacting Strategy Implementation have instead been attributed to modifying Strategy Choice. In particular, this thesis examines how two factors currently known to impact behaviour in Binary Prediction Tasks (feedback on the accuracy of participants' predictions and implementation effort) provide evidence for the distinction between Strategy Choice and Implementation. The first three experiments test the impact that feedback on the accuracy of predictions has on choice. I replicate earlier findings in the literature that feedback increases selection of the more probable outcome, but provide evidence that this increase is the result of factors influencing Strategy Implementation and not Strategy Choice. The final of these 3 experiments also examines the differential impacts of describing the contingencies in the problem versus having participants learn these contingencies through the aforementioned feedback. Here I find a large benefit to optimal performance from having the contingencies described.

An additional three experiments examine the impact of working memory load on participants' behaviour in a Binary Prediction Task. I find evidence that working memory load increases selection of the more probable outcome when Probability Matching is more difficult to

implement than Maximizing. I also find preliminary evidence that the opposite may be true when the implementation effort associated with each strategy is equal. In both cases, I find evidence that improvements in optimal responding, traditionally attributed to changes in Strategy Choice, are more likely the result of factors influencing participants' ability to implement their chosen strategy. Finally, I discuss what ramifications the distinction between Strategy Choice and Implementation has for future research on Binary Prediction Tasks.

Acknowledgements

The work included in this thesis only happened because of the extensive support I received. I would like to take this opportunity to thank my exemplary Advisors, Derek Koehler and Dan Smilek, for all of the guidance they provided me throughout my studies. You have been a tremendous support. I would also like to thank Priya Thakker and Lee Follis for their help collecting the data and running participants through the studies included here. I would like to thank my thesis committee, Jonathan Fugelsang, Mike Dixon, Michael Barnett-Cowan and John Minda for their valuable feedback on my work. Finally, I would like to recognize the National Science and Engineering Research Council of Canada for providing me with the funding necessary to complete my studies and this work.

Last I would like to thank my father, Lee James, for all the time he spent discussing numerous research ideas with me and my partner, Gus Gutoski, for listening to me talk for so long about Binary Prediction Tasks!

Table of Contents

List of Figures.....	xi
List of Tables.....	xii
Chapter 1: Introduction.....	1
Chapter 2: The Role of Feedback in Maximizing Behaviour.....	10
2.1. Introduction.....	10
2.2. Experiment 2.1: Methods.....	14
Participants.....	14
Procedure.....	15
2.3. Experiment 2.1: Results.....	16
2.4. Experiment 2.1: Discussion.....	24
2.5. Experiment 2.2: Methods.....	28
Participants.....	28
Procedure.....	28
2.6. Experiment 2.2: Results.....	29
2.7. Experiment 2.2: Discussion.....	35
2.8. Experiment 2.3: Methods.....	40
Participants.....	40
Procedure.....	40
2.9. Experiment 2.3: Results.....	42
2.10. Experiment 2.3: Discussion.....	55
Chapter 3: The role of Implementation Effort in Maximizing Behaviour.....	61
3.1. Introduction.....	61
3.2. Experiment 3.1: Methods.....	68
Participants.....	68

Procedure.....	68
3.3. Experiment 3.1: Results.....	71
3.4. Experiment 3.1: Discussion.....	74
3.5. Experiment 3.2: Methods.....	78
Participants.....	78
Procedure.....	79
3.6. Experiment 3.2: Results.....	79
3.7. Experiment 3.2: Discussion.....	81
3.8. Experiment 3.3: Methods.....	83
Participants.....	83
Procedure.....	83
3.9. Experiment 3.3: Results.....	86
3.10. Experiment 3.3: Discussion.....	90
Chapter 4: General Discussion.....	93
References.....	105
Appendix A: Experiment Materials.....	113
Experiment 2.1.....	113
Experiment Flow.....	113
Transfer Task.....	114
Questionnaire.....	115
Experiment 2.2.....	116
Experiment Flow.....	116
Questionnaire.....	116
Experiment 2.3.....	117
Experiment Flow.....	117

Instructions.....	118
Questionnaire.....	119
Experiment 3.1.....	121
Experiment Flow.....	121
Instructions.....	122
Experiment 3.3.....	124
Instructions.....	124
Questionnaire.....	125
Appendix B: Non-parametric Tests.....	128
Experiment 2.1.....	128
Binary Prediction Task Performance.....	128
Participant Estimates.....	128
Transfer Task.....	129
Experiment 2.2.....	129
Binary Prediction Task Performance.....	129
Transfer Task.....	130
Experiment 2.3.....	130
Binary Prediction Task Performance.....	130
Participant Estimates.....	131
Experiment 3.1.....	132
Binary Prediction Task Performance.....	132
Experiment 3.2.....	133
Probe Accuracy.....	133
Binary Prediction Task Performance.....	133

Experiment 3.3.....	134
Completion Time Data.....	134
Binary Prediction Task Performance.....	135
CRT2 Performance.....	135
Belief Bias Performance.....	135

List of Figures

Fig. 2.1. Experiment 2.1 Distribution of Total Greens Guessed.....	19
Figure 2.2. Experiment 2.1 Actual versus Estimated number of Green Guesses Across Conditions.....	22
Figure 2.3. Experiment 2.1 Proportion of Greens Guessed across Tasks and Conditions.....	24
Figure 2.4. Experiment 2.2 Distribution of Total Greens Guessed.....	31
Figure 2.5. Experiment 2.2 Total Number of Greens Guessed by Feedback and Order.....	33
Figure 2.6. Experiment 2.3 Distribution of Total Greens Guessed.....	44
Figure 2.7. Experiment 2.3 Mean number of Greens Guessed by Condition, Strict Maximizers included.....	46
Figure 2.8. Experiment 2.3 Mean number of Greens Guessed by Condition, Strict Maximizers removed.....	47
Figure 2.9. Experiment 2.3 Mean Difference Scores across Conditions.....	49
Figure 2.10. Experiment 2.3. Median Difference Scores by Condition.....	50
Figure 2.11 Experiment 2.3 Mean Estimates and Green Guesses by Condition.....	51
Figure 2.12 Experiment 2.3 Median Estimates and Green Guesses by Condition.....	52
Figure 2.13. Experiment 2.3. Mean Rank Index Scores across Conditions.....	54
Figure 3.1. Experiment 3.1 Number of Greens Guessed by Condition.....	74
Figure 3.2. Experiment 3.2 Mean number of Green Guesses by Condition.....	81
Figure 3.3. Experiment 3.3 Distribution of Responses to the Die problem Question by Condition.....	87

List of Tables

Table 2.1. Experiment 2.1 Central Tendency Statistics.....	20
Table 2.2. Experiment 2.2 Central Tendency Statistics.....	33
Table 2.3. Experiment 2.3 Strategy Choice by Condition.....	44
Table 2.4. Experiment 2.3 Central Tendency Statistics.....	47
Table 2.5. Experiment 2.3 Rates of Self-Reported Strategy Change by Condition.....	55
Table 3.1. Experiment 3.1 Mean and SD Scores by Condition.....	73
Table 3.2 Experiment 3.2 Mean and Median Accuracy Scores on the 3-back Task by Assigned Strategy.....	80
Table 3.3 Experiment 3.3 Central Tendency Statistics by Condition for Responses time (in seconds) to the Die Problem.....	88
Table 3.4 Experiment 3.3 Central Tendency Statistics by Condition for Responses to the Die Problem.....	89

Chapter 1: Introduction

Imagine I offered you a very simple gamble. I show you a 10-sided die on which 7 sides are painted green and 3 sides are painted red. I then tell you I'm going to roll this die 10 times and before every roll I am going to ask you which colour will turn up. How many times will you guess green?

This problem is an example of a Binary Prediction Task. As the name implies, a Binary Prediction Task is simply choosing between two predictions of an event. Typically, one event is more likely than the other (as in the example above), so you stand to maximize your winnings by always predicting which ever option occurs more often. In the die example above, you will make the most money by always predicting green (a strategy called Probability Maximizing, or just Maximizing) no matter how many times I roll the die.

In spite of its trivial simplicity, Binary Prediction Tasks have attracted quite a bit of attention over the past 70 years. The reason for this is that adult humans often do not predict which ever outcome occurs more often and, therefore, do not maximize their winnings. Rather, many people predict each option as often as it is likely to occur, a strategy called Probability Matching. So, in the example above, the probability matcher would choose green 7 times (because it should occur 70% of the time) and red 3 times (because it should occur 30% of the time). However, this strategy is obviously suboptimal, because every time you choose red, you are choosing a bet that pays out with 30% probability over a bet that pays out with 70% probability. Nevertheless, Probability Matching is a very common Strategy in Binary Prediction Tasks.

Probability Matching was initially interesting because of the challenge it raised for behaviourism, the leading theoretical account at the time of its inception. Probability Matching was first formalized by Grant, Hake & Hornseth in 1951, in an experiment designed to test conditions of partial reinforcement. Participants in their experiment attempted to repeatedly predict whether or not a light would turn on. For different groups, the light turned on at different frequencies (0%, 25%, 50%, 75% or 100% of the trials). What Grant, Hake and Hornseth observed was that participants tended to make guesses proportionate to the actual rate that the light turned on. In other words, they observed Probability Matching.

Under the behaviourist account probability matching shouldn't happen. As Grant, Hake and Hornseth noted in 1951, as long as one outcome is more likely than the other, reinforcement (the driving mechanism of behaviour under a behaviourist account) should happen more often when the more likely outcome is chosen. Over time, then, reinforcement should lead to maximizing. A flurry of research examining strategies on a Binary Prediction Task among children showed young children performing more optimally than older children (Derks & Paclisanu, 1967; Weir 1964; Lewis 1966) and various non-human species also performing optimally (Douglas & Pribram, 1966; Lee *et al.* 2004; Wilson 1960; Green *et al.* 1993). These findings are most easily explained by the role of reinforcement.¹

Grant, Hake and Hornseth's result combined with others in the field to challenge behaviourism and provided fuel for the debate between behaviourists and cognitivists. Despite behaviourists attempts to develop more complex reinforcement schedules and different

¹ Researchers have demonstrated that humans will also obey these laws of reinforcement after many many trials (Edwards, 1961; Bereby-Meyer & Erev 1998) and extensive feedback (Shanks *et. al* 2002) so perhaps we do eventually follow the rules of the behaviourists. Even so, we are much harder to teach than many other non-human species. This fact begs an explanation that tends to require a cognitivist account, so the challenge to behaviourism remains.

reinforcement units to explain probability matching, a review by Jones in 1971 concluded that “the evidence for conditioning theories exclusively explaining probability matching is small” (p. 159). This reflected a general shift in the discipline as a whole away from behaviourist accounts of psychology. Interest in probability matching from the 90s onward focused much more exclusively on cognitivist accounts and prior work on behaviourism virtually disappeared from discourse.

While interest in probability matching from a behaviourist perspective waned in the 80s and 90s, the problem became of great interest to those studying behavioural economics. The human tendency to probability match undermines the idea that people behave in ways that maximize their expected utility, an axiom that is fundamental to the whole of economic theory. In particular, it is a challenge to rationality because Probability Matching violates the principle of Statewise Dominance (A subclass of Stochastic Dominance). Put simply, Statewise Dominance is the assumption that people prefer more to less, or more technically, that they have monotonically increasing preferences. This assumption is a bedrock principle of rational choice. Thus, to violate it is to behave irrationally.

Probability matching became part of a suite of findings in behavioural economics that challenged the idea that people were rational. The challenge to human rationality presented by probability matching was addressed in one of two ways. Some argued that it was an artifact of poor information, poor research practices or both (for example Vulkan, 2000; Shanks et al 2002). Others argued that it reflected alternative preferences such as utility from factors other than the reward for correct predictions (for example Brackbill & Bravos, 1962). These arguments allow for humans to remain rational even in the face of Probability Matching and I will return to them shortly.

A second approach was to add it to a growing list of biases humans had that tended to systematically undermine rationality, but only in very particular circumstances. For example, one approach to understanding Probability Matching was to rely on the quickly growing literature on heuristics and biases. The idea here is that Probability Matching occurs because of the application of a heuristic or bias that is typically adaptive, but in this area results in suboptimal behaviour.

One such Heuristics and Biases approach is put forward by James and Koehler (2011) in their Expectation Matching theory. They argue that Probability Matching is the result of attribute substitution (Koehler & James, 2010; James and Koehler, 2011) a common heuristic from the Heuristic and Biases literature. Another Heuristics and Biases approach, what I will call the Pattern Search Hypothesis (Gaissmaier & Schooler, 2008; Wolford et al. 2004; Yellot, 1969; Gaissmaier et al., 2016; Unturbe & Corominas, 2007), emphasizes instead that probability matching is the result of an adaptive response to the environment. The environment encountered in a probability matching experiment is quite contrived. Typically an environment with these features would contain non-random events with patterns to be discovered. Thus, people have evolved an adaptive response to such environments that leads them to search for patterns. This pattern search is typically an adaptive approach that goes wrong when humans do a standard Binary Prediction Task in the lab.

Unlike the Pattern Search Hypothesis or Expectation Matching, some did not accept Probability Matching as an indication of our irrationality. For example, Vulkan (2000) argued that poor research practices, alternative incentives and structure of the learning environment prevented people from acting in a rational manner. He offers a review of the literature on probability matching with roughly this perspective in mind. In the end he concludes that "not all

hope is lost for the neoclassical approach to economics: 'matching is not robust, and some people do end up maximising their expected utility ' (p. 113).

One of Vulcan's conclusions, among others, are that the rate of actual probability matching is vastly overestimated because we simply aren't giving people enough time to learn the problem properly and enough incentive to respond optimally. These claims rely heavily on research demonstrating that after many trials with feedback humans move towards maximizing (Edwards, 1961; Bereby-Meyer & Erev 1998). Furthermore, Brackbill and Bravos (1962) found giving participants additional performance incentives tended to increase maximizing leading Brackbill and Bravos to conclude that under normal conditions participants get extra utility from predicting the infrequent event. These facts lead researchers such as Vulcan to conclude that humans are rational, but carrying out an irrational strategy because of poor information or incentives. However, this relies critically on the assumption that increasing incentives and providing extensive feedback modify participants' *strategy choice*.

Research after the behaviourists has assumed that whether or not you display probability matching or maximizing behaviour is the result of *only* your cognitive state and represents your choice to engage in a particular strategy. This assumption is shared regardless of whether you are responding to incentives, searching for patterns or using heuristics and is common to *all* of the Cognitivist approaches described above. They all assume, *a priori*, that increased maximizing behaviour is *always* synonymous with increasing endorsement of maximizing as a strategy.

Here I wish to argue that this assumption is misled. While changes in behaviour may be the result of changes in strategy, this may not be the case. Thus, I argue that behaviour in a

Binary Prediction Task is based on a two component process. The first component, common in the existing literature, I will refer to as Strategy Choice. Strategy Choice is assumed to be a largely top down process in which the participant consciously chooses what they consider to be the best strategy. Participants should be accurate at reporting this strategy and should consistently endorse this strategy (when explicitly asked) across tasks with varying surface features (ie: various chance devices and Binary Prediction Tasks). With regards to the argument above, traditional claims of rationality acts on Strategy Choice. In other words, a rational chooser selects a rational Strategy and that is why his/her behaviour is rational. Since this represents a conscious choice, the chooser is presumably quite aware of changes to their choice of strategy. In other words, if they choose to carry out a different strategy, they will be able to report that to you.

The second component I will call Strategy Implementation. Strategy Implementation is an attempt to carry out your top down Strategy Choice. Depending on the environment, Strategy Implementation may be more or less difficult. In difficult environments, features of the task and the environment may influence Strategy Implementation without impacting Strategy Choice (or endorsement). In other words, environmental features may change participants' behaviour by impacting their ability to implement their chosen strategy. These changes do not necessarily require participant awareness and may not actually alter their perception of the best strategy. Conditioning, for example, may be one bottom up mechanism that could change behaviour without changing participants conceptual understanding of what is the best Strategy. Simple association might work in a similar way. But it is not the mechanisms that are critical to my argument. The critical point is not *what* causes variance in strategy implementation, but that this variance exists and systematically impacts behaviour *without impacting Strategy Choice*.

The dissociation between Strategy Choice and Strategy Implementation has important implications. First it suggests that manipulations can selectively impact either one or both components. They need not impact both. Second, impacting either component is enough to move behaviour towards or away from maximizing without necessarily impacting the other component. I might maximize more simply because of something influencing my ability to carry out my chosen Strategy (Strategy Implementation) even if there is no impact on Strategy Choice. Finally, and most critically, I will argue that a change to maximizing (or matching) behaviour caused by Strategy Choice is not functionally equal to a change caused by Strategy Implementation. Changes to Strategy Choice should always be available to awareness and should be much more likely to be sustained. If you increase rational behaviour by impacting Strategy Choice, it makes sense to consider the individual to be more rational because they should continue to endorse the more rational strategy when they encounter the problem in varying formats in the future. By contrast, changes to Strategy Implementation are more likely to be transitory and less available to awareness. While it may momentarily increase rational responding, participants should be less likely to transfer these effects to new instances of the problem. So, to use the previous example, we would not expect an individual whose maximizing behaviour increases for reasons related to Strategy Implementation to endorse maximizing when they see similar tasks in the future. We would, however, expect that individual to endorse maximizing in the future if the manipulation affects Strategy Choice.

This is a critical point because much of the research on Binary Prediction Tasks that has been conducted over the last 70 years has conflated the two. While the Behaviourists relied exclusively on mechanisms related to Strategy Implementation, namely conditioning, the Cognitivists ignored the impact of conditioning (and other bottom-up influences) on behaviour

and instead assumed all changes in maximizing behaviour were associated with Strategy Choice. In the case of those making claims regarding rationality (for example Vulkan 2000, Brackbill & Bravos, 1962) or the role that Heuristics and Biases may play in behaviour (James & Koehler 2011; Gaismaier & Schooler 2008) this may lead to erroneously assuming you have changed participants' view of the problem when actually you have only modified their behaviour in an ephemeral way.

For example, if a manipulation momentarily improves rationality, researchers may conclude that they have demonstrated that humans are rational. Many researchers have demonstrated rationality improves after numerous trials with feedback (Edwards, 1961; Bereby-Meyer & Erev 1998). More recent research has demonstrated that feedback improves participants' performance relative to the performance of those not receiving feedback (Newell & Rakow, 2007). However, if we are to believe that these trials and feedback are actually improving rationality, evidence is needed to demonstrate that they impact Strategy Choice and that the effects are sustained. Chapter 2 reviews the effects of Feedback on behaviour in a Binary Prediction Task, and specifically tests whether participants are aware of this impact. Chapter 2 also examines whether the positive influence of Feedback on behaviour transfers to the same task in a different format. Both awareness and transfer are criteria needed to accept the assumption that feedback modified Strategy Choice.

Chapter 2 also examines the role that clearly describing the likelihood of each event (as opposed to learning if from feedback) has on participants' behaviour. Researchers that choose not to include feedback inevitably include such a description, but little is known about how this description influences participants' behaviour in a Binary Prediction Task. The final experiment

in Chapter 2 investigates the role of description and considers whether it impacts Strategy Choice or Strategy Implementation.

Besides rationality, existing research makes claims about the mechanisms behind non-maximizing behaviour. However, they do not specify at which level these manipulations are acting and confusing this can cause inaccuracy in causal theories. For example, increased maximizing under working memory load has been used to argue for a Pattern Search account of Probability Matching behaviour (Gaissmaier & Schooler, 2008; Wolford *et al.* 2000), but this argument requires that WML is impacting Strategy Choice, when, in fact, it seems plausible that it impacts Strategy Implementation instead. In Chapter 3, I examine how working memory load may impact both Strategy Choice and Strategy Implementation differently and suggest that the distinction must be understood in order to provide evidence in support of theories of Probability Matching Behaviour.

Finally, Chapter 4 summarizes the findings from Chapters 2 and 3 and relates them to the claim that Strategy Choice and Implementation have distinct impacts on behaviour. Chapter 4 also examines the short comings of the research presented here and makes recommendations for future work on the topic. Finally, Chapter 4 more broadly explores the interaction between top down processes such as Strategy Choice and the bottom up features impacting Strategy Implementation.

Chapter 2: The Role of Feedback in Maximizing Behaviour.

2.1. Introduction

In Grant, Hake and Hornseth's original paper in 1951 participants learned which or two lights would turn on by trial and error. They made guesses and observed outcomes and through doing so learned the contingencies associated with each light turning on. As early work focused on testing conditioning this was the standard way to approach Binary Prediction Tasks for many years. In this case, feedback was intrinsic to the task and tended not to be manipulated.

In 1996, Gal and Baron published a paper examining the intellectual and cognitive underpinnings of different strategies in a Binary Prediction Task. To do this, they fully described their Binary Prediction Tasks, including the contingencies. For example, Gal and Baron used a task similar to the one described in Chapter 1, in which a six sided die had 4 sides painted one colour and 2 painted another colour. Participants were told that they would be making predictions regarding the outcome of a series of die rolls with this die and asked to describe what strategy they would use. Notice that they needed no feedback on their guesses because the contingencies were described. Furthermore, they could provide a completely theoretical answer regarding their response, without ever seeing a single trial.

This descriptive approach is used extensively in the work reported here because it (a) allows us to manipulate the presence or absence of feedback, thereby determining its effect and (b) allows us to create versions of the Binary Prediction Task that have essentially no implementation component (participants may learn the contingencies from description and provide a similarly descriptive answer of their strategy). This latter component will be examined

more closely in the Chapter 3, while the role of feedback and description will be investigated in Chapter 2.

All of the experiments described in this thesis use the common descriptive problem from the introduction, which is based on Gal and Baron's (1996) example. To reiterate the earlier description, recall that I asked you to make predictions about the outcome of a 10-sided die. I used this same die with participants and they were also told that seven sides of the die were painted green and 3 sides were painted red². Participants were then asked to make some set of predictions about the outcome of the die either in a trial by trial fashion or through an aggregate self-report. Across the work presented here making more green guesses is synonymous with more optimal responding. For ease of exposition, I will often use guessing the more probable outcome and guessing green synonymously. However, I do not mean to imply that I think the results reported here are somehow specific to making green guesses in this particular task. Rather, I would expect them to apply to all similarly structured Binary Prediction Tasks.

By combining the above two techniques (describing the contingencies of the two events and providing outcome or accuracy feedback) researchers have been able to examine the role feedback plays in Binary Prediction Task behaviour in described tasks. A variety of work has demonstrated that feedback has a beneficial effect on rates of optimal responding (Newell & Rakow 2007, Shanks *et al.* 2002, Newel *et al.* 2013). In fact, after many trials (more than 1000) with feedback, participants start to behave almost optimally, with very high rates of maximizing (Edwards 1961; Shanks *et al.* 2002; Bereby-Meyer & Erev 1998). The implication of this work has been that feedback improves rationality. In fact, Shanks, Tunny and McCarthy (2002) go so

² In early work, we counterbalanced the colours presented on the die (Koehler & James, 2009; 2010), but since we repeatedly found no effect of colour, we have since discontinued this practice.

far as to claim that the challenge posed to rational choice theory by probability matching is merely an artifact of poor incentives and not enough performance feedback. In their investigation of the role of feedback in a Binary Prediction Task, Newell and Rakow (2007) conclude that “participants who are engaged actively in the task by making predictions, receiving feedback and presumably reflecting on that feedback are gradually pushed toward optimal responding” (p1138). This suggests that the mechanism behind the improvement seen from feedback is that it causes participants to deliberate on the effectiveness of their chosen strategy, thereby leading them to switch to a more optimal response.

The implication in both Newell and Rakow (2007) and Shanks *et al.* (2002) is that feedback is affecting participants’ Strategy Choice. To my knowledge, however, this has not been explicitly tested. While feedback does indeed increase optimal responding, it is not clear how or why. Feedback may directly prompt participants to think carefully about their strategy as suggested above, or it may instead impact their ability to implement their chosen strategy.

While little research directly investigates the role of feedback on Strategy Implementation there are some hints that it may at least affect Strategy Choice and Implementation differently. For example, Newell and Rakow (2007) find that initially (when Strategy Choice is presumably most active) the presence of feedback encourages *less* selection of a maximizing strategy, but that overtime feedback gradually increases optimal responding. However, it is not clear from their work whether the gradual benefit of feedback is also impacting Strategy Choice or simply operates on Strategy Implementation.

But how might feedback increase optimal responding purely through implementation? Perhaps the most straight forward account of the impact of feedback on implementation is the

simple Operant Conditioning assumed in the earliest work on Binary Prediction Tasks.

Thorndike's law of effect, that behaviour followed by pleasant consequences is likely to be repeated, would tend to encourage green responding. After all, green guesses are more likely to be correct (and therefore rewarded) than red guesses. Consistent with a conditioning account of feedback, previous work has found that feedback only improves responding when participants are making guesses (Newell & Rakow 2007) (i.e. emitting a behaviour – a requirement for conditioning) or imagining their guess (Tversky & Edwards, 1966). Simply observing the outcome sequence, by contrast, does not improve responding. Increasing payoffs improves performance (Brackbill & Bravos 1962; Shanks et al 2002) and while this is generally attributed to motivation, it is also consistent with conditioning account as increased reward tends to strengthen conditioned responses. Finally, conditioning is a handy mechanism to explain the existence of a predominantly maximizing strategy in young children (Derks & Paclisanu 1967; Lewis 1966; Weir 1964) and non-human mammals (Douglas & Pribram, 1966; Lee *et al.* 2004; Wilson 1960; Green *et al.* 1993).

For our purpose, whether or not conditioning is the mechanism behind the role of feedback may not be as important as whether feedback impacts, Strategy Choice, Strategy Implementation or both. Even if increased green responding is the result of conditioning and is effecting strategy implementation, participants could still be aware of that conditioning (Lovibond & Shanks 2002). This awareness might prompt the deliberation and engagement that would allow feedback to modify Strategy Choice. Unfortunately, little work that I know of has investigated directly (a) whether or not participants are aware of the impact of feedback on their choices, (b) whether that awareness leads them to modify their strategy and (c) whether the strategy modification is sustained.

Experiment 2.1 directly investigates these questions. The study seeks to replicate the finding that feedback increases selection of the optimal response (in this case, green), but I also extend current research by examining awareness of this improvement. This is done by asking people to estimate how often they guessed green. If participants are aware that they are guessing more optimally under conditions of feedback this should be reflected in higher estimates. If, on the other hand, they are not aware of feedbacks influence, they should not estimate differently than participants in the no feedback condition, but their choice behaviour should differ. We also ask participants directly whether or not their strategy changed throughout the learning phase. If participants are guessing green more often because feedback is leading them to reflect on the fact that guessing green pays off more, then participants should be more likely to say their strategy changed in the feedback condition than in the no feedback condition.

Finally, Experiment 2.1 also seeks to test whether the benefit of feedback transfers to the same task in a different format. If feedback leads participants to consciously modify their strategy (i.e. effects Strategy Choice), then participants should choose a more optimal strategy when presented with a very similar, but superficially different, task in the future. If, however, feedback effects only strategy implementation, participants should continue to endorse their original strategy when encountering a similar task in the future.

2.2. Experiment 2.1: Methods

Participants

Two hundred and three American adults completed an online experiment using Amazon Mechanical Turk. They were paid \$0.50 USD for their participation, as well as up to an

additional \$3.00 USD depending on performance. This study received ethics clearance through the University of Waterloo Ethics department and all participants provided online informed consent.

Procedure

Participants completed a computer task based around predicting the rolls of the 10-sided virtual die described previously. They were told that it was a fair die with 7 green sides and 3 red sides. The experiment consisted of 2 parts. In the first part (the Learning Phase or LP), participants were told that the die would be rolled 100 times. Before each roll of the die participants were asked to predict whether the die would come up green or red on that roll by clicking a green or red button at the bottom of the screen. The properties of the die (i.e. fair, with 7 green sides and 3 red sides) were clearly displayed at the top of each prediction screen. For each correct prediction they made, they would earn \$0.03. Participants were informed of their earnings as an aggregate sum at the very end of the study. The total earnings were displayed at the top of the digital feedback letter.³ Feedback on the accuracy of predictions was manipulated between subjects. For half of the participants their prediction was followed by a screen displaying either "Response Correct" or "Response Incorrect" for 1 second. For the other half of participants, this screen simply displayed "Response Recorded" for 1 second.

After completing 100 predictions of the die, all participants completed the second part of the study. In this part (henceforth referred to as the Transfer Task or TT), participants were

³ Note that participants do not receive any trial by trial information on their earnings. This is done to insure that the group receiving no feedback on the accuracy of their predictions does not deduce this information from earnings.

presented with the same game⁴ that they just played except they made only 10 predictions. They were asked to imagine that each accurate prediction would earn \$1 each. Participants then entered all 10 of their guesses into a grid containing 10 rolls and the option to click green or red for each roll. Note that since participants' choices were entered into a grid, previously choices were still visible to participants until all 10 choices were made. Additionally, they were told that their responses would be purely hypothetical, and that no die would actually be rolled for this task. Note that this means that no feedback was given during the Transfer Task phase of the experiment. Participants indicated their guess by clicking on their choice for each of the ten rolls. The exact wording and format for the Transfer Task is provided in the Appendix A.

At the end of the experiment all participants filled out a short questionnaire which asked them to estimate the number of greens they predicted in the learning phase and report whether their strategy changed over the course of the learning phrase. Participants also describe their strategy in the learning phase and answered a question testing their knowledge of the gambler's fallacy, but the findings from these latter two questions will not be reported here. A copy of the questionnaire can be viewed in the Appendix A.

2.3. Experiment 2.1: Results

A total of 203 participants completed the study. One participant was removed that guessed less than 50% green. Ninety participants remained in the Feedback condition and 112 in the No-

⁴ Half of participants were told that the die had 7 blue sides and 3 yellow sides during the Transfer Task, rather than the standard 7 green and 3 red sides. The justification for the inclusion of a differently coloured die was to test whether or not increased selection of green in the first part transferred only to an identical problem, or whether it would also transfer to a problem with superficial differences (i.e. the colour of the sides on the die). James & Koehler (2011) have previously found that these slight differences do differentiate the problems from one another and significantly impact participants' view of the problem. However, the variance in the presentation format of the problem proved to distinguish the problems sufficiently and colour had no impact on participants' responses. Therefore, the impact of colour is not discussed further here.

Feedback condition. Seventy- seven of these participants employed a strict maximizing strategy (guessed only green) during the 100 trial learning phase. When Strategy (Strict Maximizing vs. Non-Maximizing) was crossed with Feedback (Present vs. Absent) in a Chi-Square test of Independence, Strict Maximizing did not significantly differ across Feedback condition ($X^2(1)=1.58, p=0.21$). In other words, the presence of absence of feedback did not influence participants' tendency to adopt a strictly maximizing strategy. That said, the trend is that participants showed *more* Strict Maximizing when feedback was absent similar to the finding reported by Newell and Rakow (2007).

Participants who employ a strict maximizing strategy are not useful for studying how people learn from feedback because their performance is already at ceiling. Thus, it is useful to exclude these participants when examining what role feedback might have on learning during the learning phase. Thus, all Strict Maximizers (defined as guessing green on all 100 trials of the learning phase) have been removed from all further analyses. A strict cut off (allowing no red guesses among maximizers) is an unusual definition of maximizing. Previous work defining maximizing, has usually allowed for a few suboptimal guesses (for example Newell et al. 2013; Newell & Rakow 2007; Shanks et al. 2002). Thus, my definition bears some discussion.

I chose to use a strict cut off for two reasons. First, the number of red guesses to allow is an arbitrary cut off. For example, suppose an individual guesses 95% green. Was this someone with a maximizing strategy that guessed a few red, or an individual with an overmatching strategy that guessed predominantly green? I can't tell. Furthermore, overmatching (guessing more green than probability matching would dictate, but not guessing 100% green) is an explicit strategy that many participants endorse (Gal & Baron 1996), thus it is quite reasonable that participants guessing mostly green are not actually maximizers.

The second reason for choosing this cut off is merely a practical one. Maximizers were excluded from the data set because they were at ceiling and, therefore, exhibited no variability. Because I cannot examine the impact of my manipulations on this group, removing them makes sense. But justifying a cut off among those that guessed only a few red becomes more difficult. At what point do they exhibit sufficient variability to be included in the data set? I have decided that point is as soon as they are guessing any red.

Finally, it bears worth noting that the groups I define in the work presented here are not Maximizers and Probability Matchers, even though that is the traditional division in the literature. Instead, I refer to Strict Maximizers and Non-Maximizers. The reason for this is that my experiments revealed a breadth of strategies most of which are not Probability Matching. I will return to this point in the general discussion, but for now, I will simply point out that all of these Non-Maximizing strategies are suboptimal and present the same challenges to theories of rational decision making that Probability Matching does.

So, with Strict Maximizers removed, 125 participants remained that guessed at least 1 red during the 100 trial learning phase. Figure 2.1 shows a distribution of these participants' guesses. Note that Overmatching (guessing more green than a Probability Matching Strategy dictates) and Near Maximizing (guessing almost all green, but not exclusively green) were much more common strategies than strict Probability Matching. This is a common theme in much of the work presented here. I will return to this question in Chapter 4.

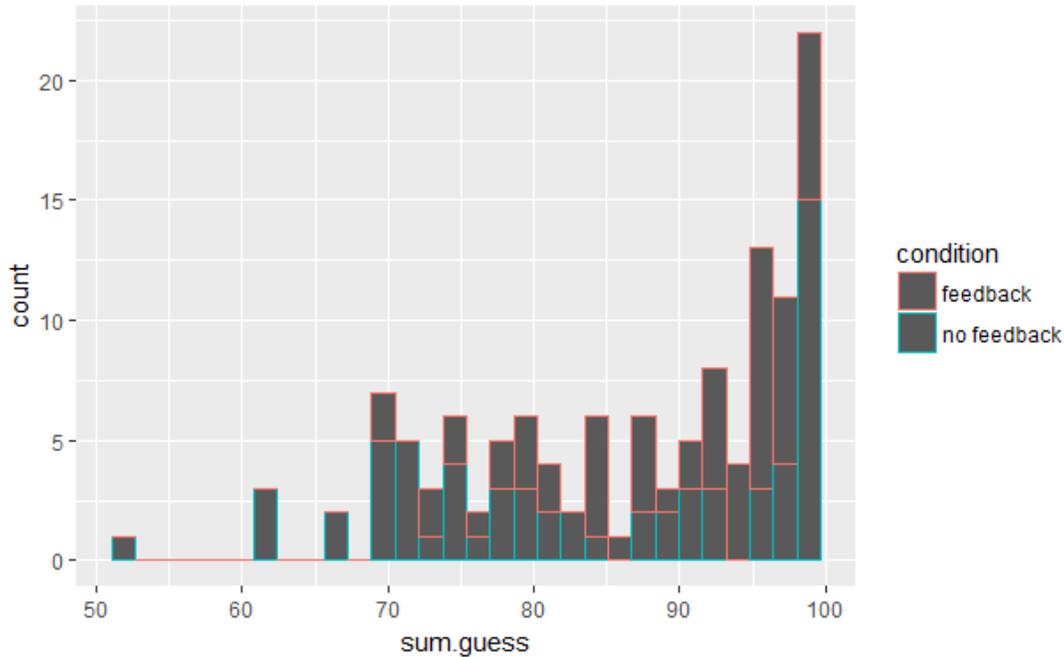


Fig. 2.1. Experiment 2.1 Distribution of Total Greens Guessed. Strict Maximizers removed.

Note also, that the data in Fig. 2.1 is not a normal distribution, but rather negatively skewed. This lack of normality is typical of all of the experiments included in this thesis. While removing Strict Maximizers from the data set tends to reduce the skew, the data remains non-normal. Furthermore, I also often violate the assumption of homogeneity of variance. One approach to dealing with this issue is to transform the data, but this would have led to considerable loss in meaning and would only help with the normality of the distribution. Another approach is to use non-parametric tests, but Zimmerman (2010) suggests that this may be less robust than using the Welch's t-test used (primarily) here. Thus, most of the tests reported here are parametric, but I have included their non-parametric counterparts in the Appendix B. In the majority of cases the pattern of results is the same regardless of whether or not I use a parametric or non-parametric test, lending credit to the claim that the parametric tests

are largely robust to the aforementioned violations. However, the non-parametric tests are done on medians, rather than means, and do tend to be less powerful. In some cases, this means that results that were significant with a parametric test are marginal or only trend in that direction with a non-parametric test. I will note in the text when this occurs. Otherwise, you can assume that the parametric and non-parametric tests produce the same result.

The remaining participants' total number of green guesses in the Learning Phase of the experiment were submitted to an independent sample Welch's⁵ T-test with Feedback (Present vs. Absent) as the between participant factor. Replicating previous work, participants who received feedback guessed significantly more green than participants who did not ($t(112.61)=2.73$, $p=0.007$, $d=0.48$)⁶⁷. Table 2.1 shows the mean and median number of green guesses by condition.

		Mean	SD	Median
Learning Phase	Feedback	89.4	8.6	92.5
	No Feedback	84.1	12.8	85
Transfer Task	Feedback	86	14.5	90
	No Feedback	82.8	14.8	80

⁵ I used the Welch's t-test (Welch, 1951) in R for all the independent t-tests presented in this thesis. Welch's t-test is an adaptation of the Student's t-test that is more reliable when two samples have unequal variance. It is not recommended to pre-test for equality of variance before using Welch's t-test, but rather to use it as the default (Zimmerman, 2004). Welch's t-test automatically corrects for violations of equality of variance by estimating the variances and accordingly adjusting the degrees of freedom used in the test. This adjustment creates degrees of freedom that may not be whole numbers or follow the standard calculation for degrees of freedom.

⁶ There was a trend in this data towards more Strict Maximizing in the No-Feedback condition. One possibility is that participants in the Feedback condition that would have been Strict Maximizers, guessed the odd red because of the presence of feedback. Since Strict Maximizers were removed from the data, but near maximizers were not, one possibility is that this effect of feedback reflects the continued presence of near Maximizers in the data. However, if the exclusion criterion for Strict Maximizers is changed to remove anyone who guesses more than 95% green, the effect of Feedback continues to exist and, in fact, becomes stronger.

⁷ Note also that the same test done on medians using a Wilcoxon rank sum test produced only a marginally significant result ($Z=2308, p=0.076$)

Table 2.1. Experiment 2.1 Central Tendency Statistics. Proportion of times participants guessed green during the Learning Phase and on the Transfer Task. (Strict Maximizers removed)

As one measure of whether participants were aware of the impact feedback had on their strategy, participants were asked to estimate how many times they guessed green in the learning phase. To index whether or not this score was accurate a difference score was calculated by subtracting how often participants guessed green from how often they estimated that they guessed green. Note that positive scores mean that participants overestimate how often they guessed green, negative scores indicate underestimation and scores near zero represent accuracy. Participants difference scores were submitted to an independent sample T-test with Feedback (Present vs. Absent) as the between participant factor. While both groups underestimated how often they guessed green, participants who received feedback made significantly worse estimates ($t(99.6)=-2.51, p=0.01, d=.45$), underestimating by an average of 8 greens⁸. Note, however, that this underestimation is driven by the aforementioned increased green guessing in the feedback condition, and that actual raw estimates made by participants did not differ by condition ($t(122.9)=0.25, p=0.81, d=.04$). Figure 2.2 shows participants mean raw estimates and actual mean number of green guesses by condition.

⁸ I did not run this as a mixed model ANOVA because it is technically not one consistent DV. However, if I combined these into a single DV and made estimate versus actual guesses into a within participant factor, the interaction did come out as significant, $F(1,123)=6.51, p=.012, \eta^2=.05$.

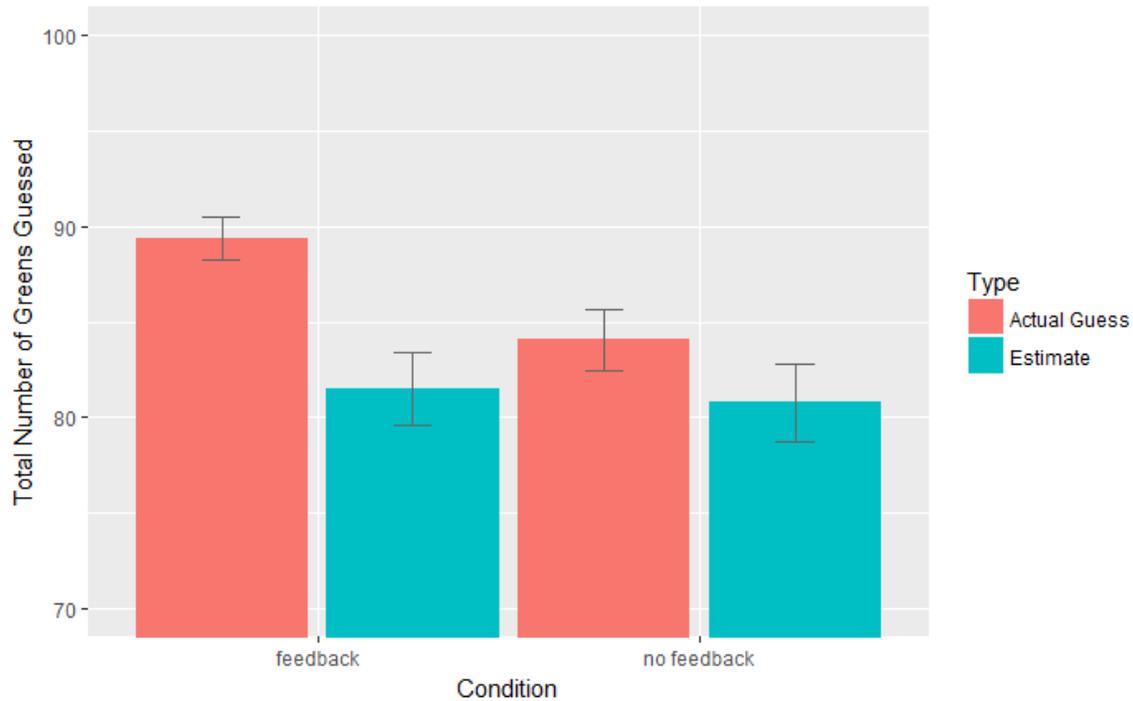


Figure 2.2. Experiment 2.1 Actual versus Estimated number of Green Guesses Across Conditions. Error bars are 1 SE of the mean.

As a further measure of awareness, we asked participants whether or not their strategy changed over the course of the learning phase⁹. Of the 127 participants who were not Strict Maximizers, 101 or 80% said their strategy did not change. Self-reported strategy change (Yes vs. No) was crossed with Feedback condition (Present vs. Absent) in a Chi-Square Test of Equivalence. Of those that did report a change in their strategy, they were not more likely to be in the feedback condition ($X^2(1)=1.2707, p=0.26$). Thus, reported strategy change was not associated with having received feedback.

⁹ Once again, we did not include Strict Maximizers in this analysis as we did not feel that they had any reason to change their strategy over time (no learning). That said, including maximizers does not change the pattern of the results.

Finally, I wished to investigate whether the improvement in optimal responding for those receiving feedback on the die task would be maintained when they completed an identical task with different surface features. Participants made 10 guesses on a grid (described in the methods section). The number of times they guessed the more likely alternative was summed to give a value between 0 and 10. I submitted these values to an independent T-test, with Feedback (absent vs. present) as the between subject factor. There was no significant effect of Feedback ($t(122.58)=1.23, p=.22, d=0.22$) on how many times participant guessed the more likely colour on the die. In other words, whatever benefit feedback has on guessing during the learning phase is no longer present when participants complete the exact same task in a grid format.

To investigate this further, I converted the number of green guesses¹⁰ during the learning phase and on the grid task to proportions so that they could be compared. I then submitted participants' proportion of green guesses during the learning phase and their proportion of green guesses on the grid task to a paired sample t-test. For those in the feedback condition, they guessed significantly more green during the learning phase than they did on the grid task ($t(59)=2.18, p=0.03, d=.27$)¹¹, but this difference was not significant for participants in the no feedback condition ($t(64)=1.02, p=0.31, d=.10$)¹². This difference is depicted in Figure 2.3. Once again, this suggests the benefit participants get from feedback during the trial by trial learning phase does not transfer to the grid task. However, this finding should be treated with caution as the median responses do not show the same pattern (see footnote).

¹⁰ This also includes blue guesses on the blue/yellow grid task problem, but I have written it as green guesses for the sake of flow.

¹¹ Note that this difference is not evident in the median responses, as a Wilcoxon signed rank test showed no difference ($V=1000.5, p=0.2629$) in the feedback condition.

¹² Note, however, that if I combine guesses on the LP with the TT into a single DV and run this through a mixed model ANOVA with LP and TT as a within participants manipulation, there is no significant interaction, $F(1,123)=1.07, p=.304, \eta^2=.01$, but there are main effects of both Feedback, $F(1,123)=4.12, p=.044, \eta^2=.03$ and type (LP or TT), ($F(1,123)=5.49, p=.021, \eta^2=.04$) .

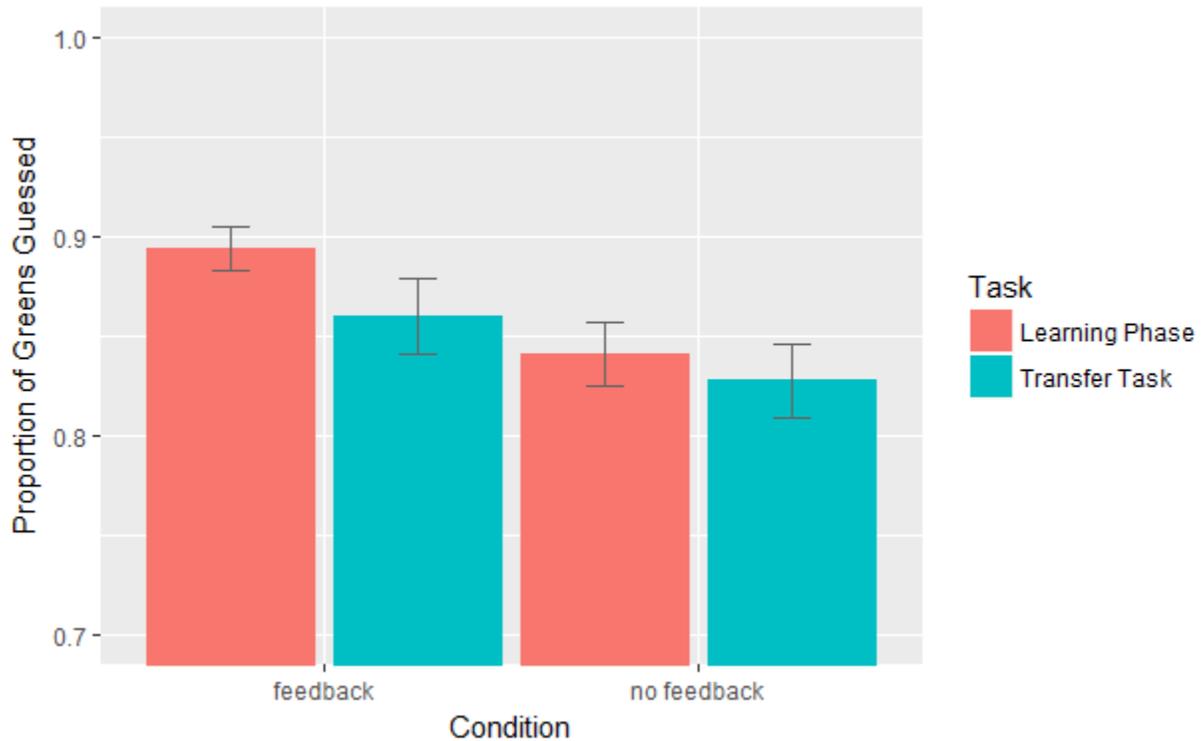


Figure 2.3. Experiment 2.1 Proportion of Greens Guessed across Tasks and Conditions. Error bars are 1 SE of the mean.

2.4. Experiment 2.1: Discussion

The results of Experiment 2.1 replicated previous work that feedback increases optimal responding for non-maximizers. However, I did not find convincing evidence that participants were aware of the impact of feedback on their choices. Participants’ estimates on how often they guessed green did not mirror their actual behaviour. Whatever mechanism participants used to generate these estimates did not seem to differ across conditions. In other words, participants who received feedback made the same estimates of their own green guessing as participants who did not receive feedback, even though they had actually guessed more green. Furthermore, most

participants did not report any awareness of a change in strategy when queried and those that did were no more likely to be in the feedback condition than the no feedback condition.

Finally, I found very little evidence that the benefit of feedback was sustained. There was no benefit of previous training with feedback on the Transfer Task, even though it was essentially identical and given directly after the training. Furthermore, for participants who did not receive feedback, the transfer task accurately reflected how they behave during the training phase (there was no difference between the two), but participants receiving feedback guessed more green during the training phase than they did on the transfer task.

In Chapter 1, I suggested that Strategy Choice was a conscious process and that alterations made to Strategy Choice would tend to be sustained across tasks. By contrast, processes that affect Strategy Implementation may operate without participants' awareness and will tend to be ephemeral. The findings from Experiment 2.1, then, support the idea that feedback increases optimal responding via a mechanism that operates on Strategy Implementation and that Strategy Choice is *not* affected by feedback. While the exact mechanism through which feedback exerts this effect is open to debate, conditioning seems to provide a plausible candidate. In particular, conditioning is highly sensitive to the specifics of the environment. While we may readily and quickly condition to emit a particular rewarded behaviour, this behaviour will not necessarily be emitted in a similar but not identical environment. The Transfer Task and Learning Phase were identical in terms of the cognitive components of the problem, but the format was different. In the Learning Phase participants made trial by trial guesses and, in the Feedback Present condition, did so with feedback. In the Transfer Task, responses were laid out on a grid simultaneously and participants received no feedback. These superficial differences would be enough to disrupt any benefit gained from

conditioning during the Learning Phase, but presumably only if the prior conditioning had not been integrated into participants' theoretical model of the problem. By contrast, if participants were actively reflecting on the Feedback they receive and were, thus, pushed towards optimal responding (as suggested by Newell and Rakow, 2007) then this would have presumably been reflected in their estimates of their own behaviour and would have transferred to their behaviour on the similar, but superficially different Transfer Task.

Experiment 2.1 provides evidence that the role of feedback in improving optimal responding may impact Strategy Implementation rather than Strategy Choice. Further, it suggests that the improvements to responding that occur as a result of feedback may not be available to participants' awareness and in particular are not sustained across problem types. These findings definitely challenge claims that humans are rational with respect to Binary Prediction Tasks as long as they are provided with proper feedback. While feedback does increase rational responding, it does not appear to increase rationality.

Experiment 2.2 further investigates maintenance of the benefits of feedback. In specific, Experiment 2.1 demonstrated that the benefits on behaviour from trials with feedback do not transfer to a similar, but superficially different task. Experiment 2.2 investigates whether they will be maintained in an identical task once the feedback is removed. If the impact of feedback disappears even when the task is identical this would be convincing evidence that participants are (a) not really aware of the impact of feedback and (b) that its impact is definitely not sustained.

In Experiment 2.2, then, feedback is manipulated within subjects with participants completing 2 phases identical to the Learning Phase in Experiment 2.1, one with feedback (Feedback Learning Phase or FLP) and one without (No Feedback Learning Phase or NLP). If

increased green responding with feedback is transient, then the order in which participants complete the FLP and NLP blocks should not matter. In other words, participants will guess more green when receiving feedback and less when receiving no feedback, even if they previously had a block with feedback. On the other hand, if feedback leads to some sustained learning that impacts Strategy Choice, then participants who benefit from feedback during their first Learning Phase should maintain that benefit during their second Learning Phase, even though feedback has been removed.

In addition to the two Learning Phases, participants in Experiment 2.2 also complete the Transfer Task from Experiment 2.1 both before and after the learning phases as a measure of strategy change. This will allow us to replicate the lack of transfer effect observed in Experiment 2.1. Furthermore, the initial Transfer Task completed before the Learning Phases will give us a metric on which to measure how participants' strategies may have changed after the Learning Phases.

Finally, as a further measure of transfer, participants are asked to describe in their own words the strategy they would use when playing a new die game involving a blue and yellow 10-sided die. These strategies are coded as Maximizing and Non-Maximizing strategies and compared to participants' behaviour during the actual Binary Prediction Tasks.

2.5 Experiment 2.2: Methods

Participants

Participants were 59 undergraduates from the University of Waterloo¹³. They completed a computer-based study in exchange for course credit. They also earned up to \$6 depending on the accuracy of their predictions in the Binary Prediction Task. This study received ethics clearance through the University of Waterloo Ethics department and all participants provided informed consent.

Procedure

Participants completed a computer task consisting of the two games they played in Experiment 1: the Learning Phase and the Transfer Task. In this experiment, however, participants began first with the Transfer Task. The Transfer Task was presented in an identical format to that of Experiment 2.1 except that only the green/red die was used (the blue/yellow die was omitted). This first presentation of the Transfer Task will henceforth be called Transfer Task 1 or just TT1. Note that no feedback was given to participants during this TT1 phase of the experiment.

Next participants completed two blocks of 100 trials each that were identical to the Learning Phase of Experiment 2.1, except that feedback was now manipulated within participants instead of between participants. For one block of 100 trials participants were

¹³ Experiment 2.2 was actually run prior to Experiment 2.1 (but reordered here because of improved coherence). Experiment 2.2 was run in lab and we were not yet aware of the appropriate sample size for this type of data. As such, Experiment 2.2's sample size is smaller than the rest of the samples provided in this work.

informed whether or not their prediction was correct (the Feedback Learning Phase or FLP block). For the other 100 trials participants were simply told that their response had been recorded (the No Feedback Learning Phase or NLP block). The order of the FLP and NLP blocks was manipulated between participants. When I refer to FLP or NLP as a condition I am referring to which of these blocks participants received first.

Finally, participants completed the Transfer Task for a second time, henceforth called TT2. This task was identical to TT1. Last, participants completed a comprehension question very similar to the original task. They were asked to consider a new ten-sided virtual die that consisted of 7 blue sides and 3 yellow sides, and described the strategy they would use to predict die rolls using this die. The colour of the die was switched because previous work has suggested that changing superficial feature of the die leads participants to consider it as distinct from the previous game (James & Koehler, 2011). Experiment 2.1 found no effect of changing the colour of the die. Nevertheless, this colour change was maintained in Experiment 2.2. The comprehension question was intended to measure whether learning during the task transferred to a theoretical understanding of the best strategy.

2.6 Experiment 2.2: Results

Fifty nine participants completed the experiment. Twenty-eight of these participants completed the FLP first followed by the NLP. An additional 31 did the reverse, completing the NLP followed by the FLP. Participants were coded as strict maximizers if they guessed green for all 100 trials in both of the two Learning Phases ($n=14$)¹⁴. All remaining participants were

¹⁴ As an aside, all of these 14 participants also guessed all green on both of the transfer tasks.

coded as using a non-maximizing strategy. Order of the NLP or FLP (First vs Second) was crossed with Strategy (Strict Maximizing vs Non-maximizing) in a Chi-Square Test of Equivalence. This revealed no significant relation between Order and Strategy ($X^2=1.02$, $df=1$, $p=0.31$). This is not a particularly good test of whether or not feedback encourages people to use a non-maximizing strategy, however, because maximizing is defined as guessing all green across both learning phases. These participants, then, guessed green both when they had feedback and when they didn't by definition. If participants get tempted into guessing red when there is feedback we would expect them to guess some red on the learning phase where they had feedback and switch to Strict Maximizing when they were not receiving feedback. Therefore, to test the impact of feedback on Strategy, it makes sense to compare the impact of feedback between participants within the first Learning Phase and then again within the second Learning Phase. Therefore, Feedback (present vs. absent) was crossed with Strategy (Strict Maximizing vs. Non-Maximizing) in the first and second Learning Phases respectively. This revealed a marginally significant difference in learning phase 1, with those that received feedback being *less* likely to engage in strict maximizing than those that did not ($X^2=3.12$, $df=1$, $p=0.08$). This effect, however, was not at all present in learning phase 2 ($X^2=0.09$, $df=1$, $p=0.76$). This suggests that either the significant effect of feedback in learning phase 1 is a type 1 error, or order is somehow relevant.

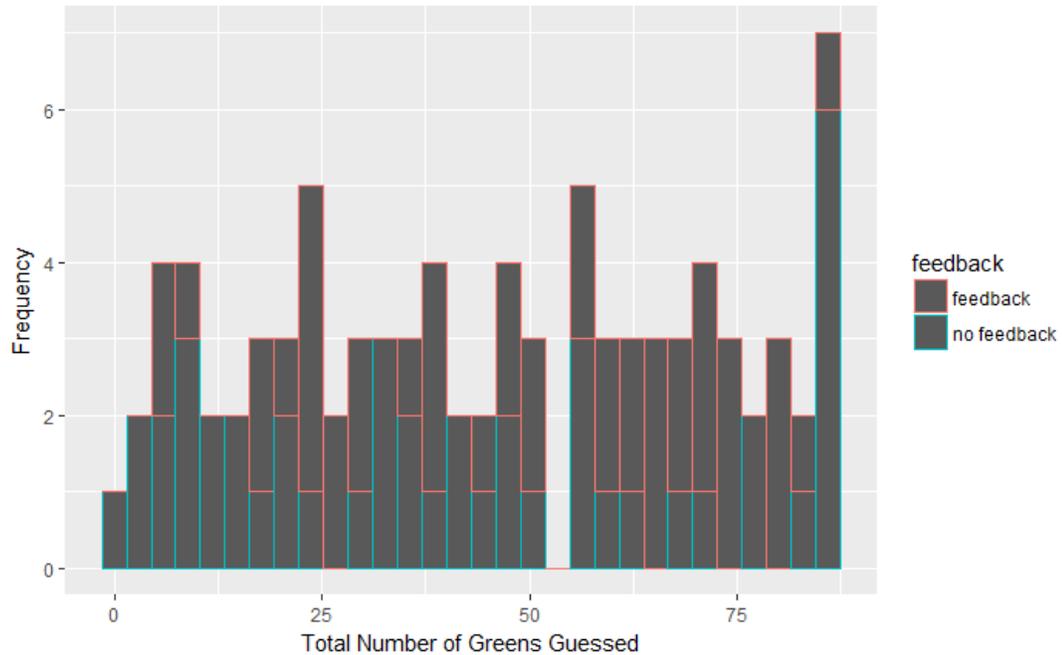


Figure 2.4. Experiment 2.2 Distribution of Total Greens Guessed. Strict Maximizers Removed.

As in Experiment 2.1, it is useful to exclude maximizers when investigating effects of learning during a Binary Prediction Task on behaviour. Thus the following analyses exclude the 14 participants who maximized across both Learning Phases of the experiment.

The remaining participants' total number of green guesses during each of the Learning Phases of the experiment were submitted to a mixed model analysis of variance with Feedback (Present vs. Absent) as the within participant factor and the order of presentation as a between participant factor. The ANOVA revealed a main effect of Feedback, $F(1,43)=5.21$, $p=.03$, $\eta^2=.11$), but no main effect of order, $F(1,57)=1.51$, $p=.23$, $\eta^2=.03$), and no interaction, $F(1,57)=1.06$, $p=.31$, $\eta^2=.02$). Figure 2.5 shows the mean rates of guessing green with and without feedback crossed with order of presentation. Table 3.2 shows the central tendency statistics for both the Learning Phase and Transfer Tasks. These results suggest that Feedback

increases how often non-maximizing participants guess green, but that this effect is not sustained. In other words, participants guessed significantly more green in the FLP than in the NLP regardless of which came first. However, closer examination of Figure 2.5 reveals that the benefit of feedback only appears to be present during the first Learning Phase. The lack of difference in Learning Phase 2 suggests there may be some sustained benefit of feedback. While the interaction is not significant, feedback is manipulated within participants, while order is manipulated between participants. I may just not have enough power to detect the latter effect. To more closely examine whether a within participants effect of feedback is present regardless of order I submitted how often each participant guessed green in LP1 and LP2 to two different paired sample t-tests: one for those that received feedback first and one for those that received it second. This revealed that there was a trend towards a difference in feedback if it came first ($t(22)=1.47, p=0.16, d=.33$) and a marginal effect of feedback if it came second ($t(21)=-1.8, p=.087, d=.39$)¹⁵.

¹⁵ Using a Wilcoxon signed rank test, this value is only a trend ($V=83, p=0.162$).

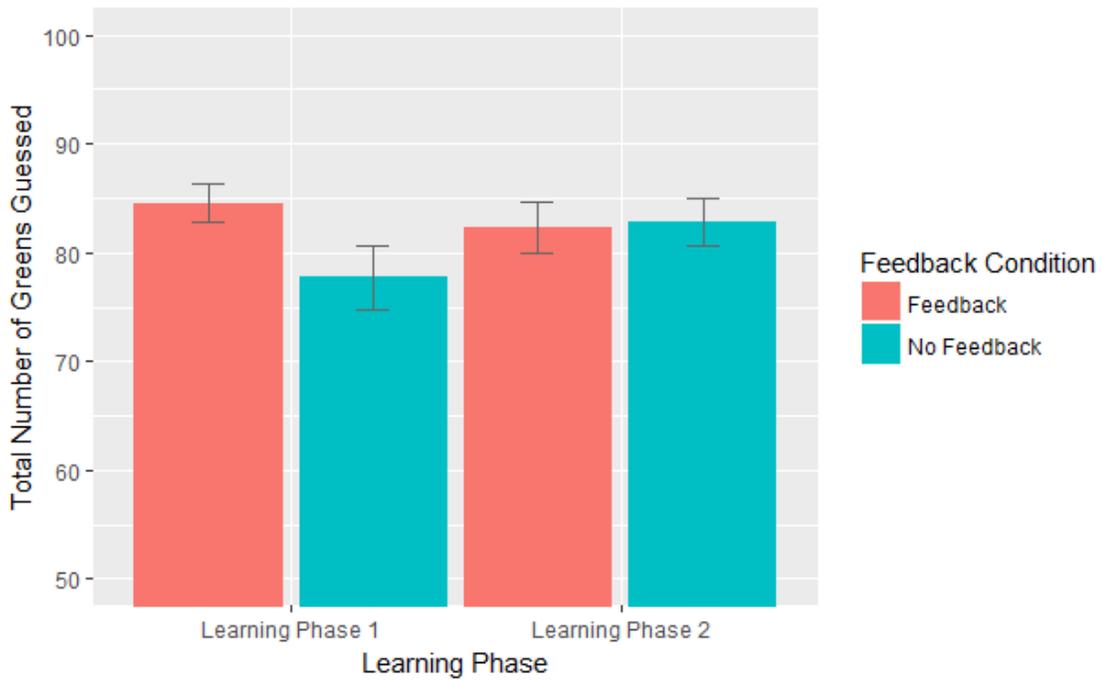


Figure 2.5. Experiment 2.2 Total Number of Greens Guessed by Feedback and Order. Error bars are 1 SE of the mean.

		Mean	SD	Median
Transfer Task 1		73.8	16.8	70
Learning Phase 1	Feedback	84.57	8.7	88
	No Feedback	77.73	13.58	74
Learning Phase 2	Feedback	82.32	10.70	83.5
	No Feedback	82.83	10.77	82
Transfer Task 2		80	13.3	75

Table 2.2. Experiment 2.2 Central Tendency Statistics. Percentage of times participants guessed green during the Learning Phase and on the Transfer Task. (Strict Maximizers removed.)

Experiment 2.2 also included two identical Transfer Tasks. The first occurred before the Learning Phases and the second afterwards. First I evaluated whether participants improved on the Transfer Task from the initial test (TT1) to the test after the Learning Phase (TT2). The number of greens each participant guessed on each test was submitted to a paired sample t-test with test time (1st versus 2nd) as the within participant factor. Participants guessed significantly more green on the second taking of the test than the first, $t(44)=-2.20, p=.03, d=0.33$. However, it is possible that this is just practice effects, as it is a pre-test, post-test design. This cannot be ruled out with the current design.

Following completion of the rest of the experiment participants were asked to describe in words the strategy they used when making their guesses. Responses were coded by two independent coders that were blind to condition. Responses were categorized as belong to a non-maximizing strategy or a maximizing strategy. In some cases, it was not possible to tell which strategy participants endorsed. These instances were coded as zeroes and excluded from analyses involving this question ($n=5$). Inter rater reliability was high ($k = .853, p<.001$). On five out of 59 cases the raters disagreed. These disagreements were resolved through discussion.

In order to compare the strategy classification above to participants initial strategy upon entering the experiment (in TT1), we needed to code the number of times they guessed green in TT1 into a strategy classification. Participants guessing between 6 and 9 green were classified as matchers and participants guessing 10 green were classified as maximizers. One participant, who guessed all red on TT1 was removed for this analysis. All other guessed 6 or more green.

When participants described what their strategy would be when encountering the same problem in the future (but with a die that had blue and yellow sides) they tended to be consistent with the

actual strategy they used at the beginning of the experiment in TT1. Only 5 of 45 participants switched from matching in TT1 to endorsing maximizing at the end of the experiment. For comparison, 3 of 45 switched the opposite way (from maximizing at the beginning to matching at the end). This suggests very little change occurred in participants' Strategy Choice over the course of the experiment.

2.7 Experiment 2.2: Discussion

Experiment 2.2 replicates the finding from Experiment 2.1 and the literature that feedback improves responding. This may be the case within participants and even when feedback precedes no feedback, although this requires replication with a larger sample.

Experiment 2.2 did find some benefit from the Learning Phases to performance on the Transfer Task, or rather that performance on the Transfer Task improved from TT1 to TT2. However, this improvement was uncorrelated with improvement during the Learning Phases and thus does not seem to reflect learning from feedback. The improvement may be better explained by a test-retest effect.

Finally, participants remained consistent with the strategy they employed at the beginning of the experiment, in TT1, when asked at the end of the experiment to describe the strategy they would use in an identical problem involving a 10-sided die with 7 blue sides and 3 yellow sides.

Together these findings support the idea that feedback's beneficial effect on responding in a Binary Prediction Task is impacting participants' Strategy Implementation rather than their Strategy Choice. The findings also provide further support for the idea that the benefits of

feedback do not transfer to future contact with the problem (as seen by participants' responses to the blue/yellow die problem). But Experiment 2.2 is less conclusive about whether or not the benefits of feedback are maintained in identical environments. While there was no significant effect of order, it did appear as though the effect of feedback was greatly diminished in Learning Phase 2, suggesting that there is some transfer to the No Feedback Learning Phase when it is preceded by the Feedback Learning Phase. However, the drop in green responding during the No Feedback Phase (especially among the median scores), even when it is second, does suggest that the effect of feedback is at least partially lost. This issue could be better clarified with further research that was suitably powered.

If the impact of feedback is actually independent of order, this raises interesting questions as to why prior benefits of feedback might not be maintained. If conditioning is the mechanism by which feedback improves responding, then one possibility is that the no feedback block serves to extinguish the conditioned response. However, I did not provide enough trials with feedback here for participants' behaviour to reach asymptote. Further, more trials without feedback may also be needed to more fully extinguish the conditioned response and may explain why there is a partial but not complete benefit to feedback in LP2. Future research more targeted towards understanding the mechanism behind the improvement from feedback should include more trials in both learning phases.

With the introduction of a described version of the Binary Prediction Task, researchers were able to investigate the role feedback played in participants' behaviour. However, little research has investigated the flip question: what role does describing the contingencies (rather than experiencing them through feedback) play in strategy choice? The existing work investigating feedback tends to provide a description of the contingencies in all cases and varies

feedback across conditions. When feedback is not being investigated, however, researchers often do not describe the contingencies, and participants must learn them solely from the feedback provided to them on the accuracy of their guesses. By contrast, some work provides only the described contingencies, and no feedback.

No one has tested the impact of describing the contingencies. The question is an interesting one because description provides a discrete moment in time under which Strategy Selection can take place and allows participants to enter the task with a clearly formulated top down strategy. I expect, then, that description would have a substantial impact on Strategy Choice, with less impact on Strategy Selection. But how would impacting Strategy Choice look different from impacting Strategy Selection?

First, I expect that impacting Strategy Choice would lead to more discretely recognizable strategies. Since participants are immediately aware of the contingencies of the die they can select *a priori* the strategy they wish to use. To the extent that choosing a maximizing strategy relies on careful deliberative consideration of the problem (as suggested by Koehler and James, 2010) we might also expect that Strict Maximizing would be more common when all information is succinctly provided in a descriptive problem. Thus more immediate selection of a Strict Maximizing strategy should occur when description is provided, if the description actually impacts Strategy Choice. With respect to Experiment 2.3, then, I would expect those who have the problem described to engage in more Strict Maximizing than those that receive no description.

To the extent that this top-down strategy inhibits bottom-up effects, I also expect description to inhibit the positive effects of feedback among non-maximizers. If, for example, a

top-down strategy choice, such as Probability Matching, requires guessing red I would expect that the better formulated this top down strategy is the better participants would resist bottom up influences encouraging green responding. Note that this does not necessarily mean there would be no impact of feedback when participants receive description, but it would be reduced.

Experiment 2.3 examines these issues by looking at the role description plays in influencing participants behaviour on the die problem used in Experiments 2.1 and 2.2. In Experiment 2.3 I vary both whether or not the problem is described *and* whether or not participants receive feedback in order to ascertain what effect both description and feedback might have on Strategy Choice and/or Implementation.

Further, in Experiment 2.1 we found that while feedback tended to increase optimal responding participants did not seem to be aware of this increase. We measured their awareness by asking them to estimate how often they had guessed green and found they tended to underestimate more with feedback. We attributed this underestimation to the increase in green guessing caused by feedback, as raw estimates between conditions did not differ. Experiment 2.3 aims to replicate this effect. Further, we can examine whether participants are aware of the impact description has on their choices. Since participants are supposed to be aware of changes to Strategy Choice I would expect participants to have more accurate estimates when they receive a description. Further, in Experiment 2.1 I claimed that inaccuracies in participants' estimates were caused by the impact of Feedback on their behaviour. If the impact of feedback is actually reduced by having a description, this would also lead to more accurate estimates.

I also investigate the question of whether feedback produces sustained benefit to optimal responding by using a ranking method for investigating transfer. In Experiment 2.3, rather than

asking participants to answer a similar problem to the one they learned on, I ask them to rank which of 4 different strategies is most likely to earn them the most money. This task has the advantage of (a) being easier because it is based on recognition rather than generation and (b) being sensitive to endorsement of hybrid strategies such as overmatching. All participants received a description of the die prior to answering this question. In hindsight, it would have been preferable to withhold the description so that participants in the feedback only group rated the strategies based on their experience with the task as the description removed the manipulation difference between conditions. However, because all participants receive a description prior to this question, I would not necessarily expect a difference between conditions.

Finally, participants were asked to report whether their strategy changes across the course of the task. If description impacts Strategy Choice, participants that receive the description should be less likely to report a change in their strategy across the task because they formulate the strategy at the time of description and do not use experience to modify it. However, similar to the results from Experiment 2.1, the additional provision of feedback, while it may improve green responding among non-maximizers, should not contribute to perceived change in strategy. In other words, those given a description should be no more likely to report strategy change than those that receive both a description and feedback. By contrast, participants in the Feedback Only condition may report change because they have to learn the contingencies from experience and develop their strategy throughout the course of the task. In other words, they do not have the benefit of description to help them formulate their strategy in advance.

2.8 Experiment 2.3: Methods

Participants

Two hundred and eighty-five participants completed the study on Mechanical Turk in exchange for \$1 USD. They also received 0.03USD per correct prediction of the die. One participant was removed from the analysis because of a failure to understand the instructions. Six participants were removed for guessing less than 50% green in block 2 of the experiment. A further 10 participants were removed because they indicated that the die had more than 10 green sides or less than 5 green sides.¹⁶

Procedure

Participants played a guessing game identical to the learning phase game presented in Experiments 2.1 and 2.2. Once again, participants predict the outcomes of 100 rolls of a 10-sided die with 7 green sides and 3 red sides. In Experiment 2.3, we varied how participants learned about the number of sides painted each colour on the die. Participants were randomly assigned to one of three conditions. One third of participants were told the die has 7 green sides and 3 red sides (Description Only condition), but receive no feedback on the accuracy of their guesses while playing the game. These participants received the standard description of the die used in the Learning Phase in Experiments 2.1 and 2.2. One third of participants were *not* told how many sides of the die are painted green or red, but they *did* receive feedback on the accuracy of their guesses (Feedback Only condition). These participants had the same

¹⁶ Interestingly, 21 participants indicated that they thought the die had an equal number of green sides and red sides. This is unusually high and must be an artifact of the instructions (when the description of the die was given, but not the contingencies, many people obviously assumed that mean the number of red and green sides were equal). However, of the 21 participants who reported the die had 5 green and 5 red sides, all of them guessed green more than red. Their mean number of green guess in block 2 of the experiment was 38 out of 50. Thus, these participants seemed to gravitate to a match strategy in spite of thinking the die had only 5 green sides.

instructions as those in the Description Only condition except that they were asked to “guess the outcome of a series of rolls of a 10-sided die that has some sides painted green and the rest painted red.” This group must guess at the number of red and green sides on the die from the feedback that they receive during the game. Note that the feedback simply informs participants whether or not their guess was correct or incorrect. They must deduce from this feedback whether a green or a red was rolled. The final group receives both the description and feedback on the accuracy of their guesses (Both condition) and serves as a comparison for the separate effects of description and feedback.

The 100 rolls of the die described above were divided into two 50 trial blocks. Since participants in the feedback only condition must spend some time learning the contingencies of the die, this allowed us to examine participants behaviour after all groups had some opportunity to learn that the die had more green sides than red sides. Thus, the analyses presented in the results section are based on participants second block of 50 trials.

After completing all 100 guesses participants filled out a short questionnaire. The questionnaire asked them to estimate how many times they guessed green in Block 2 of the experiment. Participants in the feedback only condition were also asked how many sides of the 10-sided die they thought were green. As a measure of whether our manipulations affected strategy, participants were asked to rank 4 potential strategies based on which they thought was “the best way to make the most money in the die game you just played”. Before completing this ranking question, participants in the feedback only condition were told the actual contingencies on the die to ensure all participants knew what they needed to know in order to correctly rank the strategies. The strategies described were a Maximizing strategy, a Matching strategy, an Overmatching strategy described as guessing between 80-90% green, and a 50/50 strategy which

involves guessing half green and half red. Strategies were presented in a random order across participants. Finally, participants completed a comprehension question to make sure they adequately understood the structure of the die game and were asked, as in Experiment 2.1, to report whether or not their strategy changed over time. The wording for the die game and all questionnaire questions are provided in full in the Appendix A.

2.9 Experiment 2.3: Results

The results reported here are for Block 2, when all participants had had the opportunity to learn that green was more likely than red. This was important, in particular, for participants in the Feedback Only condition who must view the die as having more green sides than red sides in order to be comparable (in terms of the correct optimal strategy) to the other conditions. Four participants reported that the die had more than 10 green sides and 6 reported that it had less than 5 green sides¹⁷. These participants were removed from all further analysis. An additional 21 participants indicated that they thought the die had an equal number of green sides and red sides. This is unusually high and must be an artifact of the instructions (when the description of the die was given, but not the contingencies, many people may have assumed that meant the number of red and green sides were equal). However, of the 21 participants who reported the die had 5 green and 5 red sides, all of them guessed green more than red. Their mean number of green guesses in block 2 of the experiment was 38 out of 50, which is quite close to the mean of the remaining non-maximizers in that condition (40 out of 50). Thus, the behaviour of these participants does not seem unusually different from the condition as a whole. As a result, their

¹⁷ Participants reporting less than 5 green sides are typically removed because if you feel the die has less than 5 green sides the optimal strategy is to guess all red. This leads to dramatic outliers and causes problems with analysis.

data has been retained in the data set. The remaining participants in the feedback condition reported a median estimate of 7 green sides on the die (mean = 7.5 green sides).

To assess what role description and feedback might have on Strategy Choice, participants were coded as having carried out a Strict Maximizing or Non-Maximizing strategy. Participants were categorized as Strict Maximizers if they guessed green on all 50 guesses in block 2 of the experiment. Overall, 114 participants out of 279 engaged in Strict Maximizing. To determine whether this varied by condition, participants strategy (Strict Maximizing versus Non-Maximizing) was submitted to a Chi Square test of Equivalence with condition as a between participants variable. Strict Maximizing (guessing green for all 50 guesses of block 2) was significantly more common in the Both and Description Only conditions than it was in the Feedback Only condition ($X^2=27.873$, $p<.001$). Table 3.3 shows the rate of strict maximizing by condition.

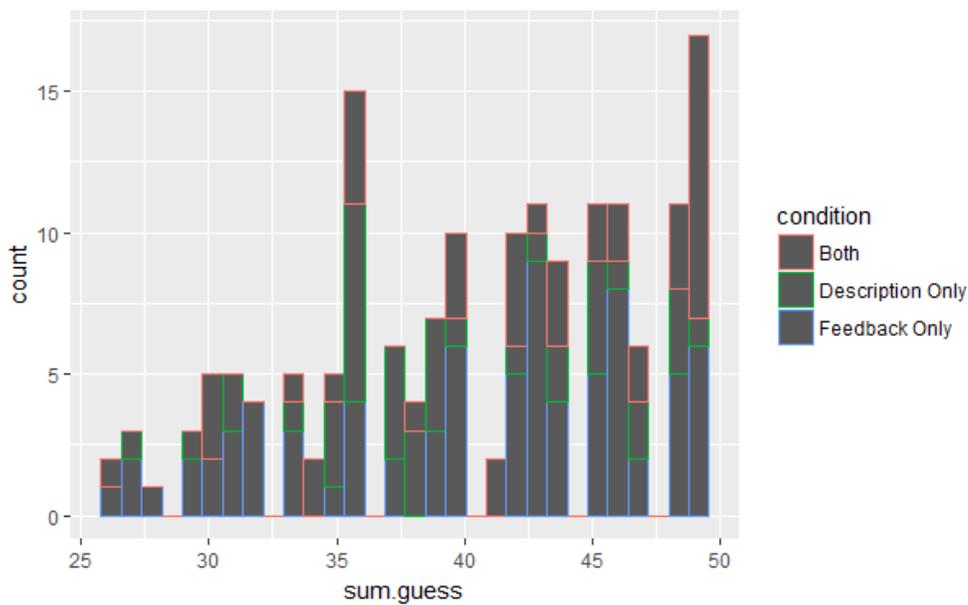


Figure 2.6. Experiment 1.3 Distribution of Total Greens Guessed. Strict Maximizers removed.

	Non-maximizers	Strict Maximizers
Both	45	48
Description Only	42	47
Feedback Only	78	19

Table 2.3. Experiment 2.3 Strategy Choice by Condition (Block 2 data only).

This comparison can also be done by looking at the number of green guesses (out of 50) made in Block 2 of the experiment. Number of greens guessed (out of 50) was submitted to a One-way Analysis of Variance (ANOVA), with condition (Both, Description Only, Feedback Only) as the between participants variable. This test revealed a significant difference between conditions, $F(1,276)=8.64, p<.001, \eta^2=.06$. To determine the nature of this effect, the number of times participants guessed green was submitted to a series of independent sample t-test with each pair of conditions as the independent variables. Participants in the Feedback Only condition guessed less green than both the Description Only condition ($t(183.95)=2.7, p=.007, d=.40$) and the Both condition ($t(185.56)=4.1, p<.001, d=.59$). The Both condition and Description Only condition were not different from one another ($t(176.91)=1.23, p=.219, d=.18$). This does not replicate the literature or the between participant effect of feedback from Experiment 2.1. However, in Experiment 2.1 we had removed maximizers because we were interested in the effect of learning. I removed Strict Maximizers and submitted the number of greens guessed to an independent T-test comparing the Both condition with the Description Only condition. This provides a direct replication of Experiment 2.1. Once again, the presence of feedback increases

green responding, $t(84.51)=2.05, p=.043, d=.44$. However, the massive difference between the Feedback Only condition and the other two conditions disappears (Feedback vs. Both condition: $t(93.58)=1.31, p=.194, d=.25$, Feedback vs. Description Only condition: $t(96.8)=-.923, p=.358, d=.17$). In other words, the large increasing in green guesses seen in the Description Only and Both conditions seems to come primarily from an increase in strict maximizing which presumably occurs by modifying strategy selection at the beginning of the task. With Strict Maximizers removed, description does not appear to increase green guessing among Non-Maximizers. Note, however, that I predicted that Non-maximizers in the Feedback condition would actually do better than those in the Both condition, because the stronger top-down strategy afforded them by the description would compete with the beneficial effects of feedback. I did not observe this. The Feedback Only group was not significantly different from the Both group, but to the extent that they differed, those received only feedback did worse than those that received both feedback and a description. The differences in mean number of greens guessed with maximizers removed can be seen in Figure 2.8. The same differences with maximizers remaining in the sample are presented in Figure 2.7

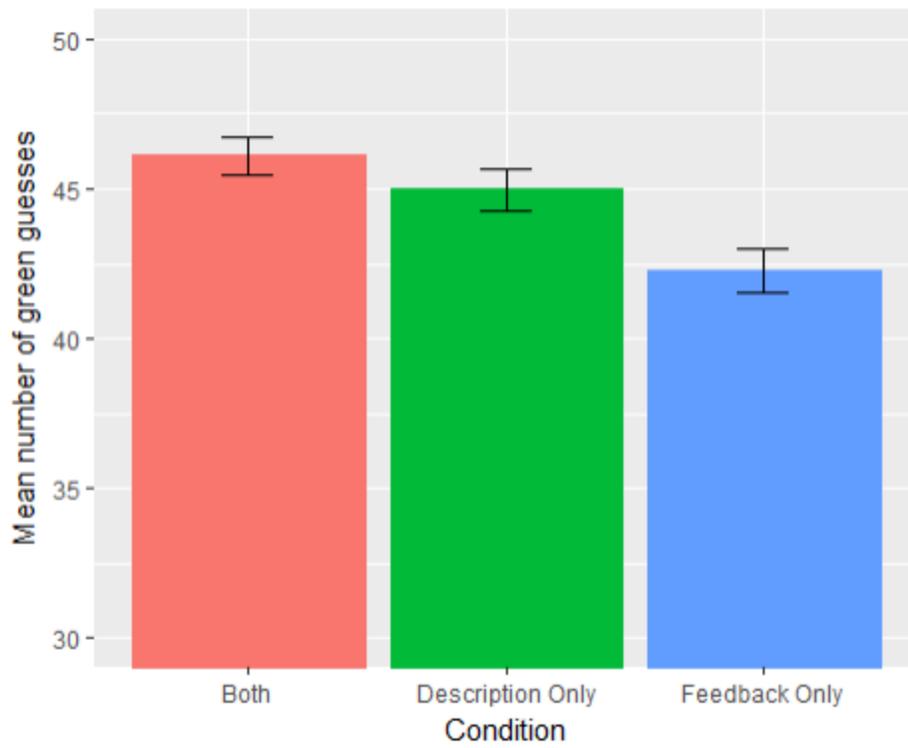


Figure 2.7. Experiment 2.3 Mean number of Greens Guessed by Condition, Strict Maximizers included. Error Bars are 1 Standard Error of the Mean.

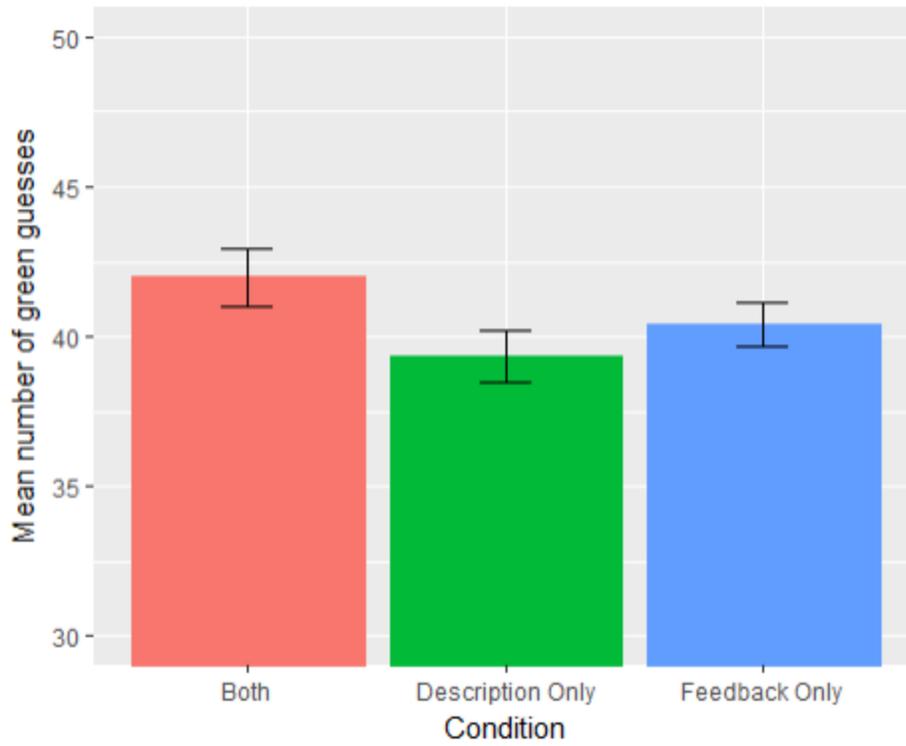


Figure 2.8. Experiment 2.3 Mean number of Greens Guessed by Condition, Strict Maximizers removed. Error Bars are 1 Standard Error of the Mean.

		Mean	SD	Median
Strict Maximizers	Description Only	44.97	6.57	50
Included	Both	46.12	6.01	50
	Feedback Only	42.27	7.04	44
Strict Maximizers	Description Only	39.33	5.57	38
Removed	Both	41.98	6.45	43
	Feedback Only	40.38	6.59	42.5

Table 2.4 Experiment 2.3 Central Tendency Statistics.

In addition to the number of greens guessed, we also tested participants' *estimates* of the number of greens guessed. This prediction is measured here using the same difference score calculated in Experiment 2.1: participants' actual number of green guesses is subtracted from their estimated number of green guesses.

I predicted that participants that received feedback on the accuracy of their behaviour would have less accurate estimates than those that do not, because feedback systematically increases green guessing without participants' awareness. As Strict Maximizers tend to have highly accurate estimates, this test was done on Non-Maximizers only (as was the case in Experiments 3.1 and 3.2). Participants' difference scores were submitted to a One-way Analysis of Variance (ANOVA) with condition (Both, Description Only, Feedback Only) as a between-participants factor. The difference in estimates between conditions was not significant ($F(2,162)=2.04, p=.133$). Once again, all participants underestimated how often they guessed green. Only the difference scores from those in the Description Only condition include zero within the 95% confidence interval (see Figure 2.9) suggesting that only this group could be considered to have accurate estimates. This is further evident when looking at the median response in each of the conditions (see Figure 2.10). The median difference scores for participants in the Description Only condition is zero, participants in the Both condition have a median score of -1 and this increases to -2.5 for the Feedback condition. To test this difference in medians, I submitted median difference scores to a Kruskal-Wallis one-way analysis of variance with condition as the between-participant factor. The test reveals that these are significantly different from one another ($X^2(2)=12.77, p=.002$). The same Kruskal-Wallis test with participants' median raw estimates as the dependent variable showed only a trend towards a difference in raw estimates across conditions ($X^2(2)=4.37, p=.112$). (The simple effects tests on

medians are present in Appendix B). This suggests, as in Experiment 2.1, that the negative difference scores are driven largely by increases in green guessing in conditions with feedback rather than differences in the estimates themselves. Note then, that the median results do support the hypotheses that (a) underestimation is driven by the presence of feedback and (b) as a result, the description condition tends to be more accurate, than the other two. However, estimates are more accurate in the both condition than in the feedback only condition (see Figures 2.11 and 2.12). This suggests that description tends to improve the accuracy of estimates. However, this should be considered weak support, as the tests done on means, while showing generally the same pattern, do not reach significance. Mean and Median raw estimates and green guesses among Non-Maximizers can be seen in Figures 2.11 and 2.12.

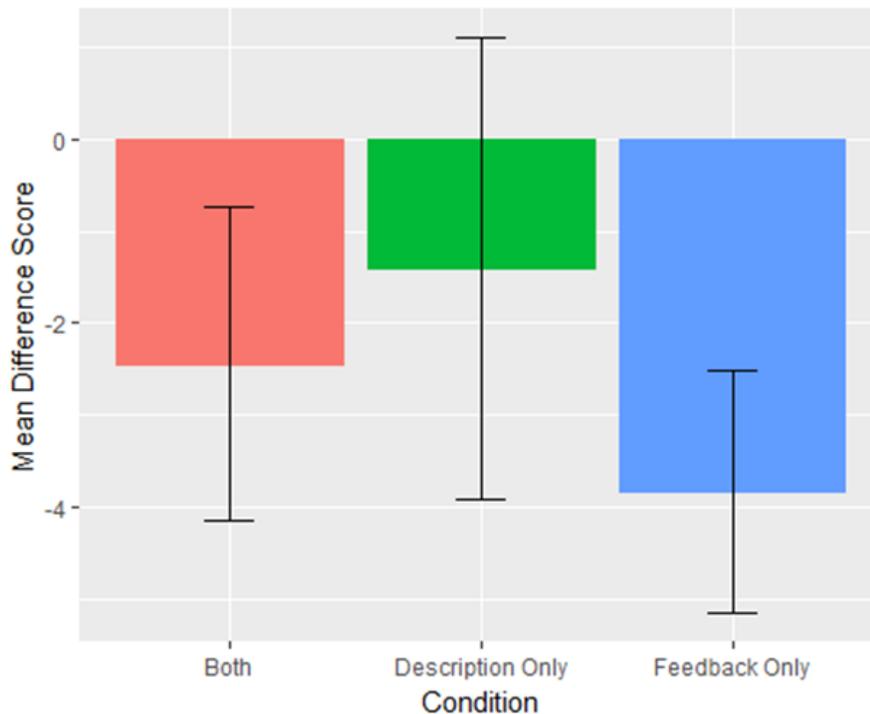


Figure 2.9. Experiment 2.3 Mean Difference Scores across Conditions (Strict Maximizers removed). Error bars are 95% Confidence Intervals.¹⁸

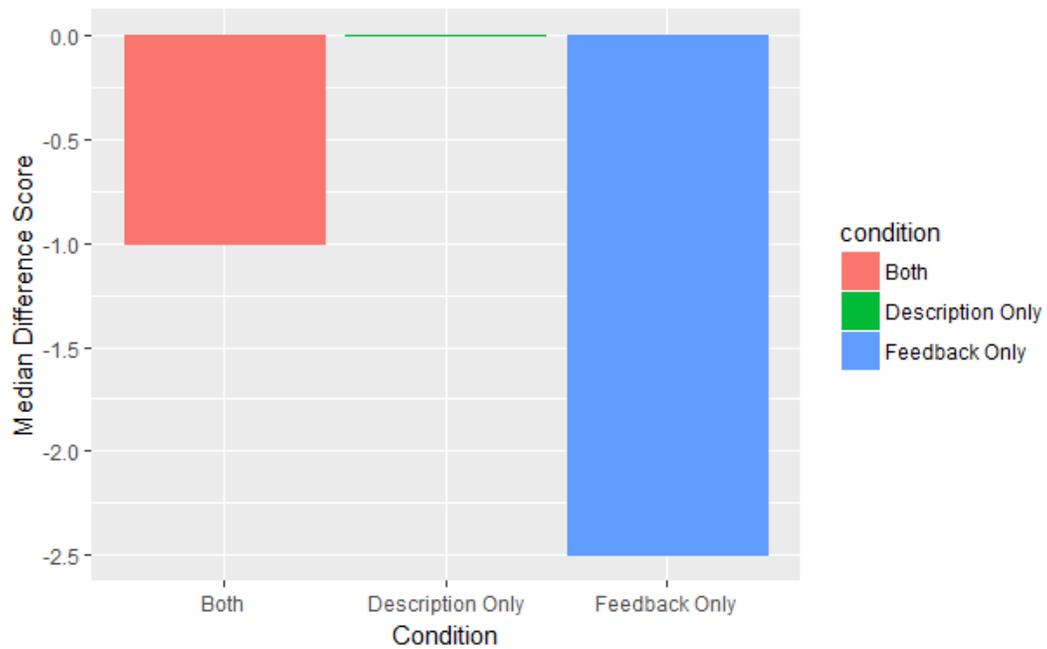


Figure 2.10. Experiment 2.3. Median Difference Scores by Condition (Strict Maximizers removed).

¹⁸ Note that I do not use the conventional 1 SE of the mean for the error bars in this graph, because it is useful to see whether or not the 95% confidence interval for the difference scores in each condition includes zero.

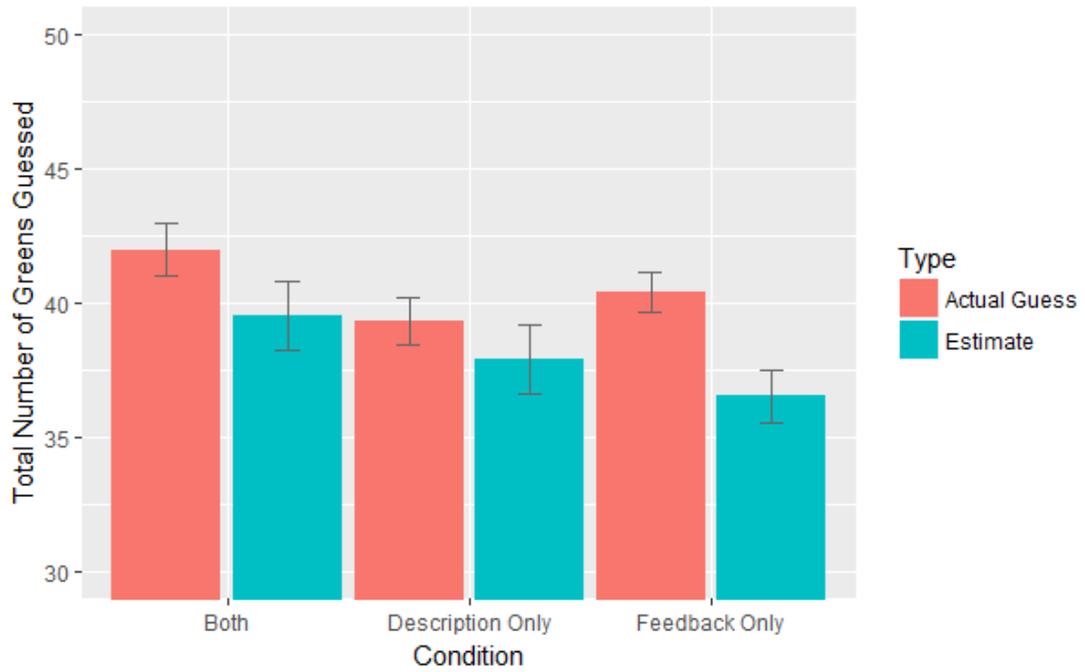


Figure 2.11 Experiment 2.3 Mean Estimates and Green Guesses by Condition (Strict Maximizers removed). Error bars are 1 standard error of the mean.

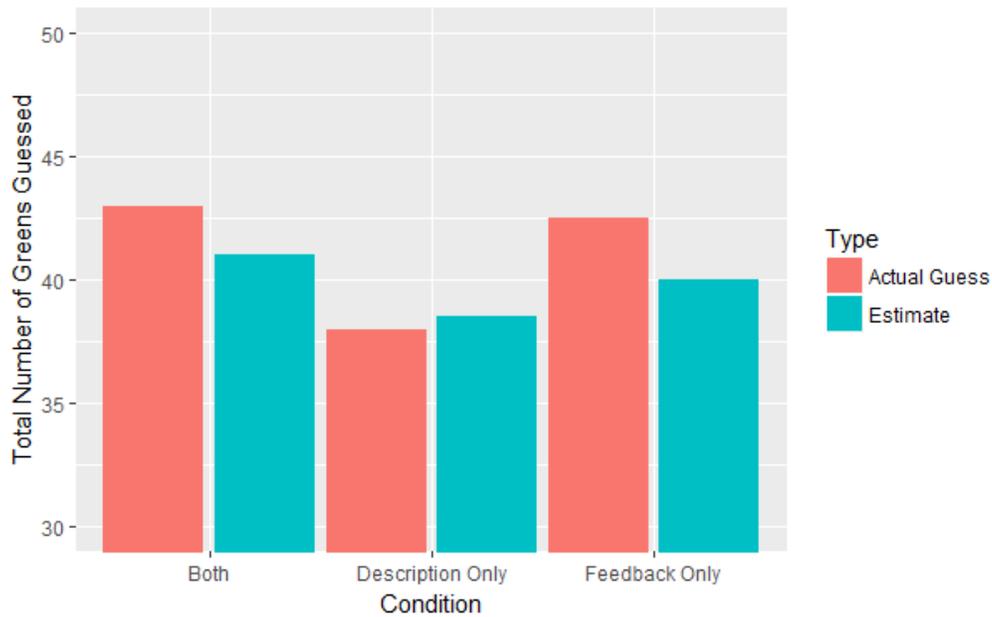


Figure 2.12. Experiment 2.3 Median Estimates and Green Guesses by Condition (Strict Maximizers removed).

Upon completing the estimation task participants were asked to rank order 4 different strategies based on which would earn them the most money. This task was included to discern whether the presence of description and/or feedback improved participants' ability to recognize the optimal strategy after completing the task. As a general metric of more optimal responding (not a measure of 1 particular strategy), I wanted a composite score to index the degree to which participants rankings favored strategies with more green guesses. Participants rank four different strategies (Maximizing, Overmatching, Matching and 50/50). I computed a Rank Index Score that starts with participants first place ranking and compares it to their second place ranking. If the first place strategy involves guessing more green than the second place strategy they receive a score of 1 for that comparison, otherwise they get a score of 0. Then I compare their second place strategy to their third place strategy, and their third place strategy to their fourth place

strategy using the same procedure. If participants ranked their strategies in descending order from most green to least green they will receive a score of 3. If their ranking was exact opposite their score will be 0. The score is supposed to index the degree to which participants favor strategies with more greens.

Rank Index Scores were submitted to a Chi Square Test of Equivalence with condition as a between participants factor.¹⁹ The test revealed no significant difference between conditions, $X^2(4) = 4.11$, $p = 0.391$. This is in spite of the large difference that participants showed in rates of maximizing during the choice task. Note, however, that the ranking task fully describes the die, thereby providing a description based problem to participants who are in the Feedback Only condition. This likely explains the lack of difference.

To examine the effects of learning from feedback on Rank Index Scores, I excluded Strict Maximizers from the analyses and reran the Chi Square test. With maximizers removed, there was a large difference between conditions, $X^2(4) = 26.37$, $p < .001$. The Feedback group had significantly higher Rank Index Scores than either of the other two groups (Feedback vs. Both condition: $X^2(2) = 10.57$, $p < .005$, Feedback vs. Description Only condition: $X^2(2) = 22.32$, $p < .001$), but the Rank Index Scores do not differ between the Both versus Description Only conditions ($X^2(2) = 3.6$, $p = .165$). The higher scores in the Feedback Only condition likely reflect the participants in the Feedback condition that switch to an optimal strategy when they are provided with a description. The lack of difference between the Both and Description Only conditions suggests that it is not the result of Feedback per se. I will discuss this effect further in the discussion. Participants mean Rank Index Scores by condition are presented in Figure 2.13.

¹⁹ Three participants were removed from all rank score analyses that had Rank Scores of zero. This was done because there were not enough participants with these scores to fill the cells necessary to run the Chi Square test. Seven participants were removed from this analysis because they did not answer the ranking question.

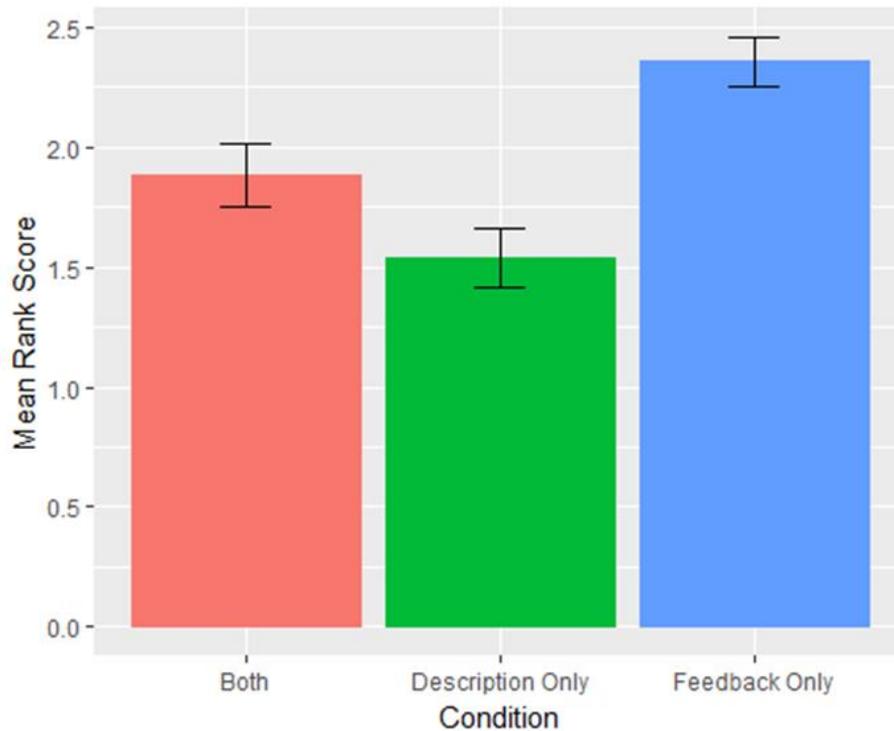


Figure 2.13. Experiment 2.3. Mean Rank Index Scores across Conditions (Strict Maximizers removed). Error bars are 1 Standard Error of the Mean.

Finally, I examined what impact both feedback and description had on participants' perception of their strategy change across the course of the 100 learning trials. Participants responded either yes or no. I hypothesized that participants that received descriptions would report less strategy change than those that had to learn the contingencies through feedback. To test this hypothesis participants yes or no responses to the question "When you were playing the die game, did your strategy change over time?" were submitted to a Chi Square Test of Equivalence with condition as a between participant factor. The test produced a significant result, $X^2(2) = 14.67, p < .001$. Further, Chi Square tests revealed that there is no significant difference between the Description Only and Both condition ($X^2(1) = 1, p = .318$), but that

participants who received Feedback Only are significantly more likely to report strategy change than either of the other two conditions (Feedback vs. Both condition: $X^2(1) = 7, p=.008$, Feedback vs. Description Only condition: $X^2(1) = 12.67, p<.001$). Similar to the rank order results, this suggests that the absence of a description leads to more strategy change across the course of the Binary Prediction Task. The number of participants reporting strategy change in each condition can be viewed in Table 2.5.

	Change Reported	No Change Reported
Both	20	73
Description Only	14	75
Feedback Only	38	59

Table 2.5. Experiment 2.3 Rates of Self-Reported Strategy Change by Condition.

2.10 Experiment 2.3: Discussion

Overall, the results from Experiment 2.3 suggest that the benefit received from feedback is minor compared to the gains participants experience from description. Experiment 2.3 found that conditions that included description lead to significantly more Strict Maximizing than the condition without description, regardless of whether or not feedback was also present.

Experiment 2.3 also replicates the findings from Experiments 2.1 and 2.3 that feedback increases selection of the green response relative to no-feedback among Non-Maximizers, but this effect is small and appears not to have a large impact on Strategy Choice. Thus, these results suggest that description tends to impact Strategy Choice leading participants to favour a Maximizing strategy.

Feedback, by contrast, seems to impact Strategy Implementation by encouraging a minimal increase in selecting the green response.

This claim is further supported by participants' responses to the question "did your strategy change over time?" Participants that received a description reported less strategy change than those that did not, suggesting that description facilitates an initial stable strategy choice. Once again, I also replicated the finding from Experiment 2.1 that among conditions with experience, feedback has no effect on self-reported strategy change. This suggests that while feedback increases green guessing, participants do not explicitly view this as a change of strategy.

Finally, I replicate the finding from Experiment 2.1 that participants underestimate how often they guessed green in the die prediction task. While this did not differ significantly across conditions, those in the Description Condition were the only ones with accurate estimates. Furthermore, among Non-Maximizers, participants' median difference scores supported the idea that feedback lead to greater underestimation while description lead to greater accuracy.

Being asked to rank four candidate strategies (Maximizing, Probability Matching, Overmatching and Guessing 50/50) seemed to level the playing field across all 3 conditions. Now those with feedback did equally well to those without feedback. This increase in performance may have been related to the presence of a description of the problem in the ranking question. This description now meant that those in the Feedback Only condition had access to the same description provided to the other two groups. With maximizers removed from the sample, the Feedback group showed a substantial advantage over the other two groups. At first glance, this appears to suggest that there is a description cost to Non-maximizers. However, this advantage is likely just an artifact of the initially low rate of Strict Maximizing in the Feedback

Only group. There are a number of participants in that group which would endorse a Strict Maximizing strategy when provided with a description. However, during the die task, they have no description and therefore do not use such a strategy. They are, therefore, classified as Non-Maximizers. However, when they encounter the question asking them to rank strategies, they now receive a description of the die and the problem. These participants now rank the strategies in an optimal fashion and this lead to better performance among those in the Feedback Only group.

The Rank question is an alternate way of measuring strategy transfer to the Transfer Task used in Experiments 2.1 and 2.2. Unfortunately, the presence of a description in all groups serves to inhibit its effectiveness as a measure of strategy transfer here. A future replication of this work should refrain from providing a description as part of the ranking question. This would be useful because if description truly impacts Strategy Choice leading to more optimal behaviour among those that receive a description, we would expect it to transfer to the Rank Question leading those participants to also have higher Rank Scores. This is because manipulations impacting Strategy Choice are available to awareness. Unfortunately, the presence of a description for all groups in the Rank question used here prevents me from accurately testing transfer.

I argued that description is an example of a manipulation that impacts Strategy Choice. Experiment 2.3 provided two pieces of evidence in support of that claim. (1) Description lead to a marked increase in Strict Maximizing and (2) that increase seemed to occur at the beginning of the task (during Strategy Choice) rather than occurring from some insight part way through (as is evident from the Strategy Change question). But this does not answer the obvious question of why description facilitates Strict Maximizing. One possibility is that description facilitates the

generation of different candidate strategies. Koehler and James (2010) argued that participants often fail to Maximize simply because they do not generate it as a candidate strategy and Probability Matching comes more readily to mind. When participants are aware of the contingencies of the die in advance, they can use that information to brainstorm strategies. When they learn the contingencies of the die through experience, they may be more myopic in their focus and may not consider which strategy would lead to the overall greatest gains. In further support of this idea, Braveman and Fischer (1968) showed more maximizing when participants were told not to think about being correct on every trial and rather to think about the task as a single problem. Description may inherently encourage a “single problem” approach, whereas trial x trial feedback in the absence of description may encourage the opposite.

Another possibility is that learning the die contingencies from experience (as required by those in the Feedback Only condition) taxes working memory. This taxation may distract participants from producing the optimal strategy. However, previous work suggests that working memory load increases maximizing (Wolford *et al*, 2004) or has no effect at all (Otto *et al* 2011). This work, however, fails to distinguish between Strategy Implementation and Strategy Choice. In the next chapter, I will argue that working memory load increases guessing green via Strategy Implementation, but that it may actually increase Selection of a Matching Strategy. If this latter claim is true, then increased working memory load among those in the Feedback Only condition, may decrease selection of a Maximizing Strategy.

In summary, Experiments 2.1 through 2.3 suggest that when participants receive feedback on the accuracy of their guesses, this tends to increase selection of the more probable option. However, this finding tends to be limited to Non-Maximizers as feedback does not seem

to encourage more uptake of a Strict Maximizing strategy. Thus, it seems likely that the beneficial impact of feedback is limited to Strategy Implementation.

Further evidence for this claim comes from the general finding that feedback does not seem to produce any sustained benefit to performance. While participants do guess more of the more probable colour when provided with feedback, this effect does not transfer to future encounters with the task, even when it is nearly identical (Experiment 2.1 and 2.2). In fact, it may disappear as soon as the Feedback is removed (Experiment 2.2). Furthermore, participants may be completely unaware that they are guessing more green when feedback is present (Experiment 2.1 and 2.3), although this finding requires more testing.

Taken together, these results suggest that the beneficial effects of feedback come largely from Strategy Implementation, at least for Studies with less than 150 trials. While longer term tests should be done, we need to remain cautious about claiming that rationality is reinstated with proper feedback (Shanks, Tunney & McCarthy, 2002) or even arguing that feedback improves performance through deliberation and reflection (Newell & Rakow, 2007). These processes suggest involvement of Strategy Choice, and I found very little evidence to suggest feedback improves Strategy Choice.

The work reviewed here does not provide conclusive evidence as to what mechanism may lead feedback to temporarily increase green guesses. While conditioning remains a good candidate, further work is needed in this area. Furthermore, the work here was limited to relatively short numbers of trials (less than 150). Newell *et al.* (2013) suggest that there may be different impacts of things like description and feedback depending on the number of trials participants receive. Perhaps the impact of outcome feedback on Strategy Choice only occurs

after a much larger number of trials. Regardless, the question of how, when and why information like outcome feedback affects Strategy Choice is a fascinating venue for future research.

Finally, the results from Experiment 2.3 suggest that description drastically increases rates of Strict Maximizing. As this effect is immediate and sustained description likely modifies participants Strategy Choice. This is further supported by the suggestion in the data that description tends to improve the accuracy of participants estimates and that participants experience little strategy change across the task when provided with an initial description. This also provides support for the claim that description influences Strategy Choice. Regardless, of what mechanism is impacted, it is certainly clear that description and feedback are *not* equivalent ways of presenting a Binary Prediction task.

Chapter 3: The Role of Implementation Effort in Maximizing Behaviour

3.1 Introduction

In Chapter 2 I investigated what role feedback might play in participants' behaviour during a Binary Prediction Task. I argued that feedback appeared to be one manipulation that impacted Strategy Implementation rather than Strategy Choice. In Chapter 3, I wish to investigate a second mechanism that appears to impact Strategy Implementation without impacting Strategy Choice: implementation effort.

If we return to our two classic strategies, Probability Matching and Maximizing, it is immediately evident that these two strategies are not equally easy to implement. Maximizing, for example, typically involves repeating the same response or action over and over. This is so trivially simple that you can imagine watching a movie while you executed this strategy and still Maximize perfectly. Probability matching, by contrast is quite a bit more effortful. First, a Probability Matcher needs to switch between 2 or more responses. Second, they need to keep track of roughly how often they have chosen each response. This latter job could be especially difficult if the Probability Matcher feels they need to alternate their responses in some semi-random way (rather than following a specific rule such as “first I guess 7 green, then I guess 3 red – repeat”).

Until very recently, differences in the effort taken to implement Maximizing as opposed to Probability Matching have been largely ignored. However, this distinction is particularly important in situations that tax resources. Under these conditions, more effortful strategies are likely to be abandoned. One obvious example of such a condition is when participants are placed under working memory load while doing a Binary Prediction Task. This is not

commonly done, but has recently been used as a useful method for determining between two potential accounts of why non-maximizing strategies (particularly Probability Matching) occur.

Specifically, these two accounts are known as the Pattern Search Account and the Expectation Matching Account (these accounts and their predictions are explained in more detail in Koehler and James (2014)). According to the Pattern Search account (Gaissmaier & Schooler, 2008; Green et al. 2010; Peterson & Ulehla 1965) Probability Matching results from a misapplication or overgeneralization of a typically useful tendency to search for patterns or predictability in outcomes in the world. When participants encounter the Binary Prediction Task, then, they tend not to believe the sequence is truly random and think, therefore, that they can improve on the predictive accuracy of a maximizing strategy by finding a predictable outcome pattern.²⁰ By the Pattern Search Account, then, Probability Matching is seen as a cognitively rich, or smart, strategy.

By contrast, the Expectation Matching Account views probability matching as a dumb strategy resulting from an intuitive error (James & Koehler, 2011; Kohler & James 2009). This theory relies on the distinction between intuitive and deliberative processes made in the influential Dual-Systems theory (Kahneman & Frederick, 2002; Sloman, 1996; Stanovich & West, 2000; for a review, see Evans, 2008). Probability Matching is suggested to result from an intuitive error that results when participants answer the question (in the standard die problem

²⁰ Although, this lack of belief in the randomness of the sequence may not even be necessary. Participants may believe that the pattern is random, but have a poor understanding of what a random pattern looks like, or indeed, what random even means (Falk & Konold, 1997). These participants, then, may interpret the true randomness they observe during the task as pattern information and believe that even random events can be predicted with some accuracy. This is supported by the strong correlation between a tendency to Probability Match and endorsement of the Gambler's fallacy (West & Stanovich 2003).

presented here) “How many greens and reds do I expect?” instead of the actual question “How many greens and reds should I predict?” This is an example of the more general heuristic attribute substitution in which participants substitute the answer to a hard question with the answer to an easy question (see Griffin et al. 2012 for a review of Judgemental Heuristics and a discussion of attribute substitution). Participants’ reliance on their expectation of the number of greens and reds gives rise to the name of the account: Expectation Matching.

Both the Pattern Search and Expectation Matching accounts are backed by a variety of evidence. For example, the Pattern Search account receives support from the finding that Probability Matching tends to be associated with complex self-reported strategies and rules (Unturbe & Corominas 2007, McMahon & Scheel 2010) and that Probability Matchers are more likely to detect the existence of actual patterns embedded in the sequence they were told was random (Gaissmaier & Schooler 2008). Expectation Matching, by contrast, is supported by the finding that measures of intelligence (West & Stanovich, 2003; Stanovich & West 2008) and deliberation (Koehler & James 2010) tend to be correlated with selecting a Maximizing strategy. Further, when participants play the game in conditions designed to decrease the generation of an expectation (each trial of the game is individuated) rates of maximizing dramatically increase (James & Koehler 2011).

Of particular interest here, however, is the directly opposing predictions that the two theories make about a third piece of evidence: the impact of working memory load (WML) on behaviour in the Binary Prediction Task. Since the Pattern Search Account assumes that Probability Matching arises from an effortful search for patterns, it assumes that WML will *decrease* the rate of Probability Matching. By contrast, the Expectation Matching account characterizes Probability Matching as an intuitive error. As the default in the dual systems

literature is that WML ought to increase reliance on intuition (Evans & Stanovich, 2013), then Expectation Matching predicts that Probability Matching should *increase* under WML.

The current literature is somewhat mixed on what impact WML actually has on strategy in the Binary Choice Task. Wolford, Newman, Miller and Wig (2004) demonstrated that giving participants a dual task that competed for left-hemisphere resources resulted in a decrease in probability matching, consistent with the Pattern Search Account. However, we failed to replicate this effect in our own lab (unpublished). Otto, Taylor and Markman (2011) also failed to find any difference in strategies at all. Neither set of data support the Expectation Matching Account and only the former supports the Pattern Search Account.

Note, however, that suggesting that WML impacts Strategy Choice is somewhat odd when we dissociate Strategy Choice from Strategy Implementation. As an imperfect rule of thumb, WML tends to have its impact while a participant attempts to implement their strategy, rather than when they select it, so the impact of WML on Strategy Choice would seem to be nil. Of course, this distinction is not at all perfect. If participants are learning the contingencies from experience, as in Experiment 2.3, Strategy Choice must occur concurrently with Strategy Implementation. Further, even if the contingencies are described, participants may amend their strategy during implementation. Nevertheless, the observation that WML ought to impact Strategy Implementation remains valid and it is particularly relevant when one of the strategies is more difficult to implement than the other, as is the case with Probability Matching relative to Maximizing.

So what predictions would an Implementation effort account make about the impact of WML on behaviour in a Binary Prediction Task? It would predict that variability among those

attempting to implement a Probability Matching strategy should increase, while this would not be the case for Maximizers. This is because Probability Matching requires more effort to implement and is, therefore, disrupted by WML. Note that without any particular mechanism drawing participants towards Maximizing, WML should be just as likely to cause Under matching as Maximizing. However, participants' simple awareness that one option is more likely than the other, feedback, or simple habitual responding may all be enough to garner support for the rule "err on the side of Maximizing".

The assumption in Wolford, Newman, Miller and Wig (2004) is that WML increases maximizing behaviour because it has encouraged selection of a Maximizing Strategy (by inhibiting Pattern Search). Similarly, Koehler and James (2014) argue that WML would increase reliance on intuition thereby decreasing selection of a Maximizing Strategy. However, neither inference is justified without ensuring that observed behavioural effects are not due to implementation effort. WML may have no impact on Strategy Choice and still increase maximizing through Strategy Implementation, simply because Probability Matching is harder to do.

In fact, there is already one study in the literature to suggest that the additional implementation effort associated with Matching may increase Maximizing when under load. Schulze and Newell (2016) varied how difficult Probability Matching and Maximizing were to implement by modifying which keys participants needed to press in order to implement each strategy. They used a die prediction task similar to the one described earlier in this thesis, in which green was the more probable colour. They then manipulated which key participants needed to press to select green. In the fixed colour-key mapping condition, the up arrow key was always assigned to guessing green. In this condition, implementing a Maximizing strategy

required pressing only the up arrow, while Probability Matching involved switching between the up and down arrow. In the variable colour-key mapping condition, the key assigned to green varied so that pressing the up arrow indicated a green guess 70% of the time and participants' must press the down arrow to choose a green guess for the remaining 30% of trials. In this condition, implementing a Maximizing strategy required monitoring your guesses and the key mapping, making it high in implementation effort, while Probability Matching now only required continually pressing the up arrow. Schulze and Newell hypothesized that whatever strategy required key switching would be the strategy that decreased under load and that is exactly what they found. When participants were in the fixed colour-key mapping condition, Maximizing increased under load, but when they were in the variable colour-key mapping condition, it decreased under load.

Schulze and Newell's study offers convincing evidence that WML may increase Maximizing because it is easier to implement than Probability Matching. While it is likely that their manipulation impacts only Strategy Implementation, it is still possible that it somehow influences participants' Strategy Choice. For example, participants may evaluate the effort involved in implementing each strategy and select their strategy accordingly, prior to beginning the task.

Similar to Schulze and Newell's work, Experiment 3.1 tests what impact WML has on participants' guesses in a Binary Prediction Task. However, I *assign* participants to carry out both Maximizing and Matching strategies. By assigning participants to carry out a particular

strategy, I can ensure²¹ that any variance in behaviour during the task is the result of the impact that WML has on participants' ability to implement the strategy assigned to them and not because of some effect on Strategy Choice.

Furthermore, I also assign WML to vary within participant (rather than between), which provides an especially strong test of the impact of load on behaviour. As mentioned above, there is some disagreement in the literature regarding the existence and nature of the impact of WML on rates of guessing the more probable outcome in a Binary Prediction Task. Because WML is manipulated within-participants (something not typically possible without strategy assignment), I have a particularly powerful test of any impact it might have, at least on Strategy Implementation.

The design of Experiment 3.2 gives us the opportunity to test (a) whether WML increases rates of guessing green (as suggested by some of the previous literature) and (b) whether the impact of load varies depending on the strategy participants are attempting to implement. In addition to load, I also manipulated whether or not participants receive feedback on the accuracy of their guesses. Feedback is manipulated between participants, so each individual either completes all trials in all conditions of the experiment with feedback or they do so without it. Including a manipulation of feedback serves two functions. First, it provides an interesting test of whether the effects of Feedback reported in Chapter 2 are acting on Strategy Choice or Strategy implementation. Since strategy is assigned, any impact feedback has on guesses must be on Strategy Implementation, rather than Strategy Choice. Thus, if participants guess more green with Feedback than they do without it we can surmise that Feedback acts, at least partially, on

²¹ Assuming participants are following instructions. Note, however, that even if they do not like the strategy they are assigned to, there is no reason to assume that they would willfully deviate from it more so under load than not under load.

Strategy Implementation. Second, as mentioned earlier, WML on its own should not necessarily move responding towards maximizing. Rather, it should make those assigned to a Probability Matching strategy less accurate in either direction. The addition of Feedback (a common practice in many Binary Prediction Tasks), however, should encourage movement in the direction of Maximizing. Thus, we would predict an interaction where when assigned to Probability Match, the most green guessing will occur for those receiving feedback *and* while they are under WML.

3.2 Experiment 3.1²²: Method

Participants

Two hundred and three adult participants completed the study online using Amazon Mechanical Turk. All participants were located in the United States and had a HIT approval rate of at least 97%. Participants received \$4 USD added to their Mechanical Turk account upon completion of the study. This study received ethics clearance through the University of Waterloo Ethics department and all participants provided online informed consent.

Procedure

Participants played the same die game used in Experiments 2.1 through 2.3. Once again the die had 10 sides with 7 sides painted green and 3 sides painted red. Participants predicted 240 rolls

²² This study is a replication of a previous study conducted by Priya Thakker for her undergraduate thesis, under the supervision of myself and Dr. Derek Koehler. The data presented here includes small modifications to improve on Priya's original work, but otherwise the method is identical.

of the die divided into 4 blocks of 60 trials each. The die was programmed so that all roles were truly random. All the materials and instructions mentioned below can be viewed in the Appendix A.

Experiment 3.1 used a 2x2x2 design in which Feedback (present vs. absent) was manipulated between-participants in an identical fashion to Experiments 2.1 through 2.3. In addition to Feedback, I manipulated 2 within-participants variables: Strategy Assignment (Matching vs. Maximizing) and Working Memory Load (Load vs. No Load). Strategy Assignment involved asking participants to carry out a particular strategy (Probability Matching or Maximizing) for a given block of 60 trials. Participants carried out each of the two strategies for 2 of the 4 blocks. The strategy a participant was asked to follow in a given block was displayed at the top of the screen during that block.

Strategies were explained to participants using the following 2 vignettes:

“Josh played this game before you. Josh always guessed the more likely event, even though he knew he would sometimes be wrong.”

“Carl tried to correctly predict each die roll. Overall the proportion of red and green Carl chose matched the die.”

Each vignette was accompanied by a pictorial description of the strategy which is available, along with the exact instructions, in the Appendix A.

My other within-participant variable was Working Memory Load, which I manipulated using a 3-back task (similar to that used in Schulze & Newell 2016). Participants received 2 blocks with the concurrent 3-back task and 2 blocks without. These were crossed with Strategy

Assignment so that participants completed a Probability matching strategy once under load and once without load and a Maximizing strategy once under load and once without load. The order participants viewed these blocks was counterbalanced across participants.

During a block in which participants were asked to do the 3-back task, participants saw a number between 0 and 9 presented between each roll of the die. Each number was presented for 0.75 seconds. Participants were given the following instructions regarding the numbers presented during these blocks:

“For this block of trials after you make a prediction, you will see a number between 0 and 9 appear in the center of the screen. At random, a probe will pop up asking you to recall the last 3 numbers that came up on the screen. Please enter the numbers in the same order they were presented to you.” After this, participants were provided with an example and asked to be as accurate as possible.

Participants saw four probes at random throughout the experiment. Participants typed in the 3 numbers they thought they had last seen and pressed enter to continue to a new roll of the die.

Finally, participants were asked at the end of the experiment whether or not they had played a similar die prediction task before. This is a standard question I include so that we can remove duplicate participants, but it was not relevant here as participants are assigned to strategy, so I did not have to worry about previous exposure to the task. It will not be discussed further here.

3.3 Experiment 3.1: Results

Four participants were removed because they guessed less than 50% green²³ in at least one of the four blocks (each of which represented a distinct combination of Strategy Assignment and WML). This left a remaining 105 participants in the feedback condition and 94 in the no feedback condition.

Due to an error in the program, I did not collect probe accuracy data, so I was unable to (a) remove participants who appeared not to be doing the dual task or (b) examine the effect of Strategy Assignment on accuracy at the 3-back task. This error is corrected in Experiment 3.2, but for now it is worth noting that in Experiment 3.2 participants were highly accurate at the 3-back task. This was also the case in Thakker's original undergraduate thesis on the topic. Given this, it is reasonable to assume that participants were doing the 3-back task.²⁴

Total number of greens guessed in each of the four blocks were submitted to a mixed model analysis of variance with Feedback (feedback vs. no feedback) as the between participant variable and WML (load vs. no load) and Strategy Assignment (Maximizing vs. Matching) as the within participant variables. The means and SD for all conditions are provided in Table 3.1. The main effect of Strategy Assignment was highly significant and very large

²³ Guessing less than 50% green is a very unusual strategy and typically represents a misunderstanding of the problem or a confusion of error (ex: participants confuse the contingencies of the die). Furthermore, these typically represent such large outliers that they can be problematic for data analysis.

²⁴ It is worth considering the value of the probe accuracy data. It is helpful to (a) ensure participants did the dual task and (b) allow us to detect a dual task cost in the probes (if participants prioritize the Binary Prediction Task). If participants ignored the probe task (the first point) we would expect no impact of our manipulation on Binary Prediction Task performance. If participants prioritize the Binary Prediction Task, and the cost occurred in the probes we would also expect no effect of our manipulation on Binary Prediction Task performance. Thus, if we do find an effect of our manipulation on performance, that suggests that neither of the above was an issue. This, combined with the results of previous and subsequent work suggesting participants do attend to probes, prompted me to include this work in the thesis in spite of the error.

($F(1,197)=1011.38, p<.000, \eta^2=.84$) with participants assigned to a Maximizing strategy guessing significantly more greens than those that were not. This means that my Strategy Assignment manipulation was generally effective in getting participants to carry out the Strategy I wished them to use. There were no main effects of Feedback($F(1,197)=1.75, p=.187, \eta^2=.01$) or WM($F(1,197)=.003, p=.954, \eta^2=0$). There was a significant Feedback by Strategy Assignment interaction ($F(1,197)=11.55, p=.008, \eta^2=.05$) and a marginal Feedback by WM ($F(1,197)=2.95, p=.088, \eta^2=.02$) interaction, but no Strategy Assignment by WM interactions ($F(1,197)=.06, p=.802, \eta^2=0$). There was a significant 3-way interaction ($F(1,197)=4.36, p=.038, \eta^2=.02$).

To examine these interactions more closely, I divided the file based on Strategy Assignment. I then submitted the total number of greens guessed for Maximizing blocks and Probability Matching blocks separately to a mixed model ANOVA, with Feedback as a between (present vs. absent) participants variable and WML (present vs. absent) as a within participants variable. When participants were assigned to a Maximizing Strategy, the number of greens they guessed did not differ based on WML ($F(1,197)=.03, p=.854, \eta^2=0$) or Feedback ($F(1,197)=3.19, p=.076, \eta^2=.02$), and the two did not interact($F(1,197)=2.73, p=.602, \eta^2=0$). That said, the effect of Feedback was marginally significant and it is worth noting its effect is in the opposite direction from typical (see Figure 3.1). In other words, when participants are asked to Maximize, they guess less green with feedback than they do without feedback. This mirrors the results reported earlier in which feedback seems to reduce rates of Strict Maximizing. The pattern of results changes when participants are asked to Probability Match. Here we see a significant effect of feedback ($F(1,197)=8.14, p=.005, \eta^2=.04$) and a significant interaction ($F(1,197)=5.45, p=.021, \eta^2=.03$), but no main effect of WML ($F(1,197)=.038, p=.845, \eta^2=0$).

To break down this two-way interaction further (and the 3-way interaction from above) I took only the Probability Matching trials and divided the file based on Feedback. I then compared trials where participants were under load to trials where they were not underload for both the feedback and no feedback conditions using a paired-samples t-test. This revealed a marginally significant difference ($t(104)=-1.71, p=.091, d=.16$) between the different levels of load for the Feedback group, where being under working memory load increased how often participants guessed green²⁵. For the no feedback group, the effect was not significant ($t(93)=1.63, p=.107, d=.16$), but trended in the opposite direction. Participants who were under load actually guessed less green than those not under load.

	Assigned to Probability Match		Assigned to Maximize	
	Single Task	Dual Task	Single Task	Dual Task
Feedback	45.15(6.04)	46.17(6.95)	58.47(4.39)	58.27(5.22)
No Feedback	43.97(5.27)	43.11(5.3)	59.23(3.03)	59.33(3.29)

Table 3.1. Experiment 3.1 Mean and SD Scores by Condition.

²⁵ This difference is not significant when using the non-parametric Wilcoxon signed rank test to compare the data ($V=1909, p=0.222$).

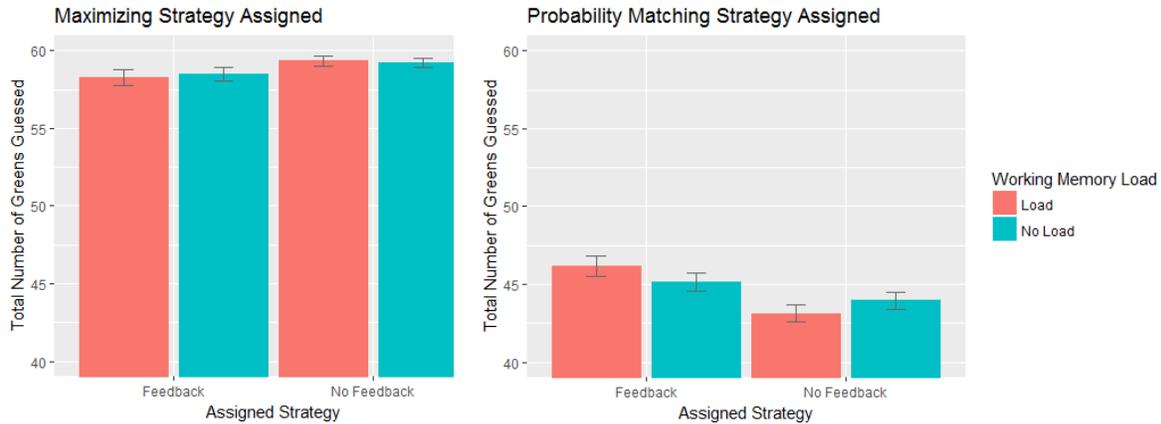


Figure 3.1. Experiment 3.1 Number of Greens Guessed by Condition.

3.4 Experiment 3.1: Discussion

Experiment 3.1 was primarily designed to test the impact of WML on Strategy Implementation. We hypothesized that load would be more disruptive to executing a Probability Matching Strategy than a Maximizing Strategy, and under conditions with Feedback, would lead to an increase in selecting green.

The results weakly confirmed this prediction, but the effect was rather small. WML had no main effects on how often participants guessed green, and its impact was only evident in the predicted 3-way interaction²⁶. Here, participants asked to implement a Probability matching strategy were marginally more likely to overmatch when under load than participants not underload, but this effect was only observed when participants were receiving Feedback. There are two possible explanations for why Feedback might moderate the effect of WML. First, as

²⁶ It is worth mentioning that this 3-way interaction is a replication of a similar but marginal 3-way interaction that was present in the initial work we did on this topic in Priya Thakker's undergraduate thesis. Unfortunately, that work included a poor manipulation of feedback that was corrected here. In Thakker's original work there was also a main effect of WML for Probability Matchers that we did not see here.

suggested in the introduction, WML in and of itself, should not encourage participants to guess the more probable option. It should simply make it harder to guess exactly 70% green.

However, the addition of Feedback does provide a mechanism for encouraging participants to guess green and would bias inaccuracies towards overmatching.

Another possibility relies on a conditioning account. Previous work has shown the impact of conditioning to be goal dependent (Verwijmeren *et al.*, 2004; Ferguson & Bargh, 2004). For example, Verwijmeren, Karremans, Stroebe & Wigboldus 2004, found that conditioning was more effective when it depends on the participants' current goal and also when the unconditioned stimulus was relevant to their goal. Assuming that the mechanism by which Feedback encourages green responding is indeed conditioning, these goal relevant factors may impact our data. Because participants were assigned to carry out a Strategy, their goal was not to correctly predict die rolls, *per se*, but to guess the correct number of greens (and reds). Thus, the feedback that they received on the accuracy of their guesses was not goal relevant. It was not telling them how closely they were adhering to their strategy, but rather whether or not they were correct in their last prediction. In fact, when assigned to Probability Match, attending to feedback actually hampers participants' performance at their goals, because Feedback encourages guessing green. Thus, participants could do better by ignoring the feedback. However, WML may tax executive function and make it difficult to keep relevant goals in mind and to ignore Feedback. The result is that, for the participants who received Feedback, green guessing was more common during the load trials than the no load trials. However, as a Maximizing Strategy is already at ceiling (in terms of guessing green), this effect was only visible when participants were assigned to Probability Match.

Testing this latter account would be somewhat complicated. As a first step, establishing that WML does moderate the occurrence of goal irrelevant conditioning would be necessary. Alternatively, research could focus on establishing the former account by using some other mechanism (besides feedback) to encourage green guessing. If the interaction between this new mechanism and WML persists, that would undermine this goal dependent account.

Besides testing the impact of WML, Experiment 3.1 once again replicated the benefit of feedback for those executing a Probability Matching Strategy. This replication, however, is of particular interest because participants were *assigned* to carry out a Probability Matching Strategy. This has two implications. First, assuming participants attempt to faithfully carry out the strategy to which they were assigned, the effect of feedback must be entirely on Strategy Implementation. While it is possible that participants willfully decide to ignore our Strategy Instructions and do, in fact, choose their own strategy, this interpretation is not supported by the massive effect of Strategy Assignment.

Second, this lends additional support to the claim that the impact of feedback on green guessing is not salient to participants. To the extent that participants are faithfully carrying out the strategy they were assigned, awareness of the impact of feedback should increase their ability to resist its impact. We don't know, however, whether they are deliberately or unintentionally deviating from their assigned strategy. Including an estimation question at the end of the experiment, similar to the one used in Experiments 2.1 and 2.3 would help to address this question. Combined with the results of these earlier experiments, however, it seems the most parsimonious explanation of the evidence is that for Non-Maximizers, Feedback impacts Strategy Implementation and it does so largely without participant awareness.

Feedback also had a marginal effect on those executing a maximizing strategy. When receiving feedback, participants tended to guess less green than when not receiving feedback. This is consistent with the finding from Experiment 2.3 (and hinted at in 2.1) that feedback tends to reduce uptake of a Strict Maximizing strategy.

When participants are assigned to a Maximizing strategy it is hard to believe that participants just don't notice that they are guessing some red and truly believed they guessed all green. Again, an estimation question would help to resolve this. Assuming participants are neglecting their assigned goal and willfully deviating from a Maximizing strategy, I think the best explanation for their behaviour is boredom and, perhaps, skilled prediction (Newell *et al.* 2013). I expect these also impacted the decreased adoption of a Strict Maximizing strategy under conditions with feedback observed in Experiment 2.3 and elsewhere in the literature (Newell & Rakow, 2007).

Finally, as somewhat of a side note, when participants were assigned to a Probability Matching Strategy, the mean response was to over match in all conditions especially those with Feedback. There has been some discussion in the literature about the prevalence of overmatching (Vulcan 2000). In particular, its existence has often been used to suggest that participants are more rational than Probability Matching would suggest (Vulcan, 2000, Shanks *et al.* 2002). Since participants were assigned to Strategy, the prevalence of overmatching here suggests that it is, at least in part, impacted by Strategy Implementation rather than being purely influenced by Strategy Choice. That said, there is also existing evidence to suggest that Overmatching is an explicit strategy chosen by participants (Gal & Barron 1996).

While the results of Experiment 3.1 are interesting, they are somewhat limited by the program error that prevented me from recording participants' probe accuracy. We are not certain, for example, that participants were actually doing the dual task. If they were not, all of our above findings would be invalid. We also cannot measure whether or not there was any effect of Strategy Assignment on accuracy at the 3-back task²⁷.

Thus, Experiment 3.2 aims to provide a partial replication of Experiment 3.1 with the program error corrected allowing accuracy of the memory probes to be recorded. To allow for a smaller sample size, in Experiment 3.2 we removed the between subject manipulation of Feedback, with all participants receiving feedback on the accuracy of their guesses. Thus, only WML and Strategy Assignment are manipulated in Experiment 3.2.

3.5. Experiment 3.2: Methods

Participants

One-hundred and five adult participants completed the study online using Amazon Mechanical Turk. All participants were located in the United States and had a HIT approval rate of at least 97%. Participants received \$4 USD added to their Mechanical Turk account upon completion of the study. This study received ethics clearance through the University of Waterloo Ethics department and all participants provided online informed consent.

²⁷ Note that in the aforementioned undergraduate thesis work I did with Priya Thakker and Derek Koehler, there was a significant impact of Strategy Assignment on probe accuracy. When participants were assigned to a Probability Matching Strategy, they were less accurate on the probes than when they were assigned to a Maximizing Strategy. We took this as evidence that Probability Matching and the 3-back task were more difficult to do concurrently than Maximizing and the 3-back task and that participants prioritized the Dual Choice Task at the expense of the 3-back task. In this work and subsequently in Experiment 3.2, we modified the instructions to emphasize equal allocation of resources.

Procedure

The procedure in Experiment 3.2 was identical to that of Experiment 3.1, with two changes. Feedback was not manipulated, so all participants received Feedback on the accuracy of their guesses. After each prediction they were told whether or not that prediction was correct. Second, participants received 6, rather than 4 memory probes to increase the sensitivity of the memory probes measure.

3.6. Experiment 3.2: Results

Two participants were removed from the data set for guessing less than half green. Eight participants were removed that got zero of the 6 memory probes correct in one or both of the 2 blocks containing the 3-back task and three were removed that did not complete the memory probes at all. Ninety-two participants remained in the data set.

First I examined how accurate participants were at the 3-back task. Participants that correctly recalled the last 3-digits they saw in the correct order when prompted were given a score of 1 for that probe. Each participant's scores were added across all 6 probes for a total score of 6 correct in each of the two blocks containing the 3-back task (one where they were assigned to Probability Match and one where they were assigned to Maximize). In general, participants were highly accurate at the probe task. The mean and median number of probes correct by condition is provided in Table 3.2. Mean number of probes correct was submitted to a paired sample t-test with Strategy Assignment as the within-participants factor. The test revealed no significant difference between conditions, $t(91)=-0.77, p=.445, d=.08$.

Next I examined the impact of Strategy Assignment and WML on participants' rate of green guessing. Total number of greens guessed in each of the four blocks was submitted to a repeated measures analysis of variance with Strategy Assignment (Probability Matching vs. Maximizing) and WML (Load vs. No Load) as within-participant factors. This revealed a large main effect of Strategy Assignment ($F(1,91)=354.28, p<.000, \eta^2=.80$), such that participants guessed more green when asked to Maximize then when asked to Probability Match (see Figure 3.2). There was also a main effect of WML ($F(1,91)=6.38, p=.013, \eta^2=.07$)²⁸. Participants guessed more green when under load than when not under load. However, this is actually only evident when they are asked to Probability Match as can be seen from the significant interaction ($F(1,91)=12.34, p<.000, \eta^2=.12$). To break down this interaction, I submitted the total number of greens guessed in each block to paired sample t-tests with WML (Load vs. No Load) as the within participant factor. This was done separately for blocks where participants were assigned to Probability Match versus blocks where they were assigned to Maximize. When asked to Probability Match, participants guessed significantly more green under load than when not under load ($t(91)=-3.64, p<.001, d=.41$), however this was not the case when they were asked to Maximize ($t(91)=1.04, p=.301, d=.11$).

Assigned Strategy	Mean	SD	Median	Min value	Max Value
Probability Matching	5.05	1.30	5	1	6
Maximizing	5.15	1.19	6	2	6

Table 3.2 Experiment 3.2 Mean and Median Accuracy Scores on the 3-back Task by Assigned Strategy

²⁸ Using a rank transformation to correct for lack of normality, this effect becomes marginal ($F(1,91)=3.46, p=.066, \eta^2=.04$)

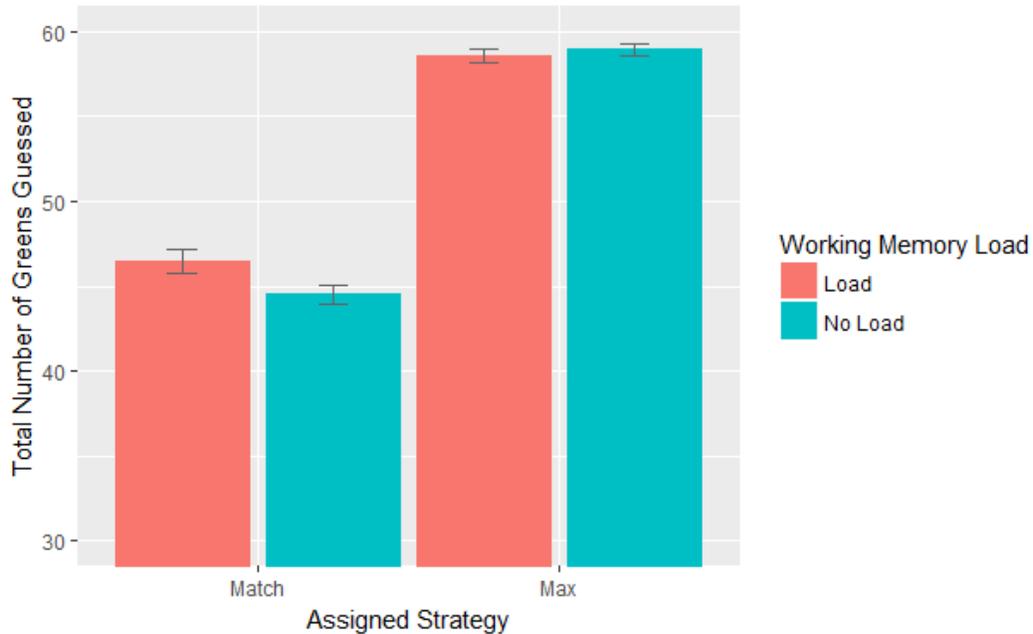


Figure 3.2. Experiment 3.2 Mean number of Green Guesses by Condition.

3.7. Experiment 3.2: Discussion

In Experiment 3.2, I confirmed that participants were doing the 3-back task and were highly accurate. I found no effect of Strategy Assignment on accuracy on the 3-back task, but I did replicate the finding that, under conditions with feedback, when participants are assigned to a Probability Matching strategy they guess more green when under working memory load than when not under load. As in Experiment 3.2, this same difference does not exist when participants are assigned to Maximize.

While the results Experiments 3.1 and 3.2, as well as the work from Schulze & Newell (2016), suggest that Working Memory Load increases Maximizing as a result of implementation effort, they do not test what impact WML might have on Strategy Choice. It is possible the

WML may affect both components in different ways. Indeed, according to the Expectation Matching Account WML should increase selection of a matching strategy because it increases reliance on intuition. However, in most paradigms, this effect on Strategy Choice would occur simultaneously with the opposite effect on Strategy Implementation. This contradiction may make it hard to detect the impact of WML on choice behaviour in a Binary Prediction Task, and may explain some of the null results in the literature.

Experiment 3.3 attempts to investigate what impact a concurrent load might have on Strategy Choice by present participants with a Binary Task that has no implementation component. This is done by providing the standard die problem used in the rest of this work, but in a word problem format. After reading the word problem, participants are simply asked to indicate how many times (overall) they would guess green. Since they supply only a single number and do not make any trial by trial guesses, both a maximizing and matching strategy are trivially easy to implement.

Unfortunately, this word problem format makes it difficult to use the 3-back Working Memory Load task used in Experiments 3.1 and 3.2. Time pressure is used as an alternative. While time pressure may not cause working memory load, per se, it is used to increase reliance on intuition (Rand 2016). As such, if the Expectation Matching Account is correct, Time pressure should decrease green guessing because Probability Matching and Maximizing are equally easy to implement, but the former should be more intuitive.

In addition to the above question, Experiment 3.3 includes a couple of standard measures of tendency towards reliance on deliberation versus intuition. One such measure, the CRT (Frederick, 2005) has been shown to correlate positively with Maximizing behaviour. In other

words, more deliberative people (as measured by the CRT) tend to Maximize. Unfortunately, the CRT has seen a great deal of use and suffers from over exposure (Haigh, 2016; Stieger & Reips, 2016). As such, we use two alternatives, the CRT2 and a set of Belief Bias Problems both of which have been shown to correlate highly with and in similar ways to the old CRT (Thomson & Oppenheimer, 2016). If time pressure truly encourages reliance on intuition, we should see that intuitive errors increase when people are put under time pressure.

3.8. Experiment 3.3: Methods

Participants

One-hundred and thirty nine University of Waterloo undergraduate students completed the study online. Participants received a half credit towards one of the courses as remuneration for completing the study. This study received ethics clearance through the University of Waterloo Ethics department and all participants provided online informed consent. Five participants were removed because of incomplete data. One additional participant was removed for guessing less than 5 green leaving a remaining 133 participants in the data set.

Procedure

Participants' completed a series of thought problems online. The first of these problems was the standard die problem presented in all of the previous studies reported here, except that it was written as a word problem and participants gave an aggregate numeric response.

Participants were asked to consider the following game:

“The game involves rolling a fair, 10-sided die. The die is “fair” in the sense that, when it is rolled, each of the 10 sides has an equal chance of turning up.

On this particular die:

7 sides are GREEN

3 sides are RED

Suppose you are to play the following guessing game: the die will be rolled ten times, and before each roll you are to guess whether a GREEN side or a RED side will turn up. Imagine that you will receive \$1 for each correct guess.”

Participants were then asked to report “Across all 10 rolls of the die, how many times would you guess green?” After providing their response participants answered two comprehension check questions (available in full in the Appendix A). These questions were included to ensure that participants did not answer the previous Die Question less optimally simply because they did not correctly understand it. In particular, we were worried that if people were reading the problem quickly this would hamper their understanding. This could lead to less Maximizing in the speeded condition, but not because of reliance on intuition. The comprehension questions help to ensure this is not an issue.

Both comprehension questions were multiple choice. Question 1 asked participants to report what they had been asked to do in the Die Problem. They then selected between the following 3 options:

“To report how many times you would guess green.”

“To report how many green would be rolled overall.”

“To report how many sides on the die were green.”

Question 2 presented participants with a hypothetical player, Mary, who had just played the Die Game. They saw a grid containing one column with Mary’s guesses and a second column with the die outcome for each roll. They were asked to indicate how much money Mary made.

In addition to the comprehension questions, participants were also asked 10 belief bias problems and an alternate version of the CRT known as the CRT2 (Thomson & Oppenheimer, 2016). The order of presentation of the CRT2 and Belief Bias problems were counterbalanced. The questions within each of these measures were also randomized. These problems were included as measures of tendency to rely on intuition and deliberation. Both problem sets are available in full in Appendix A. Finally, participants were asked whether they had seen any of the problems presented in the experiment before, in order to assess pre-exposure. While participants reported some prior exposure to the CRT2, participants reported very little exposure to the central Die problem considered here, so this will not be considered further.

Participants completed the thought problems above in one of two conditions: Speeded versus Slow. Participants in the speeded condition were asked to “be as fast and accurate as you can. Speed is important, so please answer as quickly as you can.” Participants in the slow condition read “we would like you to try and be as accurate as you can. Speed is not important, so take your time and think carefully about your answers!” Participants’ response times on each of the thought problem questions were measured from the moment the question was presented until the moment they submitted their response. Timing data was used as a manipulation check.

3.9. Experiment 3.3: Results

First, in order to determine if my Speed Instructions manipulation was effective, I recorded how long participants spent on each of the questions in the questionnaire. How long they spent answering the die problem was of particular import. How long participants spent on the die problem was submitted to an Independent Samples T-test, with Speed Instructions (fast versus slow) as the between participants variable. While participants did spend longer in the slow condition (median time of 45.96 seconds) than the fast condition (median time of 36.81 seconds), the difference was not significant ($t(94.74)=-1.10$, $p=0.273$, $d=.19$). This lack of significance likely arises from the large difference in variance between the conditions (see Table 3.3) leading to a large decrement in degrees of freedom. I return to this point in the discussion. Nonetheless, the difference is at least in the correct direction.

Next, Experiment 3.3 included 2 comprehension questions to ensure that speeding participants did not alter their responses to the die problem simply by causing poorer comprehension. Sixteen participants incorrectly answered Question 1 (asking them to identify what they had been asked to do on the die problem question). Twenty-seven participants incorrectly answered Question 2 (where they were asked to tabulate the earnings of a theoretical player that played the same guessing game). Participants' accuracy (Correct versus Incorrect) on the comprehension questions was crossed with condition (Fast vs. Slow) and submitted to a Chi-Square Test of Equivalence. Performance on these questions did not vary as a function of condition (Question 1: $X^2(1) = .001, p=.974$, Question 2: $X^2(1)= 1.07, p=.301$), suggesting that any difference between conditions on the die problem was not do to poorer comprehension in the speeded condition.

Next I examined whether the speeding manipulation had any effect on participants response to the die problem. Overall, 42 out of 134 participants indicated that they would guess green for all 10 rolls of the die in the hypothetical die problem. The full distribution of guesses by condition is provided in Figure 3.3.

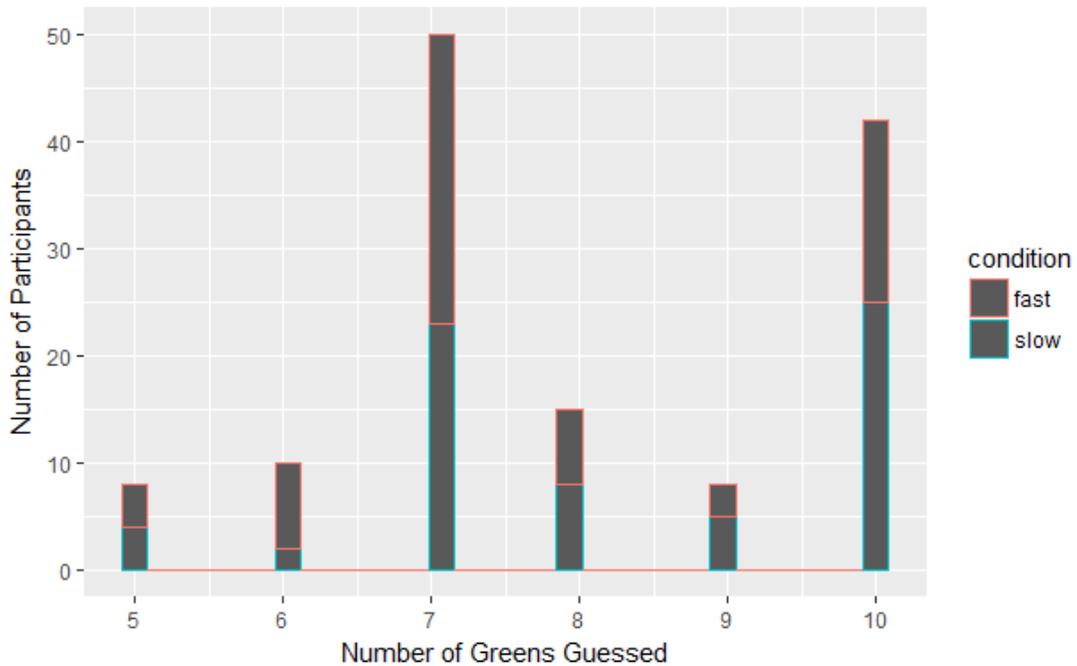


Figure 3.3. Experiment 3.3 Distribution of Responses to the Die problem Question by Condition.

Participants' answers to the die problem were classified into two different strategies. Participants classified as Strict Maximizers said they would guess green on all 10 rolls of the die. Participants classified as Non-Maximizers said they would guess between 5 and 9 greens across the 10 rolls of the die²⁹. As a final test to ensure misunderstanding of the question did not systematically influence strategy, I tested whether strategies differed as a function of accuracy on

²⁹ The one participant that said they would guess less than 5 green was removed from the analysis (see participants section).

the comprehension questions. I ran 2 separate Chi-Square Test of Equivalences in which I crossed Accuracy (Correct vs Incorrect) on each of the comprehension questions with strategy (Strict Maximizing vs Non-maximizing). There was no significant relation between strategy and accuracy in either test ($X^2(1) = 1.39, p = 0.239$ and $X^2(1) = 1.37, p = 0.241$). Given the lack of relation between accuracy on the comprehension question and performance I did not filter or remove participants based on their accuracy on the comprehension question³⁰.

To determine whether my manipulation of time spent had an impact on how participants answered the die problem the number of greens participants indicated they would guess in the Die Problem was submitted to an Independent Samples T-test with Speed Instructions (Fast vs. Slow) as the between participants variable. In spite of the weak manipulation of time pressure, the difference between conditions was marginally significant ($t(131) = -1.86, p = .065, d = 0.32$). Participants who received slow instructions engaged in higher rates of guessing green than participants who received speeded instructions. Central tendency statistics for answers to this question are provided in Table 3.3.

	Mean	Median	SD
Fast	46.74	36.81	34.84
Slow	57.7	45.96	73.22

Table 3.3 Experiment 3.3 Central Tendency Statistics by Condition for Responses time (in seconds) to the Die Problem.

³⁰ Filtering of these participants does not alter the pattern of data reported here.

	Mean	Median	SD	Percentage engaging in Strict Maximizing
Fast	7.73	7.00	1.57	26%
Slow	8.24	8.00	1.60	37%

Table 3.4 Experiment 3.3 Central Tendency Statistics by Condition for Responses to the Die Problem.

Another way of examining how participants answered the die problem was to classify their responses according to strategy. Here the same classification was used as the one described above for the comprehension question. Participants' strategy (Strict Maximizing vs. Non-Maximizing) was crossed with Speed Instructions (Fast vs. Slow) and compared using a Chi-Square Test of Equivalence. No significant difference was found, but participants answers trended in the direction of those that were speeded doing less Strict Maximizing than those that took their time ($X^2(1) = 2.05, p = 0.152$).

I also investigated whether our Speed Instructions manipulation was effective at changing the time it took participants to complete the Belief Bias and CRT2 tasks. If participants were significantly faster in one condition than the other, and if speeding does increase reliance on intuition, then we would expect less items correct in both the Belief Bias and the CRT2 problems in the Fast condition. Time taken to complete the problem set for both the Belief Bias and the CRT2 problems was submitted to a Welch's independent sample T-test, with Speed Instructions (Fast vs. Slow) as the between participants variable. Participants were significantly faster at

completing the Belief Bias problems ($t(101.84)=-2.90$, $p=.005$, $d=0.50$) and marginally faster at completing the CRT2 problems ($t(75.96)=-1.73$, $p=.088$, $d=0.30$) in the Speeded condition³¹.

Given this speed difference, participants should get more items correct in the Slow condition than in the Fast condition. Number of responses correct for each set of problems was submitted to an independent sample T-test, with Speed Instructions (Fast vs. Slow) as the between participants variable. Neither number of Belief Bias questions correct nor number of CRT2 questions correct varied as a function of condition ($t(130.74)=-1.00$, $p=0.315$, $d=0.17$ and $t(130.98)=-1.22$, $p=.226$, $d=.20$ respectively) although, once again, participants tended to get less items correct in the Fast conditions than in the Slow condition.

3.10. Experiment 3.3: Discussion

The conclusions that we can draw from Experiment 3.3 are somewhat limited by the weak time manipulation. The results suggest that with respect to the Binary Prediction Task participants were not significantly faster in the Fast condition than in the Slow condition. However, the difference between the groups was in the right direction and approached significance. The general lack of significance may be related to different variance among the two conditions (see Table 3.3). In general, the default seems to be for all participants to go fast. When asked to slow down only some of the participants in the slow condition actually did. This leads to larger variance in the slow condition than in the fast condition. This problem needs to be addressed before we can say anything definitive about what is going on here.

³¹ This marginal difference was significant ($p=.03$) when using a non-parametric test (see Appendix B).

Similarly, responses to the die problem were also marginal, but they are suggestive. Conservatively, I would argue that at the least there is no effect of speeding on green guessing, suggesting that load does not increase Maximizing when there is no difference in the effort required to implement each strategy. Less conservatively, Experiment 3.3 suggests an effect may exist in the opposite direction, with load actually increasing Non-Maximizing strategies in zero implementation effort environments. This effect is particularly evident in the median response in each condition, where participants in the slow condition were likely to guess 8 green while those in the fast condition only guessed 7. More to the point, the direction of this difference is congruent with the predictions of Expectation Matching and opposite in direction to the findings from Experiments 3.1 and 3.2.

However, this effect isn't evident when participants' guesses are classified according to strategy. Given that Maximizing and Probability Matching are distinct Strategies (and the most common selected in Experiment 3.3) we would expect slowing down would move participants from Probability Matching to Maximizing. It isn't as obvious to me why it would move them from Probability Matching to Over Matching, which seems to be at least partly what is happening.

In general, however, the findings from Experiment 3.3 are weak. First, the speeding manipulation was not as effective as desired. This could be improved with a manipulation that produces less variance in the slow condition. Second, speeding did not seem to increase reliance on intuition on the CRT2 and belief bias problems. This would suggest that it is not increasing reliance on intuition in general. However, it isn't obvious to me what the mechanism by which speeding increased Probability Matching might be if it is not reliance on intuition. It is possible, that the Belief Bias and CRT2 tasks did not show an effect simply because they are relatively

new measures of intuition and may be inaccurate. Finally, given the frequency of marginal effects, Experiment 3.3 was likely underpowered and should be replicated with a larger sample. Given that all differences were in the expected direction, many of the insignificant results may correct themselves with an adequate sample size.

Besides addressing the limitations of this particular study, the comparability of Experiments 3.1 and 3.2 with Experiment 3.3 could be improved by running the former with some sort of time pressure manipulation. If time pressure increase Maximizing under conditions with differential implementation effort then that would provide further support for the claim that load increases Maximizing because of its impact on Strategy Implementation.

Those concerns aside, if you accept that putting participants under time pressure is a valid manipulation of load and the marginal effects that slower responding did legitimately increase green guessing, then we find that the impact load has on behaviour depends on the structure of the task. In particular, when the effort associated with implementing a Probability Matching Strategy is larger than that for Maximizing, load tends to encourage more Maximizing. By contrast, when the implementation effort required for each strategy is equivalent, load tends to encourage Matching or, at the very least, did not encourage Maximizing. The opposite direction of these effects may explain why some researchers fail to find any impact of WML on behaviour during a Binary Prediction Task. The results from Experiments 3.1 through 3.3, when taken together, do support the idea that Strategy Choice and Implementation are distinct processes that may be affected differently by manipulations such as load. Researchers should consider this when making theoretical claims about how their manipulation impacts behaviour in Binary Prediction Tasks in the future.

Chapter 4: General Discussion

In Chapter 1 I proposed that behaviour during a Binary Prediction Task is influenced by two separate processes. The first of these I called Strategy Choice and I defined it as a persons' top down or conscious intention to carry out a given strategy when presented with a Binary Prediction Task. This intention may evolve over the course of the task, but the intention should be available to participants' awareness. The second process I named Strategy Implementation. This is simply an individuals' ability to carry out their intended strategy. Conditions that impact Implementation may change behaviour away from what was intended by Strategy Choice and they may do so systematically. The latter is an issue for some of the current work on Binary Prediction Tasks because this systematic movement can be interpreted as a change in Strategy Choice when it is not. This may lead researchers to make inferences about mechanisms related to Strategy Choice when their data is, in fact, the results of Strategy Implementation. In Chapters 2 and 3 I examined two such instances.

Chapter 2 investigated what impact feedback on prediction accuracy had on behaviour and whether this change was the result of Feedback's influence on Strategy Choice or Implementation. Previous work has investigated the role that Feedback has on participant behaviour in a Binary Prediction Task (Shanks, Tunney & McCarthy 2002; Newell & Rakow 2007) and found, similar to the results reported here, that feedback tends to increase selection of the more likely alternative. This effect of feedback has been used to argue that irrational behaviour at Binary Prediction Tasks are simply the results of insufficient feedback for learning and that with the help of appropriate feedback most participants behave rationally (Shanks, Tunney & McCarthy 2002). Experiments 2.1 through 2.3 collectively provided evidence that feedback (a) does encourage selecting the more probable outcome among Non-Maximizers as

suggested in previous work, (b) that this modification of behaviour is mostly unknown to the participants, as indicated by systematic underestimation of green guessing and (c) that it results in little change to behaviour in very similar future Binary Prediction Tasks. The latter two points suggest that Feedback impacts Strategy Implementation rather than Strategy Choice. Further evidence for this claim came from the continued impact of feedback even when participants were assigned a strategy (Experiment 3.1). If the impact of Feedback is entirely or largely on Strategy Implementation, it seems unwarranted to argue that Feedback improves rationality as, I would argue, that requires modification of Strategy Choice.

This is not meant to imply that Feedback cannot ever impact Strategy Choice. Indeed, Newell and Rakow 2007 suggested that feedback may have a small impact on Strategy Choice by encouraging people away from Strict Maximizing, and I found some minor support for that here as well (Experiment 2.1 and 2.3). This could be because of the allure of correctly predicting reds, which has been discussed elsewhere in the literature (Brackbill & Bravos, 1962). Furthermore, different types of feedback (ex: Shanks *et al*, 2002) or feedback for longer durations (ex: Edwards 1961; Bereby-Myer & Erev 1998, Newell *et al*. 2013) than done here may have different effects. After many trials participants may become aware of the impact of feedback and it may come to impact Strategy Choice. I have represented Strategy Choice and Implementation as two separate forces, but in this sense, they might be thought of as more of a continuum. At one end is a well understood and fully conscious intention to carry out a specific strategy. At the other end are forces modifying behaviour completely without intention or the participant's awareness. There may be a middle ground in which participants are vaguely aware of something shaping their strategy and develop a hunch. The work of Bechara, Damasio, Tranel and Damasio (1997) supports this idea. Participants chose card from four different decks. Each

card offered some monetary payoff and some loss. The cards in each deck varied such that some of the decks had an overall poorer payoff than others. Participants very quickly developed elevated skin conductance responses when choosing from the poor decks and shortly after developed a hunch about which decks were better. These things occurred long before participants could explicitly report which deck was better and why. Similarly, feedback in Binary Prediction Tasks may lead participants to develop a similar hunch that they should be guessing more green. This could occur without them realizing that they should guess all green, or the logic behind why guessing all green is a superior strategy. Whether this middle ground exists is a question for future research to explore. However, regardless of whether Strategy Choice versus Implementation is better thought of as a continuum or distinct categories it is important to clarify which is affected and how before making inferences about the role that feedback (or other manipulations) might play in shaping rational choosers in Binary Prediction Tasks.

While I am discussing rational responding, it also bears worth mentioning that one of the common themes in the experiments presented here was a relatively high rate of Strict-Maximizing. In my work, Strict Maximizing rates were often as high as 50% and many of the Non-Maximizers still guessed predominantly green (see Experiment 2.1). This seems favorable for those wishing to claim that irrationality at Binary Prediction Tasks were a quirk of experimental circumstances. Furthermore, rates of maximizing in my work are higher than rates in much of the literature.

I must rely mostly on anecdote to explain this difference, but I believe it depends largely on two main factors. First, most of the data used here was collected on Mturk. I have noticed in a variety of Binary Prediction Tasks I have done with Mturkers that they seem to produce

relatively high rates of Strict Maximizing relative the student population we have typically tested. This can be seen in the experiments presented here also. Experiments 2.2 and 3.3 both tested undergraduate students from the University of Waterloo. In these experiments the rate of Strict Maximizing was 24% and 31% respectively. By contrast, in Experiments 2.1 and 2.3, which tested Mturkers, 45% and 41% used a Strict Maximizing Strategy. (Participants do not select strategies in Experiments 3.1 and 3.2.) The 41% in Experiment 2.3 was in spite of the very low rates of Strict Maximizing (20%) when no description was provided. While more rigorous tests are required to say anything conclusively, it is my impression that Mturkers are more likely to engage in Strict-Maximizing than University of Waterloo undergraduate students. I think that the reason for this difference is that Mturkers self-select into my study in a way that undergraduate students are not able to. The Mturkers that feel most confident that they will do well at a game involving predicting the outcome of die rolls (as the study is described) are more likely to sign up for the study. These same individuals might be more intelligent, more mathematically inclined and more deliberative, all of which have been found to impact rates of Maximizing behaviour (West & Stanovich 2003; Stanovich & West 2008; Koehler & James 2010). This self-selection may increase rates of Maximizing in my Mturk samples.

The second factor impacting rates of maximizing, at least relative to the earliest work on the phenomenon and much of the work on which challenges to rationality are based, is the presence of a description of the contingencies. Early work tended to have participants learn contingencies from feedback of some sort (either their accuracy or the actual outcome). As is evident from Experiment 2.3, providing participants with a description of the contingencies dramatically increases rates of Strict Maximizing. The majority of the work presented here

provides descriptions of the contingencies which may lead to higher rates of Strict Maximizing overall.

As I mentioned earlier in the thesis, the rates of strict Probability Matching observed here are actually low, with Overmatching being a much more common strategy. For my purposes, this distinction is not critical as I am mostly concerned whether participants' behaviour (whatever it is) reflects only Strategy Choice or also Strategy Implementation. Furthermore, Overmatching is still sub optimal relative to Maximizing, albeit to a lesser degree. Nevertheless, variance in rates of Probability Matching based on experimental design is interesting and relevant to theories examining why participants engage in particular strategies such as the Pattern Search Account and the Expectation Matching Hypothesis. The work here does not directly investigate when and how Strict Probability Matching occurs, but it does seem that it is more common in situations that do not involve Strategy Implementation. For example, in Experiment 3.3, where participants simply entered in the number of green guesses they would make, Strict Probability Matching is the most common Strategy used by participants (see Figure 3.3). This is also observable in previous work that does not use trial by trial guessing (James & Koehler 2011; West & Stanovich 2003). Thus, the degree to which Strict Probability Matching occurs, like Strict Maximizing, is likely dependent on a variety of factors, but regardless, if participants do not Maximize, they are behaving sub optimally.

In addition to investigating the role of Feedback in Binary Prediction Tasks, I also examined the impact of Working Memory Load (WML) on behaviour. The impact of WML on behaviour has been used as evidence in various mechanisms hypothesized to lead to generation of a Probability Matching Strategy. For example, Koehler and James (2011) suggest that selection of a probability Matching Strategy is intuitive and argue this strategy should be more

prevalent under WML. By contrast, Wolford, Newman, Miller and Wig (2004) argue that adoption of a Probability Matching Strategy results from belief in a pattern in the data and that Probability Matching should decrease under WML. Both of these claims assume that WML will affect only Strategy Choice.

In Chapter 3, however, I found some evidence that WML may affect both Strategy Choice and Implementation and may do so in different ways. For example, when load impacts Strategy Implementation (Experiments 3.1 and 3.2), it tends to increase green guessing when the participant is receiving feedback and attempting to implement a Probability Matching Strategy. This was the circumstance of the original paper documenting a WML effect (Wolford *et al* 2004). Thus, the observed increase in Maximizing observed by Wolford, Newman, Miller and Wig may be due to the difficulty of implementing a Probability Matching Strategy. By contrast, when load impacts Strategy Choice, I found some evidence that it encourages the opposite response (Experiment 3.3) and decreases how often participants guess green. While this effect was only marginal and requires a replication with more power, it is consistent with James and Koehler's claim that Probability Matching is an intuitively generated strategy.

A further issue for research making theoretical claims about why Probability Matching exists is the large variety of methods used in Binary Prediction Tasks. Current research uses Binary Prediction Tasks that vary from fully described (for example West & Stanovich 2003; Stanovich & West 2008; Gal & Baron 1996; Koehler & James 2010; James & Koehler 2011) to only experienced (for example Shanks et al. 2002; Bereby-Meyer & Erev 1998; Wolford et al. 2004). Some tasks involve trial by trial guessing (for example Gaissmaier & Schooler 2008; Koehler & James 2009). Others do not (for example Gal & Baron 1996; Koehler & James 2009). In some of the work participants are provided with feedback (for example Rakow &

Newell 2007; Otto et al. 2011) and in others they are not (for example Koehler & James 2010; James & Koehler 2011). These varying formats are treated as though they are interchangeable and roughly equal. However, the work presented here suggests these features have significant impacts on how participants respond to the study. This presents a challenge for those wishing to compare the impacts of different manipulations across studies, a common practice when arguing for theoretical causes of Probability Matching Behaviour.

The issue is further compounded by the fact that some of these design features may interact with one another in significant ways. For example, Experiment 3.1 demonstrated that feedback interacts meaningfully with WML. If the presence or absence of feedback is manipulated without awareness of its relevance this interaction could complicate interpretation of the data across studies. In particular, designs without feedback provided would be less likely to observe a WML effect than those that did include feedback.

The distinction between Strategy Choice and Implementation and its implications for the effects of both Feedback and WML suggest that future research needs to clarify which process the target manipulation effects and how. In particular, testing whether the effects of a given manipulation transfer to similar but superficially distinct problems may be a reasonable test of whether or not Strategy Choice has been impacted. Other methods might include testing for awareness on the part of the participant that their behaviour has changed. At the very least, considering whether or not there is a plausible implementation based account of the data is worthwhile.

While the work presented here focused on behaviour in Binary Prediction Tasks, it seems plausible that a similar distinction might apply to other simple decision problems. Any time a

problem requires extended implementation (as is the case here) there is room for Strategy Implementation to masquerade as Strategy Choice. Furthermore, the results from Experiment 2.3 suggest that there may sometimes be a cost to learning contingencies from experience. At least for Binary Prediction Tasks, providing a succinct overall description of the problem seems to improve decision making. It would be interesting to investigate further real world applications of such a description and learn whether the improvement to decision making continues to apply in those contexts.

The work presented here also raised some questions less related specifically to the distinction between Strategy Choice and Implementation. One such question relates to my measure of participants' awareness of their own behaviour change. To test whether participants were aware of the impact of feedback on their behaviour I asked them to estimate how often they guessed green. However, this measure warrants some further discussion. Two questions seem of particular import: (1) how do participants generate their estimate and (2) is this estimate an accurate measure of their awareness?

In Chapter 2 I discussed the possibility that participants generate their estimates based solely on their initial Strategy Choice. Put another way, participants may base their estimates on what they intended to do, not what they actually did. Rather than attempt to keep track of exactly how often they guessed green, which is computationally expensive, it may be easier to simply report their intended strategy and assume that they carried it out relatively accurately. Thus, estimates may be a better reflection of intention than awareness.

With respect to the work on Feedback reported here, participants may have been aware that green was more often correct, and even that they were erring on the side of guessing green,

but simply not bothered to update their estimate. This could be because they considered it to be a small and inconsequential modification or simply because it required too much effort to do so. However, even if participants are aware of the impact of Feedback on their behaviour, the fact that (a) learning from feedback did not transfer to similar subsequent tasks and (b) the effect of Feedback existed even when participants were assigned to a Strategy, suggests that Strategy Implementation, and not Selection, was the locus of Feedback's effect.

Nevertheless, how participants generate estimates of their own guessing behaviour is an interesting avenue for further research. One potentially informative study would be to compare participants estimations of their own behaviour with estimations of others behaviour. Since participants do not have access to the intentions of others, they cannot use intention information to generate estimates of others behaviour. If estimates made for others look similar to estimates made for oneself, and in particular if participants continue to underestimate the frequency of green guesses, this would suggest that intentions are not how participants generate their frequency estimates.

Finally, whenever a participant goes to select red (as dictated by a top down strategy such as Probability Matching or Over Matching) their top down response is in conflict with any bottom-up mechanism (such as Feedback) that encourages green responding. This raises the interesting question of how this conflict between top down Strategy Choice and bottom-up processes is resolved. The work presented here offers a couple of interesting insights into this issue.

In this case, it seems as though top down Strategy Choice and the bottom up effects of Feedback in fact blend. In this task, participants seem to largely adhere to their Strategy Choice,

but their behaviour is subtly influenced by bottom up features, suggesting that behaviour is not solely determined by top down intentions to carry out a particular strategy. Further, participants' lack of awareness of the influence of feedback on their green guesses suggests that the influences of bottom up features operated entirely separately from their intentions.

This begs the question, when and how would participants notice these bottom up influences on their behaviour? One could think of their Strategy Choice as a reflection of the participants' mental model of the optimal strategy in the die game. The outcome feedback they receive is information they can use to update that model. Filipowicz (2017) argued that mental models tend to be updated when an event is surprising. In other words, he found that the more different an event was from that predicted by a model the more likely it was to change the model. In the context of Binary Prediction Tasks this suggests that feedback needs to be surprising in order to cause change at the Strategy Choice level. Since participants expect a certain degree of inaccuracy, most outcomes do not meet these criteria.

However, Strategy Implementation seems to be more immediately malleable. Put another way, bottom-up processes such as conditioning can have effects on behaviour without modifying Strategy Choice and this impact does not require an event to be surprising. At some point, presumably, behaviour would become deviant enough from the participants' intention that the behaviour itself acts as the surprising event capable of changing the participants' model of the problem. This, of course, is conjecture, but would serve as an interesting avenue for future research.

Some work finds much stronger impacts of Feedback depending on how many trials participants do (Edwards 1961; Bereby-Myer & Erev 1998, Newell *et al.* 2013) or how feedback

is presented (Shanks et al. 2002). This work did not test whether participants were aware of changes in their behaviour. It also did not examine whether the benefits of feedback continued on a similar but distinct task. Assuming, however, that these larger Feedback effects did represent a change to Strategy Choice, a carefully controlled study could examine when and how these changes occurred.

Another interesting avenue for future research would be to investigate whether the impacts of bottom-up control of behaviour can be increased. Thompson-Schill, Ramscar & Chrysikou (2009) argued that cognitive control, under the command of the Prefrontal Cortex, actually hampers probabilistic bottom-up learning. They suggest that young children, typically found to perform more optimally at Binary Prediction Tasks (Derks and Paclisanu 1967; Weir 1964; Lewis 1966), do so because they are more influenced by bottom up features such as conditioning. This influence exists precisely because of their under developed executive function or cognitive control. They argue that this is adaptive for young children as they need to learn large amounts of information from their environment in a very short period of time. Thus, allowing too much top down control over behaviour might inhibit learning of subtle statistical information in the environment. If this observation is true, we would expect that manipulations (such as WML) that hamper top down control of behaviour would increase bottom up impacts in adults. For example, we would expect the impact of Feedback to increase when adult participants were under WML. This view point is consistent with the interaction between Feedback and WML reported in Experiment 3.1. Feedback had the largest influence when Probability Matchers were under WML. This could be interpreted as resulting because WML hampered participants' cognitive control allowing bottom-up processes to have a larger impact on behaviour. It is possible Feedback might improve performance faster when under WML.

Participants' behaviour (rather than the die outcomes themselves) might become discrepant enough from their intentions to serve as the surprise that initiates changes in participants' models. Thus, they may (a) become aware of the change and (b) update their Strategy Choices as a result. This would have the ironic effect of leading those who were most cognitively taxed to be most rational.

In summary, this thesis provided evidence for a distinction between two processes influencing behaviour on a Binary Prediction Task: Strategy Choice and Strategy Implementation. I argued that this distinction is critical in particular when behavioural changes related to Implementation masquerade as changes in Strategy Choice. This leads researchers to make incorrect assumptions about the role of target manipulations on Strategy Choice. I provided two such examples of this phenomenon: Feedback and Working Memory Load.

Moving forward the distinction between what we think we ought to do (our Strategy Choice) and what we actually are capable of doing (our Strategy Implementation) seems important in understanding behaviour in decision making problems. In particular, it is an important distinction in Binary Prediction Tasks and deserves due consideration in future research. If we are to unravel why some people Probability Match and when and how they learn to Maximize, it is critical that we tell the difference between what they thought and what they actually did.

References

- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, 275(5304), 1293-1295.
- Bereby-Meyer, Y., & Erev, I. (1998). On learning to become a successful loser: A comparison of alternative abstractions of learning processes in the loss domain. *Journal of Mathematical Psychology*, 42(2-3), 266-286.
- Brackbill, Y., & Bravos, A. (1962). Supplementary report: The utility of correctly predicting infrequent events. *Journal of Experimental Psychology*, 64(6), 648-649.
- Braveman, N. S., & Fischer, G. J. (1968). Instructionally induced strategy and sequential information in probability learning. *Journal of Experimental Psychology*, 76(4, Pt.1), 674-676.
- Conover, W. J., & Iman, R. L. (1981). Rank transformations as a bridge between parametric and nonparametric statistics. *The American Statistician*, 35(3), 124-129.
- Derks, P. L., & Paclisanu, M. I. (1967). Simple strategies in binary prediction by children and adults. *Journal of Experimental Psychology*, 73(2), 278-285.
- Douglas, R. J., & Pribram, K. H. (1966). Learning and limbic lesions. *Neuropsychologia*, 4(3), 197-220.

- Edwards, W. (1961). Probability learning in 1000 trials. *Journal of Experimental Psychology*, 62(4), 385-394.
- Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278.
- Evans, J. S. B. T., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3), 223-241.
- Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, 104(2), 301-301-318.
- Fantino, E., & Esfandiari, A. (2002). Probability matching: Encouraging optimal responding in humans. *Canadian Journal of Experimental Psychology/Revue Canadienne De Psychologie Expérimentale*, 56(1), 58-63.
- Ferguson, M. J., & Bargh, J. A. (2004). Liking is for doing: The effects of goal pursuit on automatic evaluation. *Journal of Personality and Social Psychology*, 87(5), 557.
- Filipowicz, A. (2017). Adapting to Change: The Role of Priors, Surprise and Brain Damage on Mental Model Updating. *UWSpace*. <http://hdl.handle.net/10012/11724>
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives*, 19, 24-42.
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition, In Press*

- Gaissmaier, W., Wilke, A., Scheibehenne, B., McCanney, P., & Barrett, H. C. (2016). Betting on illusory patterns: Probability matching in habitual gamblers. *Journal of Gambling Studies*, 32(1), 143-156.
- Gal, I., & Baron, J. (1996). Understanding repeated simple choices. *Thinking & Reasoning*, 2(1), 81-98.
- Green, C. S., Benson, C., Kersten, D., & Schrater, P. (2010). Alterations in choice behavior by manipulations of world model. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, 107(37), 16401-16406.
- Green, L., Rachlin, H., & Hanson, J. (1983). Matching and maximizing with concurrent ratio-interval schedules. *Journal of the Experimental Analysis of Behavior*, 40(3), 217-224.
- Haigh, M. (2016). Has the standard cognitive reflection test become a victim of its own success? *Advances in Cognitive Psychology*, 12(3), 145.
- James, G., & Koehler, D. J. (2011). Banking on a bad bet: Probability matching in risky choice is linked to expectation generation. *Psychological Science*, 22(6), 707-707-711.
- Jones, M. R. (1971). From probability learning to sequential processing: A critical review. *Psychological Bulletin*, 76(3), 153.
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin & D. Kahneman (Eds.), (pp. 49-49-81). New York, NY, US: Cambridge University Press.

- Kahneman, D., & Klein, G. (2009). Conditions for intuitive expertise: A failure to disagree. *American Psychologist*, 64(6), 515-526.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3), 430-454.
- Koehler, D. J., & James, G. (2014). Probability matching, fast and slow. *The Psychology of Learning and Motivation*, 61, 103-131.
- Koehler, D. J., & James, G. (2009). Probability matching in choice under uncertainty: Intuition versus deliberation. *Cognition*, 113(1), 123-127.
- Koehler, D. J., & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, 38(6), 667-676.
- Lee, D., Conroy, M. L., McGreevy, B. P., & Barraclough, D. J. (2004). Reinforcement learning and decision making in monkeys during a competitive game. *Cognitive Brain Research*, 22(1), 45-58.
- Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, 28(1), 3.
- McMahon, A. J., & Scheel, M. H. (2010). Glucose promotes controlled processing: Matching, maximizing, and root beer. *Judgment and Decision Making*, 5(6), 450-457.

- Newell, B. R., Koehler, D. J., James, G., Rakow, T., & Van Ravenzwaaij, D. (2013). Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition, 41*(3), 329-338.
- Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review, 14*(6), 1133-1133-1139.
- Otto, A. R., Taylor, E. G., & Markman, A. B. (2011). There are at least two kinds of probability matching: Evidence from a secondary task. *Cognition, 118*(2), 274-279.
- Peterson, C. R., & Ulehla, Z. J. (1965). Sequential patterns and maximizing. *Journal of Experimental Psychology, 69*(1), 1-4.
- Rand, D. G. (2016). Cooperation, fast and slow: Meta-analytic evidence for a theory of social heuristics and self-interested deliberation. *Psychological Science, 27*(9), 1192-1206.
- Schulze, C., & Newell, B. R. (2016). Taking the easy way out? increasing implementation effort reduces probability maximizing under cognitive load. *Memory & Cognition, 44*(5), 806-818.
- Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making, 15*(3), 233-233-250.
- Slovic, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*(1), 3-22.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences, 23*(5), 645-665.

- Stanovich, K. E., & West, R. F. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology, 94*(4), 672-672.
- Stieger, S., & Reips, U. (2016). A limitation of the cognitive reflection test: Familiarity. *Peerj, 4*, e2395.
- Thompson-Schill, S. L., Ramscar, M., & Chrysikou, E. G. (2009). Cognition without control: When a little frontal lobe goes a long way. *Current Directions in Psychological Science, 18*(5), 259-263.
- Thomson, K. S., & Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision Making, 11*(1), 99.
- Toplak, M. E., West, R. F., & Stanovich, K. E. (2011). The cognitive reflection test as a predictor of performance on heuristics-and-biases tasks. *Memory & Cognition, 39*(7), 1275-1289.
- Tversky, A., & Edwards, W. (1966). Information versus reward in binary choices. *Journal of Experimental Psychology, 71*(5), 680-683.
- Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin, 76*(2), 105-110.
- Unturbe, J., & Corominas, J. (2007). Probability matching involves rule-generating ability: A neuropsychological mechanism dealing with probabilities. *Neuropsychology, 21*(5), 621-630.

- Verwijmeren, T., Karremans, J. C., Stroebe, W., & Wigboldus, D. H. (2012). Goal relevance moderates evaluative conditioning effects. *Learning and Motivation*, 43(3), 107-115.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, 14(1), 101-118.
- Weir, M. W. (1964). Developmental changes in problem-solving strategies. *Psychological Review*, 71(6), 473.
- Welch, B. L. (1951). On the comparison of several mean values: an alternative approach. *Biometrika*, 38(3/4), 330-336.
- West, R. F., & Stanovich, K. E. (2003). Is probability matching smart? associations between probabilistic choices and cognitive ability. *Memory & Cognition*, 31(2), 243-251.
- Wilson Jr., W. A. (1960). Supplementary report: Two-choice behavior of monkeys. *Journal of Experimental Psychology*, 59(3), 207-208.
- Wolford, G., Newman, S. E., Miller, M. B., & Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology/Revue Canadienne De Psychologie Expérimentale*, 58(4), 221-221.
- Yellot, J. I. (1969). Probability learning with non-contingent success. *Journal of Mathematical Psychology*, 6, 541-575.
- Zimmerman, D. W. (2004). A note on preliminary tests of equality of variances. *British Journal of Mathematical and Statistical Psychology*, 57(1), 173-181.

Zimmerman, D. W. (1998). Invalidation of parametric and nonparametric statistical tests by concurrent violation of two assumptions. *The Journal of experimental education*, 67(1), 55-68.

Zimmerman, D. W., & Zumbo, B. D. (1993). Relative power of the Wilcoxon test, the Friedman test, and repeated-measures ANOVA on ranks. *The Journal of Experimental Education*, 62(1), 75-86.

Appendix A – Experiment Materials

Experiment 2.1

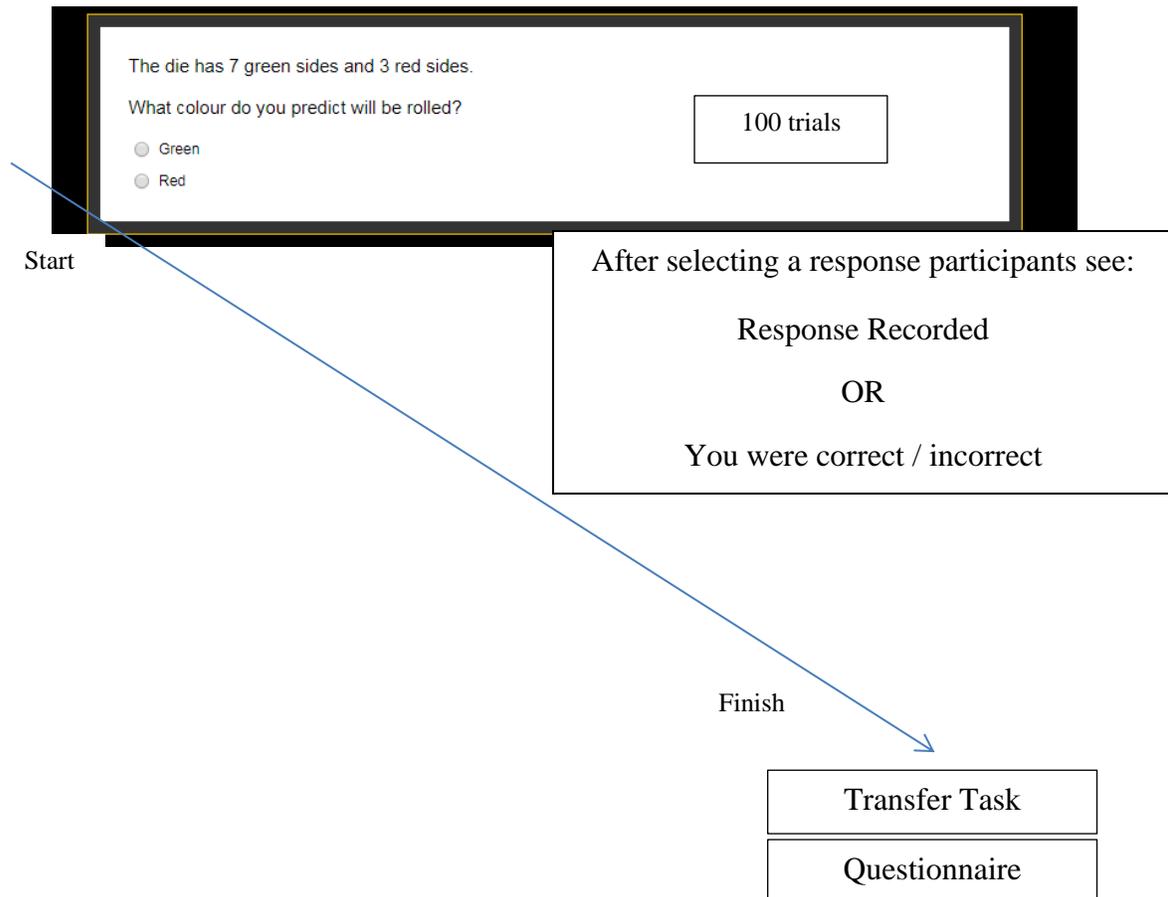
Experiment flow

Instructions:

In this portion of the study you will try to guess the outcome of a series of rolls of a 10-sided die. The die used has 7 green sides and 3 red sides. For each roll, you will be asked to predict whether the die will come up green or red. Every time your prediction is correct you will earn \$0.03.

Click on the radio button corresponding to your colour choice in order to make your prediction. You will automatically be transferred to the next page, where you will be told whether your response was correct or incorrect [alternate text for no feedback condition: “notified that your response has been recorded”]. After 1 second the page will automatically advance to the next prediction.

There will be 100 trials. Click ">>" to begin.



Transfer Task

Consider a game that involves rolling a fair, 10-sided die. The die is “fair” in the sense that, when it is rolled, each of the 10 sides has an equal chance of turning up.

On this particular 10 sided die:

7 sides are red [blue in alternate version]

3 sides are green [yellow in alternate version]

Suppose you are to play the following guessing game: The die will be rolled 10 times and before each roll you will guess whether a blue side or a yellow side will turn up. Imagine that you will receive \$1 for each correct guess.

Please indicate your guess for each of the ten rolls of the die

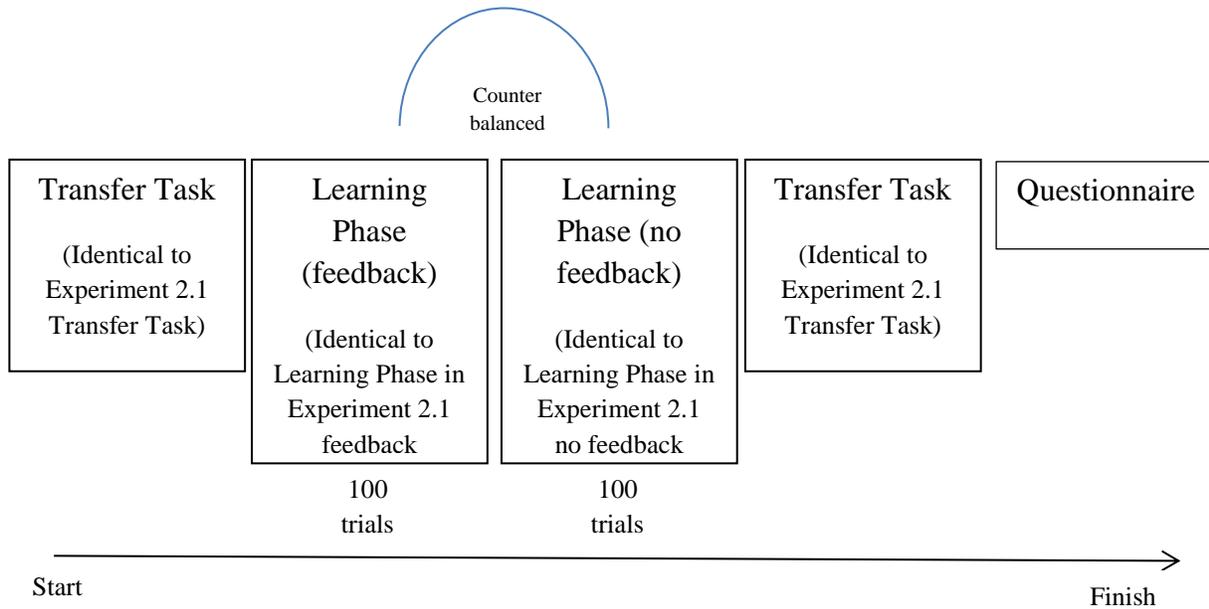
	Select your guess	
	Blue	Yellow
Roll 1	<input type="radio"/>	<input type="radio"/>
Roll 2	<input type="radio"/>	<input type="radio"/>
Roll 3	<input type="radio"/>	<input type="radio"/>
Roll 4	<input type="radio"/>	<input type="radio"/>
Roll 5	<input type="radio"/>	<input type="radio"/>
Roll 6	<input type="radio"/>	<input type="radio"/>
Roll 7	<input type="radio"/>	<input type="radio"/>
Roll 8	<input type="radio"/>	<input type="radio"/>
Roll 9	<input type="radio"/>	<input type="radio"/>
Roll 10	<input type="radio"/>	<input type="radio"/>

Questionnaire Questions

1. Please describe why you chose to make the guesses you did on the previous page [the transfer task]? What was your strategy and why did you choose that strategy? Please be as specific about your reasoning as possible.
2. Out of the 100 guesses in the first guessing game[refers to the Learning Phase], how many times do you think you guessed green?
3. Did your strategy change over time? If yes, how did your strategy change? Why did it change that way?
4. Imagine you were watching the outcome of 10 die rolls of a die with 7 green sides and 3 red sides. Imagine the die had already been rolled 7 times and all 7 times the die had come up green. On the 8th roll of the die, which colour is more likely?

Experiment 2.2

Experiment Flow

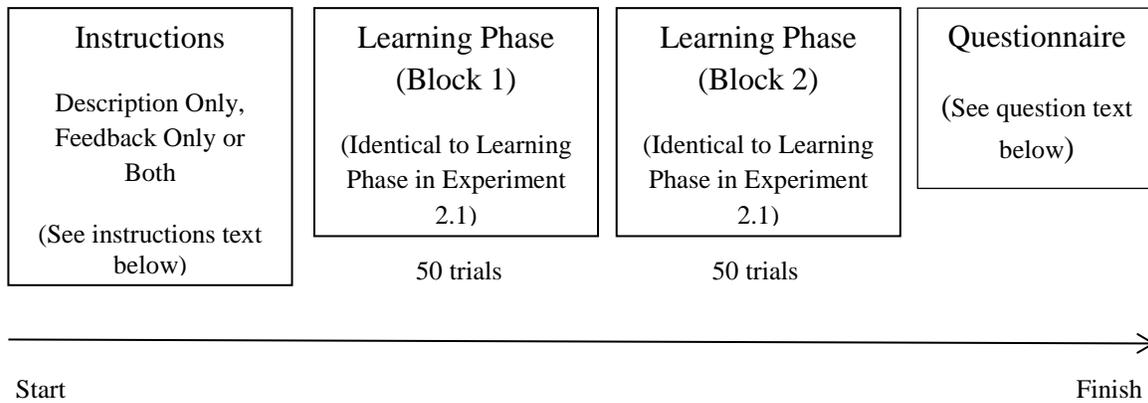


Questionnaire

1. Please describe the strategy you would use if you were asked to predict 10 rolls of a die with 7 blue sides and 3 yellow sides.

Experiment 2.3

Experiment Flow



Instructions

Both Condition: In this portion of the study you will try to guess the outcome of a series of rolls of a 10-sided die. The die used has 7 green sides and 3 reds sides. For each roll, you will be asked to predict whether the die will come up green or red. Try to use a strategy that you think gives you the best chance of getting all 100 guesses correct.

Click on the radio button corresponding to your colour choice in order to make your prediction. You will automatically be transferred to the next page, where you will be told whether your response was correct or incorrect. After 1 second the page will automatically advance to the next prediction.

There will be 50 trials in Block 1. Click ">>" to begin.

Description Only Condition: In this portion of the study you will try to guess the outcome of a series of rolls of a 10-sided die. The die used has 7 green sides and 3 reds sides. For each roll, you will be asked to predict whether the die will come up green or red. Try to use a strategy that you think gives you the best chance of getting all 100 guesses correct.

Click on the radio button corresponding to your colour choice in order to make your prediction. You will automatically be transferred to the next page, where you will be notified that your response has been recorded. After 1 second the page will automatically advance to the next prediction.

There will be 50 trials in Block 1. Click ">>" to begin.

Feedback Only Condition: In this portion of the study you will try to guess the outcome of a series of rolls of a 10-sided die that has some sides painted green and the rest painted red. For each roll, you will be asked to predict whether the die will come up green or red. Try to use a strategy that you think gives you the best chance of getting all 100 guesses correct.

Click on the radio button corresponding to your colour choice in order to make your prediction. You will automatically be transferred to the next page, where you will be told whether your response was correct or incorrect. Note that you will need to use this to figure out how many of the sides on the die are painted green and how many are painted red. After 1 second of viewing your feedback the page will automatically advance to the next prediction.

There will be 50 trials in Block 1. Click ">>" to begin.

Questionnaire

1. Out of the 50 guesses you made after the break, in **Block 2**, how many times do you think you guessed green?
2. How many sides of the 10-sided die do you think were green? (Displayed only to the Feedback Only condition)
- 3.

Mary just played the guessing game you played except the die was rolled only 10 times and Mary was paid \$1 every time she correctly predicted a die roll. The following chart shows you what Mary guessed and the die outcome for each of the ten rolls. Overall, how many dollars did Mary make?

Roll	Mary's Guess	Actual outcome
1	green	green
2	green	green
3	red	green
4	green	red
5	red	red
6	green	green
7	green	green
8	green	green
9	red	green
10	green	red

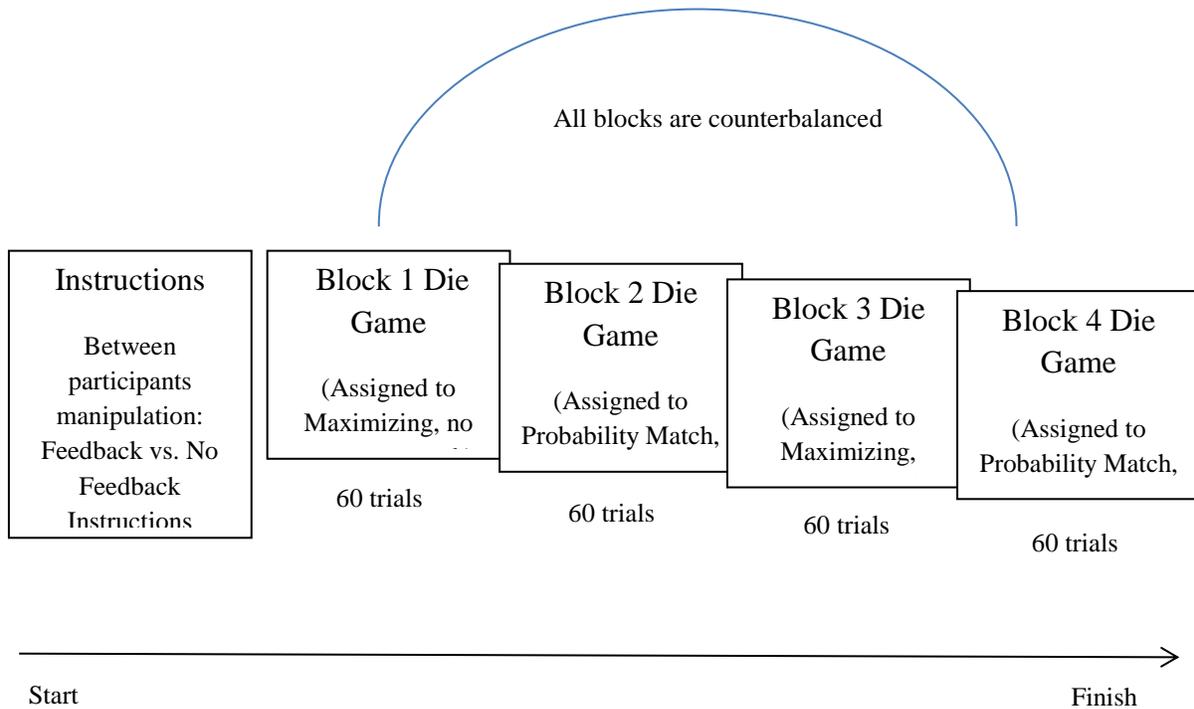
- \$10
- \$6
- \$7
- \$5

4. Please rank the following strategies based on which you think is the best way to make the most money in the die game you just played. Give the highest rank to the best strategy and the lowest rank to the worst one. You can change the strategy ranking by clicking on the strategy and dragging it to a new spot. Strategies higher up in the order are more highly ranked. (Note that these strategies were presented in random order to each participant.)

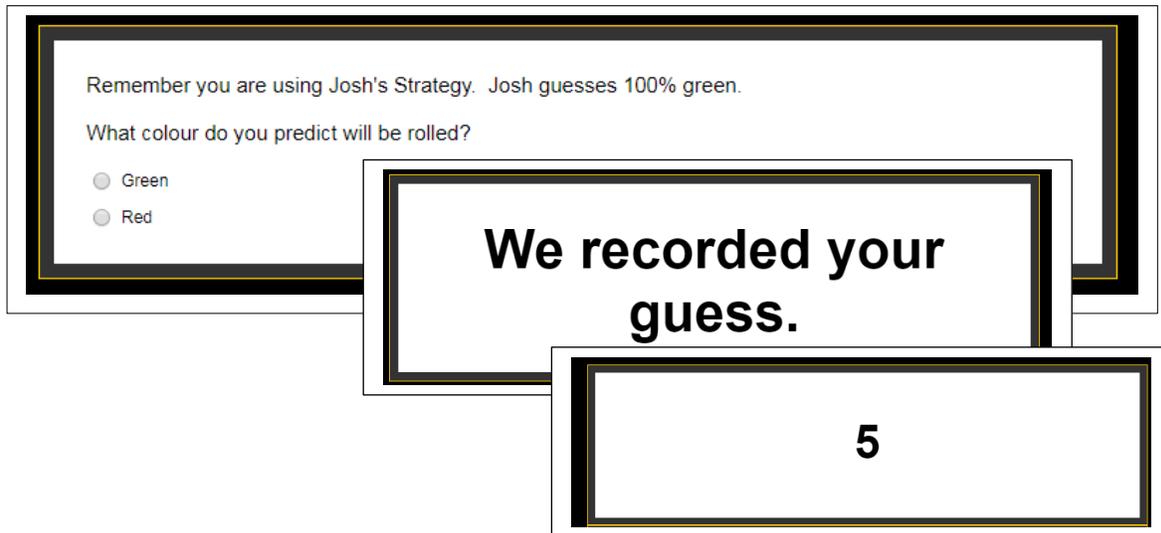
- I should always guess green and never guess red, because green is more likely to occur on any individual role of the die
 - I should choose green about 70% of the time and red about 30% of the time, because I know that overall the die will come up green about 70% of the time and red 30% of the time.
 - I should choose green half the time and red half the time, because there are two colours that could occur on any given role of the die.
 - I should guess green 80-90% of the time, because green is more likely on any given die roll, but I should still guess some red because the die will come up red some of the time.
5. When you were playing the die game, did your strategy change over time? (If participants selected yes, they saw question number 6)
6. How and why did your strategy change? Please be as specific as you can.

Experiment 3.1

Experiment Flow



Sample Die Game Trial – No feedback condition, Assigned to Maximize, concurrent task



Instructions

All participants viewed these instructions:

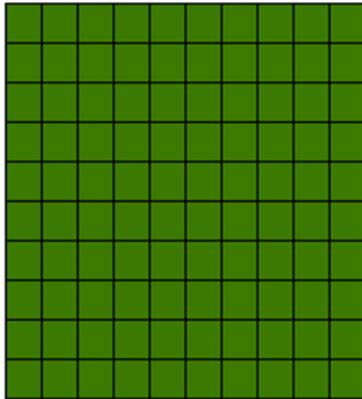
Imagine a virtual ten-sided die. Seven sides are green and three sides are red. The computer will roll this virtual die 240 times. Before each roll you will be asked to predict whether the die will come up green or red. After you make your prediction you will see a new screen telling you whether or not your prediction was correct [No feedback condition: After you make your prediction you will see a new screen telling you that your response was recorded].

You are going to make predictions in 4 blocks of 60 trials each. For each of these blocks, we will ask you to follow one of two strategies: Josh's strategy or Carl's strategy. Before each block you will be told which strategy to use and the details of that strategy will be explained.

For half of the blocks you see you will also be asked to remember some numbers while you are making your predictions. You will be given specific instructions on how to do this before completing those blocks. When you are asked to remember the numbers it is very important you do your best to remember the numbers accurately. Please advance to the next page when you are ready to begin.

Maximizing Instructions:

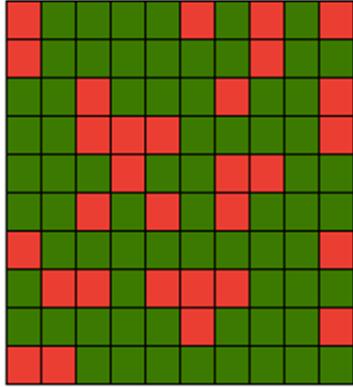
In this block, you will try to predict 60 rolls of the die. Please try to use Josh's strategy when making your predictions. Josh always guesses the more likely event, even though he knows he will sometimes be wrong. If we were to depict Josh's guesses in a grid, they might look something like this:



Please use Josh's strategy when making your guesses so that after all 60 guesses you guessed green (and red) the same amount as Josh did. Advance to the next screen when you are ready to begin.

Probability Matching Instructions:

In this block, you will try to predict 60 rolls of the die. Please try to use Carl's strategy when making your predictions. Carl's strategy was to try to correctly predict each and every die roll. Overall the proportion of red and green Carl chose matched the die. If we were to depict Carl's guesses in a grid, they might look something like this:



Please use Carl's strategy when making your guesses so that after all 60 guesses you guessed green (and red) the same amount as Carl did. Advance to the next screen when you are ready to begin.

Dual Task Instructions:

In this block you will be asked to remember some numbers while you make predictions about the outcome of the die. After each of your predictions, you will see a number between 0 and 9 appear in the center of the screen. Try to remember the last 3 of these numbers you have seen. At random, a probe will pop up asking you to enter the last 3 numbers that came up on the screen. Please be as accurate as possible in recalling the numbers. It is important that you do well on this task.

Experiment 3.3

Instructions

Slow Condition Instructions: The following questionnaire contains a series of thought problems. When solving these thought problems we would like you to try and be as accurate as you can. Speed is not important, so take your time and think carefully about your answers! Hit the advance button at the bottom of the screen when you are ready to start.

Fast Condition Instructions: The following questionnaire contains a series of thought problems. When solving these thought problems we would like you to try and be as fast and accurate as you can. Speed is important, so please answer as quickly as you can. Hit the advance button at the bottom of the screen when you are ready to start.

Questionnaire

Fully Described Die Problem:

For questions 1-3, please consider the following game:

The game involves rolling a fair, 10-sided die. The die is "fair" in the sense that, when it is rolled, each of the 10 sides has an equal chance of turning up.

On this particular 10-sided die:

7 sides are GREEN

3 sides are RED

Suppose you are to play the following guessing game: The die will be rolled ten times, and before each roll you are to guess whether a GREEN side or a RED side will turn up. Imagine that you will receive \$1 for each correct guess.

1. Across all 10 rolls of the die, how many times would you guess green?

Comprehension Questions

Which of the following is correct? In the game involving the 10-sided die, you were asked:

- To report how many times you would guess green.
- To report how many sides on the die were green.
- To report how many greens would be rolled overall.

Mary just played the game involving the 10-sided die. The following chart shows you what Mary guessed and the die outcome for each of the ten rolls. Overall, how many dollars did Mary make?

Roll	Mary's Guess	Actual outcome
1	green	green
2	green	green
3	red	green
4	green	red
5	red	red
6	green	green
7	green	green
8	green	green
9	red	green
10	green	red

- \$10
- \$6
- \$7
- It is impossible to calculate how much Mary made

CRT 2:

1. If you're running a race and you pass the person in second place, what place are you in?
2. A farmer had 15 sheep and all but 8 died. How many are left?
3. Emily's father has three daughters. The first two are named April and May. What is the third daughter's name?
4. How many cubic feet of dirt are there in a hole that is 3'deep x3'wide x 3'long?

Belief Bias Problems:

For each of the following problems, decide if the given conclusion follows logically from the premises. Select YES if, and only if, you judge that the conclusion can be derived unequivocally from the given premises. Otherwise, select NO.

Premise 1: All things that are smoked are good for the health.

Premise 2: Cigarettes are smoked.

Conclusion: Cigarettes are good for the health.

Premise 1: All unemployed people are poor.

Premise 2: Rockefeller is not unemployed.

Conclusion: Rockefeller is not poor.

Premise 1: All flowers have petals.

Premise 2: Roses have petals.

Conclusion: Roses are flowers.

Premise 1: All animals with four legs are dangerous.

Premise 2: Poodles are animals that aren't dangerous.

Conclusion: Poodles do not have four legs.

Premise 1: All mammals walk.

Premise 2: Whales are mammals.

Conclusion: Whales walk.

Premise 1: All Eastern countries are communist.

Premise 2: Canada is not an Eastern country.

Conclusion: Canada is not communist.

Premise 1: All animals like water.

Premise 2: Cats do not like water.

Conclusion: Cats are not animals.

Premise 1: All things that have a motor need oil.

Premise 2: Automobiles need oil.

Conclusion: Automobiles have motors.

Appendix B- Non-parametric Tests

Experiment 2.1

Binary Prediction Task performance

The data from 125 participants that selected at least 1 green during the Learning Phase of Experiment 2.1 were submitted to a Wilcoxon rank sum test. The dependent variable was participants total number of green guesses and the between participant factor was whether or not they received feedback on the accuracy of their guesses. The Wilcoxon test revealed a marginally significant difference between the median number of greens guessed in the feedback versus no feedback condition. Participants guessed marginally more green when receiving feedback ($Z=2308, p=0.076$).

Participant Estimates

Participants estimated how often they guessed green in the Learning Phase of Experiment 2.1. A difference score was calculated using the same method explained in the main text of this document. These difference scores were submitted to a Wilcoxon rank sum test with Feedback (Absent versus Present) as the between participant factor. This revealed a significant difference between the median difference scores in the feedback versus no feedback conditions. Participants were less accurate when receiving feedback ($Z=1448.5, p=0.012$). The same test done on participants raw estimates instead of difference scores revealed that there is no difference in raw estimates between conditions ($Z=1925.5, p=0.905$).

Transfer Task

Participants completed the Transfer Task detailed in the methods section of Experiment 2.1 and also in Appendix A. The number of times they guessed the more likely colour was summed to give a value between 0 and 10. These scores were submitted to a Wilcoxon rank sum test with Feedback (Present vs. Absent) as the between participant factor. This revealed a no difference between the median scores on the transfer task as a function of feedback ($Z=2192.5$, $p=0.199$).

Participants total number of green guesses during the Learning Phase and their total score on the Transfer Task were converted to proportions so that they could be compared. Participants proportions were submitted to a Wilcoxon signed rank test with task (Learning Phase or Transfer Task) as a within participant factor. This was done for both the Feedback condition and the No Feedback condition and revealed no difference between the proportion of green guessed across tasks in either the feedback condition ($V=1000.5$, $p=0.263$) or the no feedback condition ($V=976.5$, $p=0.652$).

Experiment 2.2

Binary Prediction Task Performance

Unfortunately, I do not know of a non-parametric equivalent to a mixed model ANOVA, so I am unable to provide a non-parametric version of the omnibus ANOVA done in Experiment 2.2. However, I can examine the simple effects using non-parametric tests. For example, a significant interaction existed in the parametric data and this interaction was examined with a t-

test. To provide a non-parametric equivalent of the t-test, I used two Wilcoxon signed rank tests examining the effect of Feedback on choice behaviour in LP1 versus LP2. These revealed that there was only a trend towards improved performance with feedback regardless of the order it was received (Feedback first: $V=182$, $p=0.183$, no feedback first $V=83$, $p=0.162$)

Transfer Task

I evaluated whether participants improved on the Transfer Task from the initial test (TT1) to the test after the Learning Phase (TT2). The number of greens each participant guessed on each test was submitted to a Wilcoxon signed rank test with test time (1st versus 2nd) as the within participant factor. Participants guessed significantly more green on the second taking of the test than the first, ($V=57$, $p=0.023$).

Experiment 2.3

Binary Prediction Task performance

A Kruskal-Wallis Ranked Sum test was used to test the difference between conditions in median greens guessed in Block 2 of Experiment 2.3. Number of greens guessed was submitted to the test with condition (Both, Description Only, Feedback Only) as the between participants variable. The test revealed a significant difference between conditions. $X^2(2, N=279)=23.12$, $p<.001$. Pair-wise comparisons between conditions done using a Wilcoxon Rank Sum test found that there was no difference in number of greens guessed between the Both and Description only

conditions ($Z=4344, p=0.533$), but that the Feedback condition was significantly different from the Both ($Z=6223.5, p<.001$) and Description only ($Z=5540, p<.001$) conditions.

As a replication of the finding from Experiments 2.1 and 2.2 that Feedback increases the number of greens guessed when both conditions have a description of the contingencies, the difference in greens guessed between the Both condition and the Description Only condition was tested. A Wilcoxon rank sum test found a significant difference between the Both and Description Only conditions, demonstrating that receiving feedback led to more green guessing ($Z=1200, p=0.03$) when both groups received a prior description of the die. Note that this test was done with Strict Maximizers removed. Note that with maximizers removed the Wilcoxon rank sum test does not find a difference between the Feedback Only and Description Only conditions ($Z=1427.5, p=0.247$) or the Feedback Only and Both conditions ($Z=2023.5, p=0.158$). The difference between these conditions (reported above), then, seems to arise largely from an increase in Strict Maximizing when a description is provided.

Participant Estimates

In addition to the number of greens guessed, we also tested participants' *estimates* of the number of greens guessed. This prediction is measured here using the same difference score calculated in Experiment 2.1: participants' actual number of green guesses is subtracted from their estimated number of green guesses. I submitted median difference scores to a Kruskal-Wallis one-way analysis of variance with condition as the between participant factor. The test reveals that estimates are significantly different from one another ($X^2(2)=12.77, p=.002$). The same Kruskal-Wallis test with participants' median raw estimates, instead of difference scores,

as the dependent variable showed only a trend towards a difference in raw estimates across conditions ($X^2(2)=4.37, p=.112$). This suggests that the difference in difference scores across conditions is driven largely by changes in green guessing rather than changes in raw estimates.

Pairwise comparisons were done using the Wilcoxon rank sum test and found that the difference across conditions in difference scores was due to a difference between the Feedback condition and the other conditions. In particular, the Feedback Only condition was drastically different from the Description Only condition ($Z=5551.5, p<.001$) and moderately different from the Both condition ($Z=5417, p=0.014$). The Description Only and Both conditions were not significantly different from one another ($Z=3658.5, p=0.147$).

Experiment 3.1

Binary Prediction Task Performance

Once again, I was unable to provide a non-parametric version of the omnibus ANOVA done in Experiment 3.1. However, I can examine the simple effects using non-parametric tests. For example, to break down the 3-way interaction from the parametric ANOVA, I took only the Probability Matching trials and divided the file based on Feedback. I then compared trials where participants were under load to trials where they were not underload for both the feedback and no feedback conditions using a Wilcoxon Signed Rank Test for participants in the Feedback condition on trials when they were assigned to probability match, with load (load vs. no load) as a within participants variable. This revealed no significant difference between conditions

($V=1909$, $p=.222$). The same test was done for participants in the No Feedback condition on trials when they were assigned to probability match, with load (load vs. no load) as a within participants variable. This revealed a trend towards a difference between conditions ($V=1819.5$, $p=.164$).

Experiment 3.2

Probe Accuracy

To determine if accuracy on the 3-back task varied as a function of Strategy Assignment the mean number of probes correct was submitted to a Wilcoxon signed rank test with Strategy Assignment as the within-participants factor. The test revealed no significant difference between conditions ($V=600.5$, $p=.543$).

Binary Prediction Task Performance

Recall that all participants in Experiment 3.2 receive feedback. Therefore, both IVs (Strategy Assignment and WML) are manipulated within participants. The typical non-parametric repeated measures ANOVA is the Friedman test, but this test does not preserve the difference between conditions and has substantially reduced power to detect effects relative to parametric tests (Zimmerman & Zumbo, 1993). As an alternative approach I rank transformed the data (in a way similar to that used in most of the non-parametric tests reported here) prior to analysis, and then ran a regular mixed model ANOVA on the transformed data. Please see Conover and Iman (1981) and Zimmerman and Zumbo (1993) for more information on this approach and its assets relative to the Friedman test.

The ranked data was submitted to a repeated measures ANOVA with number of greens guessed as the dependent variable and Strategy Assignment (Maximizing vs. Probability Matching) and WML (Present vs. Absent) as within participant factors. The ANOVA revealed a significant main effect of Strategy Assignment ($F(1,91)=368.53, p<.000, \eta^2=.80$) and a significant interaction ($F(1,91)=7.2, p=.008, \eta^2=.07$), but only a marginal effect of WML($F(1,91)=3.46, p=.066, \eta^2=.04$).

To break down the interaction, I submitted the total number of greens guessed in each block to a Wilcoxon signed rank test with WML (Load vs. No Load) as the within participant factor. This was done separately for blocks where participants were assigned to Probability Match versus blocks where they were assigned to Maximize. When asked to Probability Match, participants guessed significantly more green under load than when not under load ($V=869, p<.001$), however this was not the case when they were asked to Maximize ($V=258.5, p=.597$).

Experiment 3.3

Completion Time Data

To determine whether the speeding manipulation used in Experiment 3.3 was effective I compared the time taken to complete the die problem, the CRT2 and the Belief Bias tasks using a Wilcoxon rank sum test. For the Die problem, response time (in seconds) was submitted to this test with completion time as the dependent variable and Condition (Slow vs. Speeded) as the between participants variable. This revealed a marginally significant difference, $Z=1838, p=0.094$, with the speeded condition completing the problem faster than the slow condition. The same procedure for the CRT2 and Belief Bias problems revealed that in both cases the speeded

condition was significantly faster than the slow condition($Z=1727$, $p=0.03$ and $Z=1618$, $p=0.008$) respectively.

Binary Prediction Task Performance

A Wilcoxon rank sum test on number of green guessed for the described die problem, with number of green guessed out of 10 as the dependent variable and Condition (Slow vs. Speeded) as the between participants variable revealed a significant difference between conditions, $Z=1796.5$, $p=0.051$. The median green response in the slow condition was higher than in the speeded condition.

CRT2 Performance

A Wilcoxon rank sum test on performance for the CRT2 problems, with the number of problems correct as the dependent variable and Condition (Slow vs. Speeded) as the between participants variable revealed no significant difference between conditions $Z=1947$, $p=0.213$.

Belief Bias Performance

A Wilcoxon rank sum test on performance for the Belief Bias problems, with the number of problems correct as the dependent variable and Condition (Slow vs. Speeded) as the between participants variable revealed no significant difference between conditions, $Z=1976.5$, $p=0.288$.