

FACTORS AFFECTING LEARNING AND
DECISION-MAKING IN THE IOWA GAMBLING
TASK

Gordon Fernie

Thesis submitted to the University of Nottingham for the degree
of Doctor of Philosophy

Sept 2007

Why, it's just as if we were supplying drinks, with two fountains at our disposal: one would be of honey, standing for pleasure, the other standing for intelligence, a sobering unintoxicating fountain of plain salubrious water. We must get to work and make a really good mixture.

Philebus, Plato.

For Bob, Margaret, Harry and John

ABSTRACT

Damasio's somatic marker hypothesis (SMH; Damasio, 1994, 1996) integrates emotion with rational decision-making using evidence drawn from neurology, neuroscience and performance on a now widely cited decision-making test developed to model real-life in a laboratory setting (the Iowa Gambling Task; Bechara, Damasio, Damasio and Anderson, 1994). The SMH posits a critical input from an embodied emotional system (somatic markers) in making decisions in choice situations. But Damasio's consideration of how the undamaged brain interacts with the body has some interesting and somewhat controversial implications in the context of modern psychological research on choice behaviour. In interpreting behaviour on the IGT in accordance with the SMH three central assumptions have been made: a) that somatic markers indicate the goodness or badness of alternatives and without them decision-making cannot become optimal, b) this somatic biasing or guidance can occur unconsciously or in the absence of explicit knowledge, and c) that the system operates so as to maximize or achieve the best outcome in the long-term. The Experiments described in this thesis have explored the validity of the second and third assumptions and found that they are not accurately reflected in behaviour on the IGT. The importance of information about the IGT in the instructions participants receive suggested that explicit knowledge about the task is a more critical factor than any somatic input. No evidence of a somatic influence prior to the emergence of explicit knowledge sufficient to guide behaviour was found. Instead there were indications that knowledge precedes somatic activity on the IGT. Novel manipulations of the reinforcement contingencies in individual decks also revealed that immediate outcomes of choices are an important determinant of subsequent behaviour. Selection does not solely depend on long-term outcomes.

ACKNOWLEDGEMENTS

I would like to thank members of the Department of Psychology at Keele University for encouraging me to pursue a research career. Keele University has my gratitude for funding the first year of this thesis work. I am also grateful to Dr Richard J. Tunney for all his help and advice in the course of supervising this project. It would not have started without his encouragement for me to apply. I thank him too for allowing me to hang on to his coat tails long enough to make it into the University of Nottingham.

The thesis would not have been finished without the support and kind words of my family and friends. I could not have done it without them. Jen, you built me up when I was down. You provided a sounding board without becoming bored, and provided constant encouragement even through the dark days.

PUBLISHED PAPERS

Experiments 1, 2 and 3 are published in Fernie, G. and Tunney, R.J. (2006). The effect of reinforcer type and task instructions on learning in the iowa gambling task. *Brain And Cognition*, 60(1), 94-102.

Experiments 4 and 5 are published in Fernie, G. and Tunney, R. J. (2008). Decision-making in the Iowa Gambling Task. In F. Columbus (Ed.). *The Psychology of Decision-Making*. New York: Nova.

LIST OF CONTENTS

<i>Abstract</i>	<i>iii</i>
<i>Acknowledgements</i>	<i>iv</i>
<i>List of Figures</i>	<i>xii</i>
<i>List of Tables</i>	<i>xvi</i>

CHAPTER ONE

DECISION-MAKING, REASON AND EMOTION	1
1.1 A PHILOSOPHICAL BEGINNING	1
1.2 THE SOMATIC MARKER HYPOTHESIS	2
1.3 THE ASSUMPTION OF THE CAUSALITY OF SOMATIC MARKERS	6
1.3.1 <i>Damasio's theory of emotion and the somatic marker hypothesis</i>	6
1.3.2 <i>The case for the SMH from anatomy and SCR activity</i>	11
1.3.3 <i>Other theoretical and philosophical explanations</i>	21
1.3.4 <i>Reversal learning</i>	25
1.3.5 <i>Summary</i>	28
1.4 THE ASSUMPTION THAT DECISION MAKING PROCEEDS IN THE ABSENCE OF KNOWLEDGE	29
1.4.1 <i>The involvement of working memory</i>	31
1.4.2 <i>Other contributory factors to IGT performance</i>	37
1.4.3 <i>Summary</i>	40
1.5 THE ASSUMPTION OF A MAXIMIZING SYSTEM	40
1.5.1 <i>Rational choice theory</i>	43
1.5.2 <i>Animal choice experiments</i>	45
1.5.3 <i>Summary</i>	51
1.6 AN INTRODUCTION TO THE EXPERIMENTS	51

CHAPTER TWO

THE EFFECT OF TASK INSTRUCTIONS AND REINFORCER TYPE ON LEARNING	54
2.1 INTRODUCTION AND OVERVIEW	54
2.2 EXPERIMENT 1 A FAILURE TO REPLICATE NORMAL PERFORMANCE ON THE IOWA GAMBLING TASK - THE BEHAVIOUR OF A NORMAL SAMPLE IS NOT OPTIMAL.	56
2.2.1 <i>Introduction</i>	56
2.2.2 <i>Method</i>	57
2.2.3 <i>Results</i>	63
2.2.4 <i>Discussion</i>	66
2.3 EXPERIMENT 2 ‘STANDARD’ BEHAVIOUR REPLICATED USING ‘HINT’ INSTRUCTIONS	67
2.3.1 <i>Introduction</i>	67
2.3.2 <i>Method</i>	68
2.3.3 <i>Results</i>	69
2.3.4 <i>Discussion</i>	71
2.4 EXPERIMENT 3 THE EFFECT OF INSTRUCTION TYPE USING REAL MONEY REINFORCERS	76
2.4.1 <i>Introduction</i>	76
2.4.2 <i>Method</i>	77
2.4.3 <i>Results</i>	79
2.4.4 <i>Discussion</i>	83
2.5 CROSS EXPERIMENT COMPARISONS INSTRUCTIONS AFFECT PERFORMANCE ON THE IGT	85
2.5.1 <i>Session 1</i>	86

2.5.2	<i>Learning rates</i>	89
2.5.3	<i>Session 2</i>	90
2.5.4	<i>Learning rates in Session 2</i>	94
2.6	GENERAL DISCUSSION	94
CHAPTER THREE		
<hr/>		
	MANIPULATIONS OF LOSS FREQUENCY AFFECT LEARNING ON THE IGT	101
3.1	INTRODUCTION AND OVERVIEW	101
3.2	EXPERIMENT 4 ADVANTAGEOUS CARD SELECTION DIFFERS BETWEEN DECK COMPARISONS WHEN CHOICE IS BETWEEN TWO IOWA GAMBLING TASK DECKS	112
	3.2.1 <i>Introduction</i>	112
	3.2.2 <i>Method</i>	113
	3.2.3 <i>Results</i>	115
	3.2.4 <i>Discussion</i>	119
3.3	EXPERIMENT 5 MANIPULATION OF FREQUENCY (AND MAGNITUDE) OF LOSS AFFECTS LEARNING ON THE IGT	122
	3.3.1 <i>Introduction</i>	122
	3.3.2 <i>Method</i>	123
	3.3.3 <i>Results</i>	125
	3.3.4 <i>Discussion</i>	138
3.4	GENERAL DISCUSSION	140

CHAPTER FOUR

THE NATURE OF THE REINFORCEMENT SCHEDULE INFLUENCES

LEARNING	145
4.1 INTRODUCTION AND OVERVIEW	145
4.2 EXPERIMENT 6	
NO LEARNING ON A PROBABILISTIC VERSION OF THE IGT	151
4.2.1 <i>Introduction</i>	151
4.2.2 <i>Method</i>	152
4.2.3 <i>Results</i>	154
4.2.4 <i>Discussion</i>	158
4.3 EXPERIMENT 7	
THE VARIANT IGT – ORDER OF LOSSES AFFECTS LEARNING	166
4.3.1 <i>Introduction</i>	166
4.3.2 <i>Method</i>	167
4.3.3 <i>Results</i>	169
4.3.4 <i>Discussion</i>	174
4.4 EXPERIMENT 8	
THE VARIANT IGT – PROBABILISTIC SCHEDULES AFFECT	
LEARNING	177
4.4.1 <i>Introduction</i>	177
4.4.2 <i>Method</i>	178
4.4.3 <i>Results</i>	179
4.4.4 <i>Discussion</i>	182
4.5 GENERAL DISCUSSION	188

CHAPTER FIVE

LEARNING ON THE IGT FOLLOWS EMERGENCE OF KNOWLEDGE BUT NOT DIFFERENTIAL SOMATIC ACTIVITY

5.1	INTRODUCTION AND OVERVIEW	194
5.1.1	<i>When does a differential aSCR to deck types emerge?</i>	195
5.1.2	<i>When does knowledge sufficient to guide IGT selection behaviour emerge?</i>	199
5.1.3	<i>Summary</i>	207
5.2	EXPERIMENT 9 DIFFERENTIAL SKIN CONDUCTANCE RESPONSES BETWEEN DECK TYPES EMERGE AFTER PARTICIPANTS DISPLAY KNOWLEDGE ON THE IOWA GAMBLING TASK.	208
5.2.1	<i>Introduction</i>	208
5.2.2	<i>Method</i>	209
5.2.3	<i>Results</i>	217
5.2.4	<i>Discussion</i>	
5.3	GENERAL DISCUSSION	249

CHAPTER SIX

	GENERAL DISCUSSION	256
6.1	SUMMARY OF EMPIRICAL RESULTS	256
6.2	IMPLICATIONS FOR FUTURE USE OF THE IGT	261
6.2.1	<i>The IGT as a test of the SMH</i>	261
6.2.2	<i>The IGT as a test of decision-making</i>	265
6.2.3	<i>Implications for the SMH</i>	269

6.3	APPLICATION OF THE IGT TO OTHER CHOICE ENVIRONMENTS	271
6.4	FUTURE EXPERIMENTS	273
6.5	CONCLUSION	276
<u>APPENDIX A</u>		
<hr/>		
	FIXED SCHEDULE OF LOSSES USED IN EXPERIMENTS 1, 2, 3, 4, 5 AND 6	279
<u>APPENDIX B</u>		
<hr/>		
	FIXED SCHEDULE OF LOSSES USED IN THE MANIPULATED DECKS OF EXPERIMENT 5	281
<hr/>		
<u>APPENDIX C</u>		
	FIXED SCHEDULE OF WINS ON THE VARIANT IGT USED IN EXPERIMENT 7	284
<u>APPENDIX D</u>		
<hr/>		
	QUESTIONS USED IN THE SPECIFIC QUESTION CONDITION OF EXPERIMENT 9	287
	REFERENCES	289

LIST OF FIGURES

Figure 1.1:	The Harvard Game environment.	48
Figure 2.1:	Screen shot from the computerized gambling task	59
Figure 2.2:	Mean net score across 20-trial blocks for each session in Experiment 1	64
Figure 2.3:	Mean net score across 20-trial blocks for each session in Experiment 2	70
Figure 2.4:	Mean number of cards selected from each deck in a) Experiment 1 and b) Experiment 2	74
Figure 2.5	Mean net score across 20-trial blocks for each session in Experiment 3	80
Figure 2.6:	Mean number of cards selected from each deck in condition A (No Hint instructions) and condition B (Hint instructions)	84
Figure 2.7	Mean number of cards selected from each deck by block for each experimental condition	92
Figure 2.8a	Mean number of card selections in each block in Experiments 1 and 2	97
Figure 2.8b	Mean number of card selections in each block in Experiment 3	98
Figure 3.1	Mean number of advantageous selections across twenty ten-trial blocks in each experimental group	117
Figure 3.2:	Mean net score across ten twenty-trial blocks in each experimental condition	127
Figure 3.3:	Mean number of cards selected from each deck across all 200 trials in each condition.	131
Figure 3.4:	Mean number of cards selected from each deck across the first 100 and second 100 trials in the A) the Decreased Frequency condition and B) the Increased Frequency condition.	133
Figure 3.5a:	Mean number of cards selected from the disadvantageous decks in each condition.	134
Figure 3.5b:	Mean number of cards selected from the advantageous decks in each condition.	135

Figure 3.6:	Mean number of card selections across block from the disadvantageous decks with probability of loss of 0.3.	137
Figure 4.1:	Mean net score across 20-trial blocks for each session in Experiment 6 (Probabilistic reinforcement schedules).	155
Figure 4.2:	Mean total cards selected from each deck in each session of Experiment 6.	157
Figure 4.3:	The range of participants' net scores in both sessions of Experiment 6.	159
Figure 4.4:	Mean received value and mean number of selections in each 20-trial block for participants with a mean net score ≥ 0 .	163
Figure 4.5:	Mean received value and mean number of selections in each 20-trial block for participants with a mean net score < 0 .	164
Figure 4.6:	Mean net score across 20-trial blocks for each session in Experiment 7 (reversed IGT with fixed schedules).	170
Figure 4.7:	The range of participants' net scores in both sessions of Experiment 7.	172
Figure 4.8:	The number of cards selected from each deck in each session of Experiment 7.	173
Figure 4.9:	Mean net score across 20-trial blocks in Experiment 8 (reversed IGT with probabilistic schedules).	180
Figure 4.10:	The range of participants' net scores in both sessions of Experiment 8.	181
Figure 4.11:	Mean number of cards selected from each deck in each session of Experiment 8.	184
Figure 4.12:	Mean received value and mean number of selections in each 20-trial block for participants with a mean net score < 0 .	186

Figure 4.13	Mean received value and mean number of selections in each 20-trial block for participants with a mean net score ≥ 0 .	187
Figure 5.1:	Maia and McClelland's (2004) decision tree for classifying participants' knowledge about the IGT.	203
Figure 5.2:	Mean net score across 10-trial blocks.	218
Figure 5.3:	The number of participants in the Specific Question Group who displayed knowledge that one of the best decks is the best deck on each measure.	229
Figure 5.4:	Calculated versus actual expected value on each deck after 100 trials for each participant.	231
Figure 5.5:	Estimated versus actual expected value on each deck after 100 trials for each participant.	232
Figure 5.6:	The number of participants at each question period who selected each deck as the one deck they would choose if forced to only pick from one.	233
Figure 5.7:	Mean rating for each deck across question period.	234
Figure 5.8:	Mean proportion of cards selected from each deck in (a) the pre- and post-knowledge periods for participants who displayed knowledge ($n = 27$), and (b) the comparable periods for participants who did not display knowledge ($n = 5$).	236
Figure 5.9:	Mean aSCRs for each deck in each group, (a) across all selections, (b) in selections prior to and following knowledge expression in those participants who displayed knowledge ($n = 27$), and (c) the equivalent figure to B for participants who did not demonstrate knowledge ($n = 5$).	240

Figure 5.10: Mean rSCRs for each deck in each group, (a) across all selections, (b) in selections prior to and following knowledge expression in those participants who displayed knowledge ($n = 27$), and (c) the equivalent figure to B for participants who did not demonstrate knowledge ($n = 5$). 243

Figure 5.11: Mean pSCRs for each deck in each group, (a) across all selections, (b) in selections prior to and following knowledge expression in those participants who displayed knowledge ($n = 27$), and (c) the equivalent figure to B for participants who did not demonstrate knowledge ($n = 5$). 247

LIST OF TABLES

Table 2.1	Summary of mean learning rates by session in each experiment	90
Table 2.2	Mean number of cards selected from each deck per session by Instruction Type and Reinforcer Type	96
Table 3.1	The deck contingencies (reward magnitude and mean loss magnitude in pence and loss frequency) in each Experimental condition	114
Table 3.2	Mean number of advantageous selections in each experimental group	116
Table 3.3:	Mean learning rate (<i>b</i>) in each experimental condition across first 100 trials, second 100 trials and all trials	118
Table 3.4:	Mean net score in each condition in the first and second hundred trials, and over the whole experiment	126
Table 3.5:	Mean learning rate (<i>b</i>) in both experimental conditions across the first 100 trials, second 100 trials and all trials.	128
Table 3.6:	Mean selection rate from each deck in the first and second 100 trials in each condition.	134
Table 4.1:	Mean probability of loss and expected value (in pence) for each deck in each session in Experiment 6.	154
Table 4.2:	Mean expected value and probability of loss for each deck in each session of Experiment 6 for participants with a positive or negative net score.	161
Table 4.3:	Mean probability of win and expected value for each deck in each session in Experiment 7.	174

Table 4.4:	Mean probability of win and expected value for each deck in each session in Experiment 8.	182
Table 4.5:	Mean expected value and probability of win for each deck in each session of Experiment 8 for participants with a positive or negative net score.	185
Table 4.6:	Deck contingencies in Peters and Slovic's (2000) modified gambling task.	192
Table 5.1:	Summary of participants' knowledge expression in Bechara et al (1997a).	200
Table 5.2:	Summary of participants' knowledge expression in Maia and McClelland's replication condition.	202
Table 5.3:	Knowledge assessment for General question group	221
Table 5.4:	Knowledge assessment for Specific question group using 'partial' grouping (either C or D had to have received the best score on each measure) or 'both' grouping (C and D had to have received the best scores on each measure).	225
Table 5.5:	Knowledge assessment for Specific question group using 'partial' grouping (either deck with the highest net value at the time of questioning received the best score on each measure) or 'both' grouping (both decks with the highest net value at the time	

CHAPTER ONE

DECISION-MAKING, REASON AND EMOTION

1.1 A PHILOSOPHICAL BEGINNING

The Dualistic tradition of splitting the human psyche into a rational and emotional or appetitive parts began in Classical Greek philosophy (Barnes, 1998a). As a general definition Greek philosophy considered the psyche to be the nature of life marking out what is animate from the inanimate (Barnes, 1998b). Later dualistic formulations have been attempts to resolve the apparent discrepancy between a representation of a mental life found through introspective enquiry with the quantifiable nature of the known physical universe found through observation and measurement (e.g. Descartes, 1965). Although this dualistic philosophy is no longer held by the majority of academic or scientific thinkers it has had a profound impact on Western culture's lay understanding of human life (Rachlin, 1989; Churchland, 1993). For example, many people still consider human behaviour to be governed by two competing forces: one reason-based that operates through careful long-term planning and the weighing of the pros and cons of a given situation; and one emotion-based that operates in the moment regardless of consequence (Rachlin, 1989; Damasio, 1994). In this conception the rational force checks the emotional one and leads to beliefs that, for example, one can "lose control" when angry.

Recently, Antonio Damasio has advanced a theory that has resurrected the challenge against dualism (Damasio, 1994, 1996, 1999). His somatic marker hypothesis (SMH) is an attempt to integrate emotion and reason using evidence drawn mainly from the behaviour of certain neurological populations and compared to that

of their healthy controls. But Damasio's consideration of how the undamaged normal brain interacts with the body has some interesting and somewhat controversial implications in the context of recent psychological research on choice behaviour. The SMH will be briefly outlined, as will the Damasio's related theory of emotion, before the implications are critically evaluated.

1.2 THE SOMATIC MARKER HYPOTHESIS

The somatic marker hypothesis (SMH) integrates emotional processing as a vital component of decision-making, a process that in dualistic tradition is the preserve of reason. In the SMH decisions in uncertain or ambiguous situations are influenced, guided or biased by bodily representations of the potential goodness or badness of the available alternatives. They direct the individual towards choices that have previously been good or beneficial, or guide them away from those that have been bad or detrimental. These bodily representations are called somatic markers (somatic from the Greek *soma* meaning body) and they are the products of visceral activity that is also a fundamental aspect of emotional processing. Indeed the SMH is the central pillar in Damasio's theory of emotion, itself an attempt to provide a holistic conception of human functioning. Although, as will be elaborated below, neither this theory of emotion nor its central hypothesis are novel (Dunn, Dalgleish and Lawrence, 2006), both have captured the attention of researchers across diverse fields (as well as the general public) and provoked much enquiry into the relationship between emotion and choice. This reason alone makes the SMH worthy of further enquiry, but it is the assumptions made about human choice, and in the experimental tools used to test them, that are of interest to cognitive and behavioural researchers.

Before discussing the assumptions of the SMH it is necessary to give a short account of a decision-making task that influenced the development of this hypothesis and is used to provide empirical support. This task is the Iowa Gambling Task (IGT; Bechara, Damasio, Damasio and Anderson, 1994). In the IGT participants must make a series of selections from four decks of cards (in its original formulation the task was manual and the decks were on a table in front of participants). The participants' task is to make as much money as possible, or avoid making a loss. They are free to select from any deck and in any order. Each selection wins money but some cards contain penalties. The decks vary along three dimensions: the immediate gain, the expected long-term gain and the schedule of losses. Each selection of a card from deck A or B gives a larger immediate reward (\$100) than selection of a card from deck C or D (\$50). But the penalties in A and B are also larger meaning that the cumulative loss in decks A and B exceeds the cumulative gain, whereas the opposite is true for C and D. After ten card selections from decks A or B the cumulative loss is -\$250 whereas after ten card selections from decks C or D the cumulative gain is \$250. So decks A and B have high immediate gain but long-term loss, whereas decks C and D have low immediate gain but long-term gain. In this way decks C and D can be considered the advantageous (or good) alternatives while A and B are the disadvantageous (or bad) alternatives. The third dimension in which the decks vary is the schedule of losses. The negative expected value (the probability of win multiplied by win amount minus the probability of loss multiplied by loss amount) for decks A and B is achieved with five losses of an average -£250 in every ten cards selected from deck A whereas in deck B there is one loss of -£1250 in every ten cards selected. Similarly, the losses on decks C and D mirror the losses on A and B, but their magnitude is reduced (-£50 and -£250 respectively) to result in a positive expected value from their selections.

The SMH has three main assumptions which will be briefly stated then explored further in the subsequent sections. The first assumption is that there is a causal role for somatic markers in decision-making. Without somatic markers decisions are not guided towards the good options. This assumption is clear from the interpretation of the behaviour of healthy controls and patients with damage to the ventromedial prefrontal cortex (VMpfc) on the IGT (Bechara, Tranel, Damasio and Damasio, 1996; Bechara, Damasio, Tranel and Damasio, 1997a). While controls learn to choose advantageously (i.e. for a profit in the long-term), and develop somatic markers (measured using skin conductance responses or SCRs) that differentiate between good and bad alternatives, patients with VMpfc damage do not. They select the alternatives that are disadvantageous in the long-term, but that offer a larger immediate gain, and they do not develop differential somatic markers. This has been interpreted as showing that without somatic markers advantageous choices are not made, or in another way, somatic markers are necessary for advantageous choice behaviour.

The second assumption comes from a strong version of the SMH where somatic markers bias decision-making covertly in the absence of knowledge about the contingencies of the choice environment or indeed knowledge sufficient to guide behaviour. The empirical result on which this claim is based stems from an earlier claim that the IGT is cognitively impenetrable, or that participants do not acquire a full understanding of its contingencies despite superior memory and IQ scores (Bechara et al, 1994). Bechara et al (1997a) recorded SCRs and asked participants questions about their knowledge of the task as they completed the IGT. The authors claimed that SCRs generated prior to deck selection differentiated between the advantageous and disadvantageous decks before participants had any conscious

knowledge about the nature of the deck types. Tranel, Bechara and Damasio (1999, p.1055) claimed:

“... in normal individuals non-conscious biases guide reasoning and decision-making behaviour before conscious knowledge does, and without the help of such biases, overt knowledge may be insufficient to ensure advantageous behaviour ... we believe that the autonomic responses detected in our experiment (especially those evident [before participants articulate any hunch]) are evidence for a non-conscious signalling process.”

This claim assumes that somatic activity precedes knowledge and indeed may help in the acquisition of knowledge. It can be interpreted as an instantiation of implicit learning, itself a controversial concept in cognitive psychology (e.g. Shanks and St. John, 1994).

The final assumption is that the somatic system somehow calculates “goodness” or “badness” of a choice based on its long-term, average overall outcome. Or in other words the somatic system maximises. This is implied in the first assumption – in healthy individuals somatic markers guide decision-making towards the options with the positive expected value. It is also implied in the standard analysis of IGT behaviour – net score. This is calculated by summing the number of cards selected from the advantageous decks (C and D) and subtracting the sum of the number of cards selected from the disadvantageous decks (A and B), regardless of one decks’ schedule of losses. Positive scores indicate a preference for the advantageous decks while negative scores indicate a preference for the disadvantageous decks. Thus, the analysis of choice on the IGT is solely in terms of expected value.

This thesis looks at the second and third assumptions in the context of normal, i.e. undamaged, decision-making behaviour in healthy participants. In particular it considers why maximizing behaviour is found and whether choice is dependent on expected value. The related claim that choices proceed in the presence of differential somatic activity but the absence of knowledge is also examined in the last chapter. But before that point the support and criticism of these three assumptions of the SMH are explored.

1.3 THE ASSUMPTION OF THE CAUSALITY OF SOMATIC MARKERS

The support for the assumption of a causal role for somatic markers in decision-making comes mainly from observations of neurological patients and the location of their brain damage in the context of the SMH. To evaluate this support it is necessary to explore Damasio's theory of emotion in more detail as well as describe how the SMH fits into it.

1.3.1 Damasio's theory of emotion and the somatic marker hypothesis

The SMH arose out of the work of Damasio and his colleagues in their attempts to quantify and explain the impairment in the real-life decision-making of people with damage to the VMpfc. These patients show deficiencies in their decision-making, the results of which often include the inability to retain pre-morbid employment, unsuccessful management of personal finances and entering unsound financial investments, the breakdown in pre-morbid relationships and new relationships that do not last (Damasio, 1994; Goel and Grafman, 1997). Damasio and colleagues found no impairments in a series of standard neuropsychological tests yet the patients

showed abnormal decision-making behaviour, especially when the decisions involved personal or social matters (Damasio, 1994). This behaviour remained enigmatic until one patient reported that when viewing stimuli that would have pre-morbidly produced an emotional reaction, he felt none. This observation was confirmed in other patients with VMpfc damage (Damasio, 1994) and led Damasio to link abnormal real-life decision-making with emotional blunting. For example, when presented with disturbing images containing scenes of violence people with VMpfc damage reported they did not experience the emotional reaction they would have done pre-injury. This was measured objectively using SCR recording (Damasio, 1994). The impairment in real-life decision-making and impaired emotional responding led Damasio to develop the SMH and his wider theory of emotion. The SMH is essentially the operation of this architecture in decision environments. But since much of human activity involves making decisions (Rachlin [1989] has argued that all behaviour is choice behaviour) it can be seen as a general theory of human functioning and so the two will be discussed synonymously.

Damasio (1994, 1996; Bechara and Damasio, 2005) conceives of an emotion as the changes in body and brain states that are triggered by the operation of a dedicated brain system that processes the meaningful events in the external or internal environment. Although changes in neural activity occur (e.g. neurotransmitter release and the modifications of connections between the neural representations of stimuli), the crucial part of this emotional system is that the perception of the meaningful event produces changes in the body's physiology (e.g. heart rate, endocrine release). When the status of the somatic state is signalled back to the brain it acts as the evaluation of the meaningful event. In this theory a feeling is the perception of this change.

In the latest revision of the SMH (Bechara, 2003; Bechara and Damasio, 2005) meaningful events are called primary or secondary inducers. Primary inducers are defined as “innate or learned stimuli that cause pleasurable or aversive states” and “concepts or knowledge that through learning automatically and obligatorily elicit emotional responses” (Bechara and Damasio, 2005, p340). In previous formulations of the theory these two features of primary inducers were separately called primary and secondary emotions (Damasio, 1994; 1996). Primary emotions were conjunctions of the somatic state and the commensurate brain state, while secondary emotions were associations between primary emotions and categories of stimuli in the environment. But the reformulation of the SMH has grouped them together presumably to keep the terminology simple and to separate them from a new concept – secondary inducers. These are essentially thoughts and memories about primary inducers.

The concepts and language used to describe these meaningful events are very close to the concepts of primary and secondary reinforcers used in the long history of associative learning and studies of animal choice. This may be coincidental but it is reasonable to assume that Damasio’s theory was influenced by these earlier descriptions. Even if this is not the case, much of the neural architecture that underlies Damasio’s theory has also been explored in animal learning but with different interpretations of systemic interactions. Before discussing these differences it is necessary to give a brief account of the brain anatomy involved in representing an emotional event in the SMH. Bechara and Damasio (2005) have detailed the series of events and important structures as follows:

1. The features of a primary inducer and the somatic state associated with it are linked in the amygdala. An explicit route for processing the features of the inducer is via early sensory and higher-order association cortices, while

processing via the thalamus offers an implicit route. The somatic state is evoked by efferent connections from the hypothalamus and autonomic brainstem nuclei to the body resulting in changes in the internal milieu (“all biochemical processes occurring in an organism at any given moment, Damasio, 1994, p. 118), and in other structures having efferent connections with the body (e.g. ventral striatum, periaqueductal grey, brainstem nuclei). The nature of these changes charges the somatic state with a degree of polarity that represents the “goodness” or “badness” of the primary inducer.

2. Once somatic states are induced, the condition of these states is signalled to the brain where the brainstem nuclei and some somatosensory cortices (e.g. insula/secondary somatosensory cortex, primary somatosensory cortex, and cingulate cortex) represent the somatic state.

3. Once this primary inducer to induced somatic state relationship has been experienced once, Bechara and Damasio (2005) contend the pattern of brain activity for that somatic state is formed. This then allows secondary inducers (thoughts or memories) associated with a primary inducer to re-activate that pattern of somatic activity (although at a weaker level). This association is dependent on VMpfc. It holds the link between temporally congruent brain activity in a) areas that represent categories of events (based in higher-order association cortices), b) the structures that generate somatic states, and c) the neural patterns that represent the feeling of the somatic state. So the VMpfc operates as the connector between brain areas where knowledge of events is represented and areas where the somatic patterns of feeling for those events or situations is represented. In effect, it is analogous to a telephone switchboard operator.

4. When any part of this system is re-activated for example by encountering a primary inducer in the environment or thinking about one (secondary induction), Bechara and Damasio (2005) envisage that the system is reactivated, in whole or in part. This happens in one of two ways, but always involves a re-activation of the somatic state and consequent modulation of brain activity. Either the somatic state re-activation occurs in the body (the body loop) or it occurs in a neural representation of the body state (the “as-if body” loop).

5. Dependent on which parts of the system are re-activated, the resulting somatic activity can operate unconsciously or be perceived consciously as a good or bad feeling, or as an incentive or alarm signal. The SCR activity found on the IGT by Bechara et al (1997a) is envisaged to reflect this unconscious processing (Bechara and Damasio, 2005, p. 341).

The SMH is an instantiation of Damasio’s theory of emotion in a decision environment. Options for action become secondary inducers when they are considered and this instils a somatic valence by linking representations of their outcomes (in the first instance), previous outcomes, or outcomes in similar situations with the somatic activity that is (in the first instance) or was invoked in their presence. In this way the system for processing emotion and for evaluating alternatives (i.e. making decisions) are fundamentally intertwined. The main support for the interactions within this neural architecture has come from examination of what happens when parts of this system are damaged. And the empirical test used for this examination has been the IGT. This evidence is explored next.

1.3.2 The case for the SMH from anatomy and SCR activity

A significant strength of the SMH is this specification of its neural architecture. One of the main reasons for the critical role of the VMpfc is that it receives input from all sensory regions including somatosensory cortices (Rolls, 2004). These regions are also interconnected and among them have access to information about the whole body (Damasio, 1994, p.180). The VMpfc also has efferent and afferent links with several bio-regulatory sectors of the brain including brainstem neurotransmitter nuclei and some in the basal forebrain, as well as the amygdala, anterior cingulate and hypothalamus. Through these connections the VMpfc has direct links to every area for motor or chemical response in the brain. This architecture has been well described by others (Ongur and Price, 2000) although not necessarily interpreted in the same way (e.g. Rolls, 2004, 2005).

The consequences of lesions to the VMpfc described by Damasio (1994, 1996) are abnormal real-life decision-making that appears not to take the long-term into account, an absence of emotional reactivity where it would be expected, and normal performance on standard neuropsychological tests. On the IGT patients with VMpfc damage preferentially select the disadvantageous decks and do not generate anticipatory SCRs (aSCRs) prior to making a selection, where healthy controls' aSCRs differentiate between the expected value of the decks (Bechara et al, 1994, 1996, 1997a).

Consistent with the SMH framework bilateral amygdala damage also impairs decision-making (Tranel and Hyman, 1990; Nahm, Tranel, Damasio and Damasio, 1993) and emotional processing. The amygdala's role in the SMH is to provide the affective link to situational stimuli. Consistent with this role, amygdala damage results in an absence of physiological activity in a fear conditioning task (Tranel,

Bechara, Damasio and Damasio, 1996) as well as prior to and following card selections on the IGT (Bechara, Damasio, Damasio and Lee, 1999). Performance measured by net score was impaired relative to controls. Coupled with the VMpfc patient data it appears that there is support for the SMH framework and that the absence of somatic markers (at least as measured by SCRs) is associated with impaired IGT performance. Indeed the claim that the IGT measures emotion-based learning comes directly from these results. In support, Peters and Slovic (2000) have also reported that self-reported measures of affective reactivity added explanatory power to predictions of the choice among the decks in a modified version of the IGT. This result suggests that affective processing can provide some contribution to IGT performance.

However, it is by no means clear that the relation between somatic markers and IGT performance is anything other than correlational. It is possible that rather than reflecting the presence of an emotional biasing system, the aSCRs in healthy participants reflect the development of understanding of the task contingencies. This possibility will be returned to in section 1.4. The causality debate revolves around precedence and can also be thought of as a restatement of the classic chicken or the egg conundrum. Do somatic markers guide choice behaviour prior to consciously available knowledge or is this knowledge achieved earlier and the somatic markers a result of it? The support for the strong version of the SMH has come from two lines of investigation. The first line, and strongest support, comes from clinical studies where patients with little SCR activity also show behavioural impairment on the IGT, in the absence of impairment on other neuropsychological tests. The second is linked and has explored individual differences in SCR activity and behavioural performance.

Here the assumption is that poor deck selection is based, in part, on low SCR activity. Both methods have support in the literature but the results are not uniform.

Bechara et al (1999) replicated their aSCR and post-selection SCR results in controls and two patient groups. An amygdala-damaged patient group displayed the same deficit in aSCRs and in IGT performance as VMpfc-damaged patients, but additionally displayed no post-selection SCRs. These results implicated the amygdala in the somatic marker circuit. Further support for the SMH was provided by the non-differential aSCRs displayed by some normal participants ($n = 3$, 23%) who did not choose advantageously. However, one patient with VMpfc damage (20% of their population) also performed advantageously, though they did pick more bad cards than healthy controls, but no mention is made of this individual's aSCRs. If they were similar to the rest of the VMpfc-damaged sample then this may suggest that IGT performance is not dependent on functioning somatic markers. In a comparison of anticipatory and post-selection SCRs from early and late periods of the task, a similar pattern was found between controls and the VM-damaged group in post-selection SCRs (smaller magnitude in later trials implying habituation), whereas only controls showed the expected increase in aSCRs. However, there are some problems with this account. No description of how early and late periods were defined is made, while the authors provided neither graphical representation of the data involved nor any statistical examination of these differences.

Bechara and Damasio (2002) also looked at aSCR in four periods while participants carried out their progressive version (A'B'C'D') of the IGT. In this version the worse decks have increasing punishment with every ten choices, while best decks have increasing reward. The periods were the "pre-punishment" (trials ~1 – 10), "pre-hunch" (trials ~11 – 20), "hunch" (trials ~21 – 60) and "conceptual" (trials

~61 – 100) stages obtained from Bechara et al's (1997a) analysis of their participants' knowledge of the IGT contingencies. Healthy controls' aSCRs, combined across the disadvantageous decks, increased between periods 1 and 2 then remained relatively constant. The large punishment on deck B is likely to fall in this second period and may form part of the explanation for this. Little change was observed in aSCRs for the advantageous decks. However, no analyses of these observations are reported making any conclusions based on them tentative. When participants' SCR were split into groups determined by their behavioural performance, a three-way interaction of deck type by group by block was found but not explored. Figure 5 in Bechara and Damasio (2002) suggests that aSCR for the bad decks increased from the first to the second block for the impaired participants too. This suggests that these behaviourally impaired individuals do not respond to the somatic marker in the same way as unimpaired individuals, which further suggests that the somatic marker alone may not be enough for successful performance on the IGT. Annoni, Ptak, Caldara-Schnetzer, Khateb and Pollerman (2003) found a similar result. They describe a patient (M.F.) with a cerebellar lesion who shares the affective blunting and preference for the disadvantageous decks reported in patients with VMpfc damage, yet who generated greater aSCRs to the disadvantageous decks as reported by Bechara et al (1996, 1997a) in healthy controls. This might suggest that the aSCRs alone are insufficient to guide behaviour.

This possibility has some support. While the majority of published studies of VMpfc-damaged patients' IGT behaviour where SCR have also been recorded have replicated the absence of aSCRs, the difference in aSCRs between deck types has not been replicated in all samples of healthy controls. Some studies have replicated it (Tomb, Hauser, Deldin and Caramazza, 2002; Suzuki, Hiota, Takasawa and

Shigemasa, 2003; Crone, Somsen, Van Beek and Van Der Molen, 2004) while others have not (Campbell, Stout and Finn, 2004; Kleeberg, Bruggiman, Annoni, Melle et al 2004). Indeed studies in which physiological measurements were recorded during IGT performance have often found different results. For example, Suzuki et al (2003) used a student sample and found a difference in aSCRs between the advantageous and disadvantageous decks. They also found a difference in post-selection SCRs between deck types, and between trials with and without punishment. In some support of the SMH a correlation was found between post-selection SCR-level in the first half with selections from A and B in the second half. However, both high and low post-selection SCR groups reduced disadvantageous selections in the second 40 trials, and the low SCR group's advantageous selections were above chance levels. There was no relationship between aSCRs and performance.

Carter and Pasqualini (2004) suggested that the SMH implies that those people with stronger somatic responses would show faster learning on the IGT. In contrast to Suzuki et al (2003), no correlation in post-selection SCRs with their performance measure (the amount of money won) was found. Instead, and again unlike Suzuki et al, aSCR level did correlate with money won. Money won also correlated with a measure of neuroticism. This relationship was reduced when aSCR level was partialled out, indicating that aSCR size may mediate this relationship. The results of these studies suggest that SCR level does relate to IGT performance in some way, but differences in which SCR-type related to the performance measure clouds understanding.

One potential factor may be the method used to measure aSCRs. Campbell et al (2004) did not find any difference in aSCRs between deck types in their healthy controls or patients with Huntingdon's disease. This may be because their

participants did not show preference for the advantageous decks until the fourth trial block and even then the preference was not large. Alternatively, unlike Bechara et al (1996, 1997a, 1999) Campbell et al (2004) used a computerized version of the IGT and the timing of SCR recording was automatic. Kleeberg et al (2004) also reported no difference in aSCRs between deck types and used the computerized IGT with automatic SCR recording. Carter and Pasqualini (2004) noted that they switched from a computerized IGT to the manual version after pilot studies found that participants stopped generating SCRs after a few selections. But if SCRs are indicative of somatic markers and if their presence is dependent on test medium it begs the question of what else influences SCRs and if this also affects somatic markers. Yet aSCR differences between deck types have been reported using a computerized task (Suzuki et al, 2003; Crone et al, 2004).

Crone et al (2004) presented a detailed analysis of physiological correlates of IGT performance. They explored the findings of Bechara and colleagues (Bechara et al 2000, Bechara and Damasio, 2002) that a proportion of healthy controls performed like VMpfc patients. Crone et al hypothesised that an absence of somatic activity was linked to poor performance on the IGT, a similar suggestion to that put forward by Peters and Slovic (2000). They split their participants into three groups based on their behavioural performance and then compared these groups on the somatic measures recorded during the IGT – heart rate and SCR. Crone et al found that mean skin conductance level (SCL) compared to individual baseline was higher preceding a choice from A and B in both the moderate and best performing group. For the best performers this interacted with punishment frequency so that mean SCL was higher prior to choosing A rather than B. Similarly, in good performers heart rate slowed more prior to choices from decks A and B (heart rate slowing is greater when

preparing for an aversive event; Somsen, Van der Molen, and Orlebeke, 1983). This was not found in the worst or moderate groups. The difference in anticipatory SCL was also positively correlated with the number of advantageous selections.

Differences in both somatic measures post-selection were found to be due to receiving punishment rather than reward, and this tended to be higher when the punishment was infrequent. However, in this case there were no differences between performance groups. These results from a large sample ($n = 96$) with two somatic responses recorded suggest that somatic activity prior to making selections on the IGT is linked to learning on the IGT.

However, although an impressive and detailed study, Crone et al's (2004) results cannot inform on whether somatic markers precede knowledge or vice versa. They provide no quantification of change in physiological arousal across time (no doubt due to the statistical nightmare it produces) and this makes identifying the direction of any mechanism for feedback harder to extract i.e. are good performers faster learners (do they understand the contingencies earlier)? Crone and van der Molen (2004) did find that more of their best performing participants had greater knowledge of the decks than participants who did not perform as well. It is also feasible that the best-performing group have higher post-selection somatic activity earlier on the task and this influences anticipatory activity. However, this possibility cannot be examined from Crone et al's data as presented. It would make sense to suggest that anticipatory somatic activity migrates following feedback, and perhaps it does this faster in good performers. But does it do it faster because their physiological arousal system facilitates that or because they develop an understanding of the task faster? Crone et al (2004, p.539) consider this question citing Bechara et al (2002):

“...primary inducers are the reward and punishment events. Secondary inducers are the thoughts, and memories of prior choices, preceding a decision. These are secondary because they can only be formed after experience with the consequences (positive or negative) of choices in that situation.”

So Crone et al state that the IGT performance in the worst group (“decision-making impairment”) “arise from a weak somatic response generated by secondary (i.e. acquired) inducers”. They cite Bechara et al (1997a) and consider whether their poor performers rely on a more explicit learning strategy (rather than listening to their gut feelings - a similar contention to that proposed by Evans, Kemish and Turnbull, 2004). But the assumption is that Bechara et al’s (1997a) interpretation was correct and aSCR changes precede knowledge expression.

While this is a possibility there is also evidence to suggest that a somatic marker is not required for successful performance on the IGT at all. This evidence comes from studies using the IGT with patients with clinical damage to various parts of their affective system. North and O’Carroll (2001) found no difference in behaviour between healthy controls and patients with spinal chord damage (a complete transverse lesion at the C6 level). In these patients no connections between the peripheral nervous system and the brain exist. One would expect the absence of peripheral somatic activity would interfere with IGT performance but it did not. This result leaves two ways out for the SMH. First, other somatic connections with the brain are more important than afferent feedback through the spinal chord e.g. connections via the bloodstream, the vagus and other cranial nerves. These routes are important within the SMH according to Bechara and Damasio (2005, p.342). Second, reliance on the “as-if body” loop of the hypothesis. This is a possibility given that North and O’Carroll’s patient participants were tested some time after their injury was

sustained. However, if the “as-if body” loop is invoked it does not explain why the body loop mechanism is utilised in learning for other non-injured participants. It is also problematic that as the IGT provides a new situation for these participants (i.e. there is no stored representation of somatic feedback) the SMH framework would presumably require activation of the body loop. But it may be the case that the time between injury and IGT exposure was long enough for compensatory activation of the “as-if body” loop in new situations.

Heims, Critchley, Dolan, Matthias et al (2004) have also reported IGT performance in patients with pure autonomic failure (PAF) that was superior to their control participants. PAF results in peripheral denervation of autonomic neurons meaning no peripheral autonomic input. Dunn et al (2006) point out prolonged PAF results in changes to the structure of the brain regions involved in the representation and regulation of body state (loss of grey matter volume in anterior cingulate and insula). Such atrophy would compromise both the body and “as-if body” loops of the SMH meaning that normal performance on the IGT is problematic for the theory. However, like North and O’Carroll (2001), Heims et al suggest that other forms of somatic feedback are still intact in these patients. Together these studies suggest that disruption of the brain’s connection with the autonomic nervous system does not affect performance on the IGT in the direction predicted by the SMH account. However, the results from both studies can still fit the SMH by virtue of the undamaged pathways connecting the viscera to the brain in both patient groups, or by incorporating the “as-if body” loop of the SMH. While the first is a reasonable and testable hypothesis, the second option is less so. If the SMH is to be maintained on the basis of the “as-if body” loop it questions whether the IGT can provide a good test of the SMH.

Although not involved in somatic activation within the SMH architecture, there are suggestions that damage to the insula and somatosensory cortex impairs IGT performance (Bechara, Tranel, Damasio and Damasio, 1997b). These areas are where representations of stimuli or situations are held. Consistent with the SMH framework, when patients with damage to either the right or left hemisphere were compared to age-matched controls only right hemisphere damage resulted in impaired performance. This right hemisphere effect has also been investigated in VMpfc (or the wider orbitofrontal cortex, OFC). Tranel, Bechara and Denburg (2002) found that right, but not left, lesions were associated with more impairment in everyday decision-making and impaired IGT selections. Bechara and Damasio (2005) have linked this laterality to the polarity of somatic state. Right VMpfc processes negative somatic states while the left VMpfc processes positive somatic states. The main distinctive variable on the IGT is the punishment on the decks so this suggestion is plausible. Clark, Manes, Nagui, Sahakian, and Robbins (2003) similarly found that the laterality of frontal lesions affected performance on the IGT relative to controls. However, damage is rarely limited to one prefrontal region (Clark and Manes, 2004) and impairment was positively correlated to lesion size, implicating prefrontal regions outside of those specified in the SMH.

In summary, the neural architecture of the SMH has been well specified and there is a large body of research that provides evidence supporting it. Many lesion studies by the Iowa group exploring these neural substrates are consistent with the predictions of the SMH. However, other studies suggest disruption of somatic systems does not have the disruptive effect predicted from the SMH. There are also ambiguous accounts of the degree to which somatic activity is necessary for advantageous performance to develop on the IGT. Together these studies question

whether the assumption that somatic markers are necessary for unimpaired decision-making is justified.

1.3.3 Other theoretical and philosophical explanations

One of the strengths of the SMH, and perhaps a contributing factor to its success, is that it has an intuitive appeal. References abound in popular culture and folklore inviting us to make decisions based on our “gut instincts” when we are unsure about the wisest course, or even if we are not. However, reintegrating body processes via emotion into mental processes is not a novel concept. James (1884) and Lange (1885) were the first to argue that emotion was the perception of bodily changes in reaction to environmental stimuli. Damasio (1994) acknowledged the contribution of these peripheral feedback theories of emotion as the basis for the SMH. In including a mechanism for cognitive input, secondary inducers, he maintains he has overcome the criticism of James-Lange theory that internal as well as external events can affect and effect emotional processing. But the addition of an “as-if body” loop is also not a new addition. James (1884) considered it and Dunn et al (2006) cite a number of examples of earlier expressions of this idea in considerations of emotion. However, Damasio has extended this discussion by specifying in greater detail how such a mechanism would work.

Despite the increased specification of the SMH relative to earlier theories, Rolls (2005) has criticised it for the same reasons James-Lange theory was rejected. The major argument against it is that empirical evidence has failed to show the causal link from peripheral responses during emotional behaviour to emotional behaviour. In fact, there is much evidence that such a link does not exist including work showing

that preventing peripheral feedback through surgical severing of the spinal chord (e.g. Cannon, 1927) does not abolish emotional responses. Neither are emotional responses evoked by artificially stimulating autonomic changes through injections of adrenaline or noradrenaline (Rolls, 2005). Indeed Rolls maintains that the peripheral nervous system does not produce changes diverse or specific enough to encompass the range of emotions that can be experienced. Perhaps the most damning criticism of James-Lange theory, and by extension the SMH, is attributed to Wittgenstein by McGinn (2003). His thought refutation imagines the horribleness of grief when some loved one dies. It cannot be explained in terms of the horribleness of bodily sensations. While they may be unpleasant they are not the object of grief. James-Lange theory fails because it does not take into account the intentionality of emotion, or what the emotions are about. These are generally things outside the body.

Rolls (1990, 1994, 2005) has developed a theory of emotion that does take this intentionality into account. He proposes that emotions are states produced by rewarding and punishing stimuli. Much of his published work has explored the architecture of this system and much of it overlaps with that proposed by Damasio to underlie the SMH (e.g. a central role for the OFC and the amygdala). However, the interpretation of how that architecture works differs between the theories. For Rolls, behaviour is produced in response to learned reinforcers that also elicit autonomic responses via the OFC and the amygdala. There is no need to place peripherally mediated changes (somatic markers) as a causal mechanism for changing the behaviour into this system. Rolls suggests that such an addition would be less efficient, especially in a rapidly changing environment. It would take longer to execute a behaviour if, in order to determine a stimulus' value, it first had to be passed through the visceral system or a cortical representation of it, rather than have a direct

link between areas representing reinforcer value (OFC, amygdala) with structures that effect behavioural responses.

Damasio et al (1991) anticipated this criticism through parsimony and argued that the somatic system is evolutionarily ancient and is very effective. But this is not a strong refutation as it does not explain why a more complex system would evolve instead of a simpler one. But Damasio and colleagues have made an articulate case for how the development of decision-making based on language (reason) can be integrated into a phylogenically and ontogenically older decision system based on bio-regulatory mechanisms involved in maintaining body homeostasis. Indeed, this is one reason why Damasio (1994) has rejected a Knowledge-Somatic hypothesis for guiding behaviour on the IGT. In the evolutionary and developmental framework of the SMH, somatic signalling precedes the emergence of language (or reason) as a behavioural guide. But Rolls' account remains a more parsimonious account of the interactions in neural architecture that overlap in his and Damasio's account. Rolls (1999, 2005) has also offered a deficit in reversal learning as an account of the impaired VMpfc performance on the IGT.

Patterson, Ungerleider and Bandettini (2002) have offered an alternative view on the role of somatic markers as indexed by SCRs in IGT function – that they have no causal function. They found that SCR activity appeared to be independent from performance on the IGT and also in a working memory task. Using functional magnetic resonance imaging (fMRI) to examine brain activity, and SCR recordings, the authors found that SCR changes were correlated with activity in a number of regions, including the VMpfc, during performance on both tasks. But these changes were not related to task performance and instead suggested that SCRs are generated during complex tasks but are not specifically related to any aspect of them. These

results do pose a problem for the SMH if SCR activity is simply a correlate of brain activity or arousal. There is a get-out clause though. In recording SCR activity during the IGT Patterson et al (2002) failed to distinguish time periods before and after card choice. By combining anticipatory and reactive SCR activity the analyses employed would have missed any correlated brain activity that differentiated between pre- and post-selection SCRs.

An interesting addition to the debate is contributed by Amiez, Procyk, Honore, Sequeira and Joseph (2003). They recorded SCRs in monkeys during selection of unequally and probabilistically reinforced targets. They found that SCRs were time-locked to the monkeys' arm-movements towards rewarded stimuli, irrespective of the choice selected and its outcome. The authors suggested that the SCRs were indicative of anticipatory appetitive behaviour, not of any cognitive process associated with the best target. The absence of any cognitive component does not trouble the SMH, the fact that SCR parameters were no different during the evaluatory period and after learning had occurred does. Dunn et al (2006, p. 251) point out that this is problematic for the SMH as it suggests that 'anticipatory' changes on the IGT relate to expectancies of reward and punishment after deck selection has occurred rather than directing their selection. Or, more simply, there is no causal role for aSCRs in decision-making. However, differences between the tasks diminish these concerns. In the monkeys' task they were always rewarded and never punished, making this task free from risk as compared with the IGT. In a similar vein, in this task the monkeys were entirely familiar with the environment having experienced it many times, whereas the IGT represents a novel environment for human control participants (however, VMpfc patients often experience multiple sessions). Additionally, in defence of the SMH the integration of somatic markers and decision-making may not

be as developed in monkeys as in humans, although according to Damasio any appeal to decision-making through purely rational, cognitive processes assumes the development of two systems built one on top the other (this is Descartes' error) neglecting to acknowledge that successful decision-making must have evolved.

1.3.4 Reversal learning

Extensive evidence suggests that a major role for the orbitofrontal cortex (OFC), of which the VMpfc forms a large part, is learning stimulus-reinforcer associations (Rolls, 2004). For example, the OFC contains secondary taste and secondary olfactory cortex in which the reward values of stimuli to these senses are represented. The tastes and smells of food are primary reinforcers. The OFC also contains neurons that code the texture of food in the mouth and also neurons with links to visual areas that link tastes to what the object is (Rolls, 2004). Animals choose alternatives that have higher value and a key discriminant for primary reinforcers is a better taste e.g. higher sugar content. Rolls and colleagues have also shown that the reward values of secondary reinforcers are represented in OFC. Activity in this area (measured with fMRI) is correlated with money won or lost on a reversal learning task (O'Doherty, Kringelbach, Rolls, Hornak et al, 2001). All this has led Rolls (2004, 2005) to suggest that part of the function of the OFC in decision-making, social behaviour and emotional processing is to represent reinforcers, detect changes in those reinforcers being received and then use these changes to rapidly reset the stimulus-reinforcer associations. In turn this would rapidly change behaviour.

Consistent with such a role damage to the ventral prefrontal cortex results in deficits in reversing stimulus-reinforcer association or reversal learning as it is better known (Rolls, Hornak, Wade and McGrath, 1994; Fellows and Farah, 2003; Hornak,

O'Doherty, Bramham, Rolls et al, 2004). Selection of the previously rewarded alternative occurs despite patients being able to report the correct response. Rolls et al (1994) also reported that impairment on this task was highly correlated with evaluations of disinhibited behaviour and post-injury emotional state.

Reversal learning is an explanation for deficient IGT performance because the arrangement of the reinforcement schedules on each deck is such that the disadvantageous decks serve no punishment until after several selections (3 in deck A and 9 in deck B). Initially, a choice from them is the better option because of their higher immediate payoff. When Fellows and Farah (2005) rearranged the losses in these decks so that they occurred on the initial selections, the behaviour of patients with VMpfc damage was comparable with that of controls whereas they were impaired on the standard version of the task. Furthermore these authors also reported impairment in patients with dorsolateral prefrontal cortex (DLpfc) damage. Their results are doubly damaging for the SMH in that they suggest that the source of IGT impairment for VMpfc-damaged patients is found in a failure to reverse early learning, and that damage out with the SMH architecture impairs IGT performance. Bechara, Damasio, Tranel and Damasio (2005) have sought to explain these results as due to the lesion locations in Fellows and Farah's (2005) patients. The damage in both VMpfc and DLpfc patients was focused in the right hemisphere and Bechara et al (2005) suggest this may implicate a working memory dysfunction as the source of their IGT impairment. However, it is not clear why. In the previous study they cite in support, Bechara, Damasio, Tranel and Anderson (1998) found a single dissociation between working memory and IGT performance such that the latter was dependent on the former. Patients with right DLpfc damage were impaired on the working memory task but not the IGT. So although right hemisphere damage impacts working

memory, this impairment was not sufficient to impair IGT performance in Bechara et al (1998). More pertinent may be that Bechara et al's (1998) VMpfc patients had bilateral damage. However, the involvement of systems out with the SMH architecture in IGT impairment does question the specificity IGT impairment has to identify prefrontal damage.

This explanation in terms of a failure to reverse early learning does not explain the behaviour of VMpfc patients on repeated exposures to the IGT, when one presumes the impulse to select from A or B has declined, and knowledge can be used. However, Dunn et al (2006) have suggested that reversal learning is best understood as a failure in response inhibition (after Rescorla, 1997). In this conception successful deck selection requires inhibition of a lose-shift pattern of responding for the advantageous decks after a loss. If VMpfc-lesioned patients cannot perform this inhibition the immediate consequences rather than the long-term outcomes will govern their selections. Of course, the reason that they may not be able to perform this inhibition may be that no somatic markers act as a guide (Bechara et al, 2005). However, Dunn et al have rejected this explanation as a parsimonious alternative exists where the inhibitory mechanisms required have been well specified (Rescorla, 1997). Otherwise they could be explained by a failure in the somatic marker system to mark the long-term outcome. Further, reversal learning has been demonstrated in the absence of an intact amygdala (Izquierdo, Suda and Murray, 2004) suggesting that emotion signals are not required for reversal.

Bechara et al (2005) have also sought to rebut an explanation by reversal learning by arguing that other components of the IGT need to be learned because the task is more complex than any reversal learning task used so far. Turnbull, Evans, Kemish, Park et al (2006) offer results consistent with such a defence. They modified

the IGT so that the contingencies of the decks shift in three phases across 100 trials – i.e. A is bad for 100 trials then good for two forty trial phases then bad again for a final 40 trial phase. Schizophrenic patients with negative symptoms who were unimpaired on the standard IGT selected at chance in each of the subsequent phase shifts. This suggests that factors other than reversal learning are involved. An obvious one is a failure to acquire conceptual knowledge of the task i.e. decks that offer the lowest immediate reward are best in the long-term. Thus the failure in reversal learning could be explained in terms of a lack of understanding of how the task fundamentally worked. It is not clear if Turnbull et al (2006) controlled for this in their task design (they only state that the decks' contingencies changed so that those that were good are bad for two out of three phases and the opposite is true for the bad decks). If they did not then it would seem that their catatonic schizophrenic participants did not acquire the knowledge possessed by their control participants and schizophrenic patients with positive symptoms. The issue of participants' knowledge is revisited in section 1.4.

1.3.5 Summary

The assumption that somatic markers are necessary for learning on the IGT has been questioned from a number of directions. While there is convincing evidence that damage to areas in the SMH architecture is correlated with the loss of somatic markers and IGT impairment, there is also strong evidence that the loss of somatic markers does not necessarily lead to this IGT impairment. Indeed, several studies have implicated regions out with the SMH framework in IGT impairment. Further, the relationship between regions in the SMH architecture have fit into other explanatory frameworks that are also well established (e.g. Rolls, 1994, 2004). Impairment on the

IGT can also be explained parsimoniously within this framework without resorting to somatic markers. In short, the assumption that somatic markers are necessary for decision-making appears to have little support. However, it is still possible that SCRs aid decision-making by biasing choice in the absence of knowledge.

1.4 THE ASSUMPTION THAT DECISION-MAKING PROCEEDS IN THE ABSCENCE OF KNOWLEDGE

The assumption that decision-making on the IGT proceeds in the absence of knowledge about the task contingencies stems from two claims made by the Iowa group. The first is that choice on the IGT is “cognitively impenetrable” (Dunn et al, 2006). This claim originated in Bechara et al’s (1994) report that healthy participants with above average IQ and memory scores were unable to report the IGT deck contingencies. The second claim is a modification of the first and originates in an empirical test of this hypothesis. Bechara et al (1997a) found 70% of their healthy controls could report “conceptual” knowledge of the IGT contingencies on average after 80 trials. All acquired “hunch” knowledge (that one of the advantageous decks was good, but not specifically why) earlier and Bechara et al (1997a) claimed that aSCR activity that differentiated between deck types was found prior to the appearance of this knowledge. The second claim is fundamental to the second assumption of the SMH that learning proceeds through an emotion-based system in the absence of knowledge sufficient to guide behaviour.

Recently, Maia and McClelland (2004) have challenged this second claim that selection is made in the absence of knowledge. They found that when they asked their participants more specific questions than used by Bechara et al (1997a) they had

consciously available knowledge sufficient to guide their choices earlier than reported by Bechara et al (1997a). Crucially, this knowledge was present prior to when Bechara et al reported the differential aSCR activity. This meant that participants' selections are made with some knowledge of the likely consequences and therefore does not require an explanation dependent on unconscious somatic markers. Bechara et al (2005) have challenged this claim by pointing out that Maia and McClelland's (2004) participants still sampled from disadvantageous decks despite the professed knowledge. They claimed this leaves room for somatic markers as an explanation for why learning continues until only advantageous selections are made. Maia and McClelland (2005) rejected this interpretation and offered the more parsimonious account that in any choice environment in which there is uncertainty participants must attempt to balance exploitation with exploration. Therefore, given that learning also develops it is not surprising that participants continue to select from disadvantageous alternatives until they reach a point at which they have gathered enough information.

Maia and McClelland's (2004) study is good evidence against both the first and second assumptions of the SMH in relation to the IGT. It also raises the possibility that while somatic markers may be anticipatory they may reflect knowledge of a more risky choice. Dunn et al (2006) have also suggested that a parsimonious account of the SCR data is that greater somatic activity leads to the behavioural impairment as aSCRs are larger prior to selection from disadvantageous decks. However, this second explanation does not fit the result of Tomb et al's (2002) study where they switched the deck contingencies such that the advantageous decks involved the larger magnitude wins and losses and found aSCRs were higher prior to selection from these decks. Damasio, Bechara and Damasio (2002) rejected the claim that aSCRs are based on reinforcer magnitude suggesting instead that the polarity of

the somatic marker has reversed and that the aSCRs being measured by Tomb et al (2002) indicate an incentive to approach. This data can still be interpreted in line with a knowledge-somatic hypothesis since in Tomb et al's modified task the advantageous decks now have the largest losses. aSCRs may reflect this knowledge. This hypothesis is tested in Chapter 5.

1.4.1 The involvement of working memory

Regardless of when participants acquire knowledge of the contingencies and its source, some knowledge is required to succeed. To profit on the IGT requires that participants are aware of the current value of the available alternatives, have some representation of the long-term outcomes of these alternatives and select on the basis of them. Full knowledge requires awareness that within each deck there is an internality that relates the former to the latter. If knowledge is required to profit on the IGT then working memory must be involved. This possibility is supported by Patterson et al's (2002) finding that SCR activity during both the IGT and a working memory task correlated with activation within the same cranial network.

The assumption that learning proceeds through unconscious somatic biasing can be tested by either loading working memory using a secondary task while participants complete the IGT or attempts can be made to dissociate the working memory from IGT performance. Bechara et al (1998) attempted to do the latter and claimed that working memory is dissociated from the decision-making required for the IGT. Participants who were impaired on the IGT were also impaired on a working memory task, but impairment on the working memory task did not also result in IGT impairment. So IGT performance was dependent on intact working memory. The

participants were patients with damage to the VMpfc, the DLpfc and healthy controls. The DLpfc has been associated with the regulation of working memory Levy and Goldman-Rakic (2000). The sub-optimal behaviour of VMpfc patients was again replicated in two groups, with more anterior or posterior lesions respectively, but only the posterior group was impaired on the working memory task (the delayed non-matching to samples task, DNMS). The DLpfc patients could also be split into two groups. Neither were impaired on the IGT but only those with right hemisphere damage were impaired on the working memory task. So although right hemisphere damage impacted on working memory, this damage was not sufficient to impair IGT performance in Bechara et al (1998) although such DLpfc damage has been reported as sufficient in others (Fellows and Farah, 2005).

Other investigations of a working memory role in IGT performance have used dual-task methodologies and shown mixed results. Hinson, Jameson and Whitney (2003) found a secondary task that loaded working memory (retaining the order of a string of digits in memory) impaired performance relative to a non-working memory task (repeating digits flashed on-screen) on a modified version of the IGT. This version had only three choices, where only one was the best, one was intermediate and one was worst. In a second study, Jameson, Hinson and Whitney (2004) restricted the impact of the working memory load to the central executive component rather than the phonological loop. Bechara and Martin (2004) also found that substance dependent participants who were impaired on the progressive IGT variant were also impaired on the delayed non-matching to samples (DNMS) task used in Bechara et al (1998). As the delay on the DNMS task was not a significant factor in performance the central executive components (switching and response inhibition) were implicated.

In contrast, Turnbull, Evans, Bunce, Carzolio and O'Connor (2005) found no difference in performance between healthy controls in three working memory manipulations. In one, random number generation was used to load the central executive; a second used non-executive articulatory suppression; while a third had no secondary task. These authors interpreted these results as evidence that IGT performance is not dependent on working memory but rather an example of emotion-based learning. However, there was a trend for superior performance in the no secondary task condition which might undermine this interpretation.

A central executive role is also suggested by another study from this group. Turnbull, Berry and Bowman (2003) developed the Firefighters Task as a descriptive analogue of the IGT. Participants must assess the quality of four trainee firefighters through examination of their daily logs. These logs are analogous to the IGT decks and contain examples of good or bad deeds (e.g. saving someone from a fire or accidentally dropping them from a ladder as they are being rescued). Like on the IGT the logs were sampled one at a time. Turnbull et al reported that healthy participants were worse on the Firefighters Task than on the IGT and claimed that as the task was more impersonal no somatic markers were generated to aid decision-making. However, no physiological recordings were made to back up this claim. Another interpretation may be that participants had descriptive information about each firefighter to remember and this memory load was greater than that required for the IGT. A test of this hypothesis would be to correlate performance on the Firefighters Task with performance on an established test of working memory.

Other manipulations of the IGT designed by this group also inform on a role for working memory. Bowman, Evans and Turnbull (2005) found no effect of imposing a 6 second time constraint for selections on IGT performance relative to no

time constraint, suggesting that increased time to think about the alternatives does not improve performance. However, six seconds is not much of a time constraint and Cella, Dymond, Cooper and Turnbull (2007) have recently shown that increasing the time constraint to 2 seconds does impair performance relative to a 4-second or unrestricted inter-trial interval. This suggests that success on the IGT is reliant on having time to think about the alternatives.

The suggestion that reasoning processes are involved in IGT performance is also suggested by the results of Reavis and Overman (2001)'s study using the California Weather Task (CWT). On the CWT participants must learn the probability with which combinations of four symbols predict the chance of rain. Male participants who completed the CWT prior to the IGT had a superior performance relative to other males who completed the tasks in the opposite order. This result suggests that participants learned something of the contingent nature of choices in the CWT that aided their IGT selections. Reavis and Overman also reported that this effect was absent in their female participants opening the possibility of a gender difference in performance. This difference has been documented during development (Overman, 2004) and infancy (Kerr and Zelazo, 2004; although Garon and Moore, 2004, found the opposite result). Reavis and Overman's (2001) gender difference has been replicated in adults (Bolla, Eldreth, London, Kiehl et al, 2003) but not in the majority of studies where such a comparison has been reported. To be on the safe side, Evans et al (2004) recruited females from two age-matched populations who had left school at sixteen or gone on to university. They found a "paradoxical" effect of education whereby the university students had significantly lower net scores than the early school-leavers in the final two IGT blocks. This was despite the student sample scoring higher on a test of intelligence (the NART; though Monterosso, Ehrman,

Napier, O'Brien and Childress [2001] found a significant correlation between IQ and IGT net score). Evans et al interpreted this behavioural difference in terms of emotion-based learning and claimed their non-student population were more reliant on this system while use of it had been discouraged in the student population. An alternative interpretation is that the students were not as motivated to do well despite the use of real rather than facsimile money reinforcers i.e. the discount curves for the students may have been shallower. They may also have been more apathetic about the task's outcome given that some of them were participating as a course requirement. This issue of reinforcer value will be explored in Chapter 2.

Gutbrod, Krouzel, Hofer, Muri et al (2006) bring evidence from two patient groups that bear directly on the issue of the importance of somatic activity and knowledge. Both patient groups were amnesic, but the location of damage differed between them. In the 'anterior' group lesions were located in the basal forebrain and orbitofrontal cortex. The 'posterior' group had more heterogeneous damage affecting temporal areas and regions adjacent to the hippocampus. Gutbrod et al recorded SCRs and probed knowledge to investigate the hypothesis that consciously available knowledge is required to make advantageous decisions on the manual version of the IGT. Conscious knowledge was assessed using Bechara et al's (1997a) general questions method every twenty trials and immediately post-test using two explicit memory tasks. The first was a recognition test where participants were shown four possible sequences of twenty gains and losses and asked to match them to each of the decks. The second task was to identify the long-term consequences of each of these sequences. In both tasks getting all answers correct gave a score of 4.

Gutbrod et al (2006) found that for both patient groups mean net score hovered around zero while it increased across twenty trial blocks in healthy controls.

All healthy controls were classified as having achieved “hunch” level knowledge and 75% (n = 6) reached “conceptual” knowledge. This is comparable with Bechara et al’s finding. Only 27% (n = 3) patients had hunch knowledge (2 from the anterior group) and only one had conceptual knowledge (the same patient from the posterior group who also displayed hunch knowledge). However, the authors do not go into any detail about when any participant arrived at either level of knowledge and this is disappointing given their simultaneous examination of SCRs. In both post-test tasks the healthy controls were nearly perfect. Patients were able to identify a schedule’s long-term consequences but not to match the schedule to the deck (the posterior group were better at this task although both were worse than controls). Physiologically, two results are of note. First, aSCRs to disadvantageous choices increased across block in controls but not in either patient group. Crucially, the difference in aSCR between deck types only reached significance in the fifth block. Second, larger punishment SCRs were found to deck B in all groups. (A third result of interest is that participants with damage to the amygdala did not show lower punishment SCRs than other patients, contrary to results presented by Bechara et al, 1999). The key point here is that behaviourally healthy participants showed a switch to advantageous behaviour in the second block, yet the difference in aSCR did not appear until the final block (although there is a statistical issue here as relatively few disadvantageous selections were made). But as the authors do not discuss the emergence of knowledge in their participants, and as they use the same methodology as Bechara et al (1997a) this result can only suggest that task knowledge preceded SCR change. An additional point is that, while the absence of an aSCR difference in amnesic patients and their non-preferential choice behaviour might support the SMH, these patients did exhibit differential punishment SCRs like the control group. The difference is that the

amnesic patients did not have consciously available knowledge about the task contingencies. Gutbrod et al (2006) claim that their results demonstrate that consciously available knowledge is required for an aSCR difference between IGT deck types to emerge. However, as the authors acknowledge, there is a problem with this account given the version of the IGT they used. In order to record post-selection SCRs that had returned to baseline following aSCR activity, Gutbrod et al (2006) delayed feedback on deck selection for 10 seconds. The authors suggest that they may have inadvertently made the task extremely difficult for their amnesic participants by making the contingency between deck choice and reward or punishment difficult to notice. This methodological modification would explain why Turnbull and Evans (2006) found that their amnesic patient learned to select advantageously on a standard version of the IGT without the large delays in feedback. But Gutbrod et al's (2006) results do suggest that being unable to explicitly learn the contingencies of choice and outcome on the IGT severely impact on learning on this task.

1.4.2 Other contributory factors to IGT performance

A knowledge-somatic hypothesis is plausible in light of Maia and McClelland's (2004) data showing that knowledge sufficient to guide behaviour occurs very early in IGT selection. However, the counter claim is that knowledge alone is insufficient to guide behaviour on the IGT and this leaves room for a weaker version of the SMH where somatic markers act as guides. VMpfc-damaged patients acquire knowledge of the contingencies yet still choose the disadvantageous options (Bechara et al, 1997a). While an explanation through a failure in reversal learning can explain IGT behaviour in one instance, it is more difficult to explain on the multiple occasions on which

these patients have been tested (some of the same VMpfc patients participated in each of Bechara et al, 1994, 1996, 1997a, 1998). Although a failure of response inhibition is a realistic possibility it is difficult to comprehend why the VMpfc patients would continue to select the worst decks each time they are exposed to the task. But such behaviour is not restricted to VMpfc patients on the IGT. Despite knowledge that some behaviours have negative consequences many people continue them. Drug addictions persist in the face of well-known health concerns and potential legal issues. Even addicts of legal drugs like tobacco persist in their habit despite knowledge of the health risks and in many cases compound their irrational behaviour by paying pension contributions. However, these behaviours can be accounted for by a failure to inhibit responses to stimuli with high short-term values. This will be returned to in section 1.5.

The behaviour of healthy controls who perform in the impaired range on the IGT may be explained by their failure to acquire knowledge of the deck contingencies. However, there are some other factors that may impact on IGT behaviour. One example is risk seeking. Rather than select the alternative that is safer participants may be lured to the disadvantageous decks by the prospect of the larger immediate gain, despite the risk of a large loss. Similarly damage to the VMpfc may lead to greater risk seeking behaviour. However, there does not appear to be strong evidence to support this hypothesis. Sanfey, Hastie, Colvin and Grafman (2003) developed a task to measure risk-taking in patients with prefrontal lesions that either excluded or were restricted to the ventromedial region. In the task participants are offered a choice between two decks of cards and must select one. There are five decks in total and the expected value is the same in each. What differs is the variance of loss magnitudes and the proportion of cards with a positive outcome (out of 25

possibilities). Healthy control participants showed a preference for safe decks, minimising selections from decks with a larger variance and therefore more risk of a large loss. The VMpfc group could be split in two, with one group no different to the controls. A second group was indeed risk-seeking but they could not be distinguished from the first group on the basis of lesion location.

Another task developed to test risk-taking behaviour is the Cambridge Gamble Task (CGT; Rogers, Everitt, Baldacchino, Blackshaw et al, 1999). On each trial participants are presented with 10 boxes in two colours, behind one of which a token has been randomly placed. The proportion of each colour varies between trials. Participants are invited to place a bet from their accumulated points on the location of the token (yellow or blue box). The proportion of bets on less likely alternatives measures risk taking. Control participants tend to select more likely outcomes and adjust their betting relative to the probability of being correct i.e. they bet more when the ratio of yellow: blue is 9:1 than 6:4. Patients with VMpfc damage did not show a preference for risky or safe options nor was their betting consistent with their choice. Indeed, compared to controls bet amounts were reduced. However, Monterosso et al (2001) found that performance on the CBT and choices in the latter half of IGT were not correlated. Clark and Manes (2004) have suggested that risky behaviour is reduced following impairment to decision-making as a compensatory strategy.

1.4.3 Summary

The results of Maia and McClelland's (2004) study and investigations into the contribution of working memory leave little need for an unconscious somatic system biasing learning on the IGT. A weaker version of the SMH would have somatic

markers influence decision-making overtly as well as covertly. Indeed Damasio (1994, p. 184, 214) has advanced this account. However, if participants possess knowledge sufficient to guide behaviour there is little need for a somatic marker system to explain learning. This claim cannot be fully dismissed until it is determined when knowledge sufficient to guide behaviour and differential somatic activity emerges. But regardless of whether the weaker SMH version is correct, it is the stronger claim made by this group (Bechara et al, 1997a; Tranel et al, 1999) that researchers have seized and assumed in experiments using the IGT e.g. as a task that requires emotion-based learning (e.g. Turnbull et al, 2005).

1.5 THE ASSUMPTION OF A MAXIMIZING SYSTEM

The standard analysis on the IGT compares selections solely in terms of the expected value of the decks regardless of the other differences within these decks. Bechara et al (1994) never discuss why they set up the IGT in the manner they did. One can speculate that variations in the schedule of losses between decks with the same expected value was intended to make the task less clear to comprehend. But their interpretation of IGT performance and somatic marker activity ignores these potential differences and only considers long-term outcomes. Thus Bechara et al (1996) interpreted the behaviour of their VMpfc patients on the IGT in terms of a “myopia for the future”. Bechara, Tranel and Damasio (2000) investigated this and two other hypotheses as explanations for the sub-optimal performance of VMpfc patients. A hypersensitivity to reward and a hyposensitivity to punishment were not supported by behaviour on a reversed version of the IGT. Here the immediate consequence of selection was an immediate loss. This was higher in the advantageous decks but subsequent gains were sufficient to make a profit, whereas the disadvantageous decks

had a lower immediate loss but lower rewards insufficient to avoid an overall loss. VMpfc patients still persisted in selecting the disadvantageous decks in contrast to controls. This behaviour did not support either hypothesis based on the sensitivity to reward or punishment. To test whether VMpfc patient performance could be normalised, the 'progressive' IGT variant was created where the losses were increased or decreased across ten card blocks within each of the disadvantageous or advantageous decks respectively. VMpfc patients were still impaired relative to controls and showed no sign of learning so expressing their behavioural deficit in terms of the long-term outcomes of their deck preferences is reasonable.

Several neuropsychological pathologies are associated with behaviours that are similar to those observed in patients with VMpfc damage and many researchers have suggested that this might imply a similar underlying deficit (e.g. VMpfc damage in drug addicts). The irrational and personally destructive behaviour of drug addicts has been characterised as myopic in that lure of the high payoff from another hit outweighs the long-term consequences of such decisions (Vuchinich and Tucker, 1988; Herrnstein and Prelec, 1992; Bickel and Marsch, 2001). Supporting this conception abnormal IGT performance has been reported in various substance dependent populations (e.g. Petry, Bickel and Arnett, 1998; Grant, Contoreggi and London, 2000; Bechara, Dolan, Denburg, Hines et al, 2001; Rotherham-Fuller, Shoptaw, Berman and London, 2004; but see Dunn et al, 2006 for an extensive review). However, in a detailed examination of the behavioural performance and physiological activity of a large sample of recovering drug addicts (n = 46), Bechara and Damasio (2002) and Bechara, Dolan and Hines (2002) found that only a subgroup (23%) resembled VMpfc patients and could be classified as myopic for future outcomes. They were impaired relative to controls on both the progressive

versions of the standard and variant IGT, and did not develop differential aSCRs between good and bad decks. Another subgroup (41%) were abnormal on the standard progressive task, but not the progressive variant, and their SCRs to reward were larger than the other groups suggesting hypersensitivity to reward as an explanation. A further subgroup displayed behaviour and SCRs consistent with the majority of healthy controls. These results suggest drug addiction cannot be classified as simply myopia for the future. Similarly, although the authors state their results suggest a VMpfc involvement in drug addiction, the similarity of behaviour on the IGT does not mean that the underlying deficit (if there is one) is anatomically the same.

Regardless of the accounting ability and involvement of an underlying somatic system, the typical analysis used on the IGT implies that expected value is the deciding factor in IGT choice. This is also the assumption of a description of VMpfc patient behaviour in terms of myopia for the future. This assumption is not unreasonable given the long history of research to find a descriptive model of human decision-making. The predominant position in this literature assumes that normatively human choice does maximize, but numerous challenges demonstrate conditions in which such an assumption is not descriptive. The following sections introduce some of this literature.

1.5.1 Rational Choice Theory

Rational Choice Theory (RCT) is a collective term for the various theories of human choice that developed following earlier mathematical explorations of probability in, for example, insurance to assess risk and in economics to investigate consumer

behaviour. To be rational a decision maker must be able to weakly order outcomes dependent on preference and, most importantly, make choices so as to maximise something (Edwards, 1954). In Rational Choice Theory this something is usually called utility . Utility is best defined as the “goodness” of an option as it is used as a measure of how consequences realise values or goals. Rational Choice Theory assumes a decision maker will choose the alternative that maximises utility (Edwards, 1954). This assumption fits into the common sense view of decision-making and has become the basis of most economic thought. Rational Choice Theory is thus a normative theory of decision-making.

But people do not behave as if maximising expected value e.g. they buy insurance and play lotteries (despite the operators of such schemes making a profit because the price of the product is greater than their expected value). The first demonstration that human behaviour may not be descriptively rational is known as the St Petersburg Paradox. Here a decision maker is invited to buy a ticket to play a game where the event of interest is when an unbiased coin lands on heads. If this event occurs on the first toss they are paid £2.00, £4.00 if it is the second toss, £8.00 if it is the third toss and so on. Despite the prospective payoff being infinite people will not pay more than a few pounds to play. Bernoulli (1738/1954) suggested that this is because the value of money is a decreasing function of amount won (or indeed possessed, so that £1000 means more to a poor man than to a rich man). He suggested that people act so as to maximise expected utility rather than expected value. As the value is relative to the amount possessed (or won) the expected utility of the St Petersburg Paradox is not infinite at all and the paradox is resolved.

In the mid-twentieth century von Neumann and Morgenstern’s (1944) Expected Utility Theory (EUT) provided a mechanism for measuring utility under

risk and it became the primary model of RCT. (In economics risk and uncertainty refer to similar situations. Under risk, the probabilities of the options are known, whereas under uncertainty they are not [Edwards, 1954]). The assumption in EUT is that a decision-maker will maximize expected utility. Challenges to EUT as a descriptive as opposed to a normative theory of human choice have primarily involved testing its axioms. The most important of these are the ordering of alternatives dependent on preference that has already been outlined above (sometimes referred to as weak ordering) and transitivity (if options can be ordered that order should be transitive, e.g. If A is preferred to B and B preferred to C, then A is preferred to C); cancellation (if the available options share identical characteristics then these should cancel each other out and be ignored, i.e. choice should only depend on the characteristics which differ. This is also known as the sure-thing principle [Savage, 1954]); dominance (an option is dominated if it is inferior to another option in at least one aspect. Dominated options should never be adopted); continuity (when faced with a gamble between the best and the worst outcome versus some intermediate outcome, the decision maker should prefer the gamble as long as the odds of the best outcome are good enough); and invariance (the decision maker should not be affected by the way in which options are presented). Violation of these principles does not result in expected utility being maximised (von Neumann and Morgenstern, 1944/1947). Many violations of the axioms of EUT have been described (e.g. Allais, 1953, demonstrated a violation of the sure-thing principle; Lichtenstein and Slovic, 1971, 1973, demonstrated preference reversals that violate the transitivity axiom) and have resulted in modifications or new models to account for them (e.g. subjective expected utility theory, prospect theory). But these models

still hold that utility is maximised and thus can be counted as examples of Rational Choice Theory.

The assumption that behaviour will maximize utility is not necessarily implied from an examination of the decision-making literature. Yet this assumption is made in relation to the influence of any somatic system and in behaviour on the IGT. It is effectively a four alternative variable ratio environment made complex by the inclusion of punishing events to predominantly rewarding schedules of reinforcement. Such choice environments have been studied extensively in the animal learning literature and are worthy of review to look at the conditions in which maximization is or is not achieved.

1.5.2 Animal choice experiments

Animal choice behaviour is an extension of the research into animal learning. Choice was initially explored in environments with two operant alternatives that differed in amount (e.g. sugar concentration in available liquids). The value of one alternative relative to another can then be inferred from the animal's behaviour e.g. using the proportion of selections, or more generally the proportion of time spent selecting it, relative to the other option (Rachlin, 1989). More complicated choices have been explored using more complex operant schedules. Variable ratio schedules require a variable number of responses around a mean before reinforcement whereas variable interval schedules require a response within an amount of time that changes around some mean time. In symmetrical choice tasks the alternatives are for the same reinforcer, although not necessarily the same amount of reinforcer. In an asymmetrical choice environment the alternatives are different (e.g. water or food

pellets). Similarly, asymmetrical choice can also refer to one-choice environments where the options are to respond or not.

The repeated choice between alternatives on such reinforcement schedules gives some insight into the processes that influence choice behaviour. Indeed Rachlin (1989, 2004) has proposed that behaviour can be viewed as the inter-temporal choice of available alternatives. In most cases animals will choose the alternative that offers them the highest amount or rate of reinforcement. But situations in which animals, including humans, do not maximize have been found. One example of such an environment is a variable ratio/variable interval schedule where the choice of the option with the greater local reinforcement rate (e.g. a smaller delay till the next choice or a larger value reward) reduces the overall rate of reinforcement by increasing the delay or reducing the value on each alternative. Choice environments such as this are examples of what is known as the Harvard Game (after its place of inception). An illustration of how this environment might appear out with the laboratory is provided by the menu problem:

You are on a packaged holiday where all your meals are prepaid. The only available restaurant serves only two meals with approximately equal nutritional value – lobster and fish cakes. The only difference between the options is hedonic value where the fish cakes are relatively bland and its hedonic value changes little despite regular or occasional consumption whereas the lobster has a higher hedonic value. Due to its richness, eating the lobster everyday would rapidly reduce its hedonic value and may even reduce your desire for seafood. Such a situation would lead to a lower overall payoff across all your meals than if you ate mainly the fish cakes with the lobster on occasion. This would lead to a maximization of payoff across your meals and is the behaviour predicted by Rational Choice Theory. Melioration, on the other

hand, predicts a choice of lobster until such time as its hedonic value falls below the fish cake, whereupon choice will switch and eventually settle into matching between the, now relatively low, hedonic values of each option.

Vaughan and Herrnstein (1987) describe an environment with a similar structure to the menu problem. Figure 1.1 displays this choice environment. It represents the value of two alternatives A and B as a proportion of the number of choices from B . Here alternative B is equivalent to the fish cakes in the menu problem. Similarly, on most occasions a higher payoff will be obtained by choosing A (lobster), but with each choice that payoff is reduced. However, by choosing B the value of option A is increased. The situation is simple because one alternative provides a higher payoff (A) than the other (B), but it is complex because there is an internality such that the proportion of choices affects the overall payoff.

The structure of the IGT therefore mirrors much of the choice environment on a Harvard Game. In both the participant is faced with a choice between a larger or smaller immediate gain, but with repeated selection of the former the overall payoff is not maximized. The participant must switch to the smaller immediate payoff in order to maximize in both cases. There is no internality between these choices on the IGT in contrast to the Harvard Game. The internalities of repeated selection come within the reinforcement schedules of each deck. This may make the task easier and be a reason for the rapid development of maximizing behaviour in a normal population.




Figure 1.1: The Harvard Game environment. The delay from each alternative is plotted as a function of the proportion of choices of alternative B . The dotted line represents the average delay per trial as a function of the proportion of choices from each alternative over the previous ten trials. The maximizing equilibrium (the distribution of choices that maximizes overall payoff) is achieved when all responses are allocated to alternative B , while the melioration equilibrium is gained when all responses are for A (this diagram was originally published in Herrnstein, 1990).

The behaviour of non-human animals in Harvard Game environments has generally resulted in a preference for the option with the higher local reinforcement rate. Such behaviour conforms to a behavioural mechanism called melioration (Herrnstein and Vaughan, 1980). This is effectively a restatement of the Law of

Effect (Thorndike, 1898) whereby an animal will devote time or effort to an option proportionate to its value at the time of the choice. If these values are unequal then behaviour on successive choices is redistributed towards the more favourable alternative. An equilibrium will result if there is a distribution of responses that equalises the reinforcement available from each alternative, or where the average reinforcement rate from the available options matches. Melioration is a molecular mechanism of a more general theory of behaviour called the Matching Law (Herrnstein, 1961, 1997).

The Harvard Game thus provides an environment where maximization and melioration make different predictions. When the options were three-minute variable interval (VI) schedules and the internality was an additional 1 minute VI schedule that incremented with selection of *B*, but paid out with selection of *A*, Vaughan, Kardish and Wilson (1982) found that pigeons spent more time on *A*. This behaviour is predicted by melioration as selection of this alternative provides the higher local rate of reinforcement whereas maximization predicts greater choice of *B* as it would increase overall reinforcement rate. Similar results have been reported in humans by Herrnstein and colleagues (Herrnstein and Prelec, 1991; Herrnstein, Loewenstein, Prelec and Vaughan, 1993) in a version of the Harvard Game where the payoffs were equal in amount but choice of *A* led to an inter-trial interval two seconds shorter than choice of *B*. However, the length of the delay increased with the proportion of choices from *A* in the previous ten trials.

Melioration as a descriptive theory of human behaviour is supported under these conditions on the Harvard Game. Yet humans do maximize in their everyday life. People are capable of saving, of dieting, of giving up addictive substances (or avoid taking them up for fear of becoming addicted). So an appropriate question

seems to be under what conditions maximization or melioration dominate. Herrnstein et al (1993) reported a second condition of their delayed reinforcer version of the Harvard Game reported above. When the internality affected reinforcer amount (i.e. *B* is always worth more than *A* but repeated selection of *B* reduces the payoff from both), participants tended towards maximization. Several other manipulations of this methodology revealed that the tendency to meliorate was attenuated with feedback on previous choices; when the reward functions were parallel (as in Figure 1.1) rather than crossed (requiring non-exclusive selection of the maximizing alternative); when the averaging window for the internality was shorter rather than longer; and when payoff values rather than delays were used.

Herrnstein et al (1993) attributed the source of meliorating behaviour to limitations on information processing. In the experiments above when more information was available on the contingencies between alternatives (the internality) people tended towards maximization. As the authors point out this is not surprising. Melioration only requires knowledge of the local reinforcement rates of available options whereas maximization also requires this knowledge, as well as knowledge of the internality when present and an understanding about how to use this information. Silberberg, Thomas and Berendzen (1991) and Tunney and Shanks (2002) have shown that given information and practice participants can learn to maximize on similar schedules. However, for maximization to develop requires a large number of trials (400 trials and 150 trials of unreinforced practice in Herrnstein et al (1993); 500–700 trials in Tunney and Shanks, 2002).

1.5.3 Summary

The assumption that behaviour on the IGT conforms to Rational Choice Theory is implicit in the analysis of performance and in descriptions of behaviour in terms of the SMH in both healthy controls and patients with brain damage. This section has shown that while Rational Choice Theory is the normative theory of human choice violations have been demonstrated. These have included simple one-shot hypothetical decisions and more complex choice environments. The Harvard Game is one such example and presents a curious contrast due to the similarity of its choice environment with the IGT if viewed solely in terms of expected value. While maximization is rapidly acquired on the IGT only some conditions of the Harvard Game result in a similar outcome (Herrnstein et al, 1993; Tunney and Shanks, 2002) and this typically involves a long series of trials. This suggests that there are features of the IGT environment that permit rapid learning and choice on the IGT may not simply be determined by expected value. If that is the case then any description of behaviour in these terms may not be completely accurate and may miss important information about behavioural influences on the IGT.

1.6 AN INTRODUCTION TO THE EXPERIMENTS

The SMH makes a number of assumptions about human decision-making that have been reviewed in this chapter. The assumption that somatic markers have a causal impact on decision-making looks hard to support given the evidence that when they are absent, decision-making as measured using the IGT, can operate normally. The weaker claim that somatic markers augment a decision-making process may be true but it is unnecessary given that more parsimonious accounts of behaviour in reinforcing situations exist. This challenge to the necessity of somatic markers is also

undermined by findings that normal performance on the IGT, the main empirical test of the SMH, can be adequately explained by the presence of knowledge sufficient to guide behaviour. However, the possibility that somatic markers precede the emergence of this knowledge has not been ruled out. Nor has the hypothesis that somatic markers result from the anticipatory expression of such knowledge. The experiment that tests these competing claims is reported in Chapter 5.

Together the evidence against these first two assumptions of the SMH suggests that performance on the IGT is not sufficient to support the SMH. It is reassuring to note that this conclusion has been independently reached by other researchers (Dunn et al, 2006). But behaviour supportive of Rational Choice Theory is found using the IGT. This is interesting in itself because it conflicts with behaviour in an equivalent choice environment from the behavioural choice literature - the Harvard Game. Although maximization is found in some conditions of the Harvard Game, unlike in the IGT it is not rapidly learned. If an underlying rational somatic marker system cannot account for this behaviour then other factors must be involved. This thesis details experiments that investigate what these might be. This is important because the IGT has become a regularly used test of decision-making in clinical populations, as well as a tool to measure the development of decision-making in a normal population. It has even been cited as a test of VMpfc functioning despite the large body of evidence implicating other factors important for IGT performance. If normal behaviour on the IGT is not solely influenced by the expected value of choices then claims that the behaviour of patients with VMpfc damage or drug addictions is myopic for the future may not be complete. While important for understanding what influences normal performance on the IGT and therefore what may go wrong in

patient populations, these experiments also help to illuminate the conditions under which maximization develops.

Experiments 1, 2 and 3 of this thesis demonstrate that the information participants receive about the task has a significant effect on learning behaviour in healthy participants, suggesting that knowledge about the choices is an important contributory factor to IGT performance. Experiments 4 to 8 explore manipulations of the short-term contingencies of the IGT decks and reveal that choice is not made solely in terms of expected value. Finally Experiment 9 tests the competing claims that knowledge or somatic activity is the important factor in influencing learning on the IGT.

CHAPTER TWO

THE EFFECT OF TASK INSTRUCTIONS AND REINFORCER TYPE ON LEARNING

2.1 INTRODUCTION & OVERVIEW

This chapter details the initial experiments that examine potential factors involved in producing optimal behaviour on the Iowa Gambling Task. Experiment 1 is an attempt to replicate the behaviour of healthy control participants reported by Bechara et al (1994), using their methodology but on a computerized version of the IGT. However, the behaviour of the participants in this experiment did not correspond to that of healthy control participants in the original study – participants in Experiment 1 did not choose optimally by the end of the experiment. Subsequently, Chapter 1 goes on to describe experiments which investigate whether aspects of the administration of the IGT (the type of instructions or reinforcers participants receive) affect behaviour. The behaviour of participants under each manipulation are described separately and in general before a more detailed comparison between experiments is described and discussed at the end of the Chapter.

In Experiment 1 participants followed the methodology of Bechara et al (1994) with some minor changes and the addition of a second session (identical to the first) to investigate whether asymptotic performance could be reached. The minor changes concerned the administration of the task in a computerized rather than a manual form. Computerized administration has now become standard with no differences in performance reported due to administration methods (cf. Bechara et al 2000, Bowman et al, 2005). Participants' behaviour under these conditions was, on

average, no different to chance i.e. they did not show a preference for the advantageous decks and so were not behaving as comparable participants did in the original studies (Bechara et al, 1994, 1997a, 1999).

Experiment 2 examined one possible explanation for this failure to replicate. It followed exactly the methodology of Experiment 1 except that participants received the more detailed instructions used by Bechara et al (2000). These instructions include a hint about how to succeed on the task. In this experiment participants developed a preference for the advantageous decks at the end of the first session which continued into the second session. Experiment 3 was designed to replicate the findings from Experiments 1 and 2 but using real money rather than facsimile money as a reinforcer. Despite facsimile reinforcers being used in all the original IGT studies (Bechara et al, 1994, 1996, 1997a) and some others, use of these reinforcers is not standard. It has been argued that real reinforcers provide a more realistic incentive (Edwards, 1954; Camerer and Hogarth, 1999; Hertwig and Ortmann, 2001; Brase, Fiddick and Harries, 2006) and this might differentially affect IGT performance. While comparable behaviour was found in both conditions in the first session, convincing evidence of learning was only found when participants received the instructions with the hint. When comparisons were made between experiments, it was participants who had received the hint instructions who learned to select increasingly from the advantageous decks. The type of reinforcer received did not affect learning. The change in participants' deck selections across blocks of trials as measured by net score (advantageous selections minus disadvantageous selections) gave an estimate of learning rate. This was greatest in the first session when participants received instructions with the hint, which resulted in advantageous performance in the second

IGT session. Without the hint, performance in the first session was more affected by Reinforcer Type and was no better than chance at the end of session two.

2.2 EXPERIMENT 1

A FAILURE TO REPLICATE NORMAL PERFORMANCE ON THE IOWA GAMBLING TASK - THE BEHAVIOUR OF A NORMAL SAMPLE IS NOT OPTIMAL.

2.2.1 Introduction

Experiment 1 was an attempt to replicate the results Bechara et al (1994) using their gambling task with healthy participants. In this experiment a computerized version of the task was used, mainly to aid in experimental administration and for ease of comparison with later, perhaps more complex, manipulations of the IGT. Bowman et al (2005) have directly compared the administration methods and found no differential performance ensues. Indeed, the computerized IGT has now become the standard version (Bechara et al, 2000). In the design of this experiment several refinements to the method were made. These are detailed in the method section and were made to improve the task. Briefly, these are randomization between participants of deck position so the order of advantageous and disadvantageous decks is not fixed; an increased number of cards in each deck so it is possible for all selections to be from the same deck; a decreased inter-trial interval (ITI); and written not spoken instructions were administered.

An additional session of the IGT was also added to the standard administration in an effort to examine the acquisition of maximising behaviour in the IGT. Over 100 trials no reported results have detailed participants reaching asymptotic maximising behaviour (a net score of 20 across sequential blocks in the task i.e. all cards are

chosen from the advantageous decks). This second session allowed behaviour in a second 100 trials to be analysed and, if necessary, compared to the first 100 trials. As participants are known to learn faster with distributed than with massed trials (Dempster, 1996) the second session was run 48 hours after the first session.

It was predicted that the results of Bechara et al (1994) would be replicated i.e. that participants would prefer to choose from the advantageous decks by the end of the first session (100 trials). A tentative second hypothesis was that asymptotic performance would be reached in a second session of the task.

2.2.2 Method

Participants

Participants were first year undergraduates studying psychology at Keele University ($n = 20$, mean age = 19.7, $SD = 2.11$). Thirteen participants were female and seven were male. All participants received course credit for taking part in this experiment.

Apparatus

Several changes were necessary in transferring the IGT from a task performed using actual decks of cards to one based on a computer. The inter-trial interval (ITI) was reduced. In the original task the ITI was 6 seconds. The original time interval may have been decided upon because there was a comparison between groups and one group included neurological patients who may have required more time to carry out the task. However, given the high functioning of the neurological patients a more important factor may have been that the task was performed manually and an ITI of 6 seconds was necessary to record the relevant data. Additionally, physiological

responses were later collected using the same procedure and a relatively large time window was needed between trials. But as none of these conditions applied in Experiment 1 the ITI was reduced to 1.5 seconds to prevent fatigue effects.

Though it is not explicitly stated in Bechara et al (1994) it is clear from later papers (Bechara et al, 2000) that participants were given oral instructions before starting the task. This may have been to ensure participants, especially the neurological patients, understood what the task involved. However, in this experiment the instructions were displayed on screen prior to the start of the experiment and participants were instructed to read them. An opportunity was given for them to ask any questions. This change in procedure was to ensure that participants were given exactly the same instructions across sessions and that experimenter effects did not influence instruction administration.

The number of cards available in each deck was also increased from 40 to 100. Bechara et al (1994) explained that since participants rarely select from one deck more than forty times this is a reasonable number of cards for each deck. However, pilot studies for this experiment indicated that some participants selected from the same deck more than forty times. Since the task explores decision-making, denying participants the opportunity to select from one deck before the end of the task merely because there are not enough cards puts a restriction on their choice, especially given that there are differences in reinforcement schedules between decks (in the frequency of losses). Reinforcement schedules for these additional cards were copied from the original 40 cards to maintain the integrity of reinforcement schedules while still preserving the unpredictability of the deck schedules.

Participants were tested individually in a sound-proofed testing laboratory. A PC controlled the experiment. A program that replicated the Iowa Gambling Task was

run on the PC. Figure 2.1 displays a representative screen shot from the experiment. Four ‘decks’ of cards were displayed horizontally on screen labelled A, B, C and D. Above the decks a written message (green font colour) informed the participant how much they had won after every card selection. A second written message (red font colour) was displayed below the first message and informed the participant when a loss was made. These messages changed depending on participants’ card choices and in line with the reinforcement schedules for each deck (Appendix A). Two bars were also displayed at the top of the screen. A green bar labelled “Cash” displayed how much the participant had won on the task so far. A red bar labelled “Borrowed” displayed the amount of money ‘loaned’ to the participant to play the game. This remained at £2000 throughout the task. Participants used the mouse to select their choice of deck.




Figure 2.1: Screenshot from the computerized Gambling Task.

Design

A repeated-measures design was used with each participant taking part in two one-hundred trial sessions. Following Bechara et al (1998) the sessions were divided into twenty trial blocks to investigate participants' learning rates. The dependent variable was net score calculated by subtracting the number of disadvantageous choices (cards selected from decks A and B) from the number of advantageous choices (cards selected from decks C and D).

Procedure

Before agreeing to participate participants were told that the task would require participation in two sessions of 15 minutes separated by at least one day. Course credit was awarded on completion of the second session. The participants were told that they were taking part in a cognitive task in which the goal was to win as much fake money as possible. These were the participants' only verbal instructions.

Participants were instructed to follow the following on-screen instructions:

“You are going to see four decks of cards on the computer screen. You must make a series of card selections, one card at a time, from any of the four decks of cards until you are told to stop. After selecting each card you will receive some money. After selecting **some** cards you will be given money but will also lose money. The amount you have won will be displayed on screen as a value and in a green bar at the top of the screen.

To start you off you have been given a loan of £2000. A red bar will display this value to remind you how much money you were loaned to play the game.

The goal of the game is to maximise profit on the loan of play money (to win as much play money as possible).

You are free to switch from any deck to another, at any time, and as often as you wish until you are instructed to stop.

It is important to know that just like in a real card game, the computer does not change the order of the cards after the game starts. You may not be able to figure out exactly when you will lose money, but the game is fair and the computer does not make you lose money at random.

Any questions?

No? Then click on “Start”.

The experimenter waited in the laboratory room until the participant had asked for clarification or until they started the experiment. If clarification was required the experimenter referred to the written instructions and quoted from them.

When the experiment started participants saw four decks of cards displayed on-screen (Figure 2.1), labelled A, B, C, and D. Participants mouse-clicked a card from any of the four decks. The face of the card then appeared on top of the deck (the colour was either red or black), and a message was displayed on the screen indicating the amount of money won or lost. At the top of the computer screen a green bar changed after each selection depending on the amount of money won or lost. A red bar indicated the ‘loan’ given by the experimenter to the participant to encourage the participant to continue playing when in the red. A gain was indicated by the sound of a man shouting “Yippee!”, the appearance of a smiley face and a proportionate increase in the length of the green bar. A loss was indicated by the sound of a man shouting “Doh!”, the appearance of a frowning face and a proportionate decrease in the length of the green bar. Once the bars were updated the face of the card disappeared, and the participant could select another card.

The backs of the cards in each deck had a uniform appearance just like a real deck of cards. Each deck had 100 cards: 50 of the cards had a black face and 50 had a red face. The colour of the cards had no meaning in the task but was maintained to make the cards more realistic (Bechara et al, 2000).

The reinforcement schedule for each deck determined the amount of money won or lost for selections from that deck. The schedule of losses for the four fixed reinforcement schedules can be found in Appendix A (Bechara et al, 1994). The on-screen decks were always labelled A, B, C and D in that order. However allocation of reinforcement schedules to on-screen deck was randomised across participants. Apart from being good experimental practice this randomisation guarded against the possibility that participants might inform other members of the sample population about how to do well on the task. Selecting from deck A or B always yielded a win £100; selecting from deck C or D always yielded a win of £50. However on some selections a loss was also made. For decks A and B the loss over ten card selections was -£1250. In deck A this loss resulted from five smaller losses every ten trials which totalled -£1250. In deck B there was one loss of -£1250 every ten trials. For decks C and D the loss over ten card selections was -£250. In deck C this loss was the result of five smaller losses which totalled -£250 in every ten trials. In deck D there was one loss of -£250 in every ten trials. Thus making ten successive choices from decks A and B resulted in an average loss of -£250 whereas making ten successive choices from decks C and D resulted in a net gain of £250. Choices from C and D are advantageous in the long-term whereas choices from A and B are disadvantageous despite decks A and B having the highest immediate gain. The experiment was always conducted with hypothetical money.

The task ended after one hundred cards (trials) were selected. Participants were not aware of this prior to the task ending. When the task ended a screen informed participants that the task had ended and displayed their 'winnings'. This figure was recorded by the experimenter and its implications (either a profit or a loss made on the loan of fake money) were communicated to the participant.

The second session followed exactly the procedure of the first with the addition of the following details. Participants were told that the task was exactly the same as it was in the first session. They were reminded of the amount of money they had ‘won’ or ‘lost’. Participants were then instructed to read the instructions again and to ensure they were familiar with them. Again, the experimenter waited in the laboratory room in case the participant had any questions. If no questions were asked the experimenter left the room and the participants completed the task for a second time. Upon completion the amount of money won or lost was recorded and each participant was debriefed. Participation was rewarded with course credit in the form of 30 minutes of research participation time.

2.2.3 Results

Net score was calculated by subtracting the number of cards selected from the disadvantageous decks (decks A and B) from the number selected from the advantageous decks (decks C and D). This is the standard measure of performance on the IGT and it ranges from a minimum of -100 to a maximum of 100. A positive net score indicates a preference for the advantageous decks. Additionally, net score was calculated in blocks of twenty trials for each participant (displayed in Figure 2.2). This allows an estimate of learning rate to be calculated by looking at the change in participants’ net scores across block, i.e. the slope b . Performance in each session is discussed separately bearing in mind that the standard design features only one session.

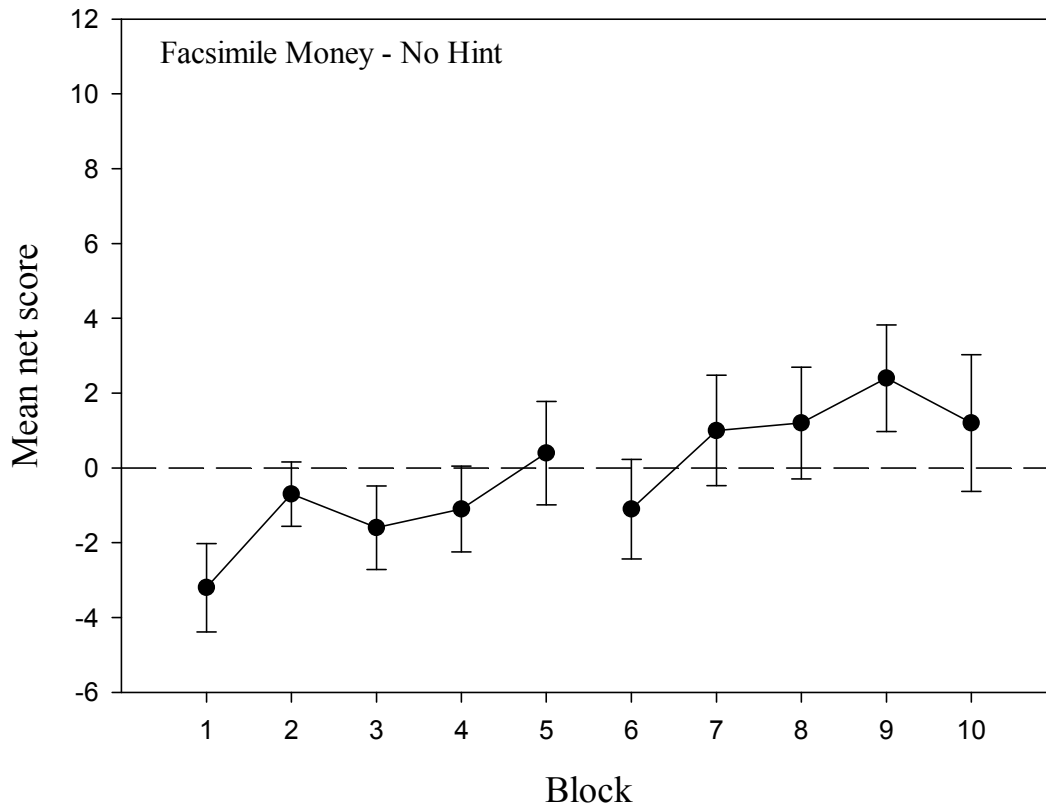


Figure 2.2: Mean net score across 20-trial blocks for each session in Experiment 1 (Facsimile reinforcers and No Hint instructions). The dashed line represents chance selection, or no preference for either advantageous or disadvantageous decks. Error bars are the standard error of the mean.

Session 1

Mean net score in session 1 was -6.2 ($SD = 14.92$). A one-sample t -test found that this was not significantly different from zero, or selection at chance, $t(19) = -1.86$, $SD = 14.92$, $p = 0.08$. However, in order to investigate whether participants' performance improved across the session two measures were used and again compared to chance (zero in both cases). First, a one-sample t -test revealed that mean net score in the final block of twenty trials was not significantly greater than zero, $t(19) < 1.0$, $SD =$

6.18, indicating that in the last twenty of one hundred trials participants had no preference for the advantageous decks. Second, it may be the case that participants' performance improved and while not preferring the advantageous decks their selection behaviour had moved away from a preference for the disadvantageous decks. In order to assess this possibility learning rate was analysed using the Lorch and Myers (1990) regression analysis for repeated measures designs. Mean learning rate in session 1 was 0.68 ($\sigma_M = 0.43$), indicating that learning rate was no different from zero, $t(19) = 1.60$, $SD = 1.90$, $p = 0.13$, i.e. participants' net scores did not increase with increased experience of the task.

Session 2

Mean net score in session 2 was 4.7 ($SD = 27.14$). A one sample t -test found that this was not significantly different from chance, $t(19) < 1.0$, $SD = 27.14$. A paired samples t -test compared net score between sessions and revealed that the increase approached significance, $t(19) = -2.03$, $SD = 24.02$ $p = 0.06$. This suggests that participants' performance improved. However, a one-sample t -test found that mean net score in the final block of session 2 was not significantly different from chance, $t(19) < 1.0$, $SD = 8.17$. These results indicate that although net score increased between sessions, at the end of the second session participants were still not showing a preference for either the advantageous or the disadvantageous decks. This was confirmed when mean learning rate in session 2 was examined, $b = 0.6$ ($\sigma_M = 0.34$). This was not significantly greater than zero, $t(19) = 1.77$, $SD = 1.52$, $p = 0.09$. These results suggest that even after two sessions on the IGT participants do not show the 'normal' behaviour reported for healthy controls on the IGT. As such these results

represent a failure to replicate the basic behaviour shown by Bechara et al (1994, 1997a, 1999).

2.2.4 Discussion

Participants in Experiment 1 did not develop a preference for the advantageous decks either after a standard administration of the IGT, or following a second session. This was reflected in the mean net score across both sessions and in the mean learning rates. Although learning rate was positive and mean net score increased between sessions, participants in Experiment 1 were still showing no preference for the advantageous decks after 200 trials. A total net score greater than 10 is indicative of 'normal' performance on the IGT as no participant with VMpfc damage had exceeded it (Bechara, Dolan, Denburg, Hindes et al, 2001). In the data reported here even in a second session mean net score was lower than this figure. Thus, Experiment 1 failed to replicate the behaviour reported as 'normal' by Bechara et al (1994, 1997a, 1999).

One possible explanation of these results is that while the instructions used in Experiment 1 were the same as those reported by Bechara et al (1994, 1997a) they were not the same as those reported by Bechara et al (1999, 2000). In these instructions participants are given more information about the task, specifically that some decks are worse than others and if these decks are avoided a winning strategy will be found. Such a suggestion has been made by Schmitt, Brinkley and Newman (1999) to explain similar behaviour in their experimental populations. However, no direct test of this hypothesis has been made. Experiment 2 was designed to test the possibility that this difference in instructions is sufficient to affect participants' behaviour on the IGT.

2.3 EXPERIMENT 2

'STANDARD' IGT BEHAVIOUR REPLICATED USING 'HINT' INSTRUCTIONS

2.3.1 Introduction

In Experiment 1 participants did not behave in the manner reported by Bechara et al (1994, 1997a) even after two sessions' experience on the IGT. At the end of these sessions participants did not show a preference for the advantageous decks. A similar result has been reported by Schmitt et al (1999) who found no difference between the performance of psychopaths and incarcerated controls. Neither group exhibited behaviour consistent with normal performance as both showed no clear preference for either advantageous or disadvantageous decks by the end of the session. Of the possibilities Schmitt et al (1999) suggested to explain the contrast between their results and those of Bechara et al (1994) the most intriguing was that although their instructions were the same as those published by Bechara et al (1994) they were not those given in the more detailed procedure published by Bechara et al (1999, 2000). In the more detailed instructions participants are given information about the nature of the decks and informed explicitly that the decks are not the same, that some are worse than others and that by staying away from the worst decks they can win. Including this "hint" in the instructions gives participants much more information about, and arguably changes, the nature of the task from one without the hint where the only information received about the nature of the decks comes from the results of one's own behaviour. Buehner and May (2004) found that less subtle changes in instructions affect human causal learning and Kudadjie-Gyamfi and Rachlin (2002) have found that including a hint about how to maximise long-term reinforcement in

the instructions of a choice task improves performance above the contingency-governed behaviour seen otherwise. Such a difference in instructions for the IGT could similarly affect behaviour.

The work of James Blair's group supports this hypothesis. Blair and Cipolotti (2000) used the Bechara et al (1994) no hint instructions and found no difference between psychopathic prisoners and incarcerated controls. But using the Bechara et al (1999, 2000) "hint" instructions Blair, Colledge, and Mitchell (2001) and Mitchell, Colledge, Leonard, and Blair (2002) found, respectively, differences in performance between boys with psychopathic tendencies and age-matched controls, and between psychopathic adults and incarcerated controls.

It may be the case that the instructions given to participants in Experiment 1 were insufficiently detailed for them to discern the nature of the task. Experiment 2 was conducted to test this possibility and as a further attempt to replicate the results Bechara et al (1994).

2.3.2 Method

Participants

Participants were first year undergraduates studying psychology at Keele University ($n = 20$, mean age = 19.2, $SD = 1.06$). Sixteen participants were female and four were male. They received course credit for taking part in this experiment.

Apparatus

As Experiment 1 but with the addition of the following instructions (addition in bold):

It is important to know that just like in a real card game, the computer does not change the order of the cards after the game starts. You may not be able to figure out exactly when you will lose money, but the game is fair and the computer does not make you lose money at random. **All I can say is that some decks are worse than others. You may find all of them bad, but some are worse than others. No matter how much you find yourself losing you can still win if you stay away from the worst decks.**

Design

As Experiment 1.

Procedure

As Experiment 1.

2.3.3 Results

As in Experiment 1 net score was calculated for each session and in blocks of twenty trials for each participant. Figure 2.3 displays the mean net scores across trial blocks for each session in Experiment 2. The same measurement from Experiment 1 has been retained to allow a visual comparison to be made. However, no cross-experimental comparisons will be made between Experiment 1 and 2. This is to avoid Type I error inflation due to multiple testing as a Factorial cross experimental comparison will be made in section 2.5 after Experiment 3 has been presented. As a result, until section 2.5 comparisons between Experiments will be general and restricted to descriptions of similarities or differences in their results.

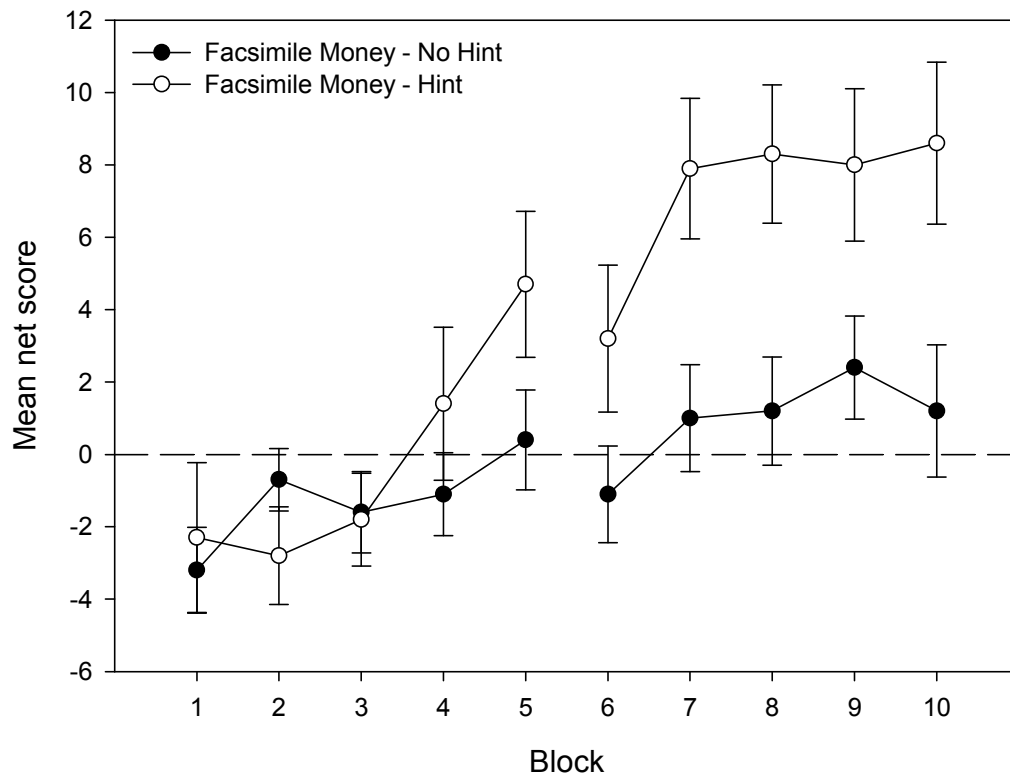


Figure 2.3: Mean net score across 20-trial blocks for each session in Experiment 2 (Facsimile reinforcers and Hint instructions). The dashed line represents chance selection, or no preference for either advantageous or disadvantageous decks. Error bars are the standard error of the mean.

Session 1

Mean total net score in session 1 was -0.8 ($SD = 22.58$). A one-sample t -test revealed that this was not significantly different from chance, $t(19) > -1.0$, $SD = 22.58$.

However, the mean net score in the final block was significantly greater than zero, $t(19) = 2.33$, $SD = 9.02$, $p = 0.03$, indicating that unlike in Experiment 1 participants finished the first session with a preference for the advantageous decks. This was

reflected in a mean learning rate, $b = 1.82$ ($\sigma_M = 0.62$), that was significantly greater than zero, $t(19) = 2.95$, $SD = 2.76$, $p < .01$.

Session 2

Mean total net score was 36 ($SD = 40.26$) and significantly greater than chance as revealed by a one-sample t -test, $t(19) = 4.00$, $SD = 40.26$, $p < 0.01$. A paired samples t -test revealed a significant difference between sessions, $t(19) = 4.59$, $SD = 35.87$, $p < 0.01$. At the end of session 2 mean net score in the final block of trials was significantly greater than chance, $t(19) = 3.84$, $SD = 10.01$, $p < 0.01$. Together these results and those from session 1 suggest that receiving the Bechara et al (1999) Hint instructions results in a preference for the advantageous decks. Indeed, the preference that emerged during session 1 continued into session 2 where learning rate was significantly greater than zero, $b = 1.09$ ($\sigma_M = 0.49$), $t(19) = 2.22$, $SD = 2.20$, $p = 0.04$. Figure 2.3 illustrates that this reflects the increase in net score from block 6 to block 7 before net score stabilises in a preference for the advantageous decks and flattens out.

2.3.4 Discussion

Participants in Experiment 2 developed a preference for the advantageous decks in the first session that continued into the second session. This was reflected in the increase in mean net score between sessions and in learning rates that were significantly greater than zero in both sessions. It appears that when the instructions include a strategy hint participants performance improves and is more like the behaviour displayed by the healthy controls in Bechara et al (1994, 1997a, 1999). However, it is of note that in the first session in Experiment 2 (and in both sessions in Experiment 1)

the mean net score is much lower than the average net score of around 39 reported as normal by Bechara et al (1994) or the mean net scores in the twenties reported by other groups (Petry, Bickel and Arnett, 1998; Grant, Contoreggi and London, 2000; Mazas, Finn and Steinmetz, 2000). According to Bechara et al (2001) a total net score greater than 10 is indicative of ‘normal’ performance on the IGT as no participant with VMpfc damage has exceeded it. In Experiment 2, despite participants’ learning rate (the change in net score) increasing significantly above chance they do not select a comparable number of cards from the advantageous decks. The failure to entirely replicate this standard normal behaviour is troublesome. The changes made to the administration of the IGT may be one explanation for these low net scores. I will examine each in turn.

In this experiment the gambling task was conducted on a computer and the lack of an experimenter interaction may have affected behaviour. However, no differences in performance were found when behaviour on the computerized task was compared to behaviour on the manual task (Bowman et al, 2005) and indeed Bechara and colleagues have not reported in any differences attributable to administration method (Bechara et al, 1999).


Another change in the methodology and one different from most previous studies is that participants were instructed to read the task instructions rather than have them read out to them. This change was made in order to control for possible confounding effects due to the experimenter and to avoid any emphasis on any of the instructions. It may be that in the reading of the instructions participants did not fully comprehend the nature of the task, or in Experiment 2 pick up on the hint as quickly as when spoken instructions were used. This may be why participants who receive the hint instructions improve so much in the second session – after some experience

on the task it is easier to spot the good decks from the bad. Participants in both experiments were after all told to re-read the instructions.

A third possibility concerns the number of cards in each deck. Bechara et al (1994) used 40 cards in each deck. This number was used as participants very rarely chose from the same deck more often than 40 times. However, because this was a possibility the number of cards in each deck was increased to 60 cards in their 2000 methodology (Bechara et al, 2000). Here there were 100 cards in each deck to ensure participants were not limited in their choices in any way. Thus it was possible for participants to select entirely from one deck. In practice this rarely happened in the first session, though in the second session this behaviour increased, presumably as participants had gained an understanding of the task. This change in the methodology is unlikely to have affected participant behaviour, especially as in the computerized version of the task participants can have no idea of the total number of cards from which they are able to select. This is not the case in the manual version where all cards are laid out in piles in front of participants.

However, focusing on individual decks does present a further possibility. Figure 2.4 shows the mean number of cards selected from each deck in each session of Experiment 1 and Experiment 2. What Figure 2.4 shows is that in session 1 of Experiment 1 participants selected most cards from deck B and deck D, although only selection from B is above chance. However, by the end of the second session the participants have no clear preference for any deck (mean card selections are at chance) except in deck A where selection is below chance. Whereas although Experiment 2 follows a similar pattern in session 1, by the end of session 2 selection

a) Experiment 1



b) Experiment 2




Figure 2.4: Mean number of cards selected from each deck between sessions in a) Experiment 1 and b) Experiment 2. Error bars are the standard error of the mean.

from B is substantially reduced and selection from C is substantially increased. These differences raise two important points. The first is that selection of the two advantageous and the two disadvantageous decks may not be as homogenous as is generally reported and assumed in analyses using net score. Secondly, if this is the case, then it is not only the expected values of the choices on the IGT that drives learning on the IGT. This has major implications for the discussion of behaviour on the IGT – a point that will be returned to and elaborated in section 2.6 when the results of the cross-experimental comparisons have been described.

The results from Experiment 2 suggest that the instructions participants receive impact on learning in the IGT. Participants' learning rates were greater than chance when the instructions carried a strategy hint. In both Experiments 1 and 2 participants' reward for participating was in the form of course credit with their on-line performance reinforced with facsimile reinforcers (though they 'won' £100 on the task, in reality they did not win anything). Bechara et al (1994, 1997a, 1999) pay their participants for taking part but use facsimile reinforcers on the gambling task. However, given that the task is named as a gambling task it might be argued that real money would be a more realistic reinforcer. Experiment 3 was run to examine whether Reinforcer Type influences performance on the IGT.

2.4 EXPERIMENT 3

THE EFFECT OF INSTRUCTION TYPE USING REAL MONEY REINFORCERS

2.4.1 Introduction

The results of Experiments 1 and 2 allowed participants' Iowa Gambling task (IGT) performance to be examined in two conditions in which different instructions had been used. Participants' performance was better when the instructions carried a strategy hint, and replicated the behaviour found in normal participants in Bechara et al (1994, 1997a, 1999). Like Bechara et al (1994, 1997a, 1999), in Experiments 1 and 2 participants were playing for facsimile reinforcers. However, given that the task has been described as emulating real-life decision environments real money would be a more realistic reinforcer. Real money reinforcers may also improve performance. Vulkan (2000, p.111) notes in his review of repeated, binary choice experiments that introducing monetary rewards increases maximising behaviour (in the case of the IGT this would be increasing advantageous card selections). Although the standard interpretation and analysis of the IGT assumes it is fundamentally a binary choice task there are four choices and so it is not clear whether real monetary rewards would have a similar effect.

Bowman and Turnbull (2003) directly compared real versus facsimile reinforcers on the IGT. The schedules of gains and losses were exactly the same in both conditions but the values in the real money conditions were one-thousandth those of the facsimile money condition. Thus £100 in facsimile money was £0.10 in real money. Bowman and Turnbull (2003) found that Reinforcer Type did not significantly affect performance. However, they report less variance in net scores

when real reinforcers are used. Less variance in net scores means any differences between groups are clearer. Their results also suggest that it is the nature of the task and not the nature of the reinforcer that is important in influencing performance. However, following the learning curves produced by each group in their experiment, they suggest that real money reinforcers may have a differential effect on performance given more exposure to the task, with performance in the real money condition superior to that of facsimile reinforcers.

What is not clear is the effect that the different instructions may have on behaviour when real reinforcers are used. With the added incentive of actual winnings the Hint instructions may be additive resulting in superior performance. The purpose of Experiment 3 was to replicate the results found in Experiments 1 and 2 using a real money reinforcer. Such a replication would indicate that it is the task instructions and not Reinforcer Type that are important in influencing performance on the IGT. Like Experiments 1 and 2 a second session was included to investigate participants' learning. This second session allows the testing of Bowman and Turnbull's (2003) hypothesis that real money reinforcers will improve performance over facsimile reinforcers in the longer term.

2.4.2 Method

Participants

Participants were recruited from the Keele University community through a poster advertisement that offered the opportunity to "Earn up to £6" by participating in a cognitive psychology experiment. The data from five participants was excluded as three participants did not attend the second experimental session, one was mistakenly

put into the wrong condition for the second session and one did not follow verbal instructions to re-read on-screen instructions prior to the start of the second session. As a result forty participants were tested and randomly allocated to one of two conditions. Participants in condition A had a mean age of 24.3 ($SD = 3.99$), thirteen were female and seven were male. Participants in Condition B had a mean age of 25.4 ($SD = 7.22$), sixteen were female and four were male.

Apparatus

The apparatus used was identical to that used in Experiments 1 and 2 with but one change due to the change in reinforcers. The reinforcement schedules used were identical to those used in Experiments 1 and 2 except that all figures used were divided by 1000, i.e. £100 in Experiments 1 and 2 equalled £0.10. As a result the amount of money 'loaned' to the participant to play the game was £2.00 and the maximum they could earn in any one hundred trials was £3.00.

Design

The design followed that of Experiments 1 and 2 with participants randomly allocated to one of two conditions equivalent to these previous experiments. Participants in condition B received the instructions containing the strategy hint.

Procedure

As Experiment 1, with the only adaptations due to the change in reinforcer used. Hence, participants were told that they were taking part in a cognitive task in which

the goal was to earn as much money as possible. They were told that they could earn up to £6 over the two sessions i.e. £3 per session. However, they were warned that there was the possibility that they would earn nothing. When the task ended after 100 trials participants were given the amount of money they had earned, reminded that they could have earned up to £3 and would have a further opportunity to do so in the second session. Upon completion of the second session any earnings were paid. If a loss had been made participants did not incur the expense. Each participant was then debriefed and reimbursed £2 for expenses incurred in traveling to the university campus (the cost of a return bus fare). This payment was not mentioned until the completion of the experiment.

2.4.3 Results

As in the previous experiments net score was calculated for each session and in blocks of twenty trials for each participant. Figure 2.5 displays the mean net scores across trial blocks for each session in each condition of Experiment 3.

2.4.3.1 No Hint condition

Session 1

When participants received the No Hint instructions, mean net score was 6.1 ($SD = 15.07$) and was not significantly greater than chance, $t(19) = 1.81$, $SD = 15.07$, $p > 0.05$. Using a one-sample t -test mean net score in block 5 was also no different from chance $t(19) < 1.0$, $SD = 8.86$. However, as in Experiment 2 it may be the case that behaviour changed away from a preference for the disadvantageous decks over the




Figure 2.5: Mean net score across 20-trial blocks in each session of each condition of Experiment 3. The dashed line represents chance selection, or no preference for either advantageous or disadvantageous decks. Error bars are the standard error of the mean.

session. To examine this possibility, learning rate was estimated and compared to zero (no learning) using the Lorch and Myers (1990) method of regression analysis. Mean learning rate approached a significant difference from zero, $b = 1.08$ ($\sigma_M = 0.54$), $t(19) = 1.99$, $SD = 2.42$, $p = 0.06$. This result reflects the increase in mean net score from block 1, where participants initially prefer the disadvantageous decks, to blocks 3 and 4 where preference is for the advantageous decks until in block 5 when mean net score dips back towards zero. These results suggest that participants do develop a preference for the advantageous decks despite not receiving the strategy

hint instructions. However, advantageous deck selections dip in the last 20 trials suggesting participants, on average, do not know which decks are the best. If participants do have a preference for the advantageous decks this should become apparent in session 2.

Session 2

Mean net score was 11.3 ($SD = 34.61$) but not significantly greater than chance, $t(19) = 1.46$, $SD = 34.61$, $p > 0.05$. Similarly, a one sample t -test found that net score in block 10 was not greater than chance, $t(19) < 1.0$, $SD = 10.61$. When session 2 net score was compared to net score in session 1 using a paired samples t -test, no significant difference was found, $t(19) < 1.0$, $SD = 34.46$. These results suggest that although participants learning rate increased in session 1, by the end of session 2 the participants in this No Hint condition were still not showing a preference for the advantageous decks. This conclusion was further supported when learning rate is examined. Unlike in session 1, learning rate is no greater than zero, $b = 0.36$ ($\sigma_M = 0.53$), $t(19) < 1.0$, $SD = 2.35$. Like in Experiment 2 when participants receive the No Hint instructions no preference for the advantageous decks developed.

2.4.3.2 Hint condition

Session 1

Mean net score was 4.1 ($SD = 24.3$) and not significantly greater than chance, $t(19) < 1.0$, $SD = 24.3$. When mean net score in block 5 was compared to chance a one sample t -test found no significant difference, $t(19) = 1.02$, $SD = 7.44$, $p > 0.05$,

indicating that after one hundred trials participants were not selecting preferentially from the advantageous decks. However, mean learning rate was positive and significantly greater than zero, $b = 1.41$ ($\sigma_M = 0.46$), $t = 3.06$, $SD = 2.06$, $p < 0.01$, indicating that mean net score was increasing with experience on the task. From Figure 2.5 it is apparent that participants have comparable behaviour in the first session between conditions. However, Figure 2.5 suggests that participants who receive the Hint instructions, unlike those who did not, develop a preference for the advantageous decks in the second session.

Session 2

Figure 2.5 shows mean net score is higher than in session 1 and increases across session 2 until flattening out at positive level. This impression was confirmed when mean net score was compared to chance; mean net score = 29.2 ($SD = 42.39$), $t(19) = 3.08$, $SD = 42.39$, $p < 0.01$. A paired samples t -test found this was significantly greater than mean net score in session 1, $t(19) = -3.54$, $SD = 31.74$, $p < 0.01$. Unlike in the No Hint condition participants also showed a preference for the advantageous decks in the final block of the session. A one-sample t -test found mean net score in block 10 to be significantly greater than chance, $t(19) = 2.39$, $SD = 11.43$, $p < 0.05$. Further evidence that participants in this condition developed a preference for the advantageous decks was found when learning rate was compared to zero and approached a significant difference, $b = 0.84$ ($\sigma_M = 0.44$), $t = 1.89$, $SD = 1.99$, $p = 0.07$.

2.4.4 Discussion

The general findings from Experiment 3 provide support for the hypothesis that the instructions participants receive affect their subsequent performance on the IGT.

Although learning rates in the first sessions in both conditions indicated that participants' net scores became more positive with experience on the task, no preference for the advantageous decks was found when mean net scores in the final block of the session was compared to chance. This result suggests that when participants are playing for real money reinforcers the instructions they receive do not influence behaviour. However, despite the similarity in session 1 behaviour (as illustrated in Figure 2.5), a clear preference for the advantageous decks only develops in session 2 in those participants who received the Hint instructions.

In Experiment 2 it was suggested that learning on the IGT may not be as simple as the development of a preference for the advantageous decks. The improvement in participants who received the Hint instructions was reflected in the change in their deck selection behaviour between sessions. These participants increased their selections from deck C and decreased their selections from deck B, with little change in selection from A or D. If this is how learning on the IGT proceeds the same pattern should be seen for the participants who received the Hint instructions in this Experiment. Figure 2.6 displays the mean number of card selections from each deck in each session for each condition in Experiment 3. A similar pattern of results is found. The major changes in deck selection in participants who received the Hint instructions come in a decrease in deck B selections from above chance in session 1 to below chance in session 2. The opposite pattern is seen for deck C while there is very little change in the other decks. This pattern can be discerned in the participants who received the No Hint instructions, but it is by no

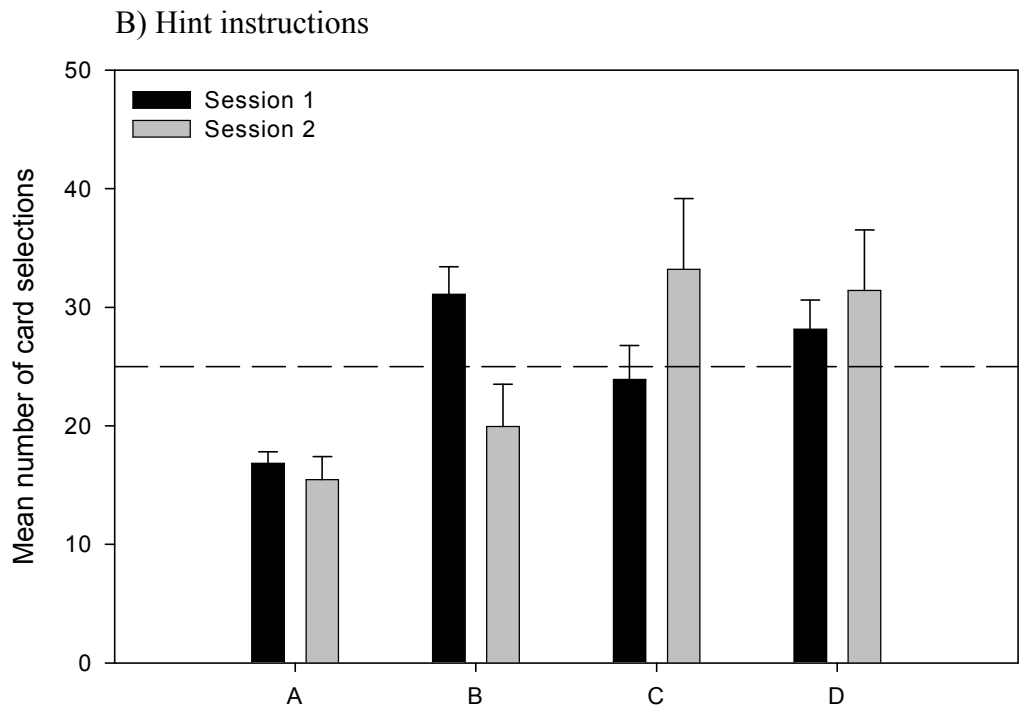
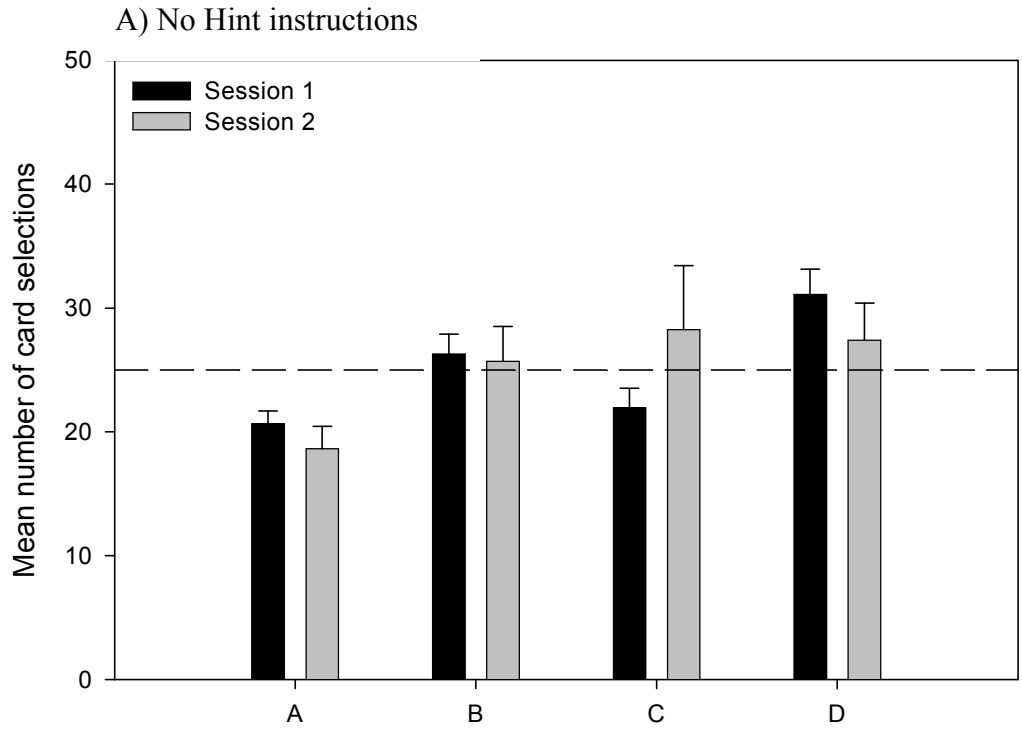


Figure 2.6: Mean number of cards selected from each deck between sessions with condition A (no hint instructions) and condition B (Hint instructions). Error bars are the standard error of the mean.

means as stark. It appears that learning on the IGT emerges as a result of a decrease in selection from deck B and an increase in selection from deck C. This point will be returned to and elaborated on in the general discussion (section 2.6).

The general findings from Experiment 3 suggest that the instructions participants receive influence their performance on the IGT. This is the same pattern of results that was found when participants received the facsimile money reinforcers in Experiments 1 and 2. However, these general findings need to be confirmed by comparing across all experiments. This will permit a comparison between Reinforcer Types to assess whether the type of money participants receive affects IGT performance. These analyses will also investigate whether there is any additive or differential effect of Instruction and Reinforcer Type. Section 2.5 goes on to describe these analyses

2.5 CROSS EXPERIMENT COMPARISONS

INSTRUCTIONS AFFECT PERFORMANCE ON THE IGT

The preceding experiments have described various manipulations in the administration of the IGT. From the general descriptions of their results the impression has formed that Instruction Type, not Reinforcer Type, influences IGT performance. In this section the results of these experiments will be compared factorially in an attempt to assess the influence of Instruction and Reinforcer Types on learning. This is an important consideration because while the Iowa group and others have used facsimile reinforcers and, presumably, the hint instructions, other combinations of each Reinforcer and Instruction Type have been used with varying results. For example, Mazas et al (2000) have replicated Bechara et al's (1994)

results using the no hint instructions and real money reinforcers. Like Mazas et al (2000), Grant et al (2000) and Petry et al (1998) paid participants for taking part (this is implied in Grant et al as participants were recruited through newspaper advertisements) and found controls performed better than substance abusers, though unlike Mazas et al facsimile reinforcers were used. In addition, Petry et al (1998) added an incentive payment of \$10 if participants finished the task in profit. Given the different methodologies used in the administration of the IGT it is not clear which factors are important in influencing IGT performance. The comparisons reported in this section will resolve what has become a hazy picture.

The preceding experiments were effectively manipulations of two factors (Instruction Type and Reinforcer Type). Experiment 1 looked at the effects of No Hint instructions with Facsimile Money reinforcers. Experiment 2 varied the instructions and included a Hint while Facsimile Money reinforcers were retained. In Experiment 3 Real Money reinforcers were used but in one condition the instructions contained No Hint while in the other they did contain the Hint. Because the standard design of the IGT features only one session performance in each session is analysed separately.

2.5.1 Session 1

A 2x2x5 (Reinforcer Type by Instruction Type by Block) mixed-design ANOVA was performed to investigate differences in net score in each experimental Block due to the experimental factors. Alpha was set at .05 for all tests unless otherwise stated. Due to a violation of the assumption of sphericity in the calculation of the repeated measures components of the ANOVA the Greenhouse Geisser correction is used where appropriate (in analyses in this and subsequent chapters).

The ANOVA revealed a main effect of Block, $F(3.25, 246.66) = 9.35$, $MSE = 43.16$, $p < 0.05$, indicating that net score differed between some blocks and reflecting the trend for net score to increase with exposure to the task. A main effect of Reinforcer Type was also found, $F(1, 76) = 3.84$, $MSE = 77.08$, $p = 0.05$; mean net score across Block was higher when participants were earning Real Money ($M = 1.02$, $\sigma_M = 0.62$) rather than Facsimile Money ($M = -0.7$, $\sigma_M = 0.62$). No other main effects or interactions were significant: Instruction Type, $F(1, 76) < 1.0$; Reinforcer Type by Instruction Type, $F(1, 76) < 1.0$; Block by Reinforcer Type, $F(3.25, 246.66) = 2.20$, $MSE = 43.16$, $p > 0.05$; Block by Instruction Type, $F(3.25, 246.66) = 1.28$, $MSE = 43.16$, $p > 0.05$; Block by Reinforcer Type by Instruction Type, $F(3.25, 246.66) > 1.0$.

Net score was greater when the reinforcer used was Real rather than Facsimile Money. This finding that monetary payoffs improve performance is in agreement with the behavioural decision-making literature (Hertwig and Ortmann, 2001; Camerer and Hogarth, 1999) but it does not replicate Bowman and Turnbull's (2003) result. They found no effect of Reinforcer Type when comparing the performance of a similar population. Bowman and Turnbull (2003) used Bechara et al's (1999) Hint instructions and the manual version of the gambling task. In contrast a computerized version of the task was used here and the effect of Reinforcer Type in this study is from data collapsed across groups who received different sets of instructions. There was no Reinforcer Type by Instruction Type interaction in the mixed-design ANOVA implying that Hint instructions did not differentially affect performance in either of the Reinforcer Type groups. However, to more clearly compare Bowman and Turnbull's (2003) results with our own, simple effects analyses¹ were conducted on

¹ Keppel (1991, p.384) recommends against pooling the error variance in mixed factor designs so separate error terms were used for the simple effects analyses.

the net scores for Session 1, which is the standard dependent variable. Comparing between the two levels of Reinforcer Type for the participants who received the Hint instructions is equivalent to Bowman and Turnbull's analysis. No effect of Reinforcer Type was found, $F(1, 38) = -1.36$, $MSE = 545.16$, $p > 0.05$, replicating Bowman and Turnbull's (2003) findings. However, analysis of the simple effects of Reinforcer Type for participants who received the No Hint instructions revealed an effect of Reinforcer Type, $F(1, 38) = 6.73$, $MSE = 224.92$, $p < 0.05$. Net score was significantly higher ($M = 6.1$, $\sigma_M = 3.37$) when the reinforcer was Real Money than when it was Facsimile Money ($M = -6.2$, $\sigma_M = 3.33$).

These findings go some way to supporting Bowman and Turnbull's (2003) result, but also provide support for the impression developed in the preceding analyses that receiving the Hint instructions changes behaviour on the IGT. The results of these further analyses indicate that Reinforcer Type does have an effect on IGT performance, but that this effect is only apparent when task instructions do not contain a Hint. These results suggest that behaviour on the IGT is influenced in at least two ways. When the only information about the contingencies of the decks is available from the results of one's own card selections (No Hint instructions) mean net score is affected by the level of incentive on offer. Real Money incentives result in more selections from the advantageous decks, consistent with the behavioural decision-making literature. However, when information on the deck contingencies is available (Hint instructions: "some decks are worse than others") the effect of Real Money reinforcers is cancelled out (Experiment 3b: Real Money – Hint) and participants who have no financial incentive to do well but information about the nature of the task (Experiment 2: Facsimile Money – Hint) are able to succeed. It may be that the level of incentive was not high enough within the Hint conditions, although it was

sufficient to facilitate a division in performance when participants had less information. This interpretation suggests that there are influences on learning in the IGT that are not additive.

2.5.2 Learning rates

While a main effect of Block is indicative of improved performance and therefore learning, more specific information is available from examination of learning rates. Learning rate has been considered in the preceding Experiments to indicate if net score increases with experience on the IGT. An advantage of using the slope b as an estimate of learning rates is it can be compared both within and across groups as well as allowing predictions of future performance to be made. Table 2.1 displays mean learning rates and shows that in each condition they are positive indicating that the number of selections from the advantageous decks increased with exposure to the task.

In order to determine if these learning rates were different, a 2x2 (Instruction Type by Reinforcer Type) independent-measures ANOVA was performed. No significant differences in learning rate were found between Reinforcer Types, $F(1, 76) < 0.01$, $MSE = 5.43$, $p > .05$; Instruction Types, $F(1, 76) = 2.02$, $MSE = 5.43$, $p > 0.05$; nor was there any interaction, $F(1, 76) = .61$, $MSE = 5.43$, $p > .05$. The effect of Reinforcer Type does not emerge when learning rate is examined implying that despite a difference in net scores between some of the experimental groups learning rate does not vary between them.

However, the highest learning rates in Session 1 followed receipt of the Hint instructions. This would suggest that with increased trials the difference between participants who had received the Hint instructions and those who had not should

Table 2.1: Summary of mean learning rates by session in each Experiment.

	Reinforcer		
	<u>Instructions</u>	<u>Facsimile</u>	<u>Real Money</u>
<u>Session 1</u>			
No Hint		0.68 (.42)	1.08 (.54)
Hint		1.82 (.62)	1.41 (.46)
<u>Session 2</u>			
No Hint		0.60 (.34)	0.36 (.52)
Hint		1.09 (.49)	0.84 (.44)

Note: n = 20 in all cases. Figures in parentheses are the standard error of the mean.

increase. As a result participants who received the Hint instructions should reach asymptotic performance sooner. Bowman and Turnbull (2003) suggested that with more trials a similar difference would be seen between Reinforcer Types. No replication of this trend was found when the conditions are the same as in Bowman and Turnbull (2003). Extrapolating from session 1 data such an effect would only be expected when No Hint instructions are used. However, by looking at participant performance in a second session on the IGT these predictions can be investigated.

2.5.3 *Session 2*

The second session on the IGT allowed an examination of when, and in what conditions, asymptotic performance is reached. The second session also allows a test

of Bowman and Turnbull's (2003) hypothesis that increased trials on the IGT would lead to differential performance between Reinforcer Types.

As in Session 1 a 2x2x5 (Reinforcer Type by Instruction Type by Block) mixed-design ANOVA was performed to investigate differences in net score due to the experimental factors. ANOVA revealed a main effect of Block, $F(3.33, 253.36) = 6.45$, $MSE = 33.01$, $p < 0.05$, indicating that net score differed between some blocks. Contrary to the Session 1 analysis a main effect of Instruction Type was also found, $F(1, 76) = 9.05$, $MSE = 267.60$, $p < 0.05$; net score was higher when participants received Hint instructions ($M = 6.52$, $\sigma_M = 1.16$) rather than No Hint instructions ($M = 1.60$, $\sigma_M = 1.16$). Also in contrast to Session 1 no effect of Reinforcer Type was found, Reinforcer Type, $F(1, 76) < 1.0$. No other main effects or interactions were significant: Reinforcer Type by Instruction Type, $F(1, 76) < 1.0$; Block by Reinforcer Type, $F(3.33, 253.36) < 1.0$; Block by Instruction Type, $F(3.33, 253.36) = 1.66$, $p > 0.05$; Block by Reinforcer Type by Instruction Type, $F(3.25, 246.66) < 1.0$.

In Session 1 participants who received the Hint instructions had the highest learning rates. These were not significantly different to those of participants who had received the No Hint instructions but did suggest that differential performance would be found with more trials on the IGT. Prior to the publication of these experiments (Fernie and Tunney, 2006) a reviewer suggested an alternative explanation – that in adding a second session after a 48-hour delay a further change was made to the standard IGT procedure. It is plausible that in re-reading the instructions after some experience on the IGT participants who received the Hint were better able to utilise the help afforded by the Hint instructions and this subsequently affected their performance. However, given that the learning rates following Hint instructions were higher in session 1, while not significantly different to those following No Hint

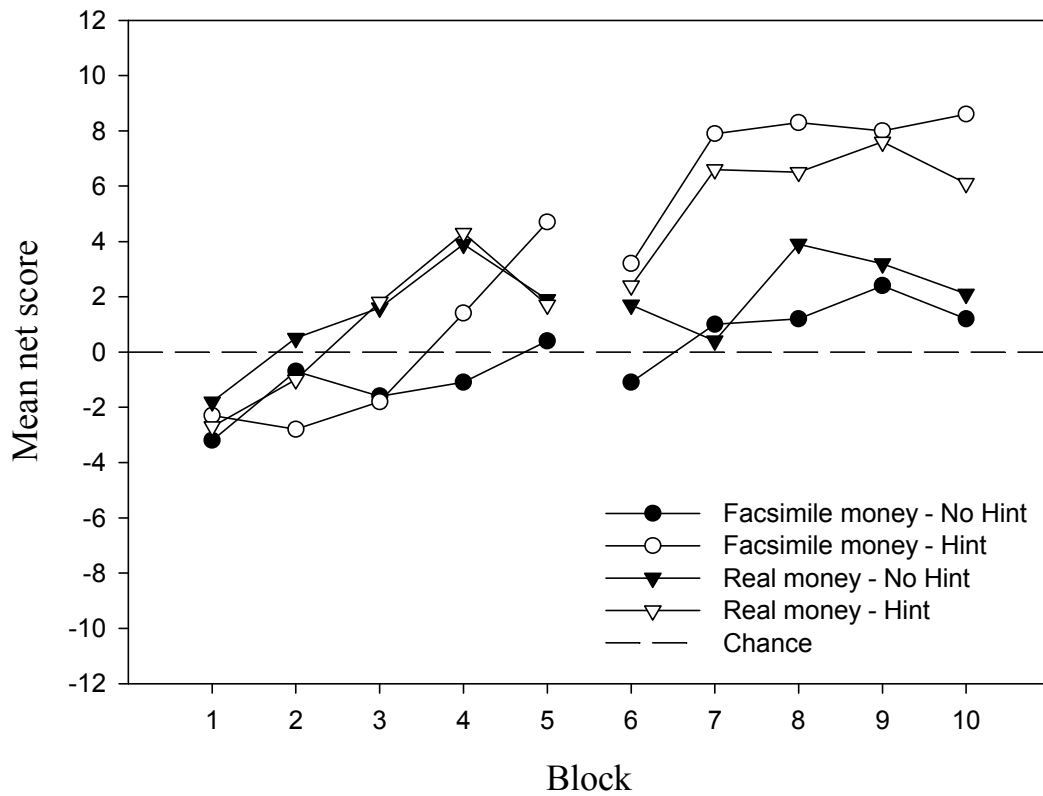


Figure 2.7: Mean number of cards selected from each deck by block for each experimental condition. Note that blocks 1 to 5 are session 1 and blocks 6 to 10 are session 2. A mean selection of 5 cards represents chance. Error bars are removed for sake of clarity.

instructions, a distinction in performance did exist between instruction groups and it seems plausible this would have occurred whether the second session followed immediately or after a delay.

It is puzzling that no effect of Reinforcer Type was found, even at one level of the Instruction Type factor. In Session 1 learning rate was higher in the No Hint - Real Money condition than in the No Hint - Facsimile Money condition, and a significant difference in net score was found between these conditions. However, this

effect disappears in the second session. Figure 2.7 displays mean net score across block in all Experiments in this Chapter. The data in this Figure suggest that there is little learning in the No Hint conditions, despite the difference in Reinforcer Type. Two explanations for this result spring to mind. Participants in the No Hint - Real Money condition may not have acquired an understanding of the task in the same way or to the same extent as those who received the Hint instructions; or the introduction of the break between sessions has affected behaviour in some way. If this were the case then a continuation of the Session 1 trend may have been found had the second session immediately followed the first.

2.5.4 Learning rates in Session 2

Learning rates were estimated as in Session 1 by calculating the slope b from the change in net score across block. A 2x2 (Instruction Type by Reinforcer Type) independent-measures ANOVA found a main effect of Instruction Type, $F(1, 76) = 10.36$, $MSE = 1.37$, $p < 0.01$; no effect of Reinforcer Type, $F(1, 76) = 1.26$, $p > 0.05$; and no interaction, $F(1, 76) = 0.29$, $p > 0.05$. These results reflect the rapid rise in net score between blocks 6 and 7 in the Hint instruction conditions despite net score reaching asymptote in the subsequent blocks. No such rise occurs following No Hint instructions and the learning rate remains flat (Figure 2.7), suggesting that with increased exposure to the IGT it is the instructions one receives that influence learning. Figure 2.7 shows that in both No Hint conditions mean net score is no different to chance in block 10 of the IGT despite the preceding 180 trials worth of experience and a gap between sessions in which to reflect on behaviour (No Hint – Facsimile Money, $t(19) < 1.0$, $SD = 8.17$, $p > 0.05$; No Hint – Real Money, $t(19) < 1.0$, $SD = 10.61$, $p > 0.05$).

2.6 GENERAL DISCUSSION

This Chapter details experiments which varied the administration of the Iowa Gambling Task in several ways that had the potential to influence performance. Participants were either given instructions that included a hint or did not, and received either real or facsimile money reinforcers. Performance was assessed by comparing net score across experimental blocks and between conditions. An effect of Reinforcer Type was found in the first experimental session. However, this effect was only apparent when participants received No Hint instructions, suggesting that Instruction Type influenced behaviour on the IGT despite no main effect emerging over the first one hundred trials. An effect of Instruction Type did emerge in the second experimental session as signalled in the first session by larger learning rates following receipt of the hint instructions. In fact, there was little evidence of learning in the participants who had received the instructions without the hint: after 200 trials mean net score in these groups was not significantly different to chance in the final 20-trial block. In addition, the effect of Reinforcer Type was not sustained into a second session in contrast to Bowman and Turnbull's (2003) prediction.

In the discussion of the general results of Experiments 2 and 3 it was suggested that selection was not uniform within the advantageous and disadvantageous decks, and in fact, learning was being driven by a reduction in selection from deck B and an increase in selection from deck C. This would be an important finding because the analysis of IGT performance using the net score measure assumes that the advantageous and disadvantageous decks are similar enough to allow the number of cards selected from them to be collapsed together. While this is true of the long-term gains, the decks also differ in the immediate gain and, crucially, the schedule of losses. Each deck is different from the others on at least one

of these three dimensions. Wilder, Weinberger and Goldberg (1998) and MacPherson, Phillips and Della Sala (2002) have reported that participants' card selection is not uniform within the advantageous and disadvantageous decks. Both studies report that their normal and clinical samples selected more cards from the decks where they received frequent gains without loss (decks B and D). This selection behaviour is consistent with what is observed in the animal operant conditioning literature and suggests that the ratio of wins to losses is important in card selection (Greenberg & Weiner, 1966).

From their published reports (Wilder et al, 1998) or personal correspondence (MacPherson et al, 2002) it is clear that both studies used Facsimile Money reinforcers and No Hint instructions. In Experiment 1 (Session 1) this preference for decks B and D was also found. In fact, this pattern was found in all Experiments (see Table 2.2). More information about the effect of Reinforcer Type reported in the cross-experimental comparisons can be found by comparing in Table 2.2 selection from decks B and D between Reinforcer Types when No Hint instructions are received. When Real Money was the reinforcer, selection was lower from deck B and higher from deck D than when Facsimile Money was the reinforcer. Comparing between Instruction Types, the mean total number of cards selected within each deck are much the same, which was reflected in the earlier cross-experimental comparisons of net score.

Figures 2.8a and 2.8b display the number of cards selected from each deck across block and within each experimental factor. This more detailed examination of deck selection than provided in Figures 2.4 and 2.6 reveals changes over time that are not found when only total cards selected are looked at. It is apparent from Figure 2.8a that selection from deck A remains consistent and below chance in all conditions,

Table 2.2: Mean number of cards selected from each deck per session by Instruction Type and Reinforcer Type.

	Session 1		Session 2		
	Instruction		Instruction		
	Reinforcer	No Hint	Hint	No Hint	Hint
A					
	Facsimile	20.45 (1.31)	18.8 (1.45)	21.85 (1.53)	14.05 (2.29)
	Real Money	20.65 (1.03)	16.85 (0.96)	18.65 (1.81)	15.45 (1.96)
B					
	Facsimile	32.65 (2.02)	31.6 (2.25)	25.8 (2.64)	17.95 (2.81)
	Real Money	26.3 (1.59)	31.1 (2.33)	25.7 (2.81)	19.95 (3.55)
C					
	Facsimile	21.65 (1.44)	21.05 (1.73)	25.55 (2.52)	34.2 (5.96)
	Real Money	21.95 (1.56)	23.9 (2.88)	28.25 (5.16)	33.2 (5.97)
D					
	Facsimile	25.25 (1.65)	28.55 (2.22)	26.8 (2.47)	33.8 (5.42)
	Real Money	31.1 (2.03)	28.15 (2.45)	27.4 (3.0)	31.4 (5.12)

Note: Figures in parentheses are the standard error of the mean.

while selection from deck D changes little despite remaining above chance. These observations do not tell us much more than what is available from looking at total card selections. However, this more local examination reveals that selection from deck B decreases while selection from deck C increases. This implies that the learning rate calculated from net score is driven by differential selection within the advantageous and disadvantageous decks. Indeed in Session 2, where the effect of

Instruction Type is found, this trend is even more apparent and shows that participants who received the Hint are correctly identifying the worst decks.

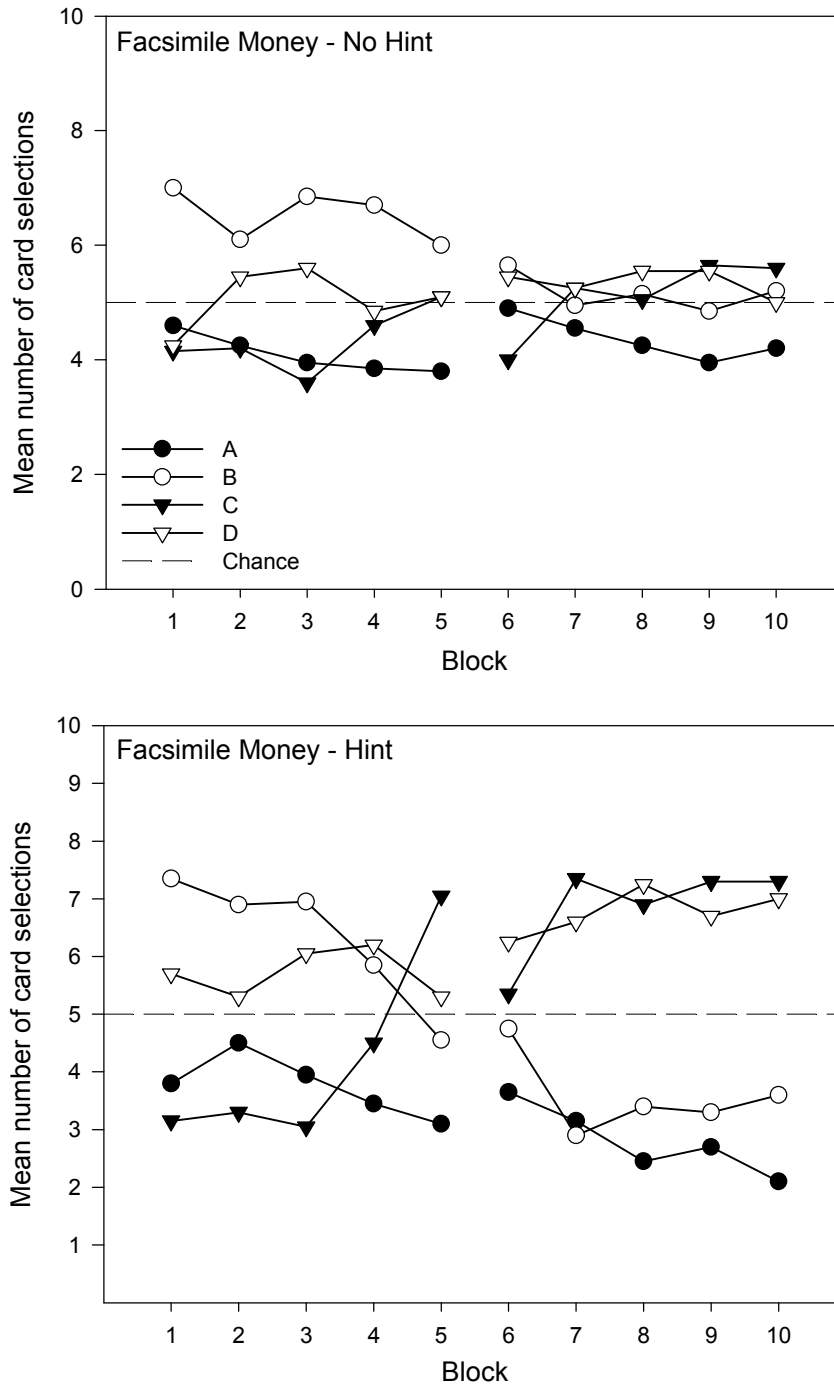


Figure 2.8a: Mean number of card selections in each block in Experiments 1 and 2.

Error bars are removed for the sake of clarity. The dashed line represents chance.





Figure 2.8b: Mean number of card selections in each block in Experiment 3. Error bars are removed for the sake of clarity. The dashed line represents chance.

The observation that there are differences in preference for decks with the same long-term expected value but different short-term contingencies implies that

learning in the IGT is not only governed by expected value. This is of interest because it may shed light on basic learning processes' sensitivity to the magnitude versus the frequency of reinforcement. This in turn has implications for researchers in all fields using the IGT. If clinical participants show differential deck selection within decks with the same expected values, then the contingencies offer a greater insight into what is influencing their behaviour on the IGT.

While these observations are interesting they must be interpreted cautiously. As the number of cards selected from one deck is dependent on the number selected from the others analysis using parametric statistics would violate the assumption of independence, making any analysis of differences between decks difficult to investigate, and this is the reason why only observations of the data have been reported here. However, the question of how the short-term deck contingencies impact on learning in the IGT is important. Therefore, Chapter 3 describes experiments where the frequency and magnitude of loss are manipulated to assess their importance on learning on the IGT.

The main finding from these Experiments is that the instructions participants receive affect their subsequent performance on the IGT. This should not come as a surprise. Being told that "some of the decks are worse than others and that if you stay away from the worst decks you will win" should affect participants' ability to distinguish between decks and thereby improve on the IGT. It is perhaps the strength of the gambling task paradigm that differences in performance between sub-populations and control groups are found without using the Hint instructions (Mazas et al, 2000; Petry et al, 1998; Grant et al, 2000). However, in altering the instructions participants are given, the nature of the task is less one of learning from behaviour as in a traditional decision-making task (or operant conditioning task), but one of

monitoring behaviour to distinguish between opposing (good and bad) options (c.f. Kudadjie-Gyamfi and Rachlin, 2002). This should certainly have an impact on participants' ability to describe what is going on as the task progresses (Bechara et al, 1997a; Maia & McClelland, 2004).

MANIPULATIONS OF LOSS FREQUENCY AFFECT LEARNING
ON THE IGT

3.1 INTRODUCTION AND OVERVIEW

In Chapter 2 it was suggested that participants' preferences within the advantageous and disadvantageous decks on the Iowa Gambling Task (IGT) were not uniform. Similar results have also been found in the few studies where deck selection from individual decks has been reported (Crone & van der Molen, 2004; MacPherson et al, 2002; Wilder et al, 1998). In these studies participants showed a preference for the two decks with infrequent losses (decks B and D) regardless of the long-term value of selecting from these decks; results that are in line with the Law of Effect (Thorndike, 1898). As these decks have a lower frequency of punishment, or a higher frequency of reward, the Law of Effect predicts that they will be chosen more often. Experiments 2 to 3 replicated and extended these findings. If only the mean number of cards selected in each session is considered, as in MacPherson et al (2002) and Wilder et al (1998), more cards were selected from decks B and D. But in a second 100-trial session on the IGT more cards were selected from the advantageous decks. This is the normal behaviour reported by Bechara et al (1994, 1997, 1999). However, the number of cards selected from any individual decks in a session does not tell the whole story. By plotting the number of selections from each deck over 20-trial blocks across the two sessions it appears that participants generally avoided the deck with high immediate reward but frequent higher losses (deck A) and an above chance

preference for the deck with low immediate reward but infrequent relatively large losses (deck D) did not change. No more information was added from this examination for these decks compared to that provided by looking at mean number of cards selected. However, in decks B and C changes across block were found. Participants initially preferred the deck with high immediate rewards and infrequent very large losses (deck B) but moved away from it with time (and with increasing losses), and developed a preference for the deck with low immediate rewards but frequent low losses (deck C). This change in selection from B to C appeared to underlie participants' learning on the IGT in Experiment 2 and the Hint condition of Experiment 3.

Very little work has examined selection from individual decks on the IGT or the relative importance of frequency or magnitude of punishment in choice behaviour. The majority of studies have considered the difference in selection from decks based on their expected values. However, some observations of differential selection within the advantageous and disadvantageous decks were summarised in the previous chapter (e.g. MacPherson et al, 2002, Wilder et al, 2002) and others have been reported in more detail.

Mintzer and Stitzer (2002) used the computerized IGT as one of a battery of neuropsychological tests to assess the cognitive performance of patients receiving methadone maintenance treatment for drug addiction. While patients had lower net scores than matched controls, their selection pattern differed within the decks with the lowest frequency (B & D), but not the highest frequency (A & C), of loss. Selection bias was measured by subtracting the number of selections of B from the number selected from D. The average net score for the patient group (-9.06, $SD = 16.25$) was significantly different from the control group (2.14, $SD = 14.13$). This measure

indicates that patients preferred deck B to deck D more than controls did. But the interesting measure is that although controls showed a slight preference for deck D (their D-B score was higher than zero), this was not different to zero indicating these participants were indifferent between the decks and suggesting that their improved performance relative to the patients was not a function of selection from deck D. The comparable measure for decks A and C, revealed no significant difference between groups, but the C-A measure was greater than zero for the control group (patients = 5.28, $SD = 14.69$; controls = 9.76, $SD = 13.94$). This shows that control participants preferred deck C to deck A, and coupled with the D-B selection measure, suggests that increased selection from C as well as decreased selection from B as compared to the patients was responsible for the higher net scores in that group. It is also of note that net scores for the control participants were not as high as reported by the Iowa group but are comparable to the lower values found in Experiments 1 to 3 and by others (e.g. Kleeberg et al, 2004; Rotherham-Fuller et al, 2004; Ritter, Meador-Woodruff and Dalack, 2004; Bowman et al, 2005; Harmsen, Bischof, Hohagen and Rumpf, 2006). Mintzer and Stitzer (2002) comment that previous studies of drug addicted populations have not reported differences as a function of loss frequency, but suggest that increasing frequency of loss may eliminate their reported effect. They go on to suggest further examination of this area is needed to understand the performance of their, and presumably other, clinical populations.

Fischer, Blommaert and Midden (2005) also noted differential activity within the disadvantageous decks. They examined performance on a manipulated version of the computerized IGT using real money reinforcers and no hint instructions. They found differential selection within the disadvantageous decks with deck A selected below chance and deck B selected above chance. Neither selection from decks C or D

were greater than chance. This contributed to an average monetary loss for participants, and occurred despite the manipulations to the original IGT design. The major changes in their task from the original IGT, aside from the computerisation and payment of real money, were fourfold. Firstly, gains and losses were reported separately in each deck (to facilitate computerisation of the experiment; but see also Peters and Slovic, 2000, discussed below). This meant that in the decks with low loss frequency there were nine gains of (10¢ in deck B; 5¢ in deck D) and one loss (115¢ in deck B; 20¢ in deck D), and in deck A there were five gains of 10¢ and five losses of 15¢. The second major change was in deck C. To avoid net yields of 0¢, the number of gains without loss was increased from five to seven. However, to maintain the expected value of +25¢, five of these gains were of 5¢, one was of 2¢ and one was of 1¢. The 3 losses were all of 1¢. These changes maintained the choice conflict in that decks A and B still had larger immediate gains but decks C and D had positive expected values. However, these manipulations make the probability of gain less certain for the participants (although the number of net gains remains the same) and magnitudes of loss are changed. This is especially relevant for selection from deck C. The third change was the randomisation of reinforcement schedules within each deck and this point is returned to in the discussion. The final major change was the addition of a money transfer aspect. In the original manual version of the IGT participants were required to give the experimenter money when they lost. This feature was restored by making participants move money from the computer dealer's account area to the participant's account area of the screen, or vice versa. This change had an interesting consequence. In a pilot study Fischer et al (2005) reported that participants developed a dislike for decks A and C compared to B and D. However, from their observations and debriefings of participants they discerned that

participants found the transfer of money with more frequent losses tedious. The authors suggested that the large number of mouse movements involved in these transfers may make avoidance of them indicative of some optimising process. However, this may be only one of the elements involved in the development of dislike since similar patterns have been found in the computerized IGT without the money exchange (Mintzer and Stitzer, 2002).

Crone, Somsen, Van Beek and Van Der Molen (2004) also report evidence of a division in selection behaviour from the four IGT decks on their computerized 'donkey' version of the task developed for children but used with adults in this study. They found an interaction between loss frequency and the net gain of the decks, such that participants selected more cards from the deck with the lower rather than the higher frequency of loss within the disadvantageous decks (deck B versus deck A), but no such division was found for the advantageous decks. However, an interaction was also found between loss frequency, net gain and performance group. Further post-hoc ANOVAs found that participants in the best performing group selected less from decks A and B, and more from deck C, with no difference in deck D selection between performance groups. Examination of Figure 2 in this study, shows that participants in the best performing group select a similar number of cards from decks B and C (~34%), at chance from deck D (~24%) and below chance from deck A (~8%). The pattern was different in the moderate and worst performing groups where above chance selection from deck B predominated (moderate group – A: ~16%, B: ~35%, C: ~24%, D: ~24%; worst group - A: ~25%, B: ~40%, C: ~16%, D: ~19%). Selection from the advantageous decks is roughly equal, but below chance in the worst group, and selection from A at chance, whereas the opposite pattern is found in the moderate group. These data suggest that differential selection from the IGT

decks may drive overall performance. However, no data were presented looking at the change in selection from individual decks across blocks. Trial block was involved in a three-way interaction with performance group and gain amount, as well as the two-way interactions involving these factors. In discussing these results the authors concluded that only the good performers improved (selected more good decks) across trial blocks. This effect could have been driven by decreasing selection from B and increasing selection from C as in the Experiments reported in Chapter 2. However, no four-way interaction between factors was revealed.

Physiological measures of heart rate and skin conductance were also recorded by Crone et al (2004) and augmented the information available from the behavioural effects. Generally, heart rate slowed more and skin conductance was increased following infrequent losses. Heart rate has been found to slow in anticipation of an aversive event (Somsen, van der Molen, & Orlebeke, 1983) and skin conductance response is a widely used measure of physiological arousal (Dawson, Shell, & Filion, 1990). Somatic activity of this sort is hypothesized to underlie decision-making on the IGT (Bechara et al, 1996; Damasio, 1996; but see Chapter 5). Crone et al (2004) performed further analyses to investigate whether heart rate and skin conductance were sensitive to the magnitude or the frequency of the losses. The physiological measures were compared following losses from decks A and D. In these decks the mean magnitude of losses is the same but the frequencies are different (losses are more frequent in deck A). No differential effects were found suggesting that the physiological measures were sensitive to the magnitude and not the frequency of losses. This result was found despite the large number of tests carried out to explore all factorial interactions and the consequent risk of Type I error inflation.

To some extent these behavioural and physiological findings fit with the behavioural results from Experiments 1 to 3 in that little change in selection from decks A and D was found, while changes in selection were apparent following experience of the high magnitude losses in deck B. However, in Experiments 1 to 3, selection from D was above chance whereas selection from A was below chance, suggesting that the frequency of loss does influence selection. The results of this study suggest that behaviourally, good performance is somewhat mediated by choices from decks B and C, much like in Experiments 1 to 3.

In a further study, Crone, Bunge, Latenstein and van der Molen (2005) investigated the performance of children and adolescents on the standard and variant versions of the 'donkey' IGT. In their experiment participants were randomly assigned within each age group (7-9 years; 10-12 years; 13-15 years) to one of three conditions: a high-complexity condition with 4 options each with punishment frequency at 50% (equivalent to two deck A's and two deck C's); a low-complexity high-punishment frequency condition (equivalent to deck A and C); and a low-complexity low-punishment frequency condition (equivalent to deck B and D). They found that the improvement in performance with increasing age reported in previous studies (e.g. Crone and van der Molen, 2004) was only apparent in the low-complexity low-punishment frequency condition where the youngest age group showed chance selection from each deck, whereas the older age groups preferred the deck with the positive expected value. Children in the middle age group were slower to develop this preference. Examining deck switching behaviour Crone et al found that the younger children switched selection more frequently in this condition. The authors argued that the younger children discounted the infrequent (large) punishment in this condition more quickly than the older age groups. These results demonstrate

that punishment frequency does influence selection behaviour, but that it should disappear into adolescence. However, even if a fourth condition with high-complexity and low-punishment frequency (BBDD) had been used it would not be as complex as the standard IGT where the frequency of punishment varies with deck types.

Peters and Slovic (2000) have criticised the design of the IGT by commenting that losses, gains and expected values are confounded as the highest magnitudes occur in the disadvantageous decks. This makes it difficult to determine if participants are selecting on the basis of size of gains, size of losses or expected value. They make no mention of the frequency of losses. Instead they created an alternative task based on the IGT. In this task, two decks have higher gains (B & D) and two have higher losses (B & C), but decks C and D have a positive expected value whereas decks A and B have a negative expected value. Unlike on the IGT gains and losses are not presented together on the same selection. In altering the IGT Peters and Slovic equalised the probability of loss in each of the decks to .5, with the exception of deck C where the probability of loss is .2 (as a result this deck also has a slightly higher expected value). While this task may allow discrimination between the magnitudes of gains, losses and expected values, arguably making it a better task than the IGT, such a manipulation does not tell us anything about the influence of loss frequency on behaviour on the IGT. However, it is of note that the average participant in the Peters and Slovic task chose deck C – the deck with the lowest frequency of loss. But this was also the deck with the highest expected value.

It was mentioned in Chapter 2 that a basic associative learning principle like the Law of Effect (Thorndike, 1898) predicts that decks with the lowest frequency of loss be preferred. This is supported by the results on Peters and Slovic's (2000)

modified task. What is not clear is why behaviour on the standard gambling task, with its arrangement of reward on every trial and infrequent losses does not conform all the time. Barron and Erev (2003, Experiment 4), have shown that in small feedback-based problems decision makers have a tendency to underweight rare outcomes. Small feedback-based problems have three features: they are repeated, each single choice is not very important overall, and there is no objective information – decision makers must rely on the immediate feedback obtained in the situation in the past. Yechiam, Stout, Busemeyer, Rock and Finn (2005a) have claimed that the choice on the IGT is an example of a small feedback-based decision situation, although they make no mention of any effects instructions have so their description is characteristic of the IGT when participants are not given the hint instructions. They investigated the influence of foregone payoffs on the selection behaviour in two manipulations of the IGT – where foregone payoffs were included or not and where reinforcer magnitude had been increased by a factor of 1.5 or not (the non-manipulated condition was equivalent to the standard task). Foregone payoffs are where the outcome of an unchosen alternative is revealed. Yechiam et al (2005a) investigated the influence of foregone payoffs on performance of a high-level drug abusing population (predominantly university students) versus controls. However, the study is of interest in the context of this Chapter because it investigated selection over 150 trials in each IGT deck. The manipulation of reinforcer magnitude produced no significant effects so results were pooled across this condition. For controls in the unmanipulated condition, selection from deck A was always below chance, and actually decreased with time. Selection from B remained above chance, and appeared to increase in the first 100 trials. Selection from C was below chance but did show some increase over all blocks, whereas in D selection increased from above chance in

the first block. One possible reason for the difference in selection patterns compared to the Experiments 1 to 3 is that reward magnitude was ten times higher – participants earned \$1 and not £0.10 per selection from decks A and B (the difference in currency is unimportant. It is the relative size within the currency that increased magnitude). Yechiam et al (2005a) found that the influence of the foregone payoffs was larger for the drug abusing group in B and D, but a different pattern of results was found – with foregone payoffs selection from B switched from learned avoidance without them to constant above chance selection with them. In A, foregone payoffs resulted in an initial attraction to this deck that reduced over time. This different results pattern lead Yechiam et al (2005a) to caution against aggregating across the disadvantageous decks and to advocate a more detailed examination of selection patterns without collapsing across the frequency difference. An examination of Figure 2 in this study reveals that the control participants have similar selection from decks A and B with and without foregone payoffs. But providing foregone payoffs actually increased selection from D and decreased selection from C. This is informative: when participants know what they missed out on they prefer the advantageous deck with the lower frequency of loss.

The reinforcement schedules on the IGT were constructed to contrast decks with positive and negative expected values, but larger or smaller immediate gains. Nowhere in the IGT literature is there any explanation for why within decks with the same expected values, the reinforcement schedules were constructed in the way they were. It can be assumed that the task was constructed with four decks as using two decks would mean that the task was too simple, whereas any more than four would increase the complexity. But it is not clear why within each deck type one deck has an infrequent, but relatively high magnitude loss, whereas the other deck has a more

frequent with consequently smaller relative magnitude losses. As detailed in Chapter 1 Bechara et al (1994) gave no indication of why they created the decks with different schedules of punishment. However, it is clear from their standard analysis and the omission of any mention of this variable from recent published descriptions (Bechara and Damasio, 2005) that they believe selection is equivalent from decks with the same expected value. This is not what was found in the experiments reported in Chapter 2, and not found in those studies that have reported selection patterns from individual decks. From these studies it is clear that there is differential selection within decks with negative expected values and possibly within the advantageous decks. The following Experiments were designed to investigate the influence of loss frequency on this choice behaviour on the IGT.

The hypothesis developed in this chapter is that deck selection on the IGT is not merely governed by the long-term consequences of the decks. The contention is that participants readily learn to avoid the disadvantageous deck with the higher frequency of loss and initially prefer the disadvantageous deck with the lower frequency of loss, while within the advantageous decks differential selection is less clear, but participants in general prefer the deck with the more frequent losses (and lower net losses). A possible explanation for this behaviour is the amount of information available to participants about the ‘goodness’ or otherwise of the decks. This information can only come from the schedule of losses, and the frequency of loss with which loss occurs. In the decks with the less frequent loss there is less information about their ‘goodness’ as there are less penalties, giving less opportunity to gather information about the long-term ‘goodness’ of the deck. Whereas in the decks with more frequent losses, there are more losses and therefore more information about the overall nature of the decks. A further aim of Experiment 5 was to test the

hypothesis that the frequency of loss provides information to participants about deck ‘goodness’ and so affects learning behaviour.

3.2 EXPERIMENT 4

ADVANTAGEOUS CARD SELECTION DIFFERS BETWEEN DECK COMPARISONS WHEN CHOICE IS BETWEEN TWO IOWA GAMBLING TASK DECKS.

3.2.1 Introduction

The results from Experiments 1-3 suggested that participants’ preferences within the advantageous and disadvantageous decks on the IGT were not uniform. The aim of the current experiment was to further investigate participants’ deck preferences and their effects on learning in the IGT by investigating behaviour when a simpler choice was offered. In four conditions, choice was examined between one advantageous and one disadvantageous deck from the IGT. This design permits the examination of learning as measured by the change in preference for the advantageous deck. As the expected values of the decks in each condition are the same, any difference in behaviour implies that the contingencies of the individual decks, and not expected values, are governing selection on this task. Any differential learning between conditions can be attributed to differences in the magnitude and frequency of losses on the individual decks. This would provide support for the observations from Experiments 1 – 3 that deck selection varies within disadvantageous and advantageous decks. Table 3.1 displays the deck contingencies in each condition.

Following the results of the previous study where the change in selection from decks B and C were hypothesized to underlie learning, it was anticipated that learning would be slowest when the choice is between these two decks. Due to the general avoidance of deck A in the earlier studies it was predicted that the fastest learning would be seen when deck A was one of the choice options. Similarly, due to the preference for deck B it was predicted that learning would be slowest in the conditions where it was one of the response options. Differences in learning between conditions where the disadvantageous decks are different and the advantageous decks are the same would suggest that selection from the disadvantageous decks is not uniform and would provide behavioural support for Crone et al's (2004) results. Differences in learning between conditions where the advantageous decks are different but the disadvantageous decks are the same would suggest that selection from the advantageous decks is not uniform. This would conflict with the results from Crone et al (2004) but provide support for the hypothesis that change in preference for deck C underlies learning on the IGT.

3.2.2 Method

Participants

Forty-eight participants (thirty-nine female) were recruited from the undergraduate and postgraduate populations at the University of Nottingham. Participants were recruited through a poster advertisement that offered the opportunity to earn up to £6 by taking part in a cognitive psychology experiment. Participants were randomly assigned to one of four conditions (see Table 3.1).

Table 3.1: The deck contingencies (reward magnitude and mean loss magnitude in pence, and loss frequency) in each experimental condition.

	Deck Comparison			
	<u>A:C</u>	<u>A:D</u>	<u>B:C</u>	<u>B:D</u>
Reward Magnitude (pence)	10:5	10:5	10:5	10:5
Mean Loss Magnitude (pence)	25:2.5	25:25	125:2.5	125:25
Loss Frequency (p [loss])	0.5:0.5	0.5:0.1	0.1:0.5	0.1:0.1

Apparatus

Participants were tested individually in a testing laboratory. A PC controlled the experiment. A 2-alternative forced choice task was created and run on the PC. The task was based on published descriptions of the Iowa Gambling Task, except that participants made choices from two rather than four decks of cards. The reinforcement schedules for each deck were the same as those published by Bechara, et al (1994) for the first 40 cards. For the remaining 160 cards in each deck the reinforcement schedules were based on the format of the first 40 cards: Deck A, five losses totalling £1.25 per ten card selections; Deck B, one loss of £1.25 per ten card selections; Deck C, five losses totalling £0.25 per ten card selections; Deck D, one loss totalling £0.25 per ten card selections. The task format was exactly the same as that reported in the preceding experiments except that choices were made from two decks and not four. All reinforcers were real money.

Design

A between-subjects design was used to compare participants' learning between four advantageous deck to disadvantageous deck comparisons. Learning when choosing between decks A and C, decks A and D, decks B and C and decks B and D was compared. The number of selections made from the advantageous decks was recorded for each of ten twenty-trial blocks. From this measure the slope, b , was calculated as an estimate of learning rate.

Procedure

Participants were randomly assigned to one of the four deck comparison conditions. The procedure was the same as in Experiment 3 except that participants took part in only one session of 200 trials and they saw on-screen, and chose between, only two decks of cards rather than four. After 100 card selections participants were invited to take a short break. The length of this break was determined by each participant and was not recorded. As there were only two choices no hint was provided.

3.2.3 Results

The number of card selections from the advantageous deck (C or D) in each condition was recorded for each participant in each of ten twenty-trial blocks. Table 3.2 displays the mean number of advantageous selections in each experimental group over the first half, the second half and the whole experiment.

To investigate whether the number of advantageous selections differed between groups, a 4x10 (Condition by Block) mixed design ANOVA was performed.

There was no interaction, $F(11.42, 167.43) < 1$, $MSE = 24.88$, $p > 0.05$. However, ANOVA revealed a main effect of Block, $F(3.81, 305.53) = 12.82$, $p < 0.001$, indicating that the number of advantageous selections differed between blocks (it increased with block). A main effect of Condition was also found, $F(3, 44) = 3.87$, $MSE = 139.33$, $p < 0.05$. Pairwise comparisons found that the number of advantageous selections was significantly greater in condition A:C versus condition B:C, $F(3,440) = 6.06$, $MSE = 23.4$, $p < 0.05$; and in condition A:D versus condition B:C, $F(3,440) = 4.70$, $MSE = 23.4$, $p < 0.05$. There were no significant differences between any other groups, $F(3,440) < 1$.

Table 3.2: Mean number of advantageous selections in each experimental group.

Trials	Deck comparison			
	A:C	A:D	B:C	B:D
1 - 100	12.75 (0.80)	12.52 (0.76)	7.43 (0.93)	11.32 (1.09)
101 - 200	15.3 (1.24)	14.73 (1.01)	11.25 (1.71)	13.38 (1.69)
1 - 200	14.03 (0.86)	13.63 (0.83)	9.34 (1.16)	12.35 (1.37)

Note: The maximum number of advantageous selections possible is 20. Figures in parentheses are the standard error of the mean.

Figure 3.1 shows the mean number of advantageous card selections in each condition across ten twenty-trial blocks. While advantageous selection appears to increase at roughly the same rate in each condition, it is always lower when the choice is between decks B and C. Learning would be indicated by an increase in the number

of advantageous selections with increased exposure to the task. As a measure of learning rate the slope, b , was calculated for each participant. Table 3.3 gives the mean learning rates for the first and second 100 trials and for all trials in each experimental condition and shows that overall learning rate is greatest when the choice is between deck C and a disadvantageous deck.




Figure 3.1: Mean number of advantageous selections across twenty ten-trial blocks in each experimental group. Error bars are the standard error of the mean.

A one-way ANOVA was run to investigate whether learning rate differed significantly between deck comparison conditions. No significant differences in learning rate between conditions were found, $F(3, 44) < 1$, $MSE = 0.44$, $p < 0.05$.

Table 3.3: Mean learning rate (*b*) in each experimental condition across first 100 trials, second 100 trials and all trials.

Trials	Deck comparison			
	A:C	A:D	B:C	B:D
1 - 100	1.04* (0.37)	0.54 (0.31)	-0.01 (0.51)	0.75 (0.40)
101 - 200	0.54* (0.24)	0.38 (0.24)	0.93* (0.39)	0.46* (0.20)
1 - 200	0.58* (0.21)	0.45* (0.13)	0.69* (0.25)	0.47* (0.17)

*Note: Figures in parentheses are the standard error of the mean. *indicates significantly greater than 0 at the .05 level using Lorch and Myers regression analyses.*

In previous experiments using the IGT selections have most commonly been examined in one hundred-trial sessions. Examining the data in a similar way reveals that there appears to be no learning in the first 100 trials in condition B:C (see Table 3.3). This result is in line with the experimental hypotheses. However, in the second 100 trials participants learn at almost twice the rate of participants in the other conditions. This mirrors the results from the previous experiments where a preference for deck C is found to develop with increased experience of the decks. Table 3.3 displays a summary of the results of Lorch and Myers (1990) regression analyses on learning rates. These regressions revealed that while learning rate was greater than zero across all 200 trials in all conditions, in the first 100 trials it was almost flat in condition B:C and was no different from zero in conditions A:D and B:D. These results were augmented by comparing mean advantageous selections made in block 5 to zero. For conditions A:D and A:C, the number of advantageous selections were

greater than zero in block 5 (A:D: $t(11) = 3.31$, $SD = 4.27$, $p < 0.05$; A:C: $t(11) = 3.83$, $SD = 4.75$, $p < 0.05$) but not in condition B:D, ($t(11) = 1.23$, $SD = 6.09$, $p < 0.05$) or condition B:C ($t(11) = -1.74$, $SD = 5.82$, $p < 0.05$). In the final block of trials advantageous selections were greater than zero in all conditions (A:D: $t(11) = 4.57$, $SD = 4.23$, $p < 0.05$; B:D: $t(11) = 2.46$, $SD = 6.11$, $p < 0.05$; A:C: $t(11) = 5.07$, $SD = 4.44$, $p < 0.05$) except condition B:C ($t(11) = 1.33$, $SD = 5.88$, $p < 0.05$).

3.2.4 Discussion

A difference was found between the number of advantageous cards selected in condition A:C and condition B:C. Fewer advantageous selections were made in condition B:C suggesting that participants found it harder to select advantageously due to a preference for deck B. This result provides support for the hypothesis that participants' preferences on the IGT distinguished between the disadvantageous decks. However, no significant differences were found in the number of advantageous selections in conditions A:D and B:D suggesting that the presence of deck C is important.

No differences were found when learning rates between conditions were examined, and learning rate was significantly greater than 0 over all 10 blocks in each condition. However, examination of Figure 3.1 suggested that condition B:C does differ from the others. Learning rate during the first hundred trial in condition B:C was flat so that by block 5 participants were not selecting advantageously. Despite a huge increase in learning rate as measured over 200 trials, in the final block participants were still not showing a preference for the advantageous deck although they were heading in that direction. These results support the experimental

hypothesis that due to changes in participants' preferences observed in the IGT this condition would be the hardest for participants to learn on. However, an alternative explanation for these findings is that identifying which option had the better long-term consequences was harder because the deck contingencies varied on both the magnitude and the frequency of loss whereas all other deck contingencies varied on only one (see Table 3.3; A:C and B:D vary on loss magnitude; A:D varies on loss frequency).

The mean number of advantageous selections was greatest when the deck comparison involved deck A. In all blocks the highest advantageous selections were seen in these two conditions. These results provide support for the hypothesis that participants generally avoid deck A, meaning that identification of the deck with better long-term consequences is easier regardless of the advantageous deck it was paired with. Unfortunately because there were no differences in the number of advantageous selections (and consequently disadvantageous deck selections) little can be inferred about the relative contribution of loss frequency or magnitude in the advantageous decks, except that they do not appear to differentially affect learning when the disadvantageous deck is A.

In condition B:D, where only the magnitude of loss differed between decks, advantageous selection did not increase in the first 100 trials and no preference for deck D had been established by block 5. Only after more exposure did a preference for deck D develop. This reflects an effect of loss magnitude. Participants learn to avoid the larger loss despite the larger gain associated with it. Table 3.3 shows that the lowest overall learning rates are found when deck D is one of the choices. However, this may reflect different processes in conditions B:D and A:D. In condition A:D little change in learning rate reflects the general and unchanging

preference for deck D. The later development of a preference for D in condition B:D may reflect the similarity in loss frequency between the decks making it harder to identify the deck selections which will result in greater long-term gains.

Overman (2004) has also identified the differences in deck contingencies as an important factor in behaviour on the IGT. As well as noting the differences in loss frequencies between the decks, Overman (2004), like Peters and Slovic (2000), has pointed out that due to the size and frequency of losses on deck C very often no overall loss is made when selecting from this deck. This may be the reason that selection from this deck changes across block: participants learn that while there are frequent losses on this deck, they rarely result in a net loss for that selection, and even when a net loss occurs it is small in comparison to all other decks. As such, deck C varies in a unique respect from the other decks on the IGT: *net loss* frequency. This net loss is lower (it varies between 1 and 4 per 10 card selections) than the frequency of losses (5 per 10 card selections). It may take the participants time to learn this, but it would mirror the apparent initial preference for the decks with infrequent losses (decks B and D) and fit into an explanation utilising the Law of Effect. Of course, the infrequent net losses also means that the magnitude of losses on Deck C are substantially lower than on the other IGT decks and it may be this that influences preference for this deck.

The results from this experiment provide some support for previous findings that participants' preferences for the disadvantageous decks are not uniform. In a two-choice environment it appears to be harder to select advantageously when deck B is one option. The results also support the hypothesis that learning on the IGT is driven by changes in selection from decks B and C. However, the experiment did not inform on whether the frequency of loss or its magnitude affect card selection on the

IGT decks, except to suggest that this relationship may be different in the advantageous and disadvantageous decks. Experiment 5 explores this relationship further.

3.3 EXPERIMENT 5

MANIPULATION OF FREQUENCY (AND MAGNITUDE) OF LOSS AFFECTS LEARNING ON THE IGT.

3.3.1 Introduction

The results from Experiment 4 provided some support for the hypothesis that the difference in deck contingencies within the disadvantageous and advantageous decks contributes to learning on the IGT. Experiment 5 was devised in order to further test this hypothesis by manipulating the reinforcement contingencies in the decks. The hypothesis tested in this chapter is that participants avoid the disadvantageous deck with the higher frequency of loss and initially prefer the disadvantageous deck with the lower frequency of loss, while within the advantageous decks differential selection is less clear, but participants in general prefer the deck with the more frequent losses (and lower net losses). A possible explanation for this behaviour is the amount of information available to participants about the ‘goodness’ or otherwise of the decks. This information can only come from the schedule of losses, and the frequency with which losses occur. In the decks with the less frequent loss there is less information about their ‘goodness’ as there are less penalties, giving fewer opportunities to gather information about the long-term ‘goodness’ of the deck. Whereas in the decks with more frequent losses, there are more losses and therefore more information about the

overall nature of the decks. Experiment 5 was designed to test the hypothesis that the frequency of loss provides information to participants about deck ‘goodness’ and so affects learning behaviour. To this end, two conditions were created by manipulating the frequency of losses on the original IGT decks. In a Decreased Frequency condition the frequency of loss in decks A and C was reduced but unchanged in decks B and D. In an Increased Frequency condition the frequency of loss in decks B and D was increased but unchanged in decks A and C. In the Decreased Frequency condition fewer losses occur across the whole task, giving participants less information about the nature of the decks, whereas in the Increased Frequency condition more information is available. Any difference in learning and in learning rate can be attributed to this difference in frequency of loss. It was predicted that if the frequency of loss is informative then a slower learning rate would be observed in the Decreased Frequency condition. A further prediction is that since participants prefer the lower loss frequency decks, selection should be higher from those decks within decks with the same expected values. A sign that loss frequency is in part informative would be that this difference is greater when loss frequency is increased rather than decreased, as participants will still be avoiding the deck with largest loss frequencies.

3.3.2 Method

Participants

Forty-two (twenty-six female) participants were recruited from the undergraduate and postgraduate populations at the University of Nottingham. Participants were recruited through a poster advertisement that offered the opportunity to earn up to £6 by taking

part in a cognitive psychology experiment. Data from two participants were excluded from the analysis due to an experiment administration error in one case, and an expression by the participant of total misunderstanding of the instructions in the other.

Apparatus

Two modified versions of the Iowa Gambling Task were created. In the Increased Frequency condition the frequency of loss was increased in the two IGT decks with low frequency (but high relative magnitude) losses (decks B and D). In the Decreased Frequency condition the frequency of loss was decreased in the IGT decks with high frequency (but lower relative magnitude) losses (decks A and C). In the original IGT the schedule of losses was fixed. This was maintained in the modified decks and the occurrence of losses was randomly determined within 10 card blocks for each deck with the caveat that the sum of losses did not change the expected value for that deck within a ten card block. Where the magnitude of losses changed (e.g. increased with the reduction in loss frequency in modified deck A and decreased with the increase in loss frequency in modified deck B) the same amounts were used in decks with the same expected value e.g. for the disadvantageous decks 5 losses of 15, 20, 25, 30, and 35 or one loss of 125 became three losses of 35, 35, and 55. The manipulated schedule of losses used in this Experiment is displayed in Appendix B. In both modified IGT versions the unmodified decks remained as they were in Experiment 3 (see Appendix A).

Design

A between-subjects design was used to compare participants' learning on the two modified IGT versions. The number of selections made from the advantageous decks minus the number of selections from the disadvantageous decks was calculated for each of ten twenty-trial blocks. From this measure the slope, b , was calculated as an estimate of learning rate. In addition, the number of cards chosen from each of the decks and the change in their selection over time was examined.

Procedure

Participants were randomly assigned to the Increased Frequency or Decreased Frequency conditions. The procedure followed that of Experiment 4. Participants took part in only one session of 200 trials and they saw on-screen, and chose between, four decks of cards. After 100 card selections participants were invited to take a short break. The length of this break was determined by each participant and was not recorded.

3.3.3 Results

Net score was calculated for each participant over the whole experiment, and for the first hundred and the second hundred trials. Mean net scores in each condition are displayed in Table 3.4. As the table shows mean net score does not differ much between groups, although contrary to the experimental hypotheses participants in the increased frequency group have a lower net score in the first 100 trials. However, an independent samples t -test found no significant difference in the overall mean net score between conditions, $t(38) < 1.0$.

Figure 3.2 displays mean net score in each of ten blocks of twenty trials for each experimental condition. Mean net score increases across blocks in both conditions, although only in the Increased Frequency condition does mean net score end above chance. However, contrary to the experimental hypothesis, mean net scores do not appear to differ between conditions. This is confirmed by the results of a 2 x 10 (Loss Frequency by Block) mixed design ANOVA. There was no main effect of Loss Frequency, $F(1, 38) < 1.0$, $MSE = 536.61$, $p > 0.05$, nor a significant interaction, $F(5.26, 199.72) = 1.57$, $MSE = 86.13$, $p > 0.05$. A significant main effect of Block was found, $F(5.26, 199.72) = 6.04$, $MSE = 86.13$, $p < 0.01$, which indicated the tendency for mean net score to increase across blocks.

Table 3.4: Mean net score in each condition in the first and second hundred trials, and over the whole experiment.

Trials	Decreased frequency	Increased frequency
1 - 100	9.5 (8.40)	3.5 (6.25)
101 - 200	20.3 (10.27) ^{††}	23.9 (10.88) [*]
1 - 200	29.8 (16.47) [†]	27.4 (16.29) [‡]

*Note: Figures in parentheses are the standard error of the mean. *indicates significantly greater than 0 at the .05 level. ‡ indicates $p = .109$. † indicates $p = 0.086$. †† indicates $p = 0.083$.*

The main effect of Block does not provide much information beyond showing that mean net score is higher in some blocks than in others. As a result the change in

mean net score across block, or the slope b , was calculated as an estimate of learning rate in each condition. Over the entire experiment learning rate was greater in the Increased Frequency condition, $b = 0.97$ ($\sigma_M = 0.26$), than in the Decreased Frequency condition, $b = 0.50$ ($\sigma_M = 0.29$). However, an independent samples t -test found that this difference was not significant, $t(38) = -1.21, p > 0.05$. This result suggests that, as with the result of the mixed-design ANOVA, there is no strong evidence to support the experimental hypothesis that increasing the frequency of loss in the low frequency decks will lead to faster learning.




Figure 3.2: Mean net score across ten twenty-trial blocks in each experimental condition. Error bars are the standard error of the mean (only negative bars are displayed for ease of viewing).

Table 3.5: Mean learning rate (b) in both experimental conditions across the first 100 trials, second 100 trials and all trials.

Trials	Decreased frequency	Increased frequency
1 - 100	2.04 (0.61)*	1.46 (0.55)*
101 - 200	-0.59 (0.56)	1.47 (0.52)*
1 - 200	0.50 (0.29)	0.97 (0.26)*

*Note: Figures in parentheses are the standard error of the mean. *indicates significantly greater than 0 at the .05 level using Lorch and Myers regression analyses.*

In Experiment 4, it was argued that as the structure of the task included a break after 100 trials and as this is the length of the standard administration of the gambling task, learning rate might be examined over the first 100 and second 100 trials. Table 3.5 displays mean learning rate over the first and second 100 trials, and over the entire experiment, in each condition. Lorch and Myers (1990) regression analyses for repeated measures designs compared these learning rates to zero. These analyses reveal that while there is a significant increase in learning rate in the first 100 trials in both conditions (Decreased Frequency: $b = 2.04$ ($\sigma_M = .61$), $t(19) = 3.35$, $p < 0.01$; Increased Frequency: $b = 1.46$ ($\sigma_M = .55$), $t(19) = 2.67$, $p < 0.02$), learning rate only continues to increase in the second 100 trials in the Increased Frequency condition, $b = 1.47$ ($\sigma_M = .52$), $t(19) = 2.82$, $p < 0.02$). Indeed, learning rate in the second 100 trials in the Decreased Frequency condition is negative, $b = -.59$ ($\sigma_M = .56$), $t(19) = -1.05$, $p > 0.05$. An independent-samples t -test found this difference to be significant, $t(38) = -2.69$, $p < 0.02$ (the same test for the first 100 trials was not

significant, $t(38) < 1.0$). This difference and the negative learning rate in the Decreased Frequency condition reflects the decline in mean net score in blocks 9 and 10 in this condition, as illustrated in Figure 3.2. This decline after 160 trials is what affects the overall learning rate in this condition, and is the reason why it is not significantly greater than zero. Whereas, in the Increased Frequency condition mean net score continues to increase until by the end of the final block, mean net score is significantly greater than chance, $t(19) = 3.70, p < 0.01$. These results give some support to the experimental prediction that learning would be greater in the Increased Frequency condition.

One reason why selection from C and D was reduced in the last two blocks in the Decreased Frequency condition may be that these trials went unpunished and participants were attempting to earn as much money as possible. This may have been an unwitting outcome of the break at halfway. However, as there was no counter for participants to keep track of the number of selections they would have had to rely on an estimate of time elapsed relative to the first one hundred trials. One demonstration that this strategy was in effect would be greater selection from deck B. However, examination of individual deck selections revealed that while seven participants had a net score less than 0 in block 9, only five showed a clear preference for the disadvantageous decks (a net score < -10). Their relative preference for each deck can be measured by subtracting selections from A from selection from B. Of these five participants, two preferred deck A (B-A score of -16 and -20), two preferred deck B (B-A score of 6 and 20) and one was indifferent between them (B-A score of 0). In block 10, nine participants in the Decreased Frequency group had a negative net score but only two had a net score less than -10 (one preferred A, B-A score of -4; one preferred B, B-A score of 4). The absence of a uniform pattern suggests that not

all these participants were following a strategy based on knowledge of the task contingencies and the time remaining, as their selections encountered sufficient losses to outweigh any short-term profit.

Individual deck selection

In Experiments 1 to 3, participants preferred to select the disadvantageous cards with the least frequent loss, whereas within the advantageous decks, although this general pattern is common early in learning, later in learning participants preferred to select from the decks with the more frequent loss. Since participants received the hint in both conditions in this Experiment a similar pattern of results should be seen, and if this were the case, it would provide additional evidence that the frequency (and magnitude) of loss does affect deck selection.

The relevant data are shown in Figure 3.3. It is clear that participants show a similar deck selection preference in both conditions. This is not surprising given the similarity in net score measures reported earlier. Unlike in previous Experiments there does not appear to be any difference in selection within the advantageous decks. However, within the disadvantageous decks participants still appear to prefer the deck with the infrequent loss. Figure 3.4 displays deck selection in the first 100 and second 100 trials in both conditions. Like the behaviour of participants who received the hint instructions in Chapter 2, deck selection from deck B decreases from the first to the second 100 trials, but unlike those earlier conditions the change in selection from deck D (increases from the first to the second 100 trials), is as large as that found in deck C. This suggests an equivalence in preference within the advantageous decks; a trend not apparent in the disadvantageous decks where deck A is always selected at a level





Figure 3.3: Mean number of cards selected from each deck across all 200 trials in each condition. Error bars are the standard error of the mean.

below chance. This implies that within the disadvantageous decks participants prefer the deck with the less frequent losses. Separate 2x2 (Deck by Time) repeated measures ANOVAs for each condition were run to investigate this claim. For the Decreased Frequency condition there was no main effect of Deck, $F(1, 19) = 3.0$, $MSE = 163.26$, $p = .1$; no main effect of Time, $F(1, 19) = 1.15$, $MSE = 104.47$, $p > 0.05$; nor was there an interaction, $F(1, 19) = 2.43$, $MSE = 29.78$, $p = .14$. For the Increased Frequency condition there was a main effect of Deck, $F(1, 19) = 26.64$, $MSE = 19.89$, $p < 0.01$; a main effect of Time, $F(1, 19) = 8.49$, $MSE = 55.40$, $p < 0.01$; but no interaction, $F(1, 19) = 1.68$, $MSE = 26.74$, $p > 0.05$. In the Increased Frequency condition the selections from B were greater than from A, and the number

of selections from these disadvantageous decks was greater in the first hundred than the second hundred trials. That this was not the case in the Decreased Frequency condition implies that participants in this condition did not discriminate between the two bad decks, whereas in the Increased Frequency condition they did.

Figure 3.5 shows the change in selections from each deck across trial blocks. In Figure 3.5a the change in selections from the disadvantageous decks follow the same trend in both conditions. Selection from deck B begins well above chance in the first block, but by block 5 selection from B is below chance. Selection from A remains below chance in both conditions. However, in the second 100 trials differences emerge between the conditions. Selection from both A and B increases in the Decreased Frequency condition, whereas they continue to decline in the Increased Frequency condition (although this is partly due to the increase in selection between blocks 5 and 6). This difference in selection between conditions would appear to be what underlies the difference in learning between these conditions. Figure 3.5b mirrors Figure 3.5a; selection in C and D increases in both conditions in the first hundred trials, but in the Decreased frequency condition selection from both declines in block 8 for deck D and block 9 for deck C. In the Increased Frequency condition selection from both decks continues to increase (although there is a dip in block 9 from D), with selection from both ending above chance.

A: Decreased Frequency



B: Increased Frequency





Figure 3.4: Mean number of cards selected from each deck across the first 100 and second 100 trials in the: A) the Decreased Frequency condition and B) the Increased Frequency condition. Error bars are the standard error of the mean.

A: Increased Frequency



B: Decreased Frequency





Figure 3.5a: Mean number of cards selected from the disadvantageous decks in each condition. Error bars are the standard error of the mean.

A: Increased Frequency



B: Decreased Frequency




Figure 3.5b: Mean number of cards selected from the advantageous decks in each condition. Error bars are the standard error of the mean.

Table 3.6 presents the change in selection (the slope, b ,) from each deck in the first and second halves of the task. The observations from Figure 3.5 are borne out. In the Decreased Frequency condition selection from the disadvantageous decks increases and selections from the advantageous decks decreases in the second hundred trials – the opposite to what happens in the Increased Frequency condition. These differences in selection support the experimental hypothesis. With a reduction in the overall frequency of losses (and a consequent increase in magnitude of loss), participants end the task with less differentiation between decks.

Table 3.6: Mean selection rate from each deck in the first and second 100 trials in each condition.

	Trial	Decreased Frequency	Increased Frequency
A			
	1 – 100	-.18 (.20)	-.28 (.17)
	101 – 200	.18 (.28)	-.28 (.16)
B			
	1 – 100	-.84 (.27)	-.53 (.18)
	101 – 200	.12 (.22)	-.46 (.17)
C			
	1 – 100	.53 (.34)	.61 (.33)
	101 – 200	-.05 (.23)	.47 (.34)
D			
	1 – 100	.49 (.28)	.20 (.27)
	101 – 200	-.25 (.17)	.27 (.40)

Note: Figures in parentheses are the standard error of the mean.




Figure 3.6: Mean number of card selections across block from the disadvantageous decks with probability of loss of 0.3. Error bars are the standard error of the mean.

It was noted that within the disadvantageous decks participants prefer the decks with the less frequent losses. This is apparent in Figure 3.5a. Figure 3.6 displays selections from deck A in the Decreased Frequency condition and selections from deck B in the Increased Frequency condition. These decks have the same probability of a loss, .3, but what differs between them is the context in which they are presented. In the Increased Frequency condition this deck is preferred initially as it has the lowest probability of loss, although selection continues to decline across block, whereas in the Decreased Frequency condition this deck has the greatest frequency of loss and is selected below chance right up until the last two blocks. Figure 3.6 illustrates that participants initially prefer the disadvantageous deck with

the lowest frequency of loss. In the Decreased Frequency condition this is not the case at the end of the task – perhaps because it is more difficult to avoid a deck that provides higher magnitude gains on seven out of ten trials.

3.3.4 Discussion

Frequency of loss was manipulated in order to test the hypothesis that learning would be affected by the amount of information participants received about how good or bad their choices were. This hypothesis received some support. Although there was no difference in learning rate across all two hundred trials, in the second half of the task, learning rate was only greater than zero in the Increased Frequency condition. In the Decreased Frequency condition it was negative.

A more detailed examination of selection from the individual decks, revealed that although selection was similar between conditions in the first one hundred trials, in the second one hundred trials participants in the Decreased Frequency condition increased their selection from the disadvantageous decks, whereas in the Increased Frequency condition selection from these decks continued to decline. At the end of the task participants in the Decreased Frequency condition were not selecting from any deck above or below chance at the end of the experiment, unlike in the Increased Frequency condition where there was preferential selection from the advantageous decks. This result supports the experimental hypothesis.

It was hypothesized from the results of Experiment 4 that changes in selection from decks B and C drives learning on the IGT. There was no sign of differential selection within the advantageous decks in either experimental condition. However, participants do appear to prefer, at least initially, the disadvantageous deck with the

lower frequency of loss. That this preference is associated with frequency of loss and not magnitude of loss was demonstrated in Figure 3.6 where selection from deck A in the Decreased Frequency condition was compared to deck B in the Increased Frequency condition. Selection patterns from these decks with the same frequency and magnitude of loss differed between groups. Participants in the Increased Frequency condition initially selected more from their deck B, while in the Decreased Frequency condition selection from their deck A remained below chance until the last two blocks. The only difference between these decks was the context in which they were presented. In the Increased Frequency condition, deck B still had the lower frequency of loss relative to the other disadvantageous deck. A possible explanation for this difference, and the higher learning of this group, is that participants in the Increased Frequency condition encountered more losses earlier than participants in the Decreased Frequency group. In the Decreased Frequency group, selections from deck B could go on unpunished for longer than selections from any other disadvantageous deck. This is because in the original task's fixed schedule of losses, the large infrequent loss in deck B occurs after nine selections of a large magnitude gain (in comparison to decks C and D). In the random order of this experiment the schedules of loss were not the same for decreased loss frequency deck A and increased loss frequency deck B. But the first loss was earlier in the modified deck A than in the modified deck B. This suggests that the number of unpunished selections before a loss in the disadvantageous decks impacts on participants' learning.

In conclusion, there was strong evidence that frequency of loss affects learning. Participants in the Increased Frequency condition were preferentially selecting from the advantageous decks by the end of the Experiment while those in the Decreased Frequency condition were not. There was evidence that participants

avoid the disadvantageous deck with the more frequent losses. However, this result may have been confounded by the fixed order of losses within decks.

3.4 GENERAL DISCUSSION

The results from Experiment 4 supported the hypothesis that differential selection within the disadvantageous and advantageous decks was driving learning on the IGT. Participants learned to select from the advantageous deck more slowly when their choice was between deck C and deck B. One possible reason for this was that participants found it easier to identify deck A as one of the worst decks because the frequency of loss was high, whereas in the advantageous decks deck C appeared better because when a frequent loss occurred it was often not a net loss.

Experiment 5 tested the hypothesis that the frequency of loss was influencing learning on the IGT in two conditions. In the Decreased Frequency condition, the identification of deck A as a bad deck and deck C as a good deck was made more difficult by reducing the frequency of loss from .5 to .3, while leaving the other decks unchanged. In the Increased Frequency condition, decks A and C were unaltered relative to the original schedules and the frequency of loss in decks B and D was increased from .1 to .3. Although there were no significant differences between overall learning rates, only learning rate in the Increased Frequency condition was significantly greater than zero, supporting the experimental hypothesis. In the Decreased Frequency condition there were no significant differences in selections from the bad decks suggesting that participants did not select preferentially from between these decks, whereas they did in the Increased Frequency condition. This non-differential selection also appeared to affect participants' selection in the

Decreased Frequency condition in the few blocks where selection from the disadvantageous decks increased, implying that these participants, on average, had not learned that decks A and B were disadvantageous. However, another possibility is that they had learned that losses were infrequent and thought they could exploit it. This issue of participants' knowledge will be returned to in Chapter 5.

That the frequency of punishment influences deck selection, at least in the disadvantageous decks, was further illustrated when selection from decks with the same frequency of loss were compared between the different conditions. There was a clear decrease in selection from deck B, but little change across blocks for deck A (until the last two blocks). The key difference between the decks was that deck B in the Increased Frequency condition had the lower frequency of loss relative to the other disadvantageous deck, whereas deck A in the Increased Frequency condition had a higher frequency of loss. A similar comparison within the altered advantageous decks did not find any differences suggesting that the manipulations to these decks made it more likely that participants would not distinguish between these decks, but gradually increase selections from them.

The issue of what exactly participants are responding to on the IGT is an important one. As Yechiam et al (2005a) found, differential selection may offer insights into what is affecting selection behaviour on the IGT. Recently, Bechara et al (2000) have described a modification of the task (the A'B'C'D' version) where the frequency of loss and the magnitude of losses and gains is altered in successive blocks of ten choices from each deck. The manipulations make the differences in expected value between the disadvantageous and advantageous decks greater. In deck A' the frequency of loss is increased 10%, but the magnitude of loss remains the same. In deck B' it is the magnitude of loss that increases every ten cards while the frequency

of loss is unchanged. The same pattern is followed in decks C' and D' except that the frequency of loss is reduced in deck C' and the magnitude of loss is reduced in deck D'. These changes would appear to make the task easier in that differentiation between what the worst decks are should be clearer. This is certainly so for deck A' where frequency of loss increases. However, patients with VMpfc damage still perform below the level of healthy controls who, if anything, asymptote at a lower level of advantageous deck selections compared to the original task. Performance was analysed using the standard net score measure and no mention was made of any differential selection within the advantageous and disadvantageous decks. However, given the results of Experiments 4 and 5, closer examination of individual deck selection will reveal more information about what is influencing selection.

Yechiam et al (2005a) found that high-level drug abusers showed differential deck selection behaviour. Bechara et al (2002) have reported that both their substance abusing participants and healthy controls could be split depending on their performance on the A'B'C'D' task. However, they do not report any individual deck selection patterns and the possibility has not been ruled out that differences exist between the groups on these measures. Bechara et al (2000) created a similar, but more complex, manipulation to their variant IGT, in that as well as gains being altered (equivalent to the A'B'C'D' task changes), gains were also increased or decreased in the advantageous and disadvantageous decks respectively. Performance on this task (E'F'G'H') allowed Bechara et al (2002) to further divide their substance abusing population into those who were not impaired on either task, a subgroup who were impaired on this task and on A'B'C'D', and those who were normal on E'F'G'H' but had large physiological responses to reward. They concluded that some substance abusers were hypersensitive to reward while others were myopic for future

consequences. However, if there was differential selection behaviour between decks then these conclusions may be extended and even supported. For example increased selection from deck B' over deck A', coupled with preference for G' over E' would support a conclusion of hypersensitivity to reward, whereas no differences in A' and B' selection and E' and G' selection would support their conclusion that participants are myopic for the future.

As mentioned in the introduction, one of the problems with manipulations of the contingencies on the IGT is that gains, loss frequencies and magnitudes and expected values are all confounded with each other. This interdependency of reinforcement magnitude, reinforcement frequency and expected value makes identifying the differential effects of each difficult. Peters and Slovic (2000) successfully removed the confound that the largest magnitudes of reinforcement were in the disadvantageous decks, but their manipulation affected loss frequencies across the decks. In support of an explanation of the importance of the frequency of loss selection was highest from the deck with the lowest losses. They also found individual differences in performance in that selection of decks with high gains correlated with extraversion whereas selection of decks with low losses correlated with high scores on the behavioral inhibition system scale (BIS; Gray, 1970). And as neither of their measures correlated with selection from the decks with the highest expected values the suggestion is that this is not the most important factor in determining deck selection. However, the participants who completed the task in this study were the forty with the most extreme scores on the each measure (less than half the total who completed the initial questionnaires, meaning that these participants were a somewhat unrepresentative sample of the normal population even if their

performance suggests individual differences are important (again this point will be returned to in Chapter 5 with regard to participants' knowledge of the task).

One of the reasons that deck B might be “preferred” so much more than deck A is that with the schedules of reinforcement in the original study, selection from deck B goes unpunished for eight consecutive card selections, whereas the first loss in deck A occurs on the third selection. The same pattern is true in the advantageous decks although the first net loss on deck C occurs later than in deck A. Thus, participants may develop a justifiable preference for the disadvantageous deck with the infrequent loss that is more difficult to overcome because of the unlikelihood of a loss from this deck, although when it comes it is massive. The order of losses in this deck has gone at least some way to explaining the performance of VMpfc patients. Fellows and Farah (2005) found patient performance as measured by net score was no different from controls when the order of losses was altered so that the first loss in deck B occurred earlier. The frequency of loss may contribute to preference for this deck but the schedule of losses is also important. The use of fixed reinforcement schedules has been criticised by Dunn et al (2006) in their recent review of the clinical use of the IGT. An implication of the order in which participants encounter losses is that the disadvantageous decks are actually the best decks up until the point the accumulated losses are greater than the accumulated gains (Maia and McClelland, 2004). Seen in this way selection from the deck with the infrequent loss is reasonable if the first loss occurs relatively late in that deck. If this were the case then it would also account for the differential selection between decks A and C. This issue will be explored further in Chapter 4.

THE NATURE OF THE REINFORCEMENT SCHEDULE
INFLUENCES LEARNING

4.1 INTRODUCTION & OVERVIEW

The behaviour of healthy control participants on the IGT has typically seen rapid learning of the correct strategy that asymptotes at a level below exclusively optimal behaviour. The results of the Experiments reported in Chapter 2 suggested that this is, in part, due to the instructions participants receive. This behaviour is interesting because in arguably simpler choice environments, learning is slower and often does not result in optimal behaviour (Herrnstein et al, 1993; Tunney and Shanks, 2002).

As an example consider a binary choice environment in which selection of option *A* is reinforced 30% of the time and option *B* is reinforced 70% of the time. This is a common scenario in probability learning tasks where the participants' task is to make as many correct (i.e. reinforced) choices as possible. The optimal strategy is, after a period of sampling, to select uniformly from the option with the high probability of reinforcement. Such behaviour in the example environment would result in reinforcement with a probability of 0.7. However, in the early days of experiments using such a task it was commonly observed that participants' learning would reach asymptote with the proportion of choices matching the probability of reinforcement from each option (e.g. Neumark and Shuford, 1959). This phenomenon is known as probability matching and is sub-optimal behaviour because the probability of reinforcement is $(0.3 * 0.3 = 0.09) + (0.7 * 0.7 = 0.49) = 0.56$,

substantially less than that obtained from exclusive choice of the high probability reinforcement (0.7). Probability matching behaviour presented a challenge to adherents to expected utility theory (von Neumann and Morgenstern, 1944) as a descriptive theory of human choice.

However, a number of criticisms of the methodology employed in the early probability matching experiments questioned whether matching behaviour really was sub-optimal. One criticism that applies directly to the nature of the IGT was that the probability of reinforcement was not independent and identically distributed (i.i.d) across all choices, i.e. they were not truly random (Myers, 1976; Vulkan, 2000). Fiorini (1971) showed that sampling without replacement within blocks of trials alters the probability of reinforcement within blocks. This means that the maximizing strategy of selecting only one option does not necessarily provide the most reinforcement. Because randomisation without replacement within blocks alters the probability of reinforcement, patterns of reinforcement can emerge between blocks. If participants are searching for a selection strategy, or if they do not believe the experimenter's claim that reinforcement is random, they may find at least enough of the pattern to encourage non-exclusive choice. And while this may look like probability matching to an experimenter who does not realise that the reinforcement environment is not random, for the subject, non-exclusive choice might actually approach an optimal strategy (Fiorina, 1971).

Jones and Myers (1966) found that people chose the reinforced response more when the trials in a probability learning task were randomised in short blocks (e.g. 20 trials) than when they were randomised in long blocks (e.g. 300 trials). The reason for this was due to the distribution of sequential dependencies - the distribution of runs of reinforcement of the same option. Randomization of reinforcement in short

blocks made the occurrence of long runs of reinforcement of one option unlikely. As a result participants exhibited more negative recency - the selection of the opposite option to the one previously reinforced with longer runs of reinforcement for that option. This indicates that with randomisation in short blocks participants are better able to predict the contingencies of the reinforcement schedule in probability matching experiments.

The findings from probability matching experiments apply to the IGT because the occurrence of reinforcements (losses) are fixed within ten-trial blocks. This means that the probability of loss is not independent and identically distributed across all selections from the decks. This may make the task easier to learn. The probability of loss fluctuates within blocks of selections from each deck contingent upon the number of choices that have been made from that deck and the number of losses encountered within that block. Consider the example of ten selections from deck A on the IGT as presented in the original Bechara et al (1994) study. The probability of loss from this deck is ostensibly 0.5 and viewed across a ten-trial block there are five losses per 10 cards. The reinforcement pattern on deck A is WWLWLWLWLL, where W is a win and L is a loss. However, the probability of loss changes within this block thus:

Loss/cards remaining:	5/10, 5/9, 5/8, 4/7, 4/6, 3/5, 3/4, 2/3, 2/2, 1/1, or
$p(\text{loss})$:	0.5, 0.55, 0.63, 0.57, 0.66, 0.60, 0.75, 1.0, 1.0

With the probability of loss contingent on the number of cards selected within a block of ten cards, on the ninth and tenth selection a loss is guaranteed. Similarly, for the decks with low frequency of loss (B and D), the probability of loss within a block of

ten cards increases with selection from that deck until the loss is encountered, and then becomes zero until that block ends. The probability matching experiments suggest that such a pattern might affect participants' behaviour although the pattern of reinforcement is complex (blocks of ten cards and four decks of cards provide a mass of permutations to mask any pattern). But the regularity with which losses are encountered within ten card blocks on decks may make their frequency more conspicuous. With such fixed schedules, in decks A and C the number of loss-free trials can be no greater than 5, and in practice are no greater than 3, while in decks B and D it is no greater than 9. This contingent learning environment results in non-normal distributions of loss-free runs in each deck and may be an important factor in the rapid learning observed on the IGT.

These observations also apply to Bechara et al's (2000) variant IGT mentioned in Chapter 1. In this version of the IGT, the task is switched from the domain of gains to the domain of losses. Bechara et al (2000) interpreted the behaviour of VMpfc patients on this task as supportive of their myopia for the future hypothesis in that they were impaired relative to the healthy controls. But the selections these patients made were not reliably different from chance on any trial and mean net score increased with block. This is not commensurate with insensitivity to future consequences and suggests learning may occur with sufficient experience. This poses the question of why the behaviour of these patients differs between tasks.

One possibility is that the variant task is easier than the standard IGT. The behaviour of healthy controls is informative. In block 1 on the variant IGT net score was not reliably different from chance, whereas it was reliably below chance in the standard IGT (Bechara et al, 2000). Net score then increased towards an asymptote above chance after 60 trials in the variant task meaning learning rate was positive but

not as large as in the standard version. The steeper slope in the standard task is due to the initial preference for the “bad” decks on this task. Given the results of the previous Experiments this probably reflects a preference for deck B. Fellows and Farah (2005) suggested that a reversal learning deficit is behind the behaviour of VMpfc patients on the standard IGT. They modified the standard IGT by changing the order in which the first losses occurred. They took the first eight cards and put them to the bottom of the deck. This resulted in the first card selected from deck B providing an immediate loss. A further change to this deck (switching card 14 with card 11) meant that on the third selection from deck B a further loss was received. Fellows and Farah (2005) reported that on the standard task, VMpfc patients’ performance was impaired compared to controls, although not at the level of that reported of VM patients in the Iowa group (6 of 9 of Fellows and Farah’s VMpfc patients selected > 50 cards from the advantageous decks). On their modified task, the behaviour of VMpfc patients was not significantly different from controls. One consequence of changing the order of losses in the fixed schedules is that the expected value of deck B becomes much more negative for the first block of ten cards. With two losses in the first three cards the frequency of loss as experienced by most participants is also not representative of the probability of loss as in the original task. Consequently, the manipulation does illustrate the importance of when the first losses occur and supports the hypothesis that a reversal learning deficit may explain VMpfc patient performance on the standard IGT.

It is notable that on the variant task the first win is earlier in the schedules for the good decks (E and G) than for the bad decks (F and H), and that compared to the standard task, the deck with the infrequent large win (deck E) provides this reinforcement much earlier (after 3 selections rather than 9). Thus the superior early

performance of healthy controls (Bechara et al, 2000) in the variant task may be an artefact of the reinforcement schedule. In Fellows and Farah's (2005) manipulated version of the standard IGT the number of advantageous cards selected by healthy controls was significantly above chance in the first block and remained stable across subsequent blocks. This suggests that the earlier placement of the events that are informative about the nature of a deck is a key determinant of the number of selections made from that deck.

The Experiments in this Chapter explore whether the nature of the reinforcement schedules, rather than their contingencies, affect learning on the IGT. In Experiment 6, the standard version of the IGT is used but the contingencies within each deck are independent and identically distributed across all cards. The results from the previous Experiments have shown the loss frequency results in differential selection between decks. As behaviour on the variant IGT has been used to support Bechara et al's (2000, 2001) conclusion that patients with VMpfc damage are myopic for future consequences it is important to understand whether what influences choice in this task varies in the same way as in the standard task. In Experiments 7 and 8 manipulations are made to the reinforcement schedules on the variant version of the IGT. Experiment 7 manipulates the fixed order of reinforcement in the variant task to match that of the standard task. While learning is still rapid, participants' pattern of selection is changed. In Experiment 8 the effect of using probabilistic reinforcement schedules on learning on the variant IGT is examined.

4.2 EXPERIMENT 6

NO LEARNING ON A PROBABILISTIC VERSION OF THE IGT

4.2.1 Introduction

Experiment 6 was designed to examine behaviour when the deck schedules on the IGT are independent and identically distributed. No previous research has used such schedules on the IGT. Some authors have randomised the occurrence of losses within the ten-trial blocks (e.g. Crone et al, 2004; Crone and Wagenmakers, personal communication). They have done this to counterbalance order effects across individuals while ensuring that the experienced payoffs are representative of the desired contingencies. However, the occurrence of losses is still constrained within blocks of ten trials. This method does mean that, like the fixed order schedules, reinforcements are not independent and identically distributed. This change to the methodology is better experimental practice as the occurrence of losses are randomised across participants and so not in the same fixed order. However, this change does not fundamentally alter the schedules and they remain deterministic because the informative events still only vary within ten-trial blocks.

The investigation of schedule type is also relevant to the claim that the IGT “simulates real-life decision-making in the way it factors uncertainty, rewards, and penalties” (Bechara et al, 1997a, p.1293). It could be argued that truly randomised fixed probability reinforcement schedules are much more realistic. Given the evidence from probability matching experiments cited in the Chapter introduction, it was predicted that when the events in the reinforcement schedules are independent and identically distributed (henceforth referred to as probabilistic), the task will be

made harder such that participants will fail to distinguish between the advantageous and disadvantageous decks.

4.2.2 Method

To determine if manipulating the schedules had any effect on learning the same experimental design employed in the Real Money – Hint condition of Experiment 3 was used. The only changes in methodology were made to the reinforcement schedules where the occurrence of losses was made probabilistic.

Participants

Twenty-two participants were recruited from the undergraduate and postgraduate populations at the University of Nottingham. Two participants failed to turn up for the second session and their data were excluded from the analysis. Participants were recruited through a poster advertisement that offered participants the opportunity to earn up to £6 by taking part in a cognitive psychology experiment.

Apparatus

A modified version of the computerised version of the IGT using real money reinforcers was created and run on a PC. The modifications made the probability of losses within each decks independent and identically distributed. This meant that for every selection of deck A, the probability of receiving a loss was 0.5 and on deck B the probability of a loss was 0.1 etc. Loss magnitudes, where they varied in the original task, were selected randomly if a loss occurred. In the fixed-schedule task ten selections from Deck A would encounter five losses of amount -£0.15, -£0.20, -£0.25,

-£0.30 and -£0.35. On the probabilistic schedule a loss with probability 0.5 would have an equal chance of being one of the five loss amounts for this deck. Similarly on deck C, the three loss amounts of -£0.025, -£0.05 and -£0.075 were equally likely if a selection resulted in a loss. Losses were always -£1.25 on deck B and -£0.25 on deck D. In all other respects the task was unchanged from the previous versions where fixed schedules of reinforcement were used. Real money reinforcers were used and have been described in the procedure for Experiment 3.

Design

A repeated measures design was utilised where participants' deck selections in each of five twenty-trial Blocks in two experimental Sessions were the repeated factors. Participants' performance was assessed by calculating a net score from the number of cards selected from the advantageous decks (A and B) minus the number selected from the disadvantageous decks (C and D). From this measure the slope, b , was calculated as an estimate of learning rate.

Procedure

The procedure followed exactly that of the Real Money – Hint condition in Experiment 3. As such participants were informed that “some decks are worse than others. You may find all of them bad, but some are worse than others. No matter how much you find yourself losing you can still win if you stay away from the worst decks” (Bechara et al, 1999; Bechara et al, 2000).

Table 4.1: Mean probability of loss and received value (in pence) for each deck in each session in Experiment 6.

	$p(\text{loss})$		Received Value	
	Session 1	Session 2	Session 1	Session 2
A	0.54 (0.04)	0.67 (0.04)	-4.12 (1.31)	-6.88 (1.30)
B	0.10 (0.01)	0.11 (0.01)	-2.87 (1.78)	-3.16 (1.58)
C	0.51 (0.03)	0.55 (0.05)	2.40 (0.14)	1.87 (0.34)
D	0.17 (0.05)	0.14 (0.02)	0.62 (1.15)	1.95 (0.51)

Note: Figures in parentheses are the standard error of the mean.

4.2.3 Results

The net score for each participant was calculated in both sessions. Additionally, net score was calculated in blocks of twenty trials for each participant allowing the change in net score across time, the slope b , to be calculated as an estimate of learning rate. To ensure that deck contingencies were similar to those of previous Experiments the probability of loss and the received value on each deck were calculated for each participant in each session. Table 4.1 displays the mean probability of loss and mean received value for each deck. With fixed schedules and uniform selection and experience of each deck, decks A and B would have an expected value of -£0.025 and decks C and D would have an expected value of £0.025. In Experiment 6 the probabilistic randomisation of losses resulted in more variability in the deck contingencies. Mean received values were more disadvantageous in decks A and B in each session, and slightly less advantageous in deck C and D. However, despite this variability, the nature of the decks remains the same as in previous experiments.

Mean loss amounts in decks A and C were similar to the mean amounts of the original task so are not considered further. As in Chapter 2, performance in each session is discussed separately.




Figure 4.1: Mean net score across 20-trial blocks for each session in Experiment 6 (Probabilistic reinforcement schedules). The dashed line represents chance selection, or no preference for either advantageous or disadvantageous decks. Error bars are the standard error of the mean.

Session 1

Figure 4.1 displays the mean net score across block in each session and across the whole experiment. Mean net score in session 1 was -6.2 ($SD = 37.81$). A one-sample

t-test found that this was not significantly different from chance, $t(19) < 1.0$, $SD = 37.81$, $p > 0.05$.

Two measures can be used to determine whether learning occurred in session 1. The mean net score in the final block of twenty trials was compared to chance and not found to be significantly greater than chance, $t(19) < 1.0$, $SD = 10.77$, $p > 0.05$, indicating that with probabilistic reinforcement schedules participants do not show a preference for the decks C and D after one hundred trials. Examination of Figure 4.1 suggests that although participants did not end session 1 with a preference for the good decks, the change in selection from them may be increasing. The possibility that learning rate increased across the session was assessed using Lorch and Myers (1990) regression analysis for repeated measures designs. Mean learning rate was estimated from the slope, b , and in session 1 was 2.07 ($SD = 3.21$). This was significantly greater than zero, $t(19) = 2.89$, $p < 0.01$. While participants did not preferentially select from the advantageous decks at the end of session 1, the increase in learning rate over the session indicates that this would occur with more exposure.

Session 2

The mean net score in session 2 was 3.7 ($SD = 41.14$). A one sample *t*-test found that this was not significantly greater than zero, $t(19) < 1.0$. A paired samples *t*-test compared net score between sessions and revealed no significant difference $t(19) < 1.0$, $SD = 49.32$, $p < 0.05$. Despite a shift between sessions from a negative to a positive mean net score, and the expectation of improved performance predicted from the increasing learning rate in session 1, mean net score for the session did not reflect a preference for the advantageous decks. Examination of Figure 4.1 shows that this




Figure 4.2: Mean total cards selected from each deck in each session of Experiment 6. Error bars are the standard error of the mean. The dashed line represents selection at chance.

may be because of the exceptionally low net score in block 6. This low net score is not predicted from the increasing learning rate found in session 1 and suggests that by the end of the first session participants have not learned which decks are advantageous. Another interpretation is that the low value in block 6 reflects exploration behaviour by participants. It is possible that they did not trust the Experimenter's assertion that the task was unchanged from the first session. However, there is a rapid increase in net score across the session, which suggests that participants do at least learn to make fewer selections from decks A and B. A

regression analysis using the Lorch and Myers (1990) method was run and determined that mean learning rate, estimated using the slope b , was 1.68 ($\sigma_M = 0.92$). This was not significantly different from zero, $t(19) = 1.83$, $SD = 4.11$, $p = 0.08$. By examining participants' selections in the final block of the experiment it can be determined if learning, as indicated by a preference for the advantageous decks, had developed. Net score in block 10 was compared to zero using a one sample t -test and was not significantly different from zero, $t(19) < 1$, $SD = 14.16$, $p > 0.05$. These results suggest that when probabilistic reinforcement schedules are used participants do not learn to select preferentially from decks C and D. Looking at the number of cards selected from each deck supports this conclusion. Figure 4.2 shows that in both sessions participants made most selections from decks B and D, with selection from B greater than chance. However, some indication that participants are learning something about the nature of the task is found in the increase in selections from deck C between sessions, while selections from the other decks decrease.

4.2.4 Discussion

The results from Experiment 6 suggest that probabilistic reinforcement schedules make it hard for participants to learn preferential selection from decks C and D. Yet there were signs that participants did increase the number of cards selected from the advantageous decks across each session. This resulted in a learning rate that was significantly greater than zero in session 1 and almost reached significance in session 2. The pattern of net score across block in session 1 is similar to that seen in the previous experiments using fixed reinforcement schedules, although in the current experiment there is more selection from A and B in the first blocks of session 1. One




Figure 4.3: The range of participants' net scores in both sessions of Experiment 6.

possible explanation for this difference in behaviour is that with the probabilistic schedules some participants do not encounter as many losses on decks A and B as do others. Until sufficient losses are encountered to outweigh the continual high reward, decks A and B are actually the more profitable alternatives (Maia and McClelland, 2004). This should be especially true of deck B where the probability of a loss is much lower and so the possibility of continued selection without loss is greater than in deck A. This is the likely location for the effect of probabilistic schedules on learning and it may account for the large change in selection between sessions that is apparent in Figure 4.1. Despite mean received values that showed decks A and B were disadvantageous and C and D were advantageous, the fluctuation in received value

across trials is greater with the probabilistic schedule than in the previous fixed schedule experiments. This may be partly why preferential selection for decks C and D does not develop across the group. In Figure 4.1 the variance in net score increases with trial block. This reflects the increased divergence in selection between participants for whom the advantageous decks are the best decks and those for whom the disadvantageous decks are not necessarily disadvantageous.

Figure 4.3 depicts the net scores of individual participants in both sessions. In session 1 55% of the participants had a negative net score ($n = 11$) while 45% had a positive net score ($n = 9$). Table 4.2 displays the mean probability of loss and mean received value for each deck for participants with a positive and negative net score in session 1 and in session 2. For participants with a preference for decks A or B, those with a negative net score, selecting from deck B was rewarding and selecting from deck D was not. The opposite is true in participants with a positive net score, those with a preference for deck C or D. Session 1 performance in this IGT with probabilistic schedules may be better than represented in Figure 4.1. Participants with a negative net score, as a result of selecting more from deck B, are doing so because it is a more rewarding deck than deck D. This is illustrated in Figures 4.4 and 4.5, where the mean received value within each block (filled circles) and the mean number of selections (open circles) from each deck is displayed for participants with a positive net score (Figure 4.4) and a negative net score (Figure 4.5). Like Figure 4.2, Figures 4.4 and 4.5 show that participants generally avoided deck A in both sessions, while selection from deck C was also generally below chance in session 1 (chance selection is presented by the dashed line). The differences between the participants split on the net score measure lie mainly in their selections from decks B and D. In the first block, when participants generally sample from all the decks, for participants with a

Table 4.2: Mean expected value and probability of loss for each deck in each session of Experiment 6 for participants with a positive or negative net score.

Group	p(loss)		Received Value	
	Session 1	Session 2	Session 1	Session 2
<u>A</u>				
Positive	0.53 (0.07)	0.65 (0.05)	-3.41 (2.16)	-6.08 (1.35)
Negative	0.55 (0.04)	0.70 (0.08)	-4.83 (1.57)	-7.77 (2.35)
<u>B</u>				
Positive	0.13 (0.05)	0.12 (0.02)	-6.43 (3.00)	-4.46 (2.41)
Negative	0.07 (0.01)	0.09 (0.02)	0.69 (1.24)	-1.56 (1.94)
<u>C</u>				
Positive	0.56 (0.03)	0.54 (0.07)	2.20 (0.14)	1.97 (0.41)
Negative	0.47 (0.05)	0.57 (0.10)	2.59 (0.23)	1.71 (0.65)
<u>D</u>				
Positive	0.10 (0.02)	0.12 (0.02)	2.11 (0.62)	2.47 (0.54)
Negative	0.24 (0.09)	0.15 (0.04)	-0.88 (2.18)	1.25 (0.96)

Note: Figures in parentheses are the standard error of the mean. n in each session (positive group, negative group): session 1 (11, 9), session 2 (9, 11).

positive net score the mean received value within block 1 from selecting deck B was negative, while it was positive for deck D (the dotted line represents a received value of 0). Participants in this group show a decrease in selection from deck B and an

increase in selection from deck D to levels below and above chance respectively. Because the mean received value from deck B generally stays below zero and the mean received value from deck D is generally greater than zero, these are advantageous choices. The increase in selection from deck C with block also contributes to these participants' positive net scores.

For participants with a negative net score, the mean received value in the first block was positive for deck B and negative for deck D. Many of these participants did not encounter a loss in their initial selections from deck B but did so on deck D. Mean received value for deck B stays above zero for blocks 1 to 3, driving selection from deck B that remains well above chance despite a decrease in the later blocks of the session. The increase in selection from deck D follows an increase in the mean received value from this deck. While, the received value is dependent on the number of selections, frequency of loss and the loss amount, if there are no losses on some decks then it is sensible to continue to pick from these decks. What Figure 4.5 shows is that participants with a negative net score have a negative net score because they make the majority of their selections from deck B. And for the majority of these participants deck B is an advantageous choice because it has both a large immediate gain and a positive received value.

Group membership changed between sessions as four participants who had a negative net score in session 1 selected more cards from decks C and D in session 2, while two participants moved from a preference for deck D to selecting more from deck B. In these two cases, the initial selections in session 2 from the previously reinforcing deck D resulted in losses. In the same block there were very few, if any, losses on deck B, although a sufficient number were encountered across the task for the deck to be a disadvantageous choice. These participants experienced a change in

the occurrence of losses between sessions on a deck they had learned to prefer. This perhaps made them doubt the experimenter's claim that nothing had changed, resulting in more exploratory behaviour. It also suggests that these participants had not learned why deck D was advantageous. Discussion of participants' knowledge will be addressed in Chapter 5.

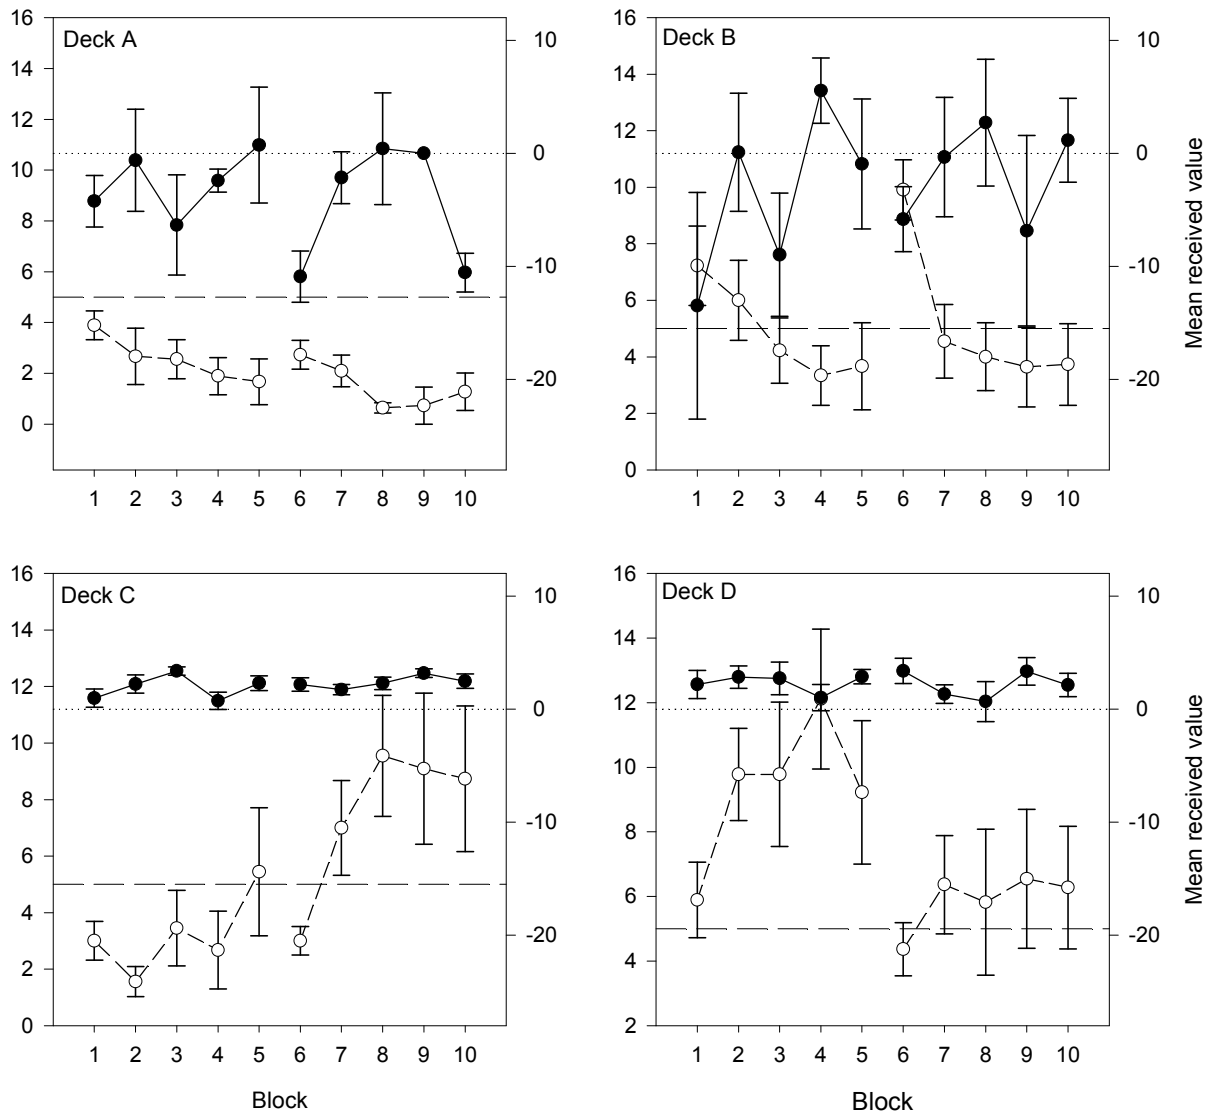


Figure 4.4: Mean received value and mean number of selections in each 20-trial block for participants with a mean net score ≥ 0 . In session 1 $n = 9$. In session 2 $n = 11$. Error bars are the standard error of the mean.

- Mean number of selections
- Selection at chance
- Mean received value
- Received value = 0

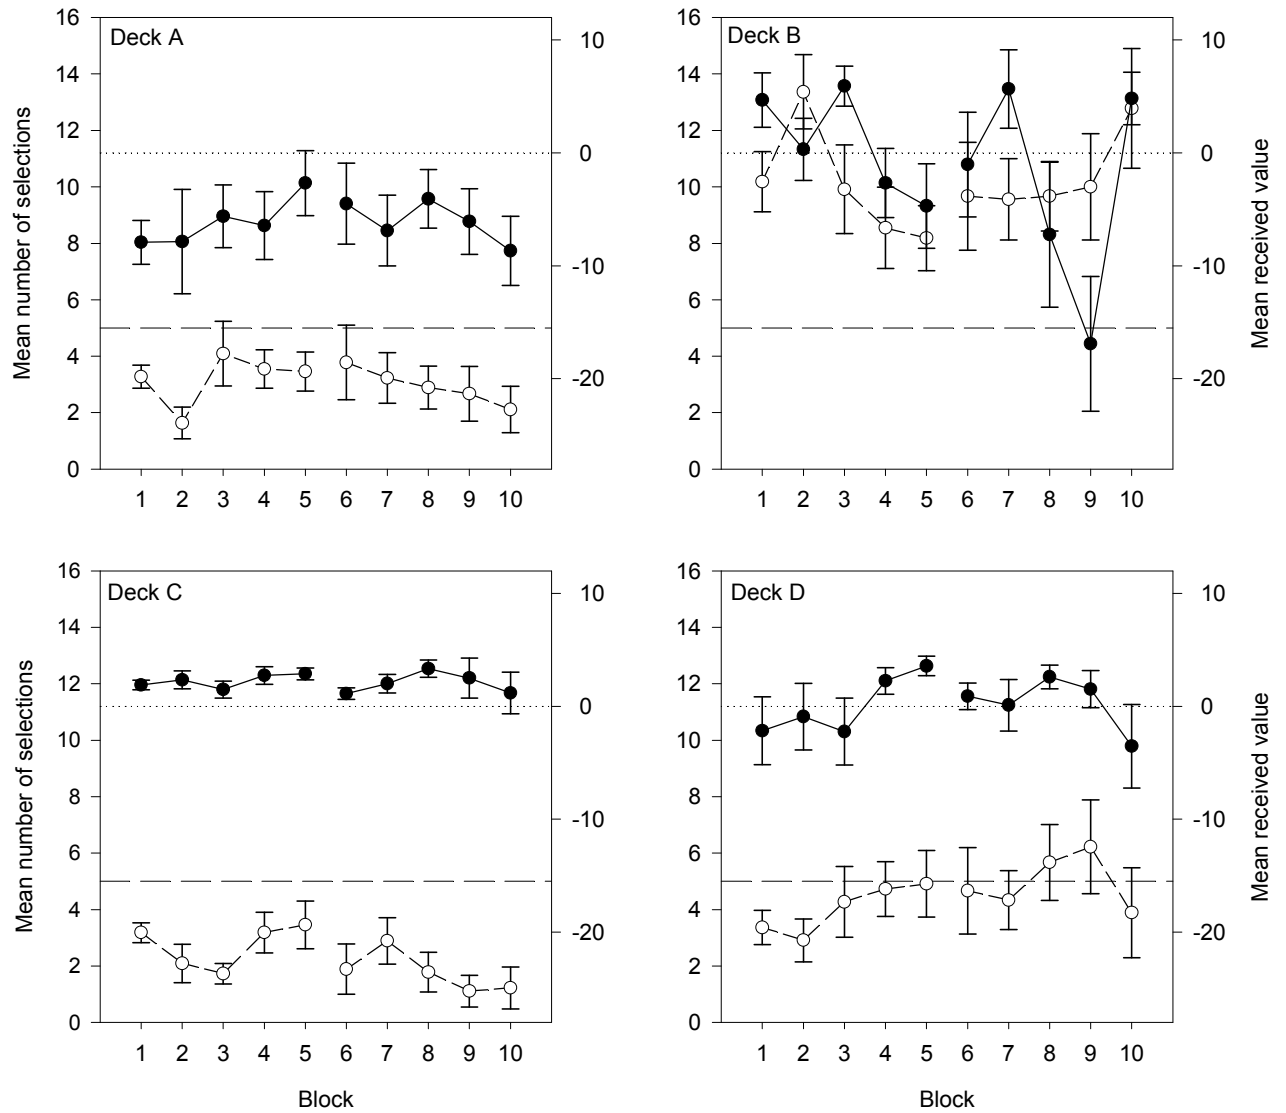


Figure 4.5: Mean received value and mean number of selections in each 20-trial block for participants with a mean net score < 0. In session 1 $n = 11$. In session 2 $n = 9$. Error bars are the standard error of the mean.

- Mean number of selections
- - - Selection at chance
- Mean received value
- Received value = 0

Table 4.2 shows that, unlike in session 1, in session 2 for both groups, decks C and D have the largest received values. These values are higher for participants with a positive net score while received value for deck B is more negative. Figures 4.4 and 4.5 show a similar pattern of reinforcement and selection behaviour as was described for session 1. For participants with negative net scores deck B is initially reinforcing

enough to maintain preferential selection above chance. It is plausible that the absence of losses that drives the initial positive received value for deck B, which was also experienced through much of session 1, makes it difficult for participants to identify deck B as a disadvantageous deck. For most of these participants, the occasional large losses mean deck B is a more favourable option than a switch to selection from deck C, as observed in participants with a positive net score.

The results from Experiment 6 suggest that deck B is the important deck in IGT selection. As with the earlier Experiments participants readily identify the disadvantageous deck with more frequent losses (A) as a bad choice. Initial preference for deck B is followed by a decline in selection as losses are encountered, as reflected in the behaviour of participants with a positive net score in this Experiment. If insufficient losses are encountered on this deck then, because of the large immediate gain from selecting it, it is clearly the best deck to select from. So much so that even with increased frequency of loss in later blocks (and decreased received value) selection persists.

Experiment 6 has shown that participants do not learn to select advantageously when, arguably more realistic, probabilistic reinforcement schedules are used on the IGT. This has some implications for learning using the standard IGT schedules that concern the occurrence of losses on deck B. This is linked to the explanation offered by Fellows and Farah (2005) for the behaviour of participants with VMpfc damage. When the original fixed schedules were altered so that participants encountered the initial losses on deck B earlier, the behaviour of patients with VMpfc damage was no different to that of healthy controls. Net scores for these controls were never near chance values indicating that they had learned all they needed to succeed on the task in the first block. That patients with VMpfc patients did not do so on the standard

version of the task, Fellows and Farah (2005) suggested was due to a deficit in reversal learning. This contention is explored further in Experiment 7 using the variant version of the IGT.

4.3 EXPERIMENT 7

THE VARIANT IGT – ORDER OF LOSSES AFFECTS LEARNING

4.3.1 Introduction

The previous experiments have shown that normal participants readily learn to avoid deck A, but it takes longer to learn to avoid deck B. The results from Experiment 6 suggested that when the first losses on a deck occur are the key determinants of further selection from that deck. This is especially true for deck B because the gains are large and the losses infrequent. When the occurrence of these losses is independent and identically distributed over all selections from decks, participants who encounter losses earlier make fewer choices from this deck. The importance of the number of selections without loss on IGT behaviour has been recognised by Fellows and Farah (2005). These authors suggested that patients with VMpfc damage have a deficit in reversal learning wherein they cannot reverse initial preference for the disadvantageous decks. This contrasts with the Iowa group's hypothesis that these patients' behaviour is the result of insensitivity to future consequences (Bechara et al, 2000), an explanation that fits in with their behaviour in their daily life (Damasio, 1994). The behaviour of healthy controls on these tasks suggests that the variant task is easier to learn and one reason may be the information received in the first selections from a deck. Experiment 7 tests the hypothesis that the order in which informative

reinforcing events are encountered within fixed reinforcement schedules influences selection behaviour, and therefore learning, on the IGT. If the order in which gains are encountered is unimportant, as implied by the assumption of a maximizing system and in the calculation of the net score measure, then selection from each deck with the same expected value (the same long term outcome) should be similar. However, if deck selection mirrors that found in the examination of the standard IGT then this assumption cannot be supported. This hypothesis is tested by altering the variant version of the IGT so that the infrequent informative events (the wins) occur in the same order in each deck as the losses occurred in each deck in the standard version of the task. If the manipulation has no effect, learning as measured by net score should still be found and above chance selection from each of the good decks should result.

4.3.2 Method

Participants

Twenty-three participants were recruited from the undergraduate and postgraduate populations of the University of Nottingham. Three participants failed to return for the second session and their data were excluded from the analysis. Participants were recruited through a poster advertisement that offered participants the opportunity to earn up to £6 by taking part in a cognitive psychology experiment.

Apparatus

The experiment took the same format as in the Real Money – Hint condition of Experiment 3 with the exception that the variant version of the IGT was used

(Bechara et al, 2000). Instead of earning money with every card selection, participants lost money. This amount was ten pence in decks E and G and five pence in decks F and H. As in Bechara et al (2000) decks E and G were advantageous as the amount won exceeded the amount lost and decks F and H were disadvantageous as the amount won was less than the amount lost. The schedules of reinforcement were altered so that they were exactly analogous to those used in Experiment 3b and in Bechara et al (1994). The major change that resulted from this was that the first win on deck E moved from the third to the ninth selection from that deck, making the occurrence of the first win on that deck synonymous with the first loss on deck B in the schedules used in the previous Experiments (see Appendix C).

Design

A repeated measures design was used where participants' deck selections in each of five twenty-trial Blocks in two experimental sessions were the repeated factors. Participants' performance was assessed by calculating a net score from the number of cards selected from the advantageous decks (E and G) minus the number selected from the disadvantageous decks (F and H). In addition, following the results of the previous Experiments individual deck selection was also examined.

Procedure

The procedure followed that of the Real Money – Hint condition of Experiment 3. The necessary alterations were made to all verbal and written instructions to reflect the change from the standard to the variant IGT.

4.3.3 Results

Net score was calculated by subtracting the number of cards selected from decks F and H from the number selected from E and G. Net score was calculated in both sessions for each participant. Additionally, net score was calculated in blocks of twenty trials for each participant allowing the change in net score across time, the slope b , to be calculated as an estimate of learning rate. Figure 4.6 displays the mean net score across block in each session and across the whole experiment. As in previous experiments performance in each session will be discussed separately.

Session 1

Mean net score in session 1 was 4.7 ($SD = 52.72$). A one sample t -test found that this was not significantly different from zero, $t(19) < 1.0$, $SD = 52.72$, $p > 0.05$. To investigate whether learning occurred across the session learning rate was estimated for each participant and compared to zero using the Lorch and Myers (1990) regression analysis for repeated measures designs. Mean learning rate was 1.70 ($SD = 2.30$) and significantly greater than zero, $t(19) = 3.31$, $SD = 2.30$, $p < 0.01$, indicating that as experience with the task increased, selection from decks F and H decreased. To investigate whether this learning had developed into a preference for decks E and G at the end of the session a one sample t -test compared mean net score in the final block of session 1 to chance. There was no significant difference, $t(19) < 1.0$, $SD = 12.16$, $p > 0.05$, indicating that while participants shifted their selection away from decks F and H, they did not prefer decks E and G after 100 card selections.




Figure 4.6: Mean net score across 20-trial blocks for each session in Experiment 7 (reversed IGT with fixed schedules). The dashed line represents chance selection, or no preference for either advantageous or disadvantageous decks. Error bars are the standard error of the mean.

Session 2

Mean net score in session 2 was 23.7 ($SD = 65.01$). A one sample t -test found that this was not significantly greater than zero $t(19) = 1.63$, $SD = 65.01$, $p > 0.05$. A paired samples t -test compared net score between sessions and revealed that net score in session 2 was significantly greater than net score in session 1, $t(19) = -2.46$, $SD = 34.55$, $p < 0.05$. Figure 4.6 shows that selection in the second session is mainly from decks E and G, and net score is above chance in all blocks except the first. Net score does not increase after block 7 and asymptotes at a value of around 6. For this reason

it is no surprise that learning rate was not significantly greater than zero, $b = 0.68$, $t(19) = 1.27$, $SD = 2.39$, $p > 0.05$. It is surprising that although mean net score in the final block is smaller than in the previous blocks, it was not significantly greater than zero as revealed by a one sample t -test, $t(19) = 1.45$, $SD = 14.19$, $p > 0.05$. These results reflect the large variability in individual selection. Figure 4.7 depicts the range of participants' net scores in both sessions. Three participants clearly show a preference for one or both of decks F and H that does not change with experience on the task. In fact, one of these participants (HF) chose from deck H on almost every selection, one (HM) chose from F on almost every selection, and one (PJ) chose mostly from F but also sampled from H. Excluding these participants the majority have a positive net score and thus a preference for decks E, G or both. The presence of these three participants increases the variance in net score, which resulted in the large mean net score for the group that was not significantly greater than zero.

Figure 4.8 depicts the number of cards selected from each deck in each session of Experiment 7. It shows that participants avoid deck E and that most selection is from the cards with frequent wins (decks F and G) with substantially greater selection from G than from any other deck. Deck G is the only deck from which selection is increased between sessions. A similar pattern was observed in Experiment 3 on the standard IGT where selection from deck C was observed to increase between sessions. These results show again the pattern of differential selection from the disadvantageous decks that was reported in Experiments 1 to 5. Additionally, selection from the advantageous decks is not uniform. As hypothesised this is influenced by the position in the fixed order of cards of the less frequent reinforcer.




Figure 4.7: The range of participants' net scores in both sessions of Experiment 7.

Table 4.3 presents the mean probability of win and mean received value for each deck experienced by participants in each session of Experiment 7. The below chance selection from decks E and H, depicted in Figure 4.8, is reflected in a probability of win that is lower than the description of these decks commonly given in the literature. These measures are related; with a fixed schedule if insufficient cards are chosen then less wins will be experienced and the probability of a win will be decreased. In the case of deck E this resulted in many participants not experiencing a win on deck E, which in turn resulted in the large negative expected value shown in Table 4.3. The manipulation of the placement of the first win in this deck thus affected the number of cards chosen from it, the deck contingencies reported here and

consequently the measure of learning determined from net score. However, the large standard errors for deck E in each session indicate that for some participants, deck E was actually an advantageous deck. This required perseverance through eight non-rewarded selections.




Figure 4.8: the number of cards selected from each deck in each session of Experiment 7. Error bars are the standard error of the mean. The dashed line represents chance selection.

Table 4.3: Mean probability of win and received value for each deck in each session in Experiment 7.

	p(win)		Received Value	
	Session 1	Session 2	Session 1	Session 2
E	0.04 (0.01)	0.05 (0.02)	-5.21 (1.52)	-2.97 (2.04)
F	0.46 (0.01)	0.40 (0.03)	-2.76 (0.06)	-3.01 (0.17)
G	0.45 (0.04)	0.48 (0.03)	1.05 (0.93)	2.03 (0.71)
H	0.05 (0.01)	0.02 (0.01)	-3.85 (0.25)	-4.39 (0.24)

Note: Figures in parentheses are the standard error of the mean.

4.3.4 Discussion

The fixed order of wins used by Bechara et al (2000) in their variant IGT was altered in Experiment 7 to match the fixed order of losses used in the equivalent decks on the standard version of the IGT. On the variant task used by Bechara et al (2000) learning is rapid and above chance selection from the advantageous decks is found from the second block onwards. In the version used in Experiment 7 the deck contingencies were exactly the same i.e. decks in each version have the same expected values and the same probability of win. All that was manipulated was the fixed order in which the informative events, the wins, occurred in each deck with specific interest on the deck with the infrequent but large magnitude event (deck E). The manipulation of the fixed order of wins resulted in selection from E that was significantly below chance as illustrated in Figure 4.8. This is very like the pattern of selection from deck A on the standard task, but instead of making the task easier, the manipulation made it harder

for participants to identify deck E as an advantageous deck. This in turn made the task more difficult. The below chance selection on deck E results because many participants stop selecting from this deck before they encounter the first win. This is reflected in the values for the mean probability of loss and the mean received value for deck E, shown in Table 4.3, that are below what would occur with selections numbering in multiples of ten. This result illustrates the importance of the order in which informative events are placed within the decks on IGT performance as measured using the change in net score. These results extend, using the variant IGT, those of Fellows and Farah (2005) using the standard IGT.

Despite below chance selection from one of the advantageous decks, evidence that participants' performance did improve both within and across sessions was found using the net score measure. This improvement was driven by selection from deck G that was significantly greater than chance and occurred despite participants showing some preference for deck F over decks E and H. It is of note that, in general, participants avoided the decks with infrequent wins, a finding also reported by Crone et al (2004). In Experiment 6 preference for the decks with infrequent losses was observed (this was also observed to some extent in the Real Money – Hint condition of Experiment 3), again similarly reported by Crone et al (2004) and others (MacPherson et al, 2002; Wilder et al, 1998). This behaviour should not be unexpected. In the context of a losing environment, participants' behaviour tends towards choices with more frequent gains, while in the context of gains, it tends away from frequent losses. This behaviour is in keeping with the Law of Effect. It also suggests that selection from the decks on the variant IGT, like on the standard IGT, is not uniform between decks with the same expected value.

Bechara et al (2000) found higher net scores earlier on the variant than the standard task, but made no comment upon this. From their Figure 2, selection from the advantageous decks in the first three trial-blocks appears to be significantly greater in the variant than in standard task. This suggests that the variant task is easier to learn than the standard task. The results of Experiment 7 suggest that this may in part be attributable to differences in the schedule of the informative events which occur later on the standard task. However, it may also be reflective of the decision-making environment. Kahneman and Tversky's (1979) value function predicts that when people are faced with certain losses, more risk-seeking behaviour results. In the context of the variant version of the gambling task, risk seeking could be defined as more selection from the decks with higher immediate losses – the behaviour that is observed in normal participants. Whereas, when people are faced with certain gains, the value function hypothesises risk aversion and in the context of the IGT this is selection from the decks with lower immediate gains. The hypothesis that the variant task is easier to learn is supported by a probability learning study that incorporated losses into the alternate options. Bereby-Meyer and Erev (1998) reported faster learning of the maximising strategy (exclusive choice for the high probability option) when the options created a loss making environment. Bereby-Meyer and Erev (1998) found that the model which best fitted the data was Erev and Roth's (1998) quantification of the Law of Effect. This learning principle also predicts the pattern of selection found on both versions of the gambling task.

4.4 EXPERIMENT 8

THE VARIANT IGT – PROBABILISTIC SCHEDULES AFFECT LEARNING

4.4.1 Introduction

The results from Experiment 7 found that changing the fixed order of losses within decks on the variant task resulted in participants avoiding the advantageous deck with infrequent but large magnitude wins. Evidence of learning was found, but it was driven by the participants' selections from the advantageous deck with frequent wins. Experiment 8 goes on to examine how the nature of the reinforcement schedules affects IGT behaviour. The same manipulation as conducted in Experiment 6 is applied to the variant IGT. In Experiment 6, as well as being probabilistic, participants' learning environment was framed in terms of gains, and this may have contributed to their failure to distinguish between the advantageous and disadvantageous decks. If participants are capable of learning on the IGT with the theoretically more difficult probabilistic schedules then their motivation, and therefore their learning, may be increased if the decision-making environment is framed in terms of losses, as in the variant task. In examining behaviour on a variant IGT with probabilistic schedules the assumption that participants' deck selection is directed solely by long-term outcomes can also be tested in an outwardly similar, but theoretically more difficult environment.

4.4.2 Method

Participants

Twenty-two participants were recruited from the undergraduate and postgraduate populations of the University of Nottingham. Two participants failed to return for the second session and their data were excluded from the analysis. Participants were recruited using the same poster as for Experiment 7.

Apparatus

The gambling task used in Experiment 7 was modified in the same way as in Experiment 6 to make the schedule of wins on all decks probabilistic. The modifications were the same as those made in Experiment 6 except that the changes were made to the schedule of wins.

Design

As in Experiment 7.

Procedure

The procedure followed exactly that of Experiment 7.

4.4.3 Results

Mean net score was calculated in the same way as in Experiment 7: the number of cards selected from decks F and H were subtracted from the number selected from decks E and G.

Session 1

Mean net score in session 1 was 11.0 ($SD = 31.08$). A one sample t -test found that this was not significantly greater than zero, $t(19) = 1.58, p < 0.05$. Figure 4.9 displays mean net score across each twenty-trial block in both sessions. It shows that net score increases across block and that by the end of the session most cards are being selected from decks E and G. This is reflected in a mean learning rate that is significantly greater than zero as revealed using Lorch and Myers (1990) regression analysis for repeated measures designs, $b = 1.31, t(19) = 2.31, SD = 2.53, p < 0.05$. Similarly, when compared to zero, mean net score in the final block of session 1 is significantly greater than zero, $t(19) = 2.10, SD = 8.72, p < 0.05$. These results show that participants experiencing the reversed probabilistic schedules are learning to select cards from the advantageous decks.

Session 2

Mean net score in session 2 was 24.10 ($SD = 38.12$). A one-sample t -test found that this was significantly greater than zero, $t(19) = 2.83, p < 0.05$. A paired sample t -test found that mean net score was not significantly greater in session 2 than in session 1, $t(19) = 1.44, SD = 40.77, p > 0.05$. However, from Figure 4.9 it is clear that the increasing learning rate found in session 1 continues in session 2. Lorch and Myers

(1990) regression analysis for repeated measures designs revealed that mean learning rate in session 2 was also significantly greater than zero, $b = 1.33$, $t(19) = 2.29$, $SD = 2.59$, $p < 0.05$. Participants who experienced the reversed IGT with probabilistic schedules develop a preference for the advantageous decks that is apparent by the end of session 1 and continues through session 2.

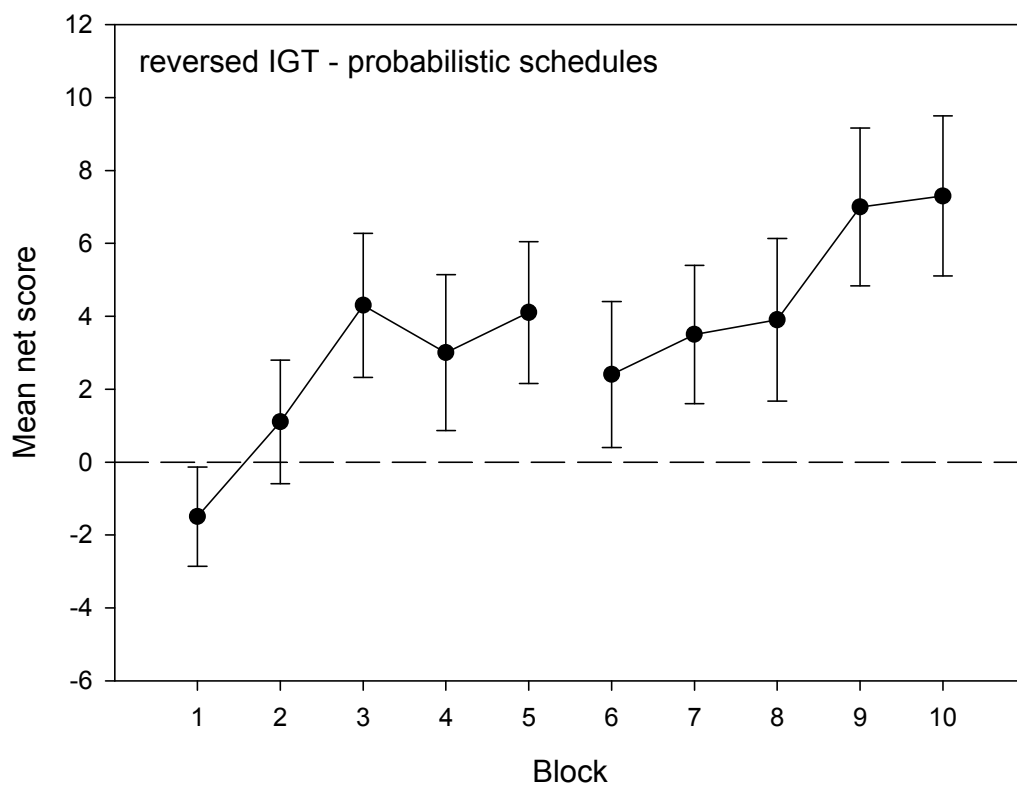


Figure 4.9: Mean net score across 20-trial blocks in Experiment 8 (reversed IGT with probabilistic schedules). The dashed line represents chance selection, or no preference for either advantageous or disadvantageous decks. Error bars are the standard error of the mean.




Figure 4.10: The range of participants' net scores in both sessions of Experiment 8.

The clear evidence of learning found in this Experiment are reflected in the change in the number of participants with a negative net score between sessions. This is depicted in Figure 4.10 and decreases from seven participants in session 1 to five in session 2. In session 1 five of these seven participants have negative net scores only marginally below zero. In session 2 all five participants had more negative scores than in the first session, but two of these were only marginally below zero (having been marginally above in session 1). Only three participants' net scores were well below zero. Participant KH chose predominantly from deck G in session 1, but in session 2 chose increasingly from deck F. One reason for this is that for KH the received value over the first selections of deck G in session 2 was negative. This is true for many of the participants with a negative net score.

Table 4.4 displays the mean received value and mean probability of loss for each deck across all participants. Due to the nature of the probabilistic schedules all decks except G have a negative expected value. Like in Experiment 7, the mean received value for deck E is negative in part because some participants did not receive a win early enough in their selections from that deck and stopped selecting from it before a win was received.

Table 4.4: Mean probability of win and received value for each deck in each session in Experiment 8.

	p(win)		Received Value	
	Session 1	Session 2	Session 1	Session 2
E	0.07 (0.01)	0.06 (0.01)	-1.82 (1.74)	-2.60 (1.71)
F	0.50 (0.02)	0.51 (0.04)	-2.57 (0.12)	-2.52 (0.23)
G	0.46 (0.02)	0.45 (0.03)	1.49 (0.52)	1.16 (0.81)
H	0.09 (0.02)	0.05 (0.02)	-2.74 (0.41)	-3.55 (0.43)

Note: Figures in parentheses are standard errors of the mean.

4.4.4 Discussion

As in Experiment 7 participants in Experiment 8 learned to select more from one of the advantageous decks with increasing exposure to the variant gambling task.

Despite the use of probabilistic schedules, performance in Experiment 8 is superior to that found with the fixed schedules used in Experiment 7. This is likely due to the

manipulation made to the fixed schedules, especially on deck E, in Experiment 7. In Experiment 7 the first win on deck E occurred after nine selections for all participants. With an equal probability of encountering a win on every selection from deck E in Experiment 8, more participants encountered a win on this deck. This resulted in greater selection from deck E as displayed in Figure 4.11. However, selection from E was still below chance as was selection from deck H. These results replicate the pattern observed in Experiment 7. On the variant task people do not sample as frequently from the decks with low frequency of wins. As in Experiment 7 learning was driven by increasing selection from deck G.

However, as Table 4.4 suggests and Table 4.5 displays in more detail, deck G was not a rewarding deck for all participants. For the five participants with a negative net score in session 2, the mean received value on deck G was almost zero. However, decks G and E were still the better decks for these participants to select despite their received values. But as Figure 4.12 illustrates these participants preferred the lower immediate losses and relatively more frequent gains on deck F in both sessions. In session 1 this was countered by increasing selection from deck G, while those few participants with a negative net score in session 2 selected at chance from G was mirrored in deck H. This contrasts with the majority of participants depicted in the corresponding Figure 4.13. Here selection from deck G clearly drives learning in both sessions with selections from deck E at chance and declining selection from decks F and H. Figure 4.13 shows that despite an initial low probability of win (reflected in mean received value < 0) selection from deck G never dips below chance levels.

Experiment 8 investigated behaviour on the variant version of the IGT when probabilistic reinforcement schedules were used. Participants' behaviour showed evidence of learning to select preferentially from the advantageous decks, despite the

use of these theoretically more difficult reinforcement schedules. When results are compared to participant behaviour reported in Bechara et al (2000), participants in Experiment 8 do not reach the same level as measured by net score until the end of the second session. However, unlike on the standard IGT with probabilistic schedules evidence of learning was found across both sessions. This may be in part because learning is faster in the context of losses (Bereby-Meyer and Erev, 1998) making the variant IGT easier.




Figure 4.11: Mean number of cards selected from each deck in each session of Experiment 8. Error bars are the standard error of the mean. The dashed line represents chance selection.

Table 4.5: Mean expected value and probability of win for each deck in each session of Experiment 8 for participants with a positive or negative net score.

	p(win)		Expected Value	
	Session 1	Session 2	Session 1	Session 2
<u>E</u>				
Positive	0.07 (0.02)	0.06 (0.01)	-0.14 (2.20)	-3.09 (1.76)
Negative	0.04 (0.02)	0.07 (0.04)	-4.94 (2.65)	-1.12 (4.71)
<u>F</u>				
Positive	0.49 (0.05)	0.51 (0.05)	-2.43 (0.24)	-2.57 (0.29)
Negative	0.54 (0.03)	0.51 (0.03)	-2.48 (0.19)	-2.36 (0.25)
<u>G</u>				
Positive	0.46 (0.02)	0.45 (0.04)	1.47 (0.55)	1.42 (1.05)
Negative	0.49 (0.04)	0.43 (0.04)	1.98 (1.12)	0.37 (0.83)
<u>H</u>				
Positive	0.08 (0.02)	0.04 (0.02)	-2.72 (0.54)	-3.99 (0.39)
Negative	0.09 (0.03)	0.11 (0.05)	-2.79 (0.65)	-2.20 (1.13)

Note: Figures in parentheses are standard errors of the mean.

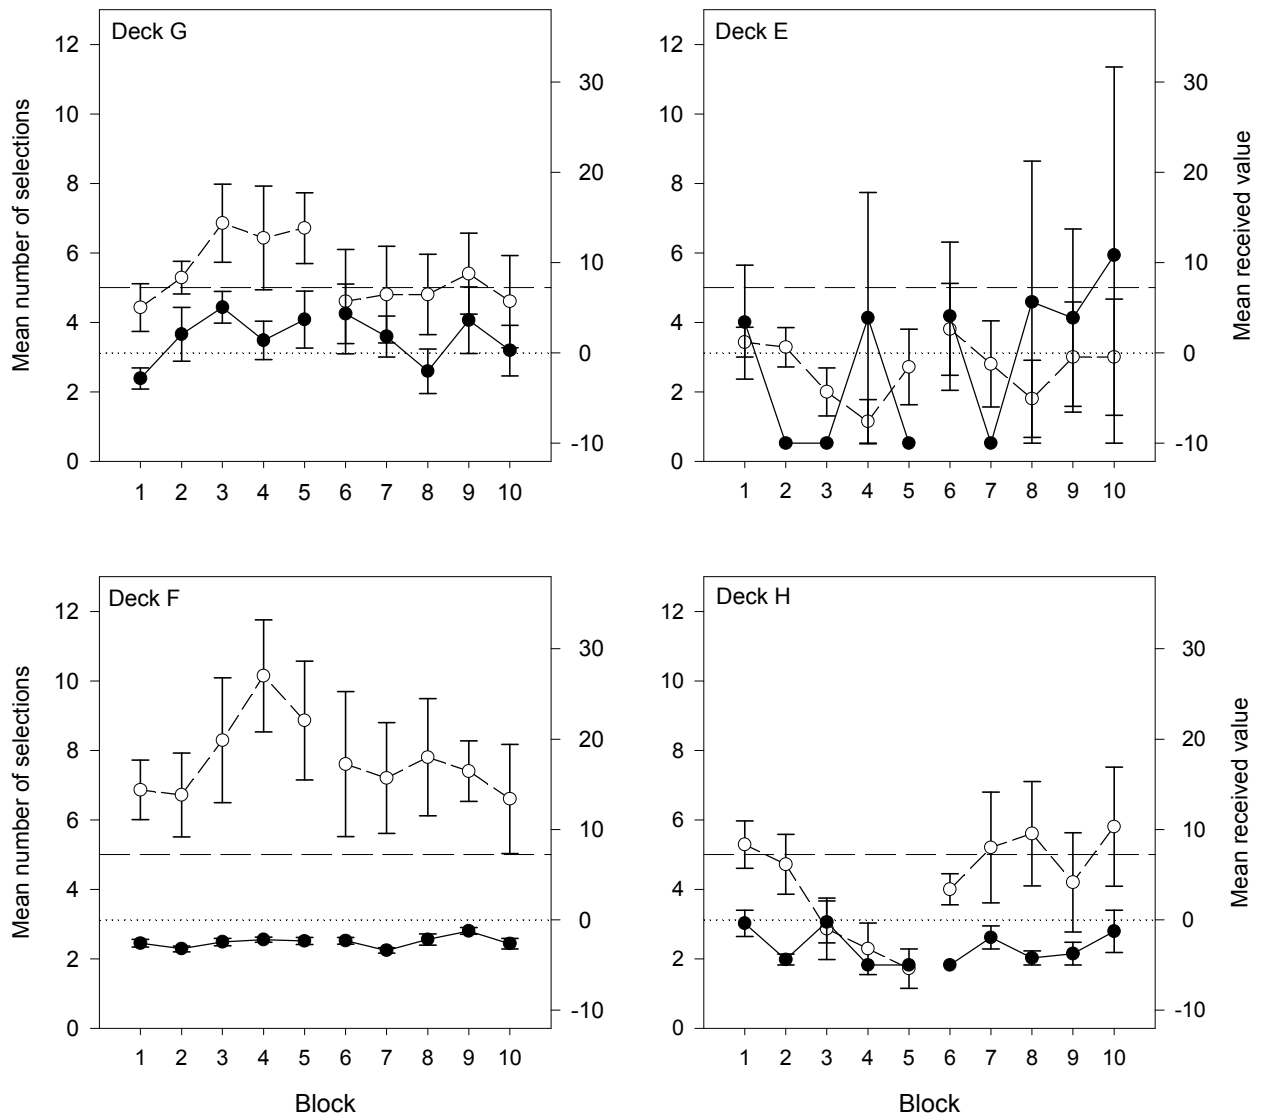


Figure 4.12: Mean received value and mean number of selections in each 20-trial block for participants with a mean net score < 0. In session 1 $n = 7$. In session 2 $n = 5$. Error bars are the standard error of the mean.

- Mean number of selections
- Selection at chance
- Mean received value
- Received value = 0

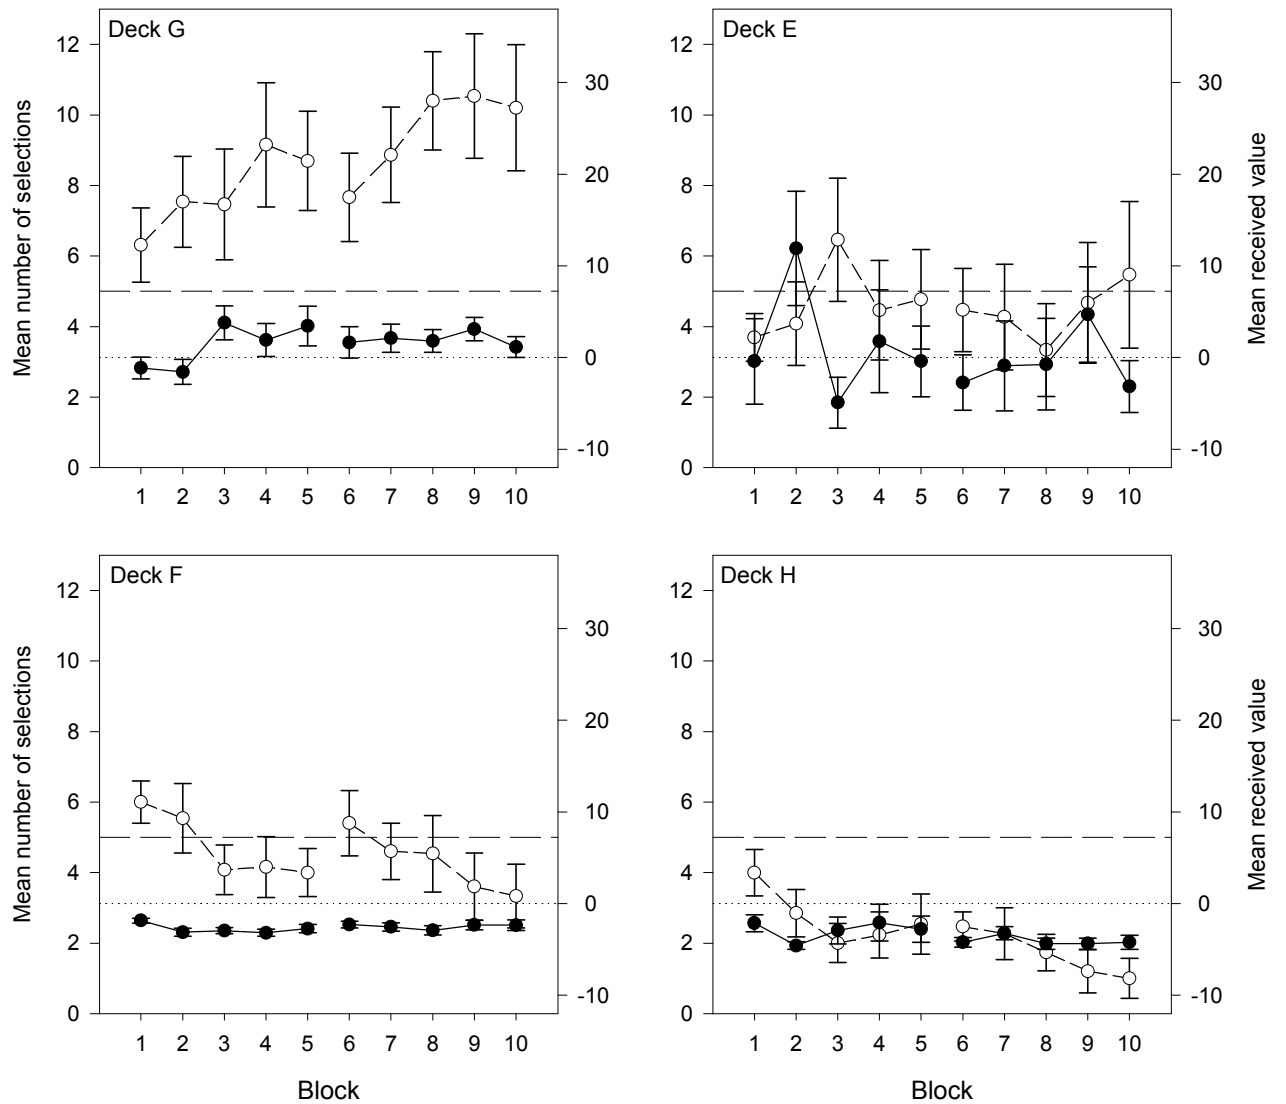


Figure 4.13: Mean received value and mean number of selections in each 20-trial block for participants with a mean net score ≥ 0 . In session 1 $n = 13$. In session 2 $n = 15$. Error bars are the standard error of the mean.

In Experiment 7 the change in win order made to the fixed reinforcement schedules particularly affected selection from one advantageous deck. This manipulation made the task more difficult for all participants as in order to obtain a win on deck E persistent selection despite mounting losses was required. So although the expected value of the deck across a block of ten trials remained the same as in the unaltered variant (Bechara et al, 2000), selection was affected. In Experiment 8 with

probabilistic reinforcement each selection from deck E carried the same probability of a win. As a result more participants encountered a win on this deck, and did so earlier in their selections from that deck than participants in Experiment 7. This resulted in selection from E being greater in Experiment 8 than in Experiment 7. Despite this, selection from deck E was still at chance and for the majority of participants deck G, with its more frequent wins, was the preferred deck. This selection behaviour shows again that selection from decks with the same (or similar) expected values is not uniform. The frequency of the informative event, the reinforcer or punishment, affects selection behaviour.

4.5 GENERAL DISCUSSION

It was hypothesised that the rapid learning of the advantageous strategy found in the gambling task literature may be due to the nature of the reinforcement schedules used. In Experiment 6, on the standard task, across participants there was no preference for the advantageous decks in the final 20-trial block. Advantageous deck preference with fixed schedule decks usually occurs after three blocks (e.g. Bechara et al, 2000; Bowman and Turnbull, 2003; Fellows and Farah, 2005). In Experiment 8, using the variant task, learning rate increased at a rate greater than chance in both sessions and using the net score measure a preference for the best decks was found in the last block of trials. However, mean net score did not rise to the asymptotic levels reported by Bechara et al (2000) until the second session. On this basis it was concluded that the probabilistic nature of the reinforcement schedules had affected learning.

On the probabilistic schedules each selection from a deck has the same probability of the infrequent informative event (the loss on the standard task, the win

on the variant task) occurring as any other. With a constrained fixed schedule as used by the Iowa group or a randomised with constraint schedule used by others (e.g. Crone et al, 2004; Crone and van der Molen, 2004) the probability of the infrequent reinforcer changes depending on the number of prior reinforcements and the number of selections remaining within that block. This means that with a probabilistic schedule predicting when the infrequent event will occur is harder. Runs of trials without the infrequent reinforcer can be longer than those that occur with a constrained schedule. One consequence of this is that in the decks with the more infrequent event (e.g. decks B or E), long sequences of cards can be selected without encountering the infrequent event. Experiment 7 found that selection from deck E was greatly below chance when the infrequent events in each deck occurred on the same fixed schedule as in the standard task. The first win from deck E was often not encountered by participants. Selection from E in Experiment 8 was close to chance, in part because more participants had received a win earlier in their selections from this deck and therefore knew such an event was possible. The results of these experiments show that the order of events within decks on the IGT is crucial to the number of cards selected from them. These figures form the units of the standard measure of learning and so support the hypothesis that learning on the IGT is influenced by the reinforcement schedules used. This claim is also supported by the increase in net score found in Fellows and Farah's (2005) manipulated version of the standard IGT.

Despite the use of probabilistic schedules increasing selection from the advantageous decks was reported on the variant task in Experiment 8. This was not found on the standard task in Experiment 6. It is possible that this may be because the variant task is easier to learn. Learning is faster on probability matching tasks where

losses are the consequences of incorrect choices (Bereby-Meyer and Erev, 1998). Kahneman and Tversky (1979) have shown that when faced with a certain loss, people prefer the option with the higher probability of winning. Translated to the gambling task this may mean that people are more readily willing to accept the larger immediate loss from selecting E or G for the chance of winning a higher amount on the variant task, than they are to accept the smaller win from selecting C or D for the chance of losing a smaller amount on the standard task. This would mean that difference in the value between loss amounts to the difference between gain amounts is larger in favour of the advantageous decks in the variant task than in the standard task. This may make the variant task easier to learn.

It has been claimed by Bechara et al (1994) that keeping track of the contingencies on individual decks is a difficult task, even for participants with higher than average IQs and memory. Part of the reason for this may be that the IGT incorporates features of various classic reinforcement schedules that combine to cloud comprehension. In turn this makes understanding what is influencing behaviour on the IGT more difficult to pin down. Peters and Slovic (2000) addressed this point by redesigning the IGT without confounding expected value with average gain and loss amounts. Table 4.10 displays the contingencies used in this task. In the standard IGT the disadvantageous decks also have the highest gains and losses associated. As many of the Figures in the preceding Experiments show, this also results in larger variances in received value for these decks. In unconfounding this relationship, Peters and Slovic were able to compare deck selections by gain (B and D versus A and C) or loss amount (A and D versus B and C) or expected value (A and B versus C and D). Their results showed that some individuals' selections were influenced by loss information. Similarly, other individuals' selections were influenced by gain information.

However, neither measure of reactivity to positive or negative events correlated with selection based on expected value and learning rate, nor improvement across trials, is not reported. Examination of Peters and Slovic's (2000) Figure 1 suggests that participants did slightly improve across trials, but performance did not reach the level of participants on the unmodified version of the IGT. In other words their modifications affected task difficulty.

Peters and Slovic's (2000) modifications did unconfound gain and loss amounts from expected value, but these modifications also altered the probability of loss among their new decks. While this feature of decks is a characteristic of the relationship between loss and gain amount and expected value, it was not entirely removed from the modified task. On three of the four decks the probability of loss was uniform ($p(\text{loss}) = 0.5$) meaning that any effect of probability loss could not be examined directly. In the fourth deck (C), one of the advantageous decks, the probability of loss was substantially lower (0.20). The hypothesis that learning is influenced by the loss probability is supported, as this was the deck most selected by participants. However, this behaviour also supports the hypothesis that expected value is important in influencing choice behaviour as the relationship between the deck contingencies mean this deck also has the highest EV (see Table 4.6). From the evidence collected in this chapter and those preceding it seems probable that the difference in probability of loss is clearer to participants than any difference in expected value, or may at least form the basis for determining any difference in expected value.

The experiments reported in this chapter and those preceding revealed differences in selection from decks with the same (or similar) expected values. This suggests that participants do not link the uniform magnitudes of immediate gains or

losses on two decks and their expected values. This may suggest that knowledge of what is happening on the task is not complete. Of course participants may possess this knowledge but are not forced to demonstrate it behaviourally because there are sufficient cards to select from only one deck. In this case they prefer the deck with infrequent losses on the standard task (deck C has infrequent actual losses) and the deck with frequent wins on the variant task. The extent of participants' knowledge has ranged from the original claim that 70% of participants had an understanding of the deck contingencies (Bechara et al, 1997a), to a more detailed examination where a minority of participants (40%) could express this level of knowledge (Maia and McClelland, 2004).

Table 4.6: Deck contingencies in Peters and Slovic's (2000) modified gambling task

	A	B	C	D
Gain amount:	50-150	150-250	50-150	150-250
Average gain amount	100	200	100	200
Probability (gain)	.5	.5	.8	.5
Loss amount	100-200	200-300	200-300	100-200
Average loss amount	150	250	250	150
EV	-25	-25	+30	+25

Another indication that participants may not fully know the contingencies of the decks is the difference in selection behaviour between sessions where net score always dips. As suggested in Chapter 2 this may reflect participants' distrust of the experimenter's claim that the task remains unchanged. This may be a more credible

position when the task is probabilistic as the appearance of infrequent events may appear to be more random. The absence of clear evidence of learning in session 2 of Experiment 6 implies that participants do not have complete knowledge of the task in session 1. This would imply that people are not learning why decks are good – i.e. developing conceptual knowledge, but reacting to the stimuli as they are presented. However, there was clear evidence of learning in both sessions of Experiment 8. The issue of what participants know about what they are learning will be addressed in Chapter 5.

LEARNING ON THE IGT FOLLOWS EMERGENCE OF
KNOWLEDGE BUT NOT DIFFERENTIAL SOMATIC ACTIVITY

5.1 INTRODUCTION & OVERVIEW

Chapter 1 introduced three assumptions the SMH makes about behaviour on the IGT. Chapters 3 and 4 demonstrated that the third assumption that decks were selected solely on the basis of their expected value was not supported by an examination of behaviour. Chapter 2 revealed that receiving more information about the decks in the form of a hint led to improved performance. This suggested that knowledge about the decks affected performance. The second assumption of the SMH about IGT behaviour is that initial selection is guided by somatic markers not available knowledge. The support for this claim came from Bechara et al's (1997) results where skin conductance responses (SCRs) were recorded as a measure of somatic activity and knowledge was periodically assessed. Bechara et al (1997a) claimed that anticipatory SCRs (aSCRs) developed in their healthy control participants before they reported any idea about a successful strategy to pursue in their deck selections on the IGT. Additionally, although 30% ($n = 10$) of controls did not develop explicit knowledge of the best strategy they still chose predominantly from the good decks by the end of the task *and* showed increased aSCRs to the bad decks. Whereas, 50% ($n = 6$) of the patients with VMPFC damage did gain an explicit understanding of the task but still chose disadvantageously and *did not* show the difference in aSCRs. Bechara et al (1997a) concluded that somatic markers, as measured by aSCRs, were necessary to choose advantageously on the IGT (assumption 1). Further, as the difference in

aSCRs preceded any “conscious knowledge” they concluded that somatic markers act as nonconscious biases that guide behaviour (assumption 2). These results arguably increased awareness of the SMH and made the IGT the widely used paradigm it is today.

Since the publication of Bechara et al (1997a) many more studies have examined SCRs during IGT performance. But few studies have questioned Bechara et al’s (1997) account of how knowledge changes during the IGT. Maia and McClelland (2004) have offered a different account and claim that healthy participants have access to knowledge sufficient to guide behaviour earlier than found by Bechara et al (1997a). However, to my knowledge, no study since Bechara et al (1997a) has investigated the relationship between somatic markers *and* knowledge during the IGT. This chapter describes such an experiment. It aims to answer two questions of importance for understanding behaviour on the IGT that also have critical implications for using the IGT as a test of the SMH: when during the IGT do SCRs that differentiate between deck types emerge, and when do participants have knowledge about the task sufficient to guide behaviour? These questions directly address the second assumption the SMH makes about IGT behaviour, or put another way they address the suitability of the IGT as a test of the SMH.

5.1.1 When does a differential aSCR to deck types emerge?

The only concerted attempt at answering this crucial question has been by the Iowa group. Bechara et al (1996) suggested that this difference emerges as healthy participants become experienced on the task. Their Figure 4 showed the mean peak SCR amplitude prior to selections from each deck. It appears that the aSCR difference between deck types emerges after approximately 10 selections each from

both decks A and B, although it is earlier for deck A. Since these decks are initially preferred these observations suggest a difference emerging possibly as early as between 16 and 20 card selections, but more realistically between 20 and 30 selections given participants' sampling from each of the four decks. However, as the observations were drawn from varying numbers of participants dependent on each participant's total selections from each deck these observations cannot provide an exact timescale for an aSCR difference. Indeed it is precisely this issue that makes any statistical analysis of differences between decks difficult. Here it is also worth bearing in mind that decks A and B are good decks to choose from until the first loss is encountered. Indeed from Bechara et al's (1996) Figure 4 it appears that the aSCR difference between deck types emerges after experience of punishment on the disadvantageous decks. As shown in Chapter 4 this is the crucial information that guides decision-making on this task.

Bechara et al (1997a) looked at change in aSCRs across the task in more detail. First they divided trials into periods depending on an individual's first encounter with a punishment (to determine the end of the "pre-punishment" period and the start of the "pre-hunch" period) and their expression of some knowledge about the task (dependent on the detail of the knowledge expressed this would end the "pre-hunch" period and start the "hunch" period or end the "hunch" period and start the "conceptual" period). These knowledge assessments were not constant but took place after the first twenty trials and then on each subsequent ten trials. Because the mean values of SCRs between deck types in each period is of great interest, the way in which these periods were created is of importance. Participants were defined as having a "hunch" if they could express the idea that decks A and B were riskier (or C and D were safer) but not articulate explicitly why. If they could detail why A and B

were riskier (or C and D were safer) they had “conceptual” knowledge. The question period when either type of knowledge was expressed determined the start of the new knowledge period and the end of the previous one (although problems with this classification will be returned to in section 5.2.3.2).

Bechara et al (1997a) reported that anticipatory SCRs for the bad decks were larger relative to the good decks and claimed that this difference emerged in normal participants approximately between trials 10 to 50. This corresponded to the “pre-hunch” period when participants did not articulate any knowledge of differences between decks. According to Bechara et al (1997a) during the pre-hunch period participants also showed no behavioural preference for either deck type. However, while healthy participants went on to show significant differences in choices from the good and bad decks, the difference in aSCR between deck types was not statistically significant in any knowledge period (although aSCRs for the bad decks, but not the good decks, were different in the hunch and conceptual periods when compared to the pre-punishment period). Despite this non-significant difference in pre-hunch aSCRs between deck types the paper generated a great deal of interest. One potential problem in the knowledge periods as determined by the method above is that the pre-punishment and pre-hunch periods may include different data points for each participant. Since the pre-punishment period ends on first encounter with a loss, when, and on which deck, this loss occurs will be different for each participant. This means that, for example, the ninth selection from deck B may be included in the pre-punishment period for one participant and the pre-hunch period for another dependent on when they make their third selection from decks A or C, or their tenth selection from deck D (see Appendix A for the fixed schedule of losses on the IGT). This effectively means that, depending on deck sampling, some participants may have

more knowledge about the rewarding contingencies of more of the decks before they encounter a loss and move to the pre-hunch period than others. It is therefore unclear what is being compared when the pre-punishment and pre-hunch period are compared. Despite these issues the core idea that emerged from Bechara et al's (1997) study was that a difference in aSCRs preceded participants' ability to express knowledge about the IGT, with all that such an assumption implies. Many other researchers have accepted this assumption as a starting point in investigations using the IGT. The validity of this conclusion will be discussed in subsequent sections.

With the importance of aSCRs on the IGT assumed, few other researchers have looked at their development across the task. Suzuki et al (2003) did consider their emergence, though not as directly as in Bechara et al's (1997) attempt. They found a difference in aSCR between deck types that did not change across their early (first 40) or late (second 40) blocks of trials. This result suggests, like that of Bechara et al (1996) that any difference in aSCRs may emerge in the first half of the IGT. The results of another study where aSCR change was plotted across blocks of trials suggested that any aSCR difference emerged in the third quarter of the task, but the task used in this study was considerably modified from the standard IGT (Jameson et al, 2004). In addition to modifying the task by using three decks (one bad, one good and one neutral), participants were also given a secondary task to load their working memory. In two of the secondary task conditions the aSCR difference between deck types resulted from a reduction in the aSCR for the good deck rather than an increase in the aSCR for the bad decks relative to the good as in Bechara et al (1997a). Kleeberg et al (2004) did examine change in both aSCR and post-punishment SCRs across trials and found them to start at a higher level and increase faster in their healthy comparison group compared to patients with MS. Faster learning (earlier

preferential selection of the advantageous decks) was also reported in the healthy controls but no correlation between the measures was reported. (Problems with such an analysis are that as the number of AB choices decreases the variance in the SCR records increases because they are drawn from fewer and fewer samples). In the comparison group the average aSCR, collapsed over all decks, increased from a flat rate between 20 and 40 card selections. This may be attributed to selections from decks A and B and the increased probability that most participants have experienced losses on both by this point in the task. Other than these studies no others chart a change in either aSCRs or post-selection SCRs across the course of the IGT. But most studies using the IGT cite and accept the version of events presented by Bechara et al (1997a).

5.1.2 When does knowledge sufficient to guide IGT selection behaviour emerge?

In addition to knowing when any aSCR difference emerges the question of when participants have knowledge about the deck contingencies sufficient to guide their choices on the IGT is of importance for any interpretation of the SMH. Damasio (1994, p. 184 and 187) has said that somatic markers can act as conscious or non-conscious biases to decision-making. But it is the stronger, non-conscious version of the SMH (Tranel et al, 1999) supported on the back of the Bechara et al (1997a) results that has caught the imagination of researchers and this interpretation that is most widely reported. In this version of the SMH somatic markers bias the decision-making environment before knowledge can have an influence.

So the issue of when enough knowledge about the task to guide behaviour emerges relative to a differential aSCR between deck types is of fundamental importance to the interpretation of the SMH and the assumptions it affords for IGT behaviour. Since Bechara et al (1997a) a number of other groups have probed participants' knowledge of IGT contingencies, although few have examined the development of that knowledge during the task. These studies are briefly reviewed below. First, more detail about Bechara et al's (1997) results is provided. Table 5.1 displays the results of their analysis of the emergence of knowledge on the IGT. Bechara et al (1997a) found that on average, healthy participants entered the "hunch" period by the fourth questioning (after trial 50, although the range was between trials 30 and 80) and the "conceptual" period by the seventh questioning (following trial 80 with a range of 60 to 90). All healthy participants achieved "hunch" knowledge, but 30% ($n = 3$) did not reach "conceptual" knowledge. This is coincidentally also the proportion of healthy participants who do not show "normal" behaviour in later studies (Bechara and Damasio, 2002).

Table 5.1: Summary of participants' knowledge expression in Bechara et al (1997a).

% participants who did not reach the hunch period:	0
Average trial number in which participants reached the hunch period:	50 (30 – 60)
% participants who did not reach the conceptual period:	30
Average trial number in which participants reached the hunch period:	80 (60 – 90)

Note: Figures in parentheses are the range of observations.

Maia and McClelland replicated Bechara et al's (1997) study and found broadly similar results when participants were questioned in the same way. However, the attempt to replicate was not straightforward due to the lack of detail about how Bechara et al assessed knowledge and categorised it into two of their four (hunch and conceptual) knowledge periods. Maia and McClelland (2004) developed a detailed solution to resolve this that resulted in a decision tree (reproduced here as Figure 5.1) to categorise each participants' knowledge at each question period into one of the six knowledge categories possible on the IGT. These are: no professed knowledge, incorrect or incomplete hunch/knowledge, partial hunch, hunch, partial conceptual and conceptual. Even with this decision tree there were still several ways knowledge could be assessed in order to integrate it into Bechara et al's knowledge periods. This integration is effectively along two axes. The first concerns whether knowledge expressed about only one of the good decks is included as conceptual knowledge (partial conceptual). In a strict interpretation of Bechara et al's criteria partial conceptual knowledge would not count as conceptual knowledge because it is not full understanding of both good decks – Maia and McClelland (2004) called this grouping “both”. In the “partial” grouping partial conceptual knowledge is included in the conceptual period.

The second axis in integration of the two knowledge assessment systems concerns when participants first show any level of knowledge. A conservative approach would only count knowledge expressed consistently throughout all question periods from the one where it was first expressed through each subsequent questioning i.e. if upon reaching one level of knowledge the participant never returned to a lower state of knowledge. An aggressive interpretation would allow an earlier expression of knowledge to be counted even if later questioning revealed that this

level of knowledge was no longer being expressed at a later question period. Table 5.2 summarises Maia and McClelland’s (2004) results from their replication condition along both these axes. Their aggressive, “partial” grouping best fit Bechara et al’s (1997) results. However, Maia and McClelland focused on the “both” grouping as it more reflected Bechara et al’s (1997) classification of conceptual knowledge. Using these figures and an aggressive approach, Maia and McClelland reported that, like Bechara et al, their participants preferentially selected the advantageous decks when they were classified as being in the hunch or conceptual knowledge periods.

Table 5.2: Summary of participants’ knowledge expression in Maia and McClelland’s replication condition.

	Approach	
	Conservative	Aggressive
% participants who did not reach hunch period:	37.5 (5.8)	12 (2.9)
Average trial number in which participants reached the hunch period:	62 (5.8)	43 (4.6)
<u>Partial Grouping</u>		
% participants who did not reach conceptual period:	47 (7.6)	25 (8.7)
Average trial number in which participants reached the conceptual period:	74 (1.6)	62 (6.6)
<u>Both Grouping</u>		
% participants who did not reach conceptual period:	77 (5.8)	60 (5)
Average trial number in which participants reached the conceptual period:	91 (5.6)	72 (4.8)

Note: Figures in parentheses are the standard deviation.

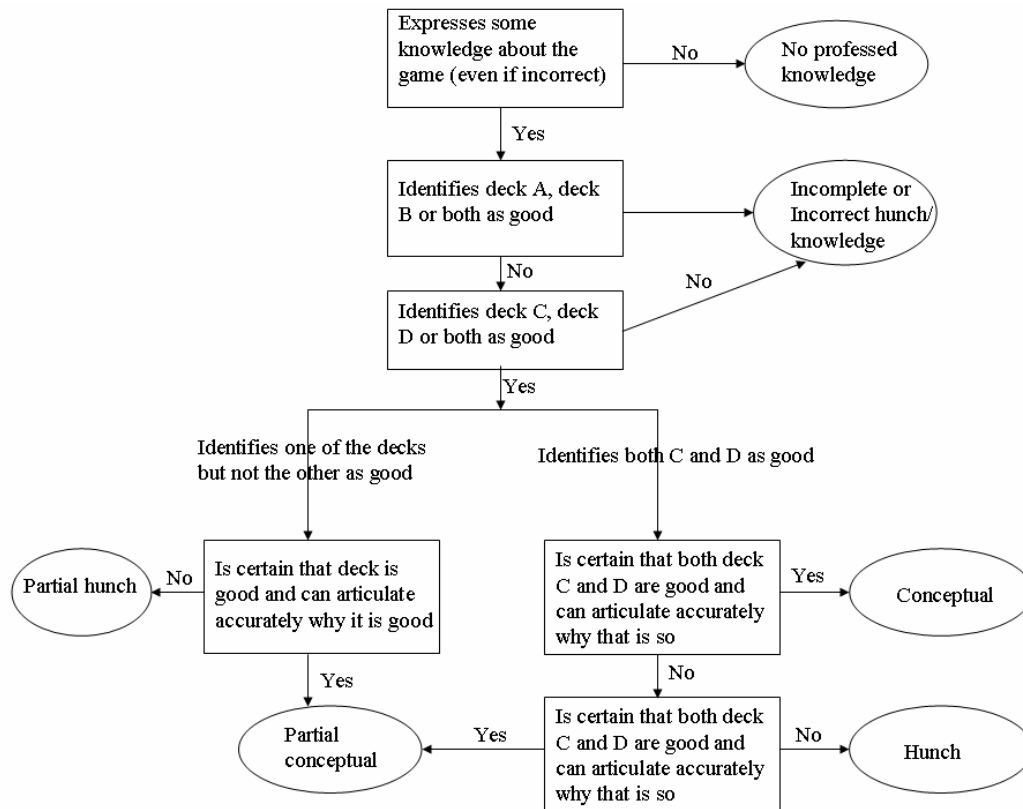


Figure 5.1: Maia and McClelland's (2004) decision tree for classifying participants' knowledge about the IGT.

The main thrust of Maia and McClelland's (2004) paper was not merely to replicate their results and specify their method in greater detail. They also challenged the ability of this method as an effective probe of participants' knowledge. This challenge stemmed from criticisms applied to the large literature on implicit learning, where the question of what knowledge participants possess is fundamental. This literature was extensively reviewed and critiqued by Shanks and St John (1994) who developed criteria necessary for determining whether learning without awareness has occurred. This was essentially the claim made by the Iowa group (Bechara et al, 1997a; Tranel et al, 1999). But Maia and McClelland (2004) pointed out that the methods used in Bechara et al (1997a) to assess knowledge (learning) failed to fulfil Shanks and St John's "Sensitivity" criterion. This concerns whether the test of

awareness reveals all the relevant knowledge that a participant consciously possesses. Maia and McClelland contended that the broad, open-ended questions used by Bechara et al (1997, “Tell me all that you know about what is going on this game” and “Tell me how you feel about this game”) may not result in participants divulging information that they still consider to be tentative, or “may not reliably cue recall of all relevant knowledge” (Maia and McClelland, 2004, p.16075). They also contended that such questioning may be influenced by factors like an individual’s personality or their level of engagement with the task. Because of these possibilities Maia and McClelland speculated that participants could have access to knowledge sufficient to guide their behaviour much earlier than revealed using Bechara et al’s questioning. As a result they developed a more focused questionnaire designed to probe all knowledge about the IGT. They found that participants were able to report a level of quantitative and qualitative knowledge that would have been sufficient to guide their behaviour from the first period of questioning. Indeed, the level of knowledge in many cases preceded advantageous selection behaviour. Despite the more detailed questioning a comparison between the behaviour of participants in this and the Bechara et al (1997a) replication condition found no significant differences. This ruled out the possibility that more detailed questioning cued participants to relevant information available in the task (otherwise learning would have been positively affected, the goal of the task being to earn as much money as possible). Thus, Maia and McClelland’s study undermines Bechara et al’s claim that the difference in aSCRs precedes knowledge on the IGT and indeed invites the suggestion that the aSCRs are a product of the conscious knowledge – what I termed in Chapter 1 a Knowledge-Somatic hypothesis to challenge the Somatic-Knowledge assumption prevalent in the IGT literature since Bechara et al (1997a). Maia and McClelland’s

findings also challenge the claim from the Iowa group that the IGT remains cognitively impenetrable even to those with above average memory and IQ.

In support of Maia and McClelland's (2004) findings, Evans, Bowman and Turnbull (2005) also found that participants could rate each deck's "goodness" and "badness" at above chance levels after 20 trials using a deck rating scale similar to one of Maia and McClelland's (2004) measures. Other researchers who have probed participants' knowledge of the deck contingencies on the IGT have primarily used post-test questioning. However, this method cannot reveal anything about *when* participants possess adequate knowledge to guide behaviour and fails to fulfil Shanks and St John's (1994) Sensitivity criterion. Nevertheless this method has been employed. Suzuki et al (2003) asked participants to rate each deck's riskiness on a 7-point scale but found no link between what they called "conscious knowledge" and two groups split on their post-selection SCR levels. This is troublesome for the Knowledge-Somatic hypothesis as it implies no relationship between knowledge and SCR levels. However, as previously pointed out any post-hoc questioning cannot inform on when awareness arises, especially in this case where group membership was also determined post-hoc. One cannot say if those in the high SCR group had awareness of deck riskiness any earlier than the low SCR group. This cannot be ruled out either and so the Knowledge-Somatic hypothesis remains viable.

Elsewhere, post-task questioning has revealed mixed accounts of participants' knowledge. Kleeberg et al (2004, p.794) informally asked one of their groups of interest (patients with MS) at the end of testing which decks it was best to avoid and state that they were generally correct. The aSCRs of less neurologically impaired patients increased across the task and they made fewer disadvantageous deck selections than less able patients. The authors concluded that since knowledge

equated between patient groups, but somatic activity did not, cognitive appraisal is not sufficient to account for advantageous IGT behaviour. However, the authors did not assess when individuals possessed any level of knowledge, or whether they acted upon it. Instead they claimed that patients with more neurological dysfunction continued to make disadvantageous choices between cards 30 to 60 because they were not guided by adequate somatic markers and so did not reach the “hunch” phase. They only began to select advantageously when “reason and comprehension became dominant”. The claim here is that the somatic markers are necessary to guide advantageous decision-making until comprehension of the task is achieved. But this claim cannot be supported on the basis of post-hoc questioning.

Fishbein, Hyde, Eldreth, London et al (2005) found that several participants from their control and substance abusing groups did not express knowledge of the correct strategy. They suggested that this absence of knowledge might explain the absence of any SCR differences between deck types or participant groups and is effectively the Knowledge-Somatic assumption. Fischer et al (2005) found overall poor behavioural performance and anecdotally reported poor knowledge of the task. They attributed the absence of learning in the majority of their participants, in part, to using participants from a technical college suggesting that they may have been too used to using reasoning processes to guide behaviour. In other words, they did not pay attention to their somatic response during the task, and this was the root of the absence of evidence of learning in their behaviour and anecdotal post-test verbal reports. This learning system hypothesis was tested by Evans, Kemish and Turnbull (2004) who found greater advantageous selection in the final two twenty-trial blocks in female participants who had left school when aged 16 compared to female age-

matched university students. This difference was attributed to better emotion-based learning in the early school leavers as a result of less time in formal education.

None of these attempts to assess knowledge can offer any replication of Bechara et al's (1997) description of the emergence of knowledge about the IGT. Nor do they challenge it and indeed most studies that do not assess knowledge repeat the Somatic-Knowledge assumption advanced by these authors. The only study that has extensively examined the progression of knowledge through the IGT is Maia and McClelland (2004). Their finding that participants possessed knowledge adequate to guide behaviour earlier than claimed by Bechara et al (1997a) seriously undermines the strong version of the SMH and the Somatic-Knowledge assumption of IGT behaviour.

The role aSCRs have, as an index of somatic markers, in decision-making on the IGT has been discussed in Chapter 1. While some studies have suggested SCR activity and IGT performance are related (Bechara et al, 1997a, 1999, 2000, 2002; Crone et al, 2004) others have failed to find a link (Tomb et al, 2002; Suzuki et al, 2003; Campbell et al, 2004; Kleeberg et al, 2004). Where knowledge has also been assessed it has generally been through post-hoc questioning and thus can offer no information about when knowledge appears. In such cases any attempt to link knowledge and somatic activity cannot differentiate between a Somatic-Knowledge hypothesis and a Knowledge-Somatic alternative.

5.1.3 Summary

This introduction addressed two questions of importance for understanding behaviour on the IGT: when during the IGT do SCRs that differentiate between deck types emerge, and when do participants have knowledge about the task sufficient to guide

behaviour? These questions also have critical implications for using the IGT as a test of the SMH. The widely cited conclusion of Bechara et al (1997a) that differential aSCR activity precedes the emergence of knowledge has some support in studies correlating aSCR activity with behavioural performance. However, such relationships have not been universally found and indeed many results offer conflicting evidence. Principally, Maia and McClelland (2004) have shown that knowledge sufficient to guide behaviour exists very early in the task. Additionally, several studies (Hinson et al, 2003; Jameson et al, 2004; Turnbull et al, 2003) have shown that impairments in executive components of working memory detrimentally impact on IGT performance suggesting that differences in aSCRs are driven by cognitive processes (implying knowledge) rather than vice versa. Experiment 9 considers these three dependent measures (behavioural performance, SCR activity and participants' knowledge) in an attempt to untangle the relationship between them.

5.2 EXPERIMENT 9

DIFFERENTIAL SKIN CONDUCTANCE RESPONSES BETWEEN DECK TYPES EMERGE AFTER PARTICIPANTS DISPLAY KNOWLEDGE ON THE IOWA GAMBLING TASK.

5.2.1 Introduction

Experiment 9 is a replication of Maia and McClelland's (2004) study with the addition of SCR recording. Maia and McClelland's study was important because participants possessed consciously available knowledge about the IGT deck contingencies at the first time of questioning (after 20 trials). This result undermined Bechara et al's (1997) claim that knowledge developed later (after ~50 trials)

allowing an explanation utilising somatic markers to explain improving performance. However, Maia and McClelland's (2004) study was not a full replication of Bechara et al (1997a). Few studies have investigated SCR change across the IGT and none since Behara et al (1997) have also attempted to assess knowledge. It therefore remains a possibility that differential SCR activity between deck types precedes consciously accessible knowledge. Such a finding would support the assumption of causality of the strong version of the SMH to explain IGT performance.

Alternatively, as Maia and McClelland's (2004) questioning was more specific than Bechara et al's (1997) it is also possible that differential SCR activity develops after consciously accessible knowledge. This would support a Knowledge-Somatic explanation for somatic activity.

5.2.2 Method

Participants

Thirty-two participants who were predominantly post-graduate students were recruited from the University of Nottingham community. They were recruited using a poster advertisement, an online advert and via direct email to members of a participant pool. Participants were told that they would be participating in a really interesting cognitive task and have the opportunity to earn up to £12. They were told that some physiological measures would be recorded and that the experiment took approximately one hour. Sixteen participants were randomly assigned to each question group (the General questions of Bechara et al, 1997a; or the Specific questions of Maia and McClelland, 2004). The mean age was 25.68 ($\sigma_M = 1.22$) in the specific question group and 24.63 ($\sigma_M = 0.92$) in the general question group.

There were nine and seven male participants in the Specific and General question group respectively.

Apparatus - Behavioural task

The computerized version of the IGT with the hint instructions and real money incentives was used. Breaks in the behavioural task occurred after the first twenty and each subsequent ten trials so that participants' knowledge could be probed using the condition-specific questions. More detail on these are provided in the next section. As a result of the addition of questionnaires and skin conductance recording the time to complete the task was increased from previous Experiments (to around one hour). To reflect this increase the reinforcers were also increased to maintain the relative incentive per hour across Experiments. As this experiment took on average four times longer than the previous purely behaviour studies (although it was longer in the specific question group), reinforcer values were also four times the amount of previous experiments. Therefore wins increased from 10p to 40p in decks A and B, and from 5p to 20p in decks C and D. All values for losses increased similarly.

Apparatus - Knowledge probes

The administration and structure of the questionnaires followed the procedure of Maia and McClelland (2004). The task was interrupted after twenty trials and thereafter after every ten trials when instructions on the computer screen informed participants that they would now be asked some questions about the task. In the Specific question group participants were given the detailed questionnaire as used in Maia and McClelland (2004; see Appendix D). The questionnaire was computer-based and

required selection of options using the mouse or entry of answers using the numerical keypad.

Participants in the General Question Group were presented on subsequent screens with the two questions used by Bechara et al (1997a): “Tell me all that you know about what is going on in this game” and “Tell me how you feel about this game”. Participants’ responses were recorded using a tape recorder operated by the experimenter who sat behind a large room-dividing screen in the same room as the participant. The presence of an experimenter during task performance was a deviation from the procedure used in the previous Experiments but was necessary in order to monitor the skin conductance record and, in this condition, to operate the tape recorder. The questions were presented on the computer in an attempt to minimise any experimenter influence and to equate the two question conditions. Interaction with the experimenter was kept to a minimum and was initially restricted to prompting participants to answer the question before them. However, some participants’ answers were so minimal that some additional prompting was occasionally required. In the main this took the form of directing participants’ answers to their knowledge of the decks.

The presentation and cessation of the questions in both conditions was accompanied by a computer beep to mark the beginning and end of the question period on the skin conductance record (more information on why this was necessary is provided in section 5.2.2.5), and to inform the experimenter when to start and end the tape recorder in the General question condition.

Apparatus – Analysis of General group transcripts

Participants' question answers were transcribed from the tape recording. Three individuals naïve to the experimental hypothesis were recruited and paid to assess the transcripts and classify the knowledge displayed at each question period using Maia and McClelland's (2004) decision tree (Figure 5.1). The assessors first undertook training on the decision tree using sample answers created to cover all possible outcomes from the tree. One hundred percent accuracy was required before the actual transcripts were assigned. When sample transcripts were not correctly rated the assessor was told and asked to try again. Most raters accurately rated each transcript on their first attempt. Rarely were three attempts required, but following correct answers the assessor had to convince the experimenter of why they had reached the assessment they had.

Once the actual transcripts had been assessed the assessors met to compare results. If there was disagreement on any participant's answer the assessors were instructed to debate their disagreement until a unanimous decision was reached. If this was not possible a majority decision for that answer was used. These final assessments of participants' answers were used to determine when knowledge was displayed in the General question group.

Apparatus - Electrodermal activity

A BIOPAC Systems MP30 system running on a Macintosh computer was used to record electrodermal activity. Skin conductance was recorded at 10Hz using two Ag/AgCl electrodes connected to the volar surfaces of the medial phalanges on participants' index and middle fingers of the left hand (all participants were right

handed). Because the MP30 system does not have the facility for a direct link between the recording computer and the task presentation computer, marking the occurrence of events was achieved by recording the sounds produced on the task presentation computer during the task. These sounds were recorded by the MP30 via an analogue input. As described in Experiment 1, during the task gains and losses were accompanied by concurrent auditory stimuli which also served as markers for events in this experiment. Additionally, the experimenter marked the skin conductance record when an event occurred. However, this measure is less reliable and not as temporally accurate. For this reason it was only referred to when the auditory record was ambiguous about when an event occurred.

Skin conductance activity analysis

Skin conductance responses were analysed using the Student Lab Pro software for the MP30 system. The first step in the analysis was the removal of the downward drift in the SCR record. A mathematical transformation provided by the Student Lab software was used to remove it prior to analysis. This “difference” transformation measures the difference in amplitude between two data samples separated by a particular number of points (in this case it was 10). The difference is then divided by the time interval between the two samples.

The SCRs were analysed using the area under the curve measurement. This measurement calculates the total area between a waveform and a baseline value within the endpoints of a selected area. In effect a line is drawn between the user defined start and end points of the waveform. For anticipatory SCRs this was the five seconds prior to deck choice as determined by the auditory signal’s mark on the analogue channel. For post-selection SCRs the start point was one second after this

marker and the end point was again five seconds later. These area under the curve measurements were then divided by the time interval to give a value in amplitude units per second ($\mu\text{S}/\text{second}$).

Design

The experiment was a replication of Maia and McClelland's (2004) study with the addition that skin conductance responses were measured. A mixed-design was used with Question Group (General or Specific) a between-subjects factor, and block of trials a within-subjects factor. Three dependent measures were obtained: participants' behaviour on the IGT, participants' knowledge of the task contingencies, and the change in participants' physiological arousal prior to card selection (aSCRs) and following card selection (r or pSCRs; reward or punishment SCRs).

Procedure

On arrival for testing participants were given a brief description of the task, an account of what was involved in the recording of electrodermal activity, and in the General Question Group, information about the recording of their answers using the tape recorder. These participants were told that questions would appear on the computer screen periodically throughout the task and they must speak their answers into the tape recorder. It was emphasised to all participants that the experimenter would not interact with them nor answer any questions about the task after the opportunity to ask them had ended (following their acknowledgement that they understood the task instructions). Informed consent was obtained from all participants.

The index and middle fingers of participants' left hands were cleaned using an alcohol free wet-wipe. Once dry an isotonic (0.5% saline) gel (Biopac Gel 101) was rubbed into the skin of the medial phalanges of the index and middle fingers of participants' left hand before the MP30 electrodes were attached. Participants were instructed that it was important to stay as still as possible throughout the experiment and to make themselves comfortable so that they only moved their right hand when controlling the mouse, and in the Specific Question Group, when they entered answers using the keyboard.

Participants then read the task instructions. These were exactly the same as for previous experiments with the addition of information about the periodic interruptions in which questions would be asked (for full task instructions see Experiment 1). A period of at least five minutes was allowed to elapse from electrode attachment to task commencement to allow the electrode gel time to be absorbed into each participant's skin. During this time participants were informed that the experimenter would be present in the room but would not be monitoring their performance. Participants were told that the purpose of the experimenter's presence was to monitor the SCR record and, in the General Question Group, to operate the tape recorder when required. They were told that there would be no interaction with the experimenter except if, in the Specific Question Group, clarification was needed on the terms used in the questionnaire. Participants were then reminded that the most important thing was to earn as much money as possible, or to avoid losing as much as possible.

SCRs were recorded without interference until the task ended. When visual inspection indicated that SCRs were present, participants were instructed to begin the task. Participants saw the IGT screen as displayed in Figure 2.1 except that a message

displayed in the middle of the screen instructed the participant to consider from which deck they would choose. The mouse pointer was not displayed and the decks could be selected while this message was on-screen. It remained for 5 seconds when it was replaced with another telling participants to “Please select a card”. The mouse pointer re-appeared and the decks became active. The five seconds prior to deck choice constituted the period during which SCRs were considered to be anticipatory.

Following the selection of a card the computer displayed the amount won accompanied by the sound of a man shouting “Yippee!” This sound was marked on an analogue channel of the SCR record and allowed the accurate pinpointing of SCR events in relation to deck choices. One second after the reward, the amount lost was displayed accompanied by the sound of a man shouting “Doh!”. The reward and loss information remained on-screen for five seconds. The instruction to “Consider your next choice” was then displayed for five seconds before participants were again instructed to choose a card. SCRs in the five seconds following deck selection were considered to be post-selection SCRs. Therefore, the inter-trial interval was at least twelve seconds but varied depending on how long participants took to choose their next card following the instruction to do so.

The experiment concluded following 100 trials on the IGT and when participants’ task knowledge had been probed nine times. The length of the Experiment differed between participants and was dependent on the speed with which they selected cards and answered the questions. As there were more questions in the specific question group these participants tended to take longer. The experiment took around one hour but this depended on the speed with which participants answered questions and made selections. So although participants were told the prospective length of the task this information could provide no hint about when it would end.

On completion of the task all electrodes were removed and participants were fully debriefed. Each participant received the amount they had earned on the task plus an additional £2. As in previous experiments participants were unaware of the additional payment and were asked not to mention it to anyone else.

5.2.3 Results

Behavioural data

Mean net score was calculated by subtracting the number of cards selected from decks A and B from the number selected from decks C and D. Additionally, net score was calculated in each block of ten trials for each participant. This differs from previous experiments where blocks of twenty trials were used. The reason for this change was twofold: to examine in more detail when the change in behaviour emerges, and as the task was interrupted after every ten trials for questioning this blocked the trials together for participants anyway.

Mean net score for the General Question group was 20.44 ($SD = 22.06$). A one sample t -test found that this was significantly greater than zero, $t(15) = 3.93$, $SD = 22.06$, $p < 0.01$ indicating that participants in this condition showed an overall preference for the advantageous decks. The same was true of participants in the Specific Question group. Their mean net score was 28.56 ($SD = 29.04$) and this was significantly greater than zero, $t(15) = 3.71$, $SD = 29.04$, $p < 0.01$.

Mean net score was calculated for each block of ten trials and compared between Question Group and across Block. Figure 5.2 displays this comparison. It can be seen that net score increases across block at a similar rate in both Question Groups. A mixed-design ANOVA revealed no main effect of Question Group, $F(1,$




Figure 5.2: Mean net score across 10-trial blocks. The closed circles represent the Specific question group and the open circles represent the General question group. Error bars are the standard error of the mean.

29) < 1. There was a main effect of Block, $F(4.31, 124.89) = 15.43$, $MSE = 44.26$, $p < 0.01$, that reflects the increase in mean net score with more trials, but no interaction, $F(4.31, 124.89) = 1.53$, $MSE = 29.0$, $p > 0.05$ indicating that learning proceeded at a similar pace in both question groups. This was confirmed by estimating learning rate, using the slope b , for each participant and comparing between groups. Mean learning rate was 0.88 ($SD = 0.72$) in the Specific question group and 0.77 ($SD = 0.63$) in the General question group. An independent-samples t -test found no significant difference between them, $t(30) = 0.44$, $p > 0.05$. However, Lorch and Myers (1990)

regression analyses revealed both learning rates were significantly greater than 0: Specific group $t(15) = 4.90, p < 0.01$; General group, $t(15) = 4.90, p < 0.01$. These results indicate that learning progressed at the same rate in both Question Groups. This is important because it shows that the nature of the questions participants received did not differentially affect their behaviour. Maia and McClelland (2004) found the same result.

Knowledge of the task: General Question group

The overall impression gained from participants in the General Question group was that most struggled to achieve any comprehension of the task. Indeed, on many occasions the experimenter had to probe participants to elicit some answer to the standard questions. This took the form of asking what each participant knew about the decks, without leading them directly to the knowledge the experimenter was seeking. This was a measure of the generally unsatisfactory nature of this questioning method, a point returned to in the discussion.

Despite these limitations the independent ratings suggested at least half the participants reached the conceptual period. This result was more like Bechara et al's (1997) than that found by Maia and McClelland (2004). However, the results do depend on the method of classifying conceptual knowledge. Like Maia and McClelland (2004) the aggressive approach provides the best fit to Bechara et al's (1997) data and the discussion that follows in this paragraph will refer to this approach only. However, unlike Maia and McClelland, in this experiment the 'partial' rather than the 'both' grouping of conceptual knowledge best matched Bechara et al's data. The summary of the raters' knowledge assessments is presented in Table 5.3. In classifying knowledge aggressively all but one participant displayed

Hunch (or in Maia and McClelland's terms level-1) knowledge and this occurred on average after 43 trials. This compares favourably with both Bechara et al (all participants had hunch knowledge on average by trial 50) and Maia and McClelland (88% of participants showed hunch knowledge on average by trial 43). However, unlike Maia and McClelland, the 'partial' grouping for conceptual knowledge fit Bechara et al's data better than the 'both' grouping. In this case only around 30% of participants (versus 62.5% using the conservative approach) failed to exhibit conceptual (or level-2) knowledge. Bechara et al's figure was also 30% and there conceptual knowledge was achieved on average by trial 80. Using either grouping methods and an aggressive approach, conceptual knowledge was achieved substantially earlier on average in the current study (by 53 or 55 trials for the 'partial' and 'both' groupings respectively, although by different numbers of participants). Maia and McClelland also found that the 'partial' grouping resulted in the majority of participants (~75%) being classified as having conceptual knowledge and on average this occurred by trial 62. However, they used the 'both' grouping when comparing their results to Bechara et al's. With the current data, the 'both' grouping would decrease the proportion of participants with conceptual knowledge to 50%. However, the 'both' grouping does seem more in keeping with the idea of conceptual knowledge. If this grouping were used and a conservative approach taken one would find that the vast majority of participants did not reach the conceptual period and if they did it was only established by the final question period. This conclusion most accurately reflects the subjective impressions of the experimenter but conflicts with the conclusions of both Bechara et al (1997a) and Maia and McClelland (2004).

Table 5.3: Knowledge assessment for General question group

	<u>Approach</u>			
	<u>Conservative</u>	<u>Aggressive</u>		
% participants who did not reach the hunch period:	50	6.25		
For participants who reached the hunch period,				
average trial number in which they did so:	73 (6.8)	43 (3.9)		
	<u>'Partial' Grouping</u>		<u>'Both' Grouping</u>	
	<u>Conservative</u>	<u>Aggressive</u>	<u>Conservative</u>	<u>Aggressive</u>
% participants who did not reach the conceptual period:	62.5	31.25	87.5	50
For participants who reached the conceptual period,				
average trial number in which they did so:	83 (9.2)	53 (6.2)	100 (0.0)	55 (8.0)

Note: Average trial values are rounded to the nearest trial. Values in parentheses are the standard error of the mean.

Knowledge of the task: Specific Question group

Three measures of knowledge were obtained for each deck at each question period: a deck rating from -10 to 10, an estimate of the average net amount won or lost on the deck and a calculated net amount based on participants' estimates of how much they would win, how often they lost, and how much that average loss was. In addition participants were also asked which deck they would choose if they only had one choice. This last measure and the qualitative deck ratings were used to assess level-1 knowledge, equivalent to Bechara et al's (1997) "hunch". Level-2 knowledge (equivalent to Bechara et al's "conceptual" knowledge) was determined by the answers to the quantitative estimate questions. Maia and McClelland (2004) pointed out that it was possible in the early question periods that one of the 'bad' decks in the long-term would be a 'good' deck at that time because the participant had not yet lost sufficient amounts on that deck to make the actual net amount from that deck negative. Thus for Maia and McClelland's analysis, knowledge of the best decks up to that question period was based on the actual net amounts won or lost from each deck at that time. Nevertheless, in order for a direct comparison to be made with the General Question condition, knowledge here is first examined assuming decks C and D are the best decks throughout the task. Table 5.4 displays the proportion of participants who attained either hunch level or conceptual level knowledge for each measure. Also displayed is the mean trial on which this knowledge was displayed.

Maia and McClelland's (2004) analysis of knowledge, at both the hunch and conceptual level, relied only on an identification that *one* of the best decks is better, or in other words the 'partial' grouping used in the analysis of conceptual knowledge in the General question condition. In calculating when knowledge first appears across participants, both an aggressive and a conservative interpretation of their knowledge

can be made. With an aggressive interpretation the earliest display of knowledge is used regardless of what later questioning revealed, whereas consistent demonstration of knowledge is required for the conservative interpretation. Maia and McClelland (2004) did not consider this distinction in their analysis as they focused on each question period alone. Such a method gives an overall picture of the change in knowledge but the interpretation is different from the more global analysis of participants' knowledge applied with the General questions. If one is seeking to establish when participants *first* show knowledge, the aggressive/conservative definition of what demonstrates knowledge is important.

The left-hand columns in Table 5.4 show this distinction with a 'partial' grouping. As with the General questions hunch level knowledge emerges earlier using the aggressive approach and all participants demonstrated knowledge at this level. This was true with deck ratings and identifying the best deck. With a conservative approach some participants did not display any hunch level knowledge (20% using the deck ratings measure; 6.25% for the best deck measure). A similar pattern is found for the conceptual level knowledge measures – using an aggressive approach resulted in more participants being categorized as displaying knowledge earlier than a conservative approach. For the expected net measure, using an aggressive approach, only 12.5% ($n = 2$) of participants failed to display conceptual level knowledge whereas using a conservative approach this proportion was more than double (31.25%). An aggressive approach also resulted in conceptual level knowledge being found earlier (after 41 trials versus 51 with a conservative approach). The same pattern occurred for the calculated net measure, although many fewer participants were classed as displaying conceptual level knowledge (25% and 50% for aggressive and conservative approaches). Those that did display knowledge

on this measure did so later than revealed by the expected net measure. With a conservative approach knowledge was present at approximately trial 64 whereas with an aggressive approach it was around trial 37. The data in Table 5.4 show that using the Specific questions and the aggressive approach that provided the best fit to previous results for the General questions, a larger percentage of participants possess hunch level and conceptual level knowledge earlier than is found when General questions are used.

In the analysis of the General questions a second grouping method was described to classify differences within participants' conceptual knowledge. To possess conceptual knowledge in the 'both' grouping required participants to display knowledge of why both of the best decks were better, whereas the 'partial' grouping classified similar knowledge for only one deck as conceptual. With the current analysis of the Specific questions it is possible to use a grouping that corresponds to this 'both' grouping. In such a grouping knowledge that the two best decks at the time of questioning are the best decks is required (decks C and D in the current examination). Thus the 'both' grouping is a more stringent criterion than the 'partial' grouping used by Maia and McClelland. The relevant data is displayed in the right-most columns of Table 5.4 and the stringency of this grouping is reflected across all measures as conceptual level knowledge was found later and in fewer participants than if the 'partial' grouping is used.

Table 5.5 explores the data when actual received values at each time of questioning are considered. In this case, decks A and B are the best decks until such time as losses on them exceed gains. The comparison between tables is informative in some important ways. First, and most importantly, more participants show evidence of knowledge of the best decks earlier when the actual received values are used to

Table 5.4: Knowledge assessment for Specific question group using ‘partial’ grouping (either C or D had to have received the best score on each measure) or ‘both’ grouping (C and D had to have received the best scores on each measure).

Question Type	‘Partial’ Grouping		‘Both’ Grouping	
	Conservative	Aggressive	Conservative	Aggressive
<u>Ratings</u>				
% participants who did not reach Hunch level:	25	0	50	12.5
average trial number in which they did so:	47 (5.7)	36 (4.4)	60 (7.1)	33 (3.8)
<u>One deck</u>				
% participants who did not reach Hunch level:	6.25	0	<i>as ‘partial’ grouping as there is</i>	
average trial number in which they did so:	49 (7.2)	29 (3.2)	<i>only one response possible</i>	
<u>Expected net</u>				
% participants who did not reach Conceptual level:	31.25	12.5	56.25	25
average trial number in which they did so:	51 (6.4)	41 (2.9)	56 (6.9)	52 (5.1)
<u>Calculated net</u>				
% participants who did not reach Hunch level:	50	25	68.75	31.25
average trial number in which they did so:	64 (9.2)	37 (3.6)	72 (12.4)	44 (4.9)

Note: Average trial values are rounded to the nearest trial. Values in parentheses are the standard error of the mean.

determine the best decks. This is certainly the case for the aggressive approach and reflects participants' early identification that decks A or B are generally profitable in the first question period. When the aggressive approach is used with a partial grouping participants, on average, demonstrate hunch level knowledge at the first question period (trial 20) and conceptual level knowledge by the second (trial 30).

Second, when a conservative approach is used the differences between Tables 5.4 and 5.5 essentially disappear. This is in part because for most participants by trial 30 decks C and D have become the decks with the best received values. That the values in each table do not greatly differ with such an approach is informative. Remember that the difference between the aggressive and conservative approach is that sustained demonstration of knowledge is required to fulfil the conservative criteria. We can assume that the similarities between the figures in the two tables demonstrates that participants' knowledge does not keep up with the change in actual received values as decks A and B become disadvantageous (generally between the second and third question period) otherwise the average trial on which knowledge was reached with a conservative approach would be lower in Table 5.5 than in Table 5.4.

A third point has already been mentioned but is worth re-emphasising here. When the 'both' grouping is used it is clear that most participants can identify the two best decks at some point in the early part of the task as evidenced by the figures generated from an aggressive approach. However, only a minority go on to demonstrate conceptual level knowledge that both of the good decks are the best. In other words few participants appear to have complete understanding of the deck contingencies. However, though important for a discussion on the extent of participants' task knowledge, it is actually not surprising given the constraints of the task. The stated goal is to earn as much money as possible but there is an unknown

Table 5.5: Knowledge assessment for Specific question group using ‘partial’ grouping (either deck with the highest net value at the time of questioning received the best score on each measure) or ‘both’ grouping (both decks with the highest net value at the time of questioning received the best scores on each measure).

Question Type	‘Partial’ Grouping		‘Both’ Grouping	
	Conservative	Aggressive	Conservative	Aggressive
<u>Ratings</u>				
% participants who did not reach Hunch level:	20	0	50	12.5
average trial number in which they did so:	39 (6.8)	22 (1.0)	59 (7.6)	33 (3.04)
<u>One deck</u>				
% participants who did not reach Hunch level:	6.25	0	<i>as ‘partial’ grouping as there is</i>	
average trial number in which they did so:	47 (7.5)	21 (0.6)	<i>only one response possible</i>	
<u>Expected net</u>				
% participants who did not reach Conceptual level:	25	0	62.5	12.5
average trial number in which they did so:	51 (7.9)	26 (2.2)	57 (9.5)	36 (4.2)
<u>Calculated net</u>				
% participants who did not reach Conceptual level:	50	0	68.75	12.5
average trial number in which they did so:	65 (8.5)	26 (2.4)	72 (12.4)	36.4 (4.3)

Note: Average trial values are rounded to the nearest trial. Values in parentheses are the standard error of the mean.

and limited time period to do this within. These constraints make the exploration necessary to achieve complete understanding and the goal of the task somewhat incompatible. It is also important to remember that complete knowledge is not required to profit on this task. As Maia and McClelland (2004) have pointed out, knowledge that one of the good decks is a good deck is sufficient to successfully guide behaviour on the IGT. A participant does not need to know that both good decks are best in order to select advantageously. As in many decision-making environments a simple heuristic that led one to choose one advantageous deck achieves the same result as total comprehension of the task structure. Given these reflections a conservative approach using a ‘partial’ grouping would seem to be the best to capture the emergence of knowledge sufficient to guide behaviour.

Indeed, in Maia and McClelland’s (2004) knowledge assessment they used the ‘partial’ approach throughout (referred to in notes to their Figures 2 and 3). Their Figure 2 displayed the proportion of participants who selected one of the two best decks for each measure at each question period. This can be considered an aggressive approach as it takes no account of a change in individual participants’ reports over time. As shown in Table 5.5 using the ‘partial’ grouping and an aggressive approach most participants in Experiment 9 have hunch level knowledge the first time they are asked about the task and a large majority show evidence for conceptual level knowledge.

Indeed if this approach is used at each question period a similar figure to Maia and McClelland’s is produced. Figure 5.3 is this figure. Like Table 5.5 it shows that the majority of participants have hunch level knowledge at the first question period. However, whereas for Maia and McClelland the number of participants displaying knowledge stayed high and relatively consistent across question period on each

knowledge measure, here the numbers fluctuate much more. As observed in the discussion of Tables 5.4 and 5.5 accurate partial knowledge dips on all measures as decks A and B become disadvantageous (between trials 20 and 40). The majority of participants recover some hunch level knowledge, but fewer reach conceptual level. This is despite the majority consistently selecting one of the best decks on each trial.

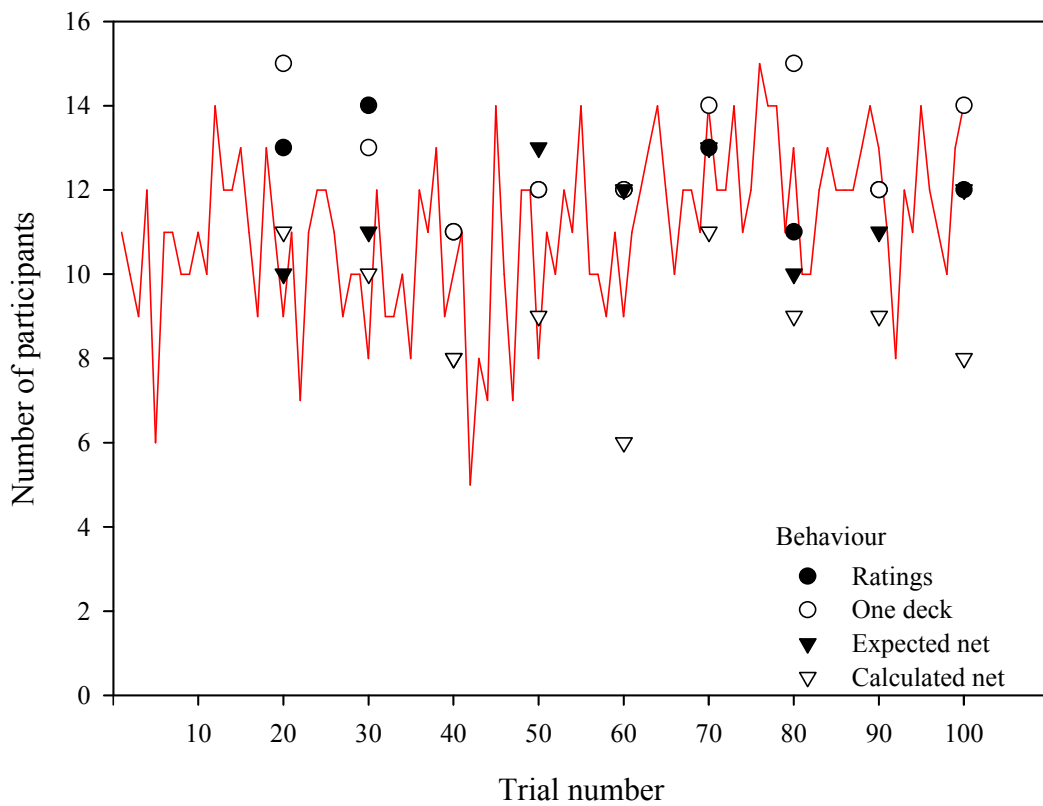


Figure 5.3: The number of participants in the Specific Question Group who displayed knowledge that one of the best decks is the best deck on each measure. The grey line represents the number of participants who chose one of the best decks available on each trial. On trial 40, 50, 60 and 90 the one deck marker covers the rating marker. On trial 40 the calculated net covers the expected net marker. On trial 70 and 100 the rating marker covers the expected net marker.

What is also informative from Figure 5.3 is that for many participants the calculated net measure (the open triangle) does not reflect one of the best decks. The vastly varying answers that participants provided in the calculation of this measure may be responsible. Figure 5.4 displays the calculated net measure for each deck from every participant in the final question period. The dashed line shows that the mean received value for each deck is close to its pre-test expected value. The figure shows that the calculated net measure does not correspond to individual participants' received values on each deck. Figure 5.5 shows that the same is true for the expected net measure. Together these Figures suggest that most participants' quantitative knowledge of the deck contingencies is not accurate. Indeed for many participants the expected or calculated nets are positive for decks A and B, and negative for decks C and D. This may indicate that participants are unable to retain quantitative knowledge about the decks or that they did not comprehend what was required in the answer for the measures themselves. In support of the latter explanation Figure 5.3 shows that on less complex measures most participants were able to select one of the better decks. Remember again that complete knowledge is not required to succeed on this task.

Figure 5.6 shows the number of times each deck was identified as the one deck participants would choose if they could only choose one for the remainder of the task. Aside from the first question period when deck B is often advantageous, most participants would choose deck C or deck D. Indeed the number of participants who would choose deck C increases with experience of the task, mirroring the behavioural data from this and previous experiments. As Figure 5.3 shows, using this measure the majority of participants were able to identify one of the best decks at each question period. However, given the choice most did not subsequently choose it exclusively.




Figure 5.4: Calculated versus actual expected value on each deck after 100 trials for each participant (except for participant 3 for whom figures are following 80 trials). The calculated expected value was calculated from a participant's estimates of the average gain, average loss and frequency of loss over ten selections from that deck. The dashed lines are the mean actual expected values.




Figure 5.5: Estimated versus actual expected value on each deck after 100 trials for each participant (except for participant 3 for whom figures are following 80 trials).

The dashed lines are the mean actual expected values.




Figure 5.6: The number of participants at each question period who selected each deck as the one deck they would choose if forced to only pick from one.

Figure 5.7 shows the change in ratings for each deck across block. The ratings are mostly negative for all decks and this may partly explain the discrepancy between the measures collected from participants and the actual received values displayed in Figures 5.4 and 5.5. It is clear that most participants do not believe any of decks are good. However, it is equally clear that decks C and D are perceived as being less bad than decks A and B. Although this indicates that participants have not comprehended the nature of the decks, and thus of the task, such knowledge would be sufficient to guide behaviour advantageously. This knowledge is present in most participants at the second question period. Participants also consistently rate deck A as one of the worst decks from the first opportunity they are given. This supports the hypothesis

developed in Chapter 2 that participants are able to identify deck A as a bad choice very early on, and this is the reason it is chosen much less than chance.




Figure 5.7: Mean rating for each deck across question period. Error bars are the standard error of the mean.

Summary of knowledge section

The focus of this section has been the extent of participants' knowledge and when that knowledge emerges. In the General question condition an aggressive approach provided the best fit to the data in the previous studies and using this approach most participants (93.75%) in that condition had hunch level knowledge by trial 43. However, using a conservative approach resulted in fewer participants being

categorized as displaying knowledge about the decks, a situation that reflected the experimenter's subjective impressions.

In the Specific question group and regardless of classification strategy most participants make a distinction between the advantageous and disadvantageous decks following trial 30. But although an aggressive, 'partial' approach applied to all measures revealed a majority of participants had knowledge at the first question period, this number was fewer than found by Maia and McClelland (2004). It was argued that in order to identify the earliest point at which knowledge sufficient to guide behaviour is reached consistently, a conservative approach using a 'partial' grouping would be best. This strategy suggests that most participants have hunch level knowledge by trial 39 using the ratings measure or trial 47 using the one deck measure. Neither are substantially different to trial 43 provided by the best-fitting strategy in the General question condition. But as deck ratings required more information from participants the figures gained from this measure will be used in the further analyses where differences pre- and post-knowledge are considered. Although, the strategies used to determine when knowledge was present were different in each group, this is appropriate because participants showed no differences in behaviour and so it can be assumed that their experience of the task was similar. We can further assume that their pre-task knowledge was similar and as their behaviour did not differ their knowledge remained similar throughout the task. All that differed between the groups then was the specificity of knowledge probe. If this is the case then an aggressive approach is appropriate for the General group because their knowledge was not probed as effectively as the Specific group participants. Ideally, a conservative partial approach would have been used throughout but this




Figure 5.8: Mean proportion of cards selected from each deck in (a) the pre- and post-knowledge periods for participants who displayed knowledge ($n = 27$), and (b) the comparable periods for participants who did not display knowledge ($n = 5$). Error bars are the standard error of the mean. The dashed line represents chance selection.

would not have been sensitive enough in the General condition to indicate when knowledge sufficient to guide behaviour appeared. The use of these two approaches results in data that is consistent between groups and with the previous literature using the General questions. It is also consistent with the behaviour shown in Figure 5.2. Mean net score first moves above chance in both groups in block 4, the block during which the above measures suggest participants can determine C and D to be the best decks.

Further support is provided by an analysis of the proportion of selections from each deck in the pre- and post-knowledge periods across all participants who were categorized as having displayed knowledge (displayed in Figure 5.8a). The proportion of selections from decks A and B declines from the pre- to post-knowledge period, whereas the proportion increases for decks C and D. This supports the supposition that participants' choices are guided by knowledge of the decks. A 4 x 2 (Deck by Time) repeated measures ANOVA examined these data. A significant interaction between Deck and Time was revealed, $F(2.28, 59.35) = 17.41, MSE = 0.03, p < 0.01$; as was a main effect of Deck, $F(3, 78) = 7.48, MSE = 0.03, p < 0.01$. There was no effect of Time, $F(1, 26) < 1$. A complex interaction comparison examined the interaction between Deck Type and Time by collapsing data across advantageous and disadvantageous decks in each knowledge period. This 2 x 2 repeated measures ANOVA found a significant interaction between Deck Type and Time, $F(1, 26) = 35.60, MSE = 0.03, p < 0.001$; a main effect of Deck Type, $F(1, 26) = 15.38, MSE = 0.03, p < 0.001$; but no main effect of Time, $F(1, 26) = 2.09, MSE < 0.01, p > 0.05$. Subsequent simple comparisons found that the proportion of advantageous choices in the pre-knowledge period was not significantly greater than the number of disadvantageous choices, $F(1, 26) = 2.41, MSE = 0.03, p > 0.05$;

whereas it was in the post-knowledge period, $F(1, 26) = 31.84$, $MSE < 0.01$, $p < 0.001$. Figure 5.8a shows that, consistent with previous experiments, this difference appears to be due to changes in selections from decks B and C. In the post-knowledge period the proportion of selections from deck B has decreased below chance and the proportion of selections from deck C has increased above chance. Similar patterns are found in decks A and D, but the major changes lie in decks B and C.

A similar pattern is shown in Figure 5.8b for the participants who displayed no knowledge. The early period shown in the Figure represents the proportion of choices from each deck up until the mean trial at which participants in the knowledge group displayed knowledge. The late period is the period from this mean trial until the end of the task. While behaviour in this group looks similar to the knowledge group, there are several differences. The proportion of selections from each deck is much closer to chance in both time periods. In the late period, unlike the participants with knowledge, selections from B are not below chance nor are selections from deck C above chance. These observations were tested in a 4 x 2 (Deck by Time) repeated measures ANOVA. It found no interaction, $F(1, 26) = 2.44$, $MSE = 0.01$, $p > 0.05$; no main effect of Deck, $F(1, 26) = 1.29$, $MSE < 0.01$, $p > 0.05$; and no main effect of Time, $F(1, 26) < 1$. These results suggest that only with knowledge sufficient to guide behaviour do participants select advantageously on the IGT. The next section will examine whether differences in physiological responses exist prior to knowledge being displayed and so leave an opportunity for an explanation of IGT behaviour incorporating somatic markers.

Physiological measures - aSCR

Anticipatory SCRs were the mean area under the curve of the SCR in the five seconds prior to selecting a card. Mean aSCRs for each deck were obtained by taking the average aSCR for that deck for each participant and dividing across participants. These mean aSCRs are displayed by Group in Figure 5.9a. The Figure shows that mean aSCRs are generally very low and that they are similar in each Group. To determine if any differences existed, a 2x4 (Group by Deck) mixed-factor ANOVA was run. Although mean aSCR was higher in the Specific Question Group than in the General Question Group no main effect of Group was found, $F(1, 30) < 1$. There was also no main effect of Deck, $F(1,30) < 1$. Despite the higher mean aSCR for deck B in the Specific Question Group, there was no interaction between Question Group and Deck, $F(3, 90) = 2.02$, $MSE < 0.01$, $p = .12$. As in the behavioural analysis no differences in aSCR were found between groups nor were any differences observed between decks. This first result supports the conclusion that the different questioning did not differentially affect participants, whereas the second contrasts with the data reported by Bechara et al (1997a).

In the previous section it was determined that most participants in each group display at least hunch level knowledge of the task between trials 40 and 50. In order to determine whether aSCR differences existed between decks prior to this period, average aSCRs before and after each participant's expression of knowledge were calculated for each deck for those participants who displayed knowledge (80% in the Specific group, 93.75% in the General group). As there were no differences in aSCR between groups in the previous analysis this factor was not included in the subsequent analyses. This also removes the problems associated with unequal sample sizes that would result with its inclusion. Some participants did not select cards from some of

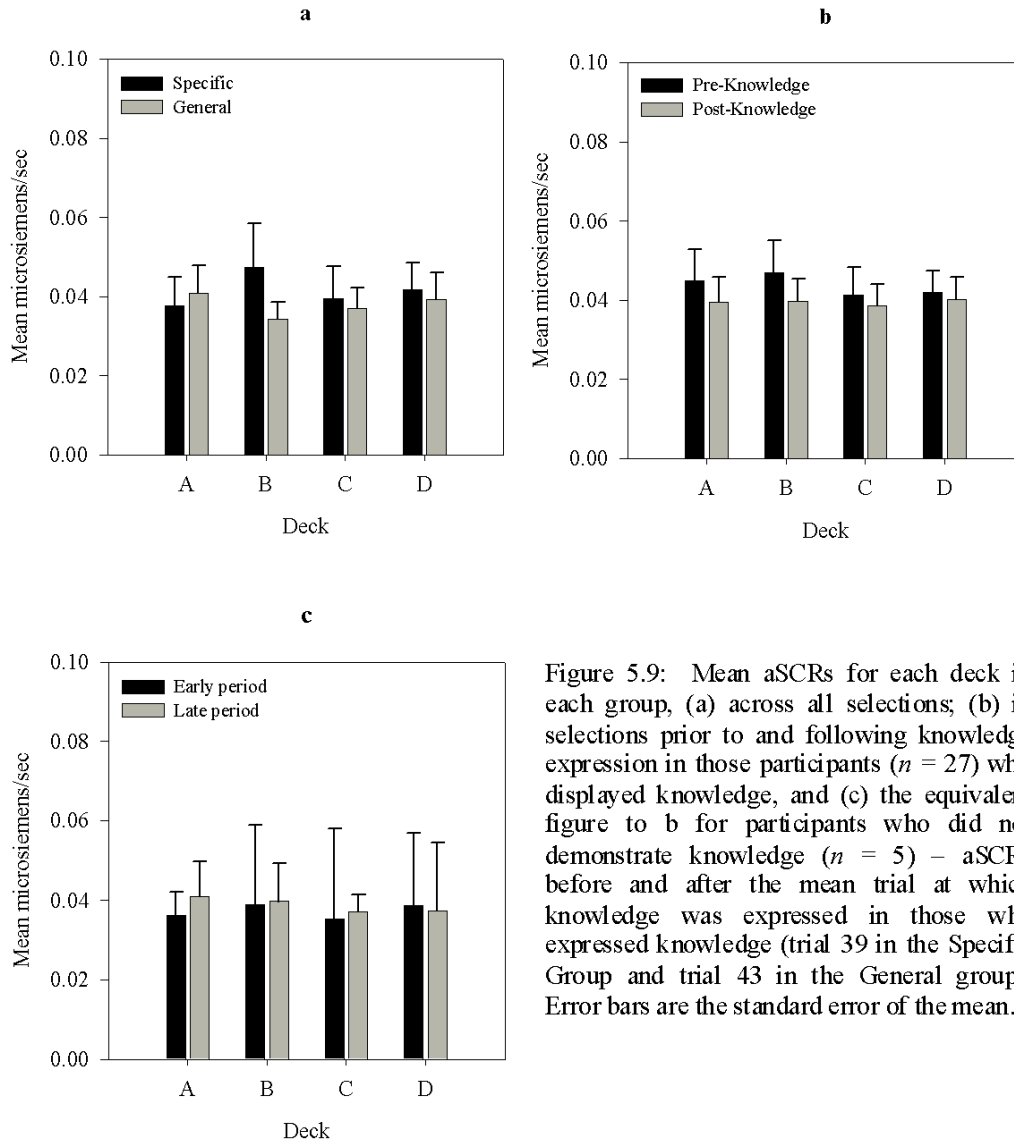


Figure 5.9: Mean aSCRs for each deck in each group, (a) across all selections; (b) in selections prior to and following knowledge expression in those participants ($n = 27$) who displayed knowledge, and (c) the equivalent figure to b for participants who did not demonstrate knowledge ($n = 5$) – aSCRs before and after the mean trial at which knowledge was expressed in those who expressed knowledge (trial 39 in the Specific Group and trial 43 in the General group). Error bars are the standard error of the mean.

the decks in the period following their expression of knowledge. As a result there is no SCR on these decks for these participants. As identifying missing values provided statistical results that did not correspond to the actual data (estimated marginal means were different to the descriptive mean) missing values were replaced by the mean

value for each level of each factor. For deck A, there were 2 missing values post-knowledge in each group. They were replaced by 0.039 in the Specific question group and 0.052 in the General question group. For deck B, there were 4 missing values post-knowledge in the Specific question group (replaced by 0.043) and 3 missing values post-knowledge in the General question group (replaced by 0.031). Finally, for deck D one post-knowledge missing value in the Specific question group was replaced by 0.037. The missing values all came from the same people who either chose only one deck in the period after they displayed knowledge (deck C in one participant in the Specific question group), or no longer chose from both deck A or B (two participants in both groups) or did not select from deck B (two participants in the Specific question group and one in the General question group). The resulting 4 x 2 (Deck by Time) repeated measures ANOVA found no significant effects: Deck by Time, $F(1.78, 46.41) = 1.25, MSE < 0.01, p > 0.05$; Deck, $F(2.04, 52.99) < 1$; Time, $F(1, 26) < 1$.

The main purpose of this experiment was to determine if any physiological responses distinguish between decks prior to participants' expression of knowledge i.e., SCR changes in the pre-hunch period of Bechara et al's (1997). No differences in aSCR were found between decks in the pre-knowledge period. This replicates Bechara et al's result, and like their data the mean values found in the present study within this period, displayed in Figure 5.9b, suggested that a difference between decks A and B and decks C and D may exist although there was no significant interaction. Therefore no evidence was found to support the hypothesis that differences in aSCRs precede knowledge expression in participants who express hunch level knowledge. Indeed, Figure 5.9c shows that in participants who did not display any knowledge mean aSCRs across the same time periods were at a similar level. Additionally, for

those participants' who did display knowledge, aSCRs did not appear to be related to their knowledge state as no differences pre- and post-knowledge in any deck were found.

Physiological measures – post-selection SCR

Post-selection SCR was the mean area under the curve of the SCR in the five seconds after a card was selected. These SCR were split into those following a reward with no punishment (reward SCR or rSCR) and those following trials on which punishment occurred (punishment SCR or pSCR). Mean rSCR and pSCR for each deck were calculated for each individual. The mean of these values provided the mean post-selection SCR displayed by Group in Figures 5.10a and 5.11a for reward and punishment SCR respectively. Figure 5.10a shows that mean rSCRs are similar in each Group but that there is a trend for rSCRs to be higher in decks A and B. A 2x4 (Group by Deck) mixed-factor ANOVA was run to examine rSCRs across all selections. There was no interaction, $F(1,30) < 1$; no main effect of Group, $F(1,30) < 1$; but a main effect of Deck was found, $F(1, 30) = 5.97, MSE < 0.01, p < 0.01$. A planned complex main comparison was performed to investigate whether rSCRs differentiated between the advantageous and disadvantageous decks. It found that rSCRs were higher for the disadvantageous decks, $F(1, 30) = 10.12, MSE < 0.01, p < 0.01$. In addition, separate pairwise comparisons revealed that mean rSCR was higher for deck A than deck C, $F(1, 30) = 11.44, MSE < 0.01, p < 0.01$; and deck D, $F(1, 30) = 8.20, MSE < 0.01, p < 0.01$. Mean rSCR was also higher for deck B than deck C, $F(1, 30) = 5.55, MSE < 0.01, p < 0.05$. The difference in rSCR between decks B and D was marginally non-significant, $F(1, 30) = 3.44, MSE < 0.01, p = 0.07$. There was no difference in rSCR between decks A and B or between decks C

and D. These results suggest that, in keeping with previous research (e.g. Tomb et al, 2002), selections that provided a larger reward result in larger rSCRs.




Figure 5.10: Mean rSCRs for each deck in each group, (a) across all selections; (b) in selections prior to and following knowledge expression in those participants ($n = 27$) who displayed knowledge; and (c) the equivalent figure to b for participants who did not demonstrate knowledge ($n = 5$) – rSCRs before and after the mean trial at which knowledge was expressed in those who expressed knowledge (trial 39 in the Specific Group and trial 43 in the General group). Error bars are the standard error of the mean.

To investigate whether rSCRs distinguished between selections prior to or following the display of knowledge a 4 x 2 (Deck by Time) repeated-measures design ANOVA was conducted. As no group differences were discovered in the initial analysis Group was removed as a factor in subsequent analyses. As with the equivalent aSCR analysis missing values due non-selection of a deck in the post-knowledge period were replaced by the group mean. For deck A, 2 missing values post-knowledge in each group were replaced by 0.093 in the Specific question group and 0.099 in the General question group. For deck B, the 4 missing values post-knowledge in the Specific question group were replaced by 0.093 and the 3 missing values post-knowledge in the General question group were replaced by 0.089. Finally, for deck D the single post-knowledge missing value in the Specific question group was replaced by 0.063.

An interaction was found between Deck and Time, $F(2.39, 62.13)^1 = 4.65$, $MSE < 0.01$, $p = 0.01$. As with the overall analysis a main effect of Deck was also found, $F(3, 78) = 5.00$, $MSE < 0.01$, $p < 0.01$, but there was no effect of Time, $F(1, 26) < 1$. Figure 5.10b displays the mean rSCRs pre- and post-knowledge in each deck. The interaction between Deck and Time appears to be the result of lower mean rSCRs in the post-knowledge period for the advantageous decks as compared to the disadvantageous decks. In order to examine this further data was collapsed across Deck to provide values for the advantageous and disadvantageous decks in each time period and an interaction contrast was performed. This is effectively a 2 x 2 (Deck Type by Time) repeated-measures ANOVA. It revealed a significant interaction between Deck Type and Time, $F(1, 26) = 11.83$, $MSE < 0.01$, $p < 0.01$; a main effect of Deck Type, $F(1, 26) = 15.74$, $MSE < 0.01$, $p < 0.001$; but no effect of Time, $F(1, 26) < 1$. Subsequent simple comparisons found a difference between Deck Types in

the post-knowledge period, $F(1, 26) = 19.56$, $MSE < 0.01$, $p < 0.001$, and not in the pre-knowledge period, $F(1, 26) < 1$. In the selections after knowledge is displayed participants' physiological reactions following reward distinguish between the good and bad decks.

This interaction should not have been affected by the replacement of missing values with group means for several participants. Nevertheless, to rule out this possibility the participants for whom there were two or more missing values in the post-knowledge period (two in the Specific Group and three in the General group) were excluded and the analysis run again. The three additional participants in the Specific group who did not select deck B in the post-knowledge period were retained as excluding them would have left only nineteen participants. The resulting ANOVA revealed identical results with the relevant interaction between Time and Deck still significant, $F(2.38, 47.56) = 2.96$, $MSE < 0.01$, $p < 0.05$. This suggests that the results from the full analysis are legitimate and do reveal that a difference in rSCRs between deck types exists only after knowledge has been displayed.

Figure 5.10c presents rSCRs for the participants who did not display knowledge. Here the pre- and post-knowledge periods are based on the mean values from the participants who did display knowledge. The early period includes the trials up to trial 39 and 43 for participants in the Specific and General groups respectively. The late period includes all the subsequent trials. The mean values depicted in this Figure are much lower than those for participants with knowledge, suggesting that knowledge and physiological activity may be linked. A similar pattern of reduced physiological activity in the post-knowledge period in decks C and D is also found in this group as in the participants with knowledge, but here it is also found for deck B. A 4 x 2 (Deck by Time) repeated-measures ANOVA was also conducted on this data.

There was no interaction between Deck and Time, $F(3, 12) = 1.31$, $MSE < 0.01$, $p > 0.05$; no main effect of Deck, $F(3, 12) = 1.54$, $MSE < 0.01$, $p > 0.05$; and no main effect of Time, $F(1, 4) < 1$. This result supports the conclusion from the analysis of the with-knowledge group that knowledge influences physiological activity. However, this conclusion is qualified by the low number of participants included in this analysis.

Figure 5.11a shows pSCRs over all selections and all participants. Mean pSCRs are higher in the decks with low frequency of punishment (B and D). Mean pSCRs are also higher than mean rSCRs. A 4 x 2 (Deck by Group) mixed-factor ANOVA revealed no interaction, $F(3, 90) < 1$ and no main effect of group, $F(1, 30) < 1$, thus replicating the other SCR data that found no group differences in SCR. A main effect of Deck was found, $F(2.12, 63.66)^1 = 4.40$, $MSE < 0.01$, $p < 0.05$. Subsequent simple comparisons revealed that pSCRs following selections from deck A were significantly lower than those from deck B, $F(1, 30) = 6.73$, $MSE < 0.01$, $p < 0.05$; as were selections from deck C, $F(1, 30) = 10.02$, $MSE < 0.01$, $p < 0.05$; while pSCRs for deck D were also significantly higher than those from deck C, $F(1, 30) = 5.73$, $MSE < 0.01$, $p < 0.05$. There was no difference in pSCRs following selections from decks B and D, $F(1, 30) = 2.96$, $MSE < 0.01$, $p > 0.05$, nor between decks A and D, $F(1, 30) = 2.96$, $MSE < 0.01$, $p = 0.10$, which replicates Crone et al (2004) and supports their conclusion that it is the magnitude of punishment and not the frequency that is influential for pSCRs.

Due to the infrequent nature of punishment relative to reward in all of the decks (far greater in decks B and D), many participants received no punishment in the post-knowledge period on some decks either as a result of not choosing them or because no punishment resulted from their choices. As this applied across so many

participants a 4 x 2 (Deck by Time) analysis became impractical with the addition of missing values reaching unacceptable levels. However, the question of interest was




Figure 5.11: Mean pSCRs for each deck in each group. (a) Across all selections. (b) Mean pSCRs for the advantageous and disadvantageous decks in selections pre- and post-knowledge expression in those participants ($n = 27$) who displayed knowledge. (c) The equivalent graph to b for the participants who did not demonstrate knowledge ($n = 5$) – pSCRs in the period before and after the mean trial at which knowledge was expressed in the majority of each group. Error bars are the standard error of the mean.

whether physiological activity distinguished between the decks prior to a display of knowledge. As such pSCRs were averaged within participants in two ways. First, the

mean pSCR for the advantageous and disadvantageous decks in the pre- and post-knowledge period were calculated for each participant. Figure 5.11b displays these means for those participants who displayed knowledge. A 2x2 (Deck Type by Time) repeated measures ANOVA, equivalent to that performed on the rSCR data, revealed a significant interaction between Deck Type and Time, $F(1, 26) = 4.44$, $MSE = 0.02$, $p < 0.05$; but no main effect of Deck Type, $F(1, 26) < 1$; nor a main effect of Time, $F(1, 26) = 1.96$, $MSE = 0.02$, $p > 0.05$. Subsequent simple comparisons revealed that pSCRs were higher for the disadvantageous decks prior to knowledge being displayed than in the period after, $F(1, 26) = 6.04$, $MSE = 0.01$, $p < 0.05$.

Second, the mean pSCRs for the decks with frequent and infrequent punishments were also calculated in each knowledge period. A 2 x 2 (Punishment Frequency x Time) repeated measures ANOVA found no interaction, $F(1, 26) < 1$; no main effect of Punishment Frequency, $F(1, 26) < 1$; and no main effect of Time, $F(1, 26) = 1.96$, $MSE = 0.02$, $p < 0.05$. This result contrasts with Crone et al (2004) who found higher pSCRs following choices from decks B & D.

Similar analyses were carried out for the participants who showed no knowledge. Figure 5.11c displays the mean values of pSCRs collapsed across the advantageous and disadvantageous decks up to and after the mean trial at which participants with knowledge displayed that knowledge. The 4 x 2 (Deck Type by Time) ANOVA revealed no interaction, $F(1, 26) = 1.42$, $MSE < 0.01$, $p > 0.05$; no main effect of Deck Type, $F(1, 26) < 1$; and no main effect of Time, $F(1, 26) = 1.11$, $MSE < 0.01$, $p > 0.05$. The Punishment Frequency x Time ANOVA also revealed no interaction, $F(1, 26) = 1.43$, $MSE < 0.01$, $p > 0.05$; no main effect of Deck Type, $F(1, 26) < 1$; and no main effect of Time, $F(1, 26) < 1$.

Summary of the physiological section

No differences in any SCR type were found between groups suggesting that the questions did not affect participants' physiological responses. There was also no difference in aSCR between decks or deck types at any point during the task.

Differences in post-selection SCRs were found. Reward SCRs distinguished between the advantageous and disadvantageous decks across the whole experiment. This effect was found to emerge only in participants who displayed knowledge and then only in later trials following their display of knowledge. No differences in rSCRs between decks were found in those who did not display knowledge. Punishment SCRs were found to be larger for the disadvantageous decks in the pre-knowledge period but only for participants who displayed knowledge.

Overall, the SCR values recorded were low in comparison to other studies. However, a consistent pattern was apparent. Anticipatory SCRs were generally low, but rSCRs were higher and pSCRs were higher still. This is consistent with previous research in the literature and argues against any methodological flaw in SCR recording.

5.3 GENERAL DISCUSSION

The results of this experiment are not definitive and provide only mixed support for the theories of IGT behaviour explored in the introduction. This is because no difference in aSCRs between any of the decks were found. The absence of this previously reported effect makes it impossible to determine whether such an effect precedes or follows participants' displays of knowledge. However, while no aSCR differences were found participants on the whole learned to select from decks C and

D. This suggests that differential somatic activity, reflected in aSCR activity, is not necessary for learning to proceed.

While it was found that aSCRs do not differentiate between deck types prior to knowledge being displayed, a difference between deck types found over all rSCRs was localised within participants who displayed knowledge in the period following that knowledge being displayed. This result provides more support for the Knowledge-Somatic than the Somatic-Knowledge hypothesis, although that support must be qualified. The absence of any difference in aSCRs is problematic as a null effect can never be evidence for any hypothesis, and the results from the pSCRs suggest physiological responses occur for larger primary punishers but only in the initial period of the task. It could be argued that pSCRs did not distinguish between decks in the post-knowledge period because participants were aware that those decks had the worst losses. Conversely, it could also be argued that the pre-knowledge pSCRs influence subsequent decisions and constitute the first stage in a process towards somatic markers. This position is supported by the absence of these effects in participants who displayed no knowledge. So the physiological results are ambiguous showing that differences in post-selection SCR activity emerge following knowledge for rewards but prior to knowledge for punishments. It could be argued that the post-knowledge difference in rSCRs indicates relief at escaping from a choice on a disadvantageous deck without a punishment. This would constitute an effect of knowledge and offer better support for the Knowledge-Somatic hypothesis than the Somatic-Knowledge hypothesis. After all, these decks are more risky than the advantageous decks. Differential SCR activity, including aSCRs, may just reflect this awareness of risk.

The absence of an aSCR difference between deck types is not without precedent (Kleeberg et al, 2004; Campbell et al, 2004). That aSCRs did not increase over time also replicates earlier results using a computerized version of the task (Suzuki et al, 2003; Carter and Pasqualini, 2004). A possible explanation for the absence of differences in the aSCRs is the automated way in which they were gathered. The experimenter controlled the length of the inter-trial interval between SCR acquisitions in Bechara et al (1997a). This was to ensure that participants' physiological activity had returned to baseline following the previous choice. It is possible that as the inter-trial interval was fixed to a greater extent in the current experiment, physiological activity following the previous choice interfered with anticipatory physiological activity on the next choice. But Crone et al (2004) employed a similarly automatic methodology ensuring that the inter-trial interval was as long as reported by Bechara et al (1997a) and found similar effects to them. The inter-trial interval in the current study was as long as the mean given in Bechara et al's report (twelve seconds). If interference between SCR types occurred then larger aSCRs would be expected following a loss than following a gain, mirroring the results with the post-selection SCRs. But an examination of aSCRs in each deck following a gain and a loss found no difference. This data was calculated for each participant and entered into a 4 x 2 (Deck by Reinforcer) repeated measures ANOVA. No main effect of Reinforcer was found, $F(1, 27) < 1$; nor was there a main effect of Deck, $F(1.98, 53.33) < 1$; nor an interaction, $F(1.74, 46.88)^1 < 1$. This suggests that automatic gathering of SCRs did not impact on the clarity of the physiological record.

The SCRs recorded were also small in comparison to previous studies, with area under the curve measures one-tenth the size of some reported in the literature (Bechara et al, 1997a). This suggests that an error in the recording or analysis

clouded the results. But differences between SCR types were found in line with the previous literature. The aSCRs were smaller than the rSCRs, which were smaller than the pSCRs. A possible explanation, that may also have a bearing on the absence of an aSCR effect relates to a difference in the methodology employed as compared to most other studies. It will be remembered that the experimenter, though present in the same room as the participant, sat behind a screen and did not interact. This was communicated to the participant prior to task commencement. This absence of an interpersonal interaction has been maintained throughout the Experiments reported in this thesis, but is a major difference between this computerized task and the manual task. How might this impact the physiological record? Carter and Pasqualini (2004) reported that they switched from a computerized to a manual task because participants SCRs decreased substantially as the computerized task progressed. The absence of an interpersonal interaction or even the presence of an actively observing person may be required for SCRs to exist. In support, Van 't Wout, Kahn, Sanfey and Aleman (2006) reported that SCRs were higher when unfair offers were made and rejected in an Ultimatum Game, but only when those offers were made by a human rather than a computer. This interpersonal interaction hypothesis does not explain the presence of differential SCRs between deck types as reported by Bechara et al (1997a) and others even using a computerized task (Hinson et al, 2003; Crone et al, 2004) but without more detail about their exact procedure (e.g. the presence and position, or absence of an experimenter) it is difficult to entirely reject it .

This experiment found that the emergence of knowledge occurred at a similar time as Bechara et al (1997a), yet found no replication of the aSCR effect. If accepted at face value this result is problematic for the SMH. Participants in this experiment improved on the IGT at levels similar to those in the previous Experiments, and

displayed knowledge of which decks were worst in the long-run, yet the results suggest aSCRs played no part in this process. It may be that participants in this experiment did not have the same physiological reaction as those in other experiments but if this is the case it suggests that like other, clinical studies (North and O'Carroll, 2001; Heims et al, 2004) the absence of autonomic activity does not preclude learning on the IGT.

The results of this experiment are not only problematic for Bechara et al's (1997) account of IGT behaviour. Knowledge sufficient to guide long-term advantageous selection emerged in the majority of participants at around the same time as Bechara et al (1997a) claimed. Participants are able to identify one of the best decks at the initial question times as Maia and McClelland (2004) claimed. But when the change over from the disadvantageous to the advantageous decks as the best decks occurs participants have a problem keeping up. This was reflected in the similarity of the data displayed in Tables 5.4 and 5.5. This overlaps with when Bechara et al (1997a) claimed the aSCR difference emerged (trials 10 – 50). Kleeberg et al (2004) reported that although they found no difference in aSCRs between deck types the increase in aSCR they observed averaged across all decks emerged between trials 20 and 40. These aSCR differences may be related to the shift in polarity of deck received values. The results from this study mean that Maia and McClelland's (2004) assertion that participants have knowledge sufficient to guide their behaviour from the first questioning has not been replicated here and this does leave open the possibility that somatic activity precedes knowledge. However, no evidence was found to support this hypothesis in this study.

As Maia and McClelland (2004) found, the assessments of participants' knowledge here often indicated that their behaviour did not reflect the knowledge they

possessed. Participants often did not select one of the best available choices despite the knowledge probes indicating that they were able to make this distinction. One explanation for this behaviour that Maia and McClelland did not consider, but which is apparent from the detailed examination of participants' knowledge, is that their knowledge is not complete and few possess accurate knowledge of the deck contingencies. This makes non-optimal deck selection a reasonable option as participants attempt to explore the decks to learn more about their contingencies. However, as Figures 5.4 and 5.5 suggest few come close to achieving this understanding.

Most participants gave all the decks a negative rating suggesting that they were unaware that either decks C or D were profitable with repeated selection. This also suggests that for participants in this experiment the times when they lost money were most influential when they made their ratings. This accords with the hypothesis emerging through this thesis that the short-term outcomes of deck choice are more important than has been previously assumed. Certainly the pattern of changing selection from decks B and C driving learning observed in previous Experiments was replicated here and was reflected in the question responses of participants given the Specific questions.

These results differ from those reported by both Maia and McClelland (2004) and Bechara et al (1997a). Both groups suggest that the majority of their participants end the experiment with conceptual knowledge of the IGT. Little evidence was found in this experiment to support this conclusion, but the results did show that conceptual knowledge is not critical for learning to occur. But the difference in degree of knowledge at the end of the task may help explain why participants in the present

Experiments make fewer advantageous choices than those of the Iowa group. Such differences in performance have been recorded in other tasks (Brase et al, 2006).

In conclusion, the results from this experiment suggest that participants do not generate anticipatory physiological activity sufficient to differentiate between deck types in the period prior to acquiring knowledge sufficient to guide their behaviour. Knowledge required to profit on the IGT emerges later than claimed by Maia and McClelland (2004) but it is not a complete understanding of the nature of the IGT. Despite this, advantageous deck selection is learned mainly through learned avoidance of B and increased selection from C.

GENERAL DISCUSSION

6.1 SUMMARY OF EMPIRICAL RESULTS

Behaviour on the Iowa Gambling Task (IGT, Bechara et al, 1994) has been widely interpreted as empirical support for Damasio's (1994, 1996) somatic marker hypothesis (SMH). The SMH posits a critical input from an embodied emotional system in making decisions in choice situations. In interpreting behaviour on the IGT in accordance with the SMH three central assumptions have been made. These are: a) that somatic markers indicate the goodness or badness of alternatives and without them decision-making cannot become optimal, b) this somatic biasing or guidance can occur unconsciously or in the absence of explicit knowledge, and c) that the system operates so as to maximize or achieve the best outcome in the long-term. The Experiments described in this thesis have explored the validity of the second and third assumptions and found that they are not accurately reflected in behaviour on the IGT.

Chapter 2 explored the nature of instructions and type of reinforcement employed on the IGT. The finding that performance, as measured using the net score, was improved by the presence of a strategy hint in the task instructions is not surprising but had remained untested, though not unremarked (Schmitt et al, 1999), in the preceding literature. This novel result demonstrated that cognitive information impacts on task performance and implies that knowledge about the IGT influences performance. A strong version of the SMH (e.g. Bechara et al, 1997a; Tranel et al, 1999) would suggest that somatic markers precede and indeed facilitate the

development of this knowledge. However, were this the case it would be predicted that learning (guided by somatic markers) would be found in the absence of the hint instructions. The Experiments in Chapter 2 revealed that this was not the case. Without the hint instructions participants were selecting the advantageous decks at chance levels even after 200 trials. This is problematic for a strong version of the SMH and for the related third assumption of IGT behaviour in the SMH framework. If somatic markers are an integral system for decision-making, and so guide behaviour on the IGT towards maximization, then it should not matter what information participants receive about the task. That learning does not occur without hint information suggests that explicit knowledge is important for performance. The possibility that somatic markers are a reflection of knowledge about deck contingencies (e.g. that A and B have the highest risk due to the magnitude of losses) was explored in Chapter 5. This Experiment was a direct test of the second IGT assumption that somatic activity in the form of anticipatory SCRs precedes knowledge. No evidence was found to support this assumption but it was found that knowledge sufficient to guide behaviour was achieved before post-reward SCRs differentiated between deck types. A Knowledge-Somatic hypothesis to rival the Somatic-Knowledge hypothesis implicit in the second assumption cannot be rejected. These results along with the absence of learning when no hint is provided support recent assertions that the IGT is not as cognitively impenetrable (Maia and McClelland, 2004) as claimed by Bechara et al (1994, 1997a).

The third assumption of IGT behaviour in the SMH framework is that deck selection proceeds solely on the basis of long-term outcome (expected value). It was explored in detail in Chapters 3 and 4 where the consequences of altering the short-term contingencies, but not expected value, were explored. These tests of the third

assumption were provoked after participants' behaviour in the Experiments in Chapter 2 was found to be inconsistent with an account of behaviour based solely on expected value. Participants' selection decreased from deck B but increased from deck C and this change (i.e. learning) was behind the improving IGT performance. Experiment 4 supported this account in a simple two-choice version of the IGT, and Experiment 5 found some evidence that the information available from the relative frequency of loss was responsible. However, the prediction that learning would differ between an Increased Frequency and Decreased Frequency version of the IGT only received marginal support. While a significant difference in learning rate and mean net score between the conditions was not found, only in the Increased Frequency condition was net score above chance by the end of the Experiment. These Experiments in Chapter 3 replicated the differential selection from IGT decks with the same expected values found in Chapter 2. The demonstration that relative frequency of loss is influential in determining preference adds new information to understanding of IGT behaviour and suggests that when the informative events are encountered is also influential. Indeed, the experiments reported in this chapter add empirical evidence to the increasing number of studies (Wilder et al, 1998; MacPherson et al, 2002; Peters and Slovic, 2000; Yechiam et al, 2005a) that have reported differential selection based on the immediate consequences of selection may be more important than has generally been assumed. This is much more apparent for the disadvantageous decks where the difference in the relative frequency of loss appears to determine which is preferred and which is avoided. The number of unpunished selections from the deck would also appear to be a critical factor in determining the number of repeated selections from these decks.

Differences in the immediate outcomes of advantageous deck selections do not appear to create as stark a disparity as in the disadvantageous decks. It was suggested that deck C is perhaps the more important deck in influencing learning. This may be because it takes participants longer to learn that the more frequent losses in this deck do not actually punish them in the longer term. Whereas losses in deck D are so infrequent (and relatively smaller than deck B) that above chance selection does not need to change to gain in the long-term. There may also be a role for individual differences in the relative contribution that loss frequency, loss magnitude, gain magnitude or expected value to individual deck preference and therefore learning on this task. The focus in Chapter 3 was on the frequency of loss but given the structure of the IGT it is impossible to alter one component of a deck without a commensurate change in another. Even where gain magnitude and expected value were unconfounded (Peters and Slovic, 2000) participants exercised preference for the deck with the lowest frequency of loss and the dependent highest expected value.

Chapter 4 continued the exploration of deck selection in the IGT. Manipulations of the underlying reinforcement schedule again showed that deck selection is not solely governed by expected value. For example, despite deck A having a higher received value than deck B for most participants in Experiment 6, participants still on the whole preferred deck B. This disadvantageous deck with the infrequent reinforcer is really the key deck on the IGT. Participants readily avoided deck A but because the first loss on deck B occurs after many unpunished selections selection generally continued until the second loss was encountered. Prior to this point (13 selections) B is still a profitable deck. Discovering that repeated selection results in an overall loss is the crucial selection issue and this can only occur once a loss has been experienced. In Chapter 4 the reinforcement schedules were

manipulated and the importance of the first occurrence of this event was demonstrated. Although the nature of the probabilistic schedule meant that deck B was not as bad for some participants as for others, its average expected value was negative. Yet, justifiably, many participants persisted in its selection despite the higher losses. This resulted in a difference in selection between A and B in the probabilistic schedule that was more stark than when fixed schedules were used (compare Figures 4.2 and 2.6).

The importance of the infrequent event was also demonstrated when the variant IGT was used in Experiment 7. Again, differential selection between decks with the same expected value was found, but this was more apparent in the advantageous decks. In the variant task these decks have the highest immediate losses but the higher gains. When the order of the infrequent events was manipulated to the same order as in the standard IGT, participants rarely persisted in selecting from deck E (large infrequent gain) long enough to encounter it. However, they readily selected deck G and it was this deck that resulted in their learning an advantageous strategy. Similar results were found when a probabilistic version of the variant task was used. Together the Experiments in Chapter 4 demonstrated that deck selection on the IGT is not solely dependent on the expected value of decks. The assumption that selections are made on this basis cannot be supported by a detailed examination of behaviour. This has implications for both an underlying somatic system should it exist and also the analysis of IGT behaviour. These are discussed in the next section.

The main results from this thesis can be summarised thus: Participants do better on the IGT if they are told what to do. Generally they learn to avoid the disadvantageous deck with infrequent large losses and learn to prefer the advantageous deck with frequent low losses (and lower net losses). Selection is not

made on the basis of expected value alone, and participants have access to knowledge that is sufficient to guide behaviour in the absence of any anticipatory somatic markers. The implications of these findings for the SMH and future studies using the IGT will now be discussed.

6.2 IMPLICATIONS FOR FUTURE USE OF THE IGT

The IGT has been predominantly used for two purposes. The first was in the development of and as a test of the SMH. The second is as a test of decision-making in disparate populations. In some cases the second use has resulted in interpretations in accordance with the first. But, in both respects a basic understanding of how healthy participants (with whom comparisons are generally made) actually decide on the task has been left behind. The Experiments in this thesis have gone a long way to remedying this situation and therefore have implications for the use of the IGT in the future.

6.2.1 The IGT as a test of the SMH

The claim that deck selection on the IGT is guided by somatic markers prior to the development of consciously accessible knowledge (Bechara et al, 1997a) is central to the assumption that somatic markers are causally related, or at least necessary for, decision-making to proceed successfully. This strong version of the SMH is the version that has been cited in the majority of published studies where the IGT has been used. Maia and McClelland (2004) challenged this assumption by showing knowledge about the decks is available earlier than assessed using Bechara et al's (1997a) method. The Experiments in Chapters 2 and 5 have extended this challenge

by demonstrating that explicit knowledge about the IGT decks is the crucial component of participants' eventual performance. Without information (necessarily interpreted in cognitive terms) that the alternatives are qualitatively different participants did not learn to select advantageously. Experiment 9 showed that participants acquire knowledge sufficient to guide behaviour on average after 40 trials, but that there was a transition period during which knowledge lagged behind the change in deck contingencies as the previously better decks became the worst. This detail was absent from Maia and McClelland's (2004) analysis where they examined knowledge assessed using detailed questioning only within individual question periods. The global approach applied here was important because it left open the possibility that differential somatic marker activity may have preceded knowledge acquisition – the transition period falls within the period when Bechara et al (1997a) reported differential somatic activity between deck types (although it was different this difference was non-significant). However, no such somatic activity was found in this Experiment 9 using SCRs to measure somatic activity.

These results are problematic for a strong version of the SMH because an explanation without somatic markers cannot fit into the framework of how decisions are made. Several explanations consistent with the SMH suggest themselves but also bring about problems. One possibility is that somatic activity did precede knowledge acquisition, or contributed to its development, but were not recorded because a computerized test was used (see also Carter and Pasqualini, 2004). This is problematic because if somatic markers guide decisions in a manual IGT (e.g. Bechara et al, 1996; Carter and Pasqualini, 2004) but not when it is computerized it suggests that somatic activity is affected by situational factors other than the choice

environment. Since the choice remains constant across mediums such an explanation would imply that somatic markers are not ubiquitous for decision-making on the IGT.

Another possible explanation is that participants somehow accessed the “as-if body” loop of the somatic system. In Bechara and Damasio’s (2005) recent reformulation of the SMH this is unlikely since description of this loop’s operation requires an initial bodily representation of the somatic response to novel situations. Since all participants were naïve to the task prior to exposure some initial somatic activity would be expected. It is possible that the similarity of the IGT environment to other card games, or “real-life” contingencies meant body loop activation was unnecessary. Ultimately, this possibility is not testable and this point will be returned to later in this section.

A third possibility is that the strong version of the SMH is not supported and a weak version where somatic activity contributes to decisions but is neither causal nor necessary may be true. The problem with this possibility is that somatic activity is no longer necessary since consciously available knowledge can explain behaviour anyway. If somatic activity does not precede knowledge then it is equally possible that it occurs as a result of knowledge and not as a bias that feeds into knowledge development. The importance of the executive component of working memory in IGT selection supports such a contention (e.g. Hinson et al, 2003; Jameson et al, 2004).

For an explanation of behaviour on the IGT to be consistent with the SMH despite evidence that the body loop is impaired (e.g. North and O’Carroll, 2001; Heims et al, 2004) requires involving the “as-if body” loop. The main problem with this account is that it is virtually impossible to test. If the SMH can always account for IGT learning by appealing to the “as-if body” loop then it is difficult to see how

such a hypothesis is falsifiable. It seems unlikely that neurological patients will be found in whom only this loop is damaged since it is fundamentally the same system as the body loop but without the bodily input. Any damage would affect both hypothetical systems. Recently, Bechara and Damasio (2005) have suggested a differentiation between the body and “as-if body” loop. The former is engaged when faced with situations under ambiguity whereas the latter is engaged in choice under risk. The main support for this claim is that aSCRs on the Cambridge Gamble Task (CBT; Rogers et al, 1999) are lower than on the IGT. The difference between the tasks is that the contingencies are explicit in the CBT (therefore choice operates under risk), but are uncertain or ambiguous on the IGT. But this distinction is contradicted by the studies cited earlier where IGT learning was unaffected by body loop damage, and does not explain why somatic markers are not always found in encounters with the IGT regardless of the environmental format.

Another implication of the results from this thesis concerns descriptions of abnormal behaviour on the IGT. It has been explained in terms of a “myopia” for future consequences (e.g. Bechara et al, 1999; 2002) and hypersensitivity to reward (Bechara and Damasio, 2002) dependent on performance on the standard, variant and progressive versions of the IGT and analysis of somatic activity using SCRs. Neither of these explanations takes into account the differential deck selection analysed in this thesis. It may be that selection from the disadvantageous decks in clinical populations is uniform and in such cases selection from A and B might constitute myopic behaviour. However, examination of individual deck selection, and how it changes over time, can provide more information about the factors influencing choice. Myopia for the future only works as an explanation for behaviour when data is examined in terms of expected value. If selection is mainly from B, or B is returned

to more frequently than A, then this may be myopia but may also reflect risk-seeking and/or an absence of understanding of the nature of the task. Equally possible, selection from B might persist as an attempt to gain the higher reward despite knowledge of the risk of the large infrequent penalty. Or selection from B might persist for the same reason but in the absence of understanding that losses outweigh gains. Similarly non-selection of C and preference for D may reflect an avoidance of frequent punishments or an inability to maintain an account of overall losses. Differences such as these may be explained by myopia for the future or hypersensitivity to reward but this thesis has demonstrated that more information is available than from simply collapsing across decks with the same expected value.

6.2.2 The IGT as a test of decision-making

One of the problems for the IGT as a test of decision-making is that many factors have been linked to abnormal performance on it. This is partly due to the need for it to be complex in the attempt to model real-life decision-making. Indeed the claim that the IGT is cognitively impenetrable received a degree of support in this thesis due to the inability of participants to accurately quantify the contingencies on any of the decks when they are asked. A similar result may have been found by Maia and McClelland (2004) but they did not report the accuracy of their participants' estimates of the calculated and expected net measures. However, the impenetrability claim is challenged because participants' answers to specific questions demonstrated their knowledge was sufficient to distinguish decks C and D as the least bad relatively early in the questioning. What they had trouble doing was keeping an accurate estimate of the contingencies.

Other processes contribute to IGT performance and deficits in any of these may result in abnormal behaviour. Working memory has already been mentioned (Hinson et al, 2003; Jameson et al, 2004) as has risk-seeking although no correlations between IGT and CBT performance have been reported when performance on the two tasks was compared (Monterosso et al, 2001). Also important as explanatory mechanisms are hypersensitivity to reward (Bechara et al, 2002; Bechara and Damasio, 2002) and possibly even hyposensitivity to punishment (Bechara et al, 1999). In some respects the possible mechanisms for abnormal choice on the IGT reflect the complexity of the system that must deal with real-life decisions. But the lack of specificity of an underlying impairment for abnormal performance makes comparisons across populations difficult. For example, Dunn et al (2006) have highlighted six studies of schizophrenic patient groups where abnormal behaviour has not been replicated. While this may reflect the debate about the nature and classification of schizophrenia (Frith, 1995) it is problematic that abnormal performance does not replicate across similar populations (this has also been reported in other clinical groups e.g. OCD: Nielen, Veltman, de Jong, Mulder and den Boer, 2002; Cavadini, Riboldi, Keller, D'Annuncci and Bellodi, 2002).

A related problem for the IGT is that sizable proportions of healthy participants select disadvantageously. Bechara et al (2001, 2002) and Bechara and Damasio (2002) reported this figure at between 30 and 40%. This has also been found by other groups (Bowman and Turnbull, 2003; Evans et al, 2004; Crone et al, 2005). Similarly performance of normal controls using the net score measure rarely reached the levels reported by Bechara et al (1994, 1996) in this thesis. These data are problematic for the IGT as it has been argued that it raises questions about its ecological validity (Dunn et al, 2006). If controls do not learn to select

advantageously but presumably have no deficit in day-to-day decision-making then the utility of the IGT as a test of such functioning is questionable.

Yechiam, Busemeyer, Stout and Bechara (2005b) have attempted to remedy this by naming three component processes that contribute to IGT performance from Busemeyer and Stout's (2002) expectancy-valence model. These are the importance of gains versus losses, the effects of recency (or forgetting) and the consistency of choices between their expectancies and actual selections. In analysing a variety of clinical populations relative to their control groups they have identified clusters of populations that vary along one or more of these parameters. The behaviour of VMpfc patients was best described by an over-attention to recent outcomes. The explanatory mechanism for this may be somatic markers, but a reversal learning deficit is equally possible. Other models have been able to account for IGT behaviour using a reinforcement learning model (Oya, Adolphs, Kawasaki, Bechara et al, 2005). To my knowledge none specifically factor somatic markers into their models leaving it as a possible contributory mechanism, and an additional free parameter, but given the problems with the IGT as a test of the SMH identified in the previous section and in 6.2.3, that its absence does not reduce the success of the models is informative.

With all the factors that contribute to IGT performance it is reasonable to ask whether the IGT adds anything to discussions of human decision-making. Other tasks better reflect components of this process than the IGT and as a result are more useful as diagnostic tools. Sanfey et al's (2003) task allows risk-seeking to be measured as does analysis of behaviour on the CBT (Rogers et al, 1999). In its standard analysis using net score the IGT cannot pinpoint the root cause of deficient performance. Examination of individual deck selection, as advocated here and by others (Yechiam et al, 2005b; Dunn et al, 2006) may help. But modifications of the

IGT have resulted in improved versions. In Bowman and Turnbull's (2004) Bangor Gambling Task (BGT) one deck is used where the probability of loss increases with selection. This task is simpler to administer yet allows measurement of decision-making in the face of increasing losses. Peters and Slovic's (2000) redesign of the IGT removed the confound of immediate gains and expected value. Examination of individual deck selection or selection across pairs of decks with similar features can better inform on the factors influencing choice. This test is a better test for the myopia for the future hypothesis as it does not require the additional administration of the variant or progressive IGT versions. However, because magnitude and probability of loss (and gain) create expected value it is possible that these factors will not allow dissociation of their influence on individual selection.

Another concern for the IGT is in its analysis. The deck selections are not truly independent observations and so measuring differences between them does violate one of the assumptions of ANOVA. Yet with only a few exceptions (Bolla et al, 2003; Harmsen et al, 2006), analyses have proceeded using parametric statistics. This concern is circumvented by using an additive measure like net score. However, this cumulative measure misses out on important differences between the alternatives it adds up. Analysis of individual deck selection in isolation is one analysis option but such multiple testing would inflate Type I error. This makes the analysis of IGT performance difficult as there is informative data within each alternative about how selection proceeds but comparisons across alternatives involves violation of a fundamental assumption of parametric data. This problem is not restricted to examination of differences between individual decks but also applies to collapsing across decks to compare between them. This has implications for the examination of physiological activity dependent on these choices. This issue was commented upon in

Chapter 5 where some participants either did not select disadvantageous cards in the period of interest (e.g. post-knowledge expression) or did but were unpunished (and therefore had no pSCR). While collapsing across deck type may miss physiological activity specific to individual decks, it is also problematic because of the nature of the IGT where learning is measured by a reduction in selections from disadvantageous decks with experience. This issue has rarely been mentioned in the literature but should receive consideration if the IGT is to continue being used as a test of decision-making. At the very least researchers should be aware of the consequences of such analyses.

6.2.3 Implications for the SMH

The evidence reviewed in Chapters 1 and 5 as well as the results reported in Chapter 2 and in Chapter 5 suggest that the IGT is not an adequate test of the SMH. This is in part because behaviour on the IGT can be explained by knowledge without recourse to somatic markers. Even if somatic markers do influence decisions it is by no means certain that anticipatory physiological activity is a somatic representation of the selection that is actually made. Equally, a sub-optimal behaviour on the IGT can be explained by a failure in the mechanisms involved in reinforcement learning (i.e. reversal learning). The neural architecture of this system and that proposed in the SMH are very similar although the interactions between some areas differ. The reinforcement learning system is well specified (e.g. Rolls, 2004) and offers a more parsimonious account for behaviour than taking a route through the periphery, or a neural representation thereof, in order to learn to avoid or select alternatives (Rolls, 2005). This is not to say that a somatic signal might not contribute, just that it is not necessary.

The very nature of SMH also makes it difficult to falsify, especially using a multi-component behavioural task like the IGT. An explanation in terms of the SMH can always fall back on the “as-if body” loop or a conscious integration of somatic signals and knowledge. The specificity and definition of the system may be one reason for its influence in the resurgent interest in the interaction between emotion and decision-making, but not being able to test it means its explanatory power is reduced.

Despite the increased specificity of the neural architecture questions about the operation of the SMH remain unanswered (Maia and McClelland, 2005). One question is the mechanism by which an option’s goodness or badness is determined. Damasio, Bechara and Damasio (2002) suggested that the differential somatic activity in aSCRs on the standard IGT indicates a warning that the bad decks are bad, whereas on Tomb et al’s (2002) modified task they signal that the good decks are good. But it has not been made clear how somatic activity signals alarm in one case and incentive in another. Indeed Schmitt et al (1999) suggested that if VMpfc damage results in an inability to link somatic activity with representations of the outcomes of alternatives then no behavioural bias should result. Bechara and Damasio (2005) have suggested the significance of an option is determined by the feeling it creates. This in turn implies that pleasure or pain are involved and takes the SMH into the territory of Rolls’ (1990, 1994) reinforcement learning. It also presents an interesting quandary when immediate pleasure and delayed pleasure are juxtaposed in a situation like the menu problem described in Chapter 5 despite the implication that the somatic system operates rationally. Bechara and Damasio (2005) suggest that when there are competing options, and their commensurate somatic states arise, the system operates by natural selection where the strongest somatic marker is the one that exerts

influence. The hypothetical arena in which somatic markers operate is further complicated by the addition of background somatic states. Bechara and Damasio (2005) suggest that when these are weak or neutral decisions are most sensitive to long-term consequences, but when the background activity is noisy or strong somatic markers that are in line with this activity are enhanced and those that are not are weakened. But this seems to be a restatement of Descartes' dualism reduced into somatic activity – when somatic (high emotion) activity is stronger it is harder to make decisions based on long-term consequences (reason). This increased detail in how the somatic system operates integrates it with so many other accounts of neural and mental operation that it becomes difficult to separate the SMH and thereby test it.

6.3 APPLICATION OF IGT TO OTHER CHOICE ENVIRONMENTS

One motivation for this thesis was to resolve an apparent paradox in behaviour on the IGT with behaviour on the Harvard Game. The two tasks are functionally equivalent, but learning of a maximization strategy was more rapid on the IGT than on the Harvard Game. Herrnstein et al (1993) suggested that the difference in when maximization or melioration dominated behaviour could be attributed to limitations in cognitive aspects of information processing. In order to maximize on the Harvard Game, decision makers must be able to represent the immediate outcome of an alternative, as well as have awareness of any internality between this outcome and long-term outcome (in the Harvard Game this is a between alternative internality; in the IGT it is a within alternative internality) and use this information to maximize. On the Harvard Game factors that made the internality easier to identify increased maximization behaviour.

On the IGT, Chapter 2 demonstrated the importance of giving participants information about the internality by telling them that some of the decks were better than others and staying away from the worst decks would lead to profit. Without this information, behaviour neither meliorated nor maximized. The behaviour on individual decks in Chapters 2 and 3 indicated that the internality is easier to acquire on decks A and D, but takes longer to acquire in deck B simply because the loss does not come early in selection with fixed reinforcement schedules. The importance of the first loss (or win) in the deck with the infrequent large magnitude reinforcer was illustrated in behaviour in the Chapter 4 when more realistic probabilistic reinforcement schedules were used. Chapter 5 demonstrated that participants have qualitative knowledge about the internality reflected in ratings of goodness or badness but that their ability to quantify it is limited. These results suggest that the choice environment of the IGT is simpler than on the Harvard Game and this is reflected in the faster learning rates.

While learning in the Experiments herein was still rapid relative to the Harvard Game, it was not as fast nor did it reach the levels reported by Bechara et al (1994, 1997). This has been true for other European research groups (e.g. Bowman and Turnbull, 2003; Van den Bos, Houx and Spruijt, 2006). This may reflect differences between using student samples versus older participants (generally middle-aged as they are matched to VMpfc patients in the Iowa group cohorts) who may be more motivated due to participating in medical research or for a large fee (larger than the incentives on offer with Real Money reinforcers). Whatever the root of this difference in population net scores it is possible to use the average slopes (learning rates) over the first session in the samples tested here to extrapolate across additional trials and predict when selection would reach a maximizing asymptote.

This assumes a linear relationship holds between net score and trial block (i.e. no exponential change in net score that may reflect conceptual understanding of the contingencies). Such extrapolations indicate that a maximising asymptote would be reached after 300 – 400 trials depending on instructions given which suggests that learning on the IGT may be closer to that in the Harvard Game than suggested from an examination of the “normal” performance observed by the Iowa group.

6.4 FUTURE EXPERIMENTS

The IGT seems likely to continue to be used as a test of decision-making in spite of the difficulty of narrowing the source of impairment from the variety of processes that contribute towards successful performance. This thesis has made important contributions to the future administration and analysis of the IGT. As well as highlighting the importance of the information participants are provided with about the task in their instructions, the administration of the IGT has been improved in a number of minor but important ways. For example, rather than presenting decks in the same order for all participants as in the standard administration deck position was counterbalanced across participants in the Experiments here. Also, the number of available selections from each deck was the same as the number of total selections meaning participants could make all their selections from one deck. This is an improvement on the initial (Bechara et al, 1994) administration with 40 cards in each deck, or later versions where 60 have been used (Bechara et al, 1999; 2000). This is important if individual deck selection, rather than selection of deck type, is of interest. Real Money reinforcers have been used as standard, despite the evidence that at their current level they do not impact on behaviour.

Contributions have also been made in the analysis of behaviour where the importance of looking at learning rate and the change in individual deck selection has been demonstrated. Similarly it is not enough to compare populations on the net score measure. These measures within groups should be compared with chance to determine if learning has occurred. In many clinical populations learning may develop at a slower pace than in the healthy controls they are compared with.

The work in this thesis has generated questions as well as answers about behaviour on the IGT and its relationship with the SMH. Future experiments that extend from it fall into the two categories explored in sections 6.2.1 and 6.2.2 – using the IGT to test the SMH and exploring decision-making using the IGT. It has been argued that the IGT is not a good test of the SMH. While this is true direct tests of the SMH may be possible by attempting to disrupt the putative somatic markers and observing the effect on decision-making as measured in IGT performance. This has also been suggested by Bechara and Damasio (2005, p. 351) who report preliminary data showing reduced selections from decks C and D when emotional imagery was invoked (by asking participants to recall personally emotional experiences) relative to neutral imagery. Another potential test would be to disrupt or augment somatic states by presentation of consistent or inconsistent emotional imagery in the form of rapidly presented visual stimuli after initial deck selection. Memory and other cognitive processes have been probed using images from the International Affective Picture System bank of emotional images. Other manipulations may be possible by simulating an anticipatory somatic response with tactile stimulation (e.g. mild electric shock or a more pleasant equivalent). Should sub-optimal performance on the IGT result it would provide some support for a somatic role in decision-making.

An important follow-up study to separate the Somatic-Knowledge and Knowledge-Somatic hypotheses more clearly would be to combine the instruction conditions from Experiment 2 and the specific questioning from Experiment 9. The development of knowledge in the No Hint condition could be important, especially if it occurs in tandem with physiological changes (not limited to SCR but including heart rate monitoring and perhaps even pupil dilation to measure arousal). The importance of when hint instructions are delivered is also worth exploring further. The effect of Hint largely occurred in the second session although it was predicted from session 1 learning rates, and may have been augmented by repetition between sessions. Emphasising the hint in the initial instructions may be one reason for the differences in performance in these experiments and in those of the Iowa group.

However, as discussed above many factors contribute to IGT behaviour. Their relative importance and relationship to personality factors should also be examined in more detail. While this may not inform on the SMH, more understanding of the factors influencing IGT behaviour in a normal population is useful when it forms the baseline to which clinical populations are compared. One such factor should be individual discounting functions. This is important when considering a participants' motivation to profit on the task when Real Money is used, especially given results from experimental economics where this has been associated with improved performance relative to other reinforcers or no reinforcer at all (Camerer and Hogarth, 1999; Hertwig and Ortman, 2001, Tunney and Shanks, 2002; Brase et al, 2006). It may be that the absence of an effect of Reinforcer Type in Chapter 2 was because, as a fraction of the minimum wage, they were not motivating enough. By examining individual discount rates it would be possible to tailor IGT contingencies to individuals and so be more certain that the task was motivating. Alternatively, values

either side of individuals' indifference points between immediate and delayed reward could be assigned to each deck type.

Effects of manipulating the rewards on offer have recently been documented. Van den Bos et al (2006) manipulated the relationship between gains in the advantageous and disadvantageous decks and found that learning was slower as the magnitude of gains in A and B were increased. Franken, Georgieva, Muris and Dijksterhuis (2006) manipulated the frame in which the IGT is approached by manipulating losses or wins in a previous task. More risky decisions followed prior losses, while safer choices followed gains. Differences in delay discounting have been related to real-life behaviour within students (Silva and Gross, 2004).

In any of these future studies additional information on participants' personality and individual differences should be collected to allow their influence to be partialled out of the relationship of interest as well as to assess their relative impact on performance. Only in this way will a full picture of the factors that are important for performance on the IGT be obtained.

6.5 CONCLUSION

A strong version of the SMH assumes that somatic markers are necessary for decision-making to be successful, that they influence choice in the absence of, or prior to, consciously available knowledge, and that the somatic system maximizes. These assumptions are integral to an interpretation of IGT behaviour in terms of the SMH, and indeed this behaviour has been used as key support for the SMH. This thesis has explored these assumptions as they relate to behaviour on the IGT and concluded that they cannot be supported from a detailed examination of this behaviour. The

importance of information about the IGT in the form of hint instructions has suggested knowledge about the task may be a more critical factor than any somatic input. When somatic activity was recorded using SCRs no evidence of a somatic influence prior to the emergence of consciously accessible knowledge sufficient to guide behaviour was found. Instead there were indications that a Knowledge-Somatic hypothesis may explain somatic activity on the IGT.

The assumption that selection on the IGT is solely determined by expected value was not supported by examining behaviour. Instead, it appeared that the crucial selection issue for participants was learning to avoid deck B due to its infrequent large loss. A similar learned preference for deck C in the advantageous decks was also suggested but support was not as strong across Experiments and variations of the standard task.

This thesis provides more evidence against the ability of the IGT to serve as a test of the SMH. It is satisfying to know that a recent review of the IGT literature has independently come to a similar conclusion (Dunn et al, 2006, p. 258). While the IGT should not be used as a definitive test of normative decision-making, it is useful within a battery of such tests, especially if behaviour is examined in detail. However, performance on it should not be used as a diagnostic tool from which to conclude deficient VMPfc functioning or myopic decision-making.

But the IGT does provide another example that normal human behaviour is rational when faced with a complex choice environment in which information about the short and long-term contingencies can be learned with experience. That learning in repeated choice environments such as the IGT proceeds on the basis of available information without recourse to an unconscious, and pre-emptive, emotional biasing system is reassuring. Full comprehension is not required for qualitative distinctions to

be made between available reinforcement contingencies. With such heuristics, learning can occur. That such environments model real-life (Rachlin, 1989) implies that humans remain the masters of their destiny.

APPENDIX A

FIXED SCHEDULE OF LOSSES USED IN EXPERIMENTS 1, 2, 3, 4, 5 AND 6

Card number	A	B	C	D
1	0	0	0	0
2	0	0	0	0
3	15	0	5	0
4	0	0	0	0
5	30	0	5	0
6	0	0	0	0
7	20	0	5	0
8	0	0	0	0
9	25	125	5	0
10	35	0	5	25
11	0	0	0	0
12	35	0	2.5	0
13	0	0	7.5	0
14	25	125	0	0
15	20	0	0	0
16	0	0	0	0
17	30	0	2.5	0
18	15	0	7.5	0
19	0	0	0	0
20	0	0	5	25
21	0	125	0	0
22	30	0	0	0
23	0	0	0	0
24	35	0	5	0
25	0	0	2.5	0
26	20	0	5	0
27	25	0	0	0
28	15	0	0	0
29	0	0	7.5	25
30	0	0	5	0
31	35	0	0	0
32	20	125	0	0
33	25	0	0	0
34	0	0	2.5	0
35	0	0	2.5	25
36	0	0	0	0
37	15	0	7.5	0
38	30	0	0	0
39	0	0	5	0
40	0	0	7.5	0
41	20	0	0	0
42	0	0	0	0
43	0	0	7.5	0
44	25	0	0	0
45	15	0	2.5	25
46	0	125	0	0

47	0	0	5	0
48	30	0	5	0
49	35	0	5	0
50	0	0	0	0
51	35	0	2.5	0
52	20	0	0	0
53	25	0	2.5	0
54	0	0	0	0
55	15	0	7.5	0
56	0	0	0	0
57	0	0	7.5	0
58	30	125	0	25
59	0	0	0	0
60	0	0	5	0
61	0	0	0	0
62	0	0	0	0
63	15	0	5	0
64	0	0	0	0
65	30	0	5	0
66	0	0	0	0
67	20	0	5	0
68	0	0	0	0
69	25	125	5	0
70	35	0	5	25
71	0	0	0	0
72	35	0	2.5	0
73	0	0	7.5	0
74	25	125	0	0
75	20	0	0	0
76	0	0	0	0
77	30	0	2.5	0
78	15	0	7.5	0
79	0	0	0	0
80	0	0	5	25
81	0	125	0	0
82	30	0	0	0
83	0	0	0	0
84	35	0	5	0
85	0	0	2.5	0
86	20	0	5	0
87	25	0	0	0
88	15	0	0	0
89	0	0	7.5	25
90	0	0	5	0
91	35	0	0	0
92	20	125	0	0
93	25	0	0	0
94	0	0	2.5	0
95	0	0	2.5	25
96	0	0	0	0
97	15	0	7.5	0
98	30	0	0	0
99	0	0	5	0
100	0	0	7.5	0

APPENDIX B

FIXED SCHEDULE OF LOSSES USED IN THE MANIPULATED DECKS OF

EXPERIMENT 5

Card number	A Decreased loss frequency	B Increased loss frequency	C Decreased loss frequency	D Increased loss frequency
1	0	0	0	0
2	0	0	0	0
3	35	0	5	0
4	0	0	0	0
5	55	0	0	0
6	0	0	0	0
7	0	35	12.5	7.5
8	0	55	0	0
9	0	0	0	12.5
10	35	35	7.5	5
11	35	0	0	0
12	0	55	0	0
13	55	0	0	0
14	35	0	0	5
15	0	0	0	5
16	0	0	7.5	0
17	0	35	0	15
18	0	0	7.5	0
19	0	0	0	0
20	0	35	10	0
21	0	0	5	0
22	55	0	0	0
23	0	0	0	0
24	0	0	0	0
25	0	0	10	0
26	35	0	0	12.5
27	0	35	10	0
28	0	0	0	7.5
29	35	35	0	5
30	0	55	0	0
31	35	0	10	0
32	0	35	7.5	5
33	35	55	0	7.5
34	0	35	0	12.5
35	0	0	0	0
36	55	0	0	0
37	0	0	0	0
38	0	0	0	0
39	0	0	0	0
40	0	0	7.5	0
41	0	0	5	0
42	55	0	0	0
43	0	0	12.5	10

44	35	55	0	10
45	0	0	7.5	5
46	0	0	0	0
47	0	35	0	0
48	0	0	0	0
49	0	0	0	0
50	35	35	0	0
51	0	0	0	0
52	0	0	0	0
53	35	35	0	0
54	0	0	0	0
55	0	0	10	0
56	35	0	0	0
57	0	55	7.5	0
58	0	35	7.5	12.5
59	0	0	0	7.5
60	55	0	0	5
61	0	0	0	0
62	0	55	0	5
63	0	35	0	7.5
64	0	0	0	0
65	0	0	0	0
66	35	0	5	0
67	55	0	0	0
68	0	0	10	0
69	0	0	10	0
70	35	35	0	12.5
71	0	35	0	0
72	0	0	0	5
73	55	0	0	10
74	0	35	12.5	0
75	0	55	0	0
76	0	0	0	0
77	0	0	7.5	10
78	0	0	5	0
79	35	0	0	0
80	35	0	0	0
81	0	0	0	0
82	35	35	0	0
83	35	0	0	0
84	0	0	0	12.5
85	0	55	0	0
86	55	35	10	7.5
87	0	0	0	0
88	0	0	0	0
89	0	0	7.5	0
90	0	0	7.5	5
91	0	0	0	0
92	35	0	0	5
93	55	0	0	10
94	35	35	5	10
95	0	0	10	0
96	0	35	0	0
97	0	0	0	0

98	0	55	0	0
99	0	0	0	0
100	0	0	10	0

APPENDIX C

FIXED SCHEDULE OF WINS ON THE VARIANT IGT USED IN EXPERIMENT

7

Card number	E	F	G	H
1	0	0	0	0
2	0	0	0	0
3	0	5	15	0
4	0	0	0	0
5	0	5	30	0
6	0	0	0	0
7	0	5	20	0
8	0	0	0	0
9	125	5	25	0
10	0	5	35	25
11	0	0	0	0
12	0	2.5	35	0
13	0	7.5	0	0
14	125	0	25	0
15	0	0	20	0
16	0	0	0	0
17	0	2.5	30	0
18	0	7.5	15	0
19	0	0	0	0
20	0	5	0	25
21	125	0	0	0
22	0	0	30	0
23	0	0	0	0
24	0	5	35	0
25	0	2.5	0	0
26	0	5	20	0
27	0	0	25	0
28	0	0	15	0
29	0	7.5	0	25
30	0	5	0	0
31	0	0	35	0
32	125	0	20	0
33	0	0	25	0
34	0	2.5	0	0
35	0	2.5	0	25
36	0	0	0	0
37	0	7.5	15	0
38	0	0	30	0
39	0	5	0	0
40	0	7.5	0	0
41	0	0	20	0
42	0	0	0	0
43	0	7.5	0	0
44	0	0	25	0

45	0	2.5	15	25
46	125	0	0	0
47	0	5	0	0
48	0	5	30	0
49	0	5	35	0
50	0	0	0	0
<hr/>				
51	0	2.5	35	0
52	0	0	20	0
53	0	2.5	25	0
54	0	0	0	0
55	0	7.5	15	0
56	0	0	0	0
57	0	7.5	0	0
58	125	0	30	25
59	0	0	0	0
60	0	5	0	0
<hr/>				
61	0	0	0	0
62	0	0	0	0
63	0	5	15	0
64	0	0	0	0
65	0	5	30	0
66	0	0	0	0
67	0	5	20	0
68	0	0	0	0
69	125	5	25	0
70	0	5	35	25
<hr/>				
71	0	0	0	0
72	0	2.5	35	0
73	0	7.5	0	0
74	125	0	25	0
75	0	0	20	0
76	0	0	0	0
77	0	2.5	30	0
78	0	7.5	15	0
79	0	0	0	0
80	0	5	0	25
<hr/>				
81	125	0	0	0
82	0	0	30	0
83	0	0	0	0
84	0	5	35	0
85	0	2.5	0	0
86	0	5	20	0
87	0	0	25	0
88	0	0	15	0
89	0	7.5	0	25
90	0	5	0	0
<hr/>				
91	0	0	35	0
92	125	0	20	0
93	0	0	25	0
94	0	2.5	0	0
95	0	2.5	0	25
96	0	0	0	0
97	0	7.5	15	0
98	0	0	30	0

99	0	5	0	0
100	0	7.5	0	0

APPENDIX D

QUESTIONS USED IN THE SPECIFIC QUESTION CONDITION OF

EXPERIMENT 9

Q1. Rate on a scale of -10 to +10, how good or bad you think deck A is, where -10 means that it is terrible and +10 means that it is excellent.

Rate on a scale of -10 to +10, how good or bad you think deck B is, where -10 means that it is terrible and +10 means that it is excellent.

Rate on a scale of -10 to +10, how good or bad you think deck C is, where -10 means that it is terrible and +10 means that it is excellent.

Rate on a scale of -10 to +10, how good or bad you think deck D is, where -10 means that it is terrible and +10 means that it is excellent.

Q2. Okay, why did you rate deck A with ...?
Okay, why did you rate deck B with ...?
Okay, why did you rate deck C with ...?
Okay, why did you rate deck D with ...?

Q3. In answering the questions that follow, consider the following definitions. Your "winning amount" for a trial is the amount you won on that trial. Your "loss" on a trial is the amount you lost on that trial. Your "net result" for a trial is the amount you won minus the amount you lost on that trial. Do you understand these definitions and the differences between the three terms? [If not, explain again using examples.]

Okay, now suppose you were to select 10 cards from deck A.

Q3.1. What would you expect your average net result to be?

Q3.2. What would you expect your average winning amount to be?

Q3.3. In how many of the 10 trials would you expect to get a loss (not necessarily a net loss)?

Q3.4. For those trials in which you would get a loss, what would you expect the average loss to be?

[Repeat question Q3 for decks B through D]

Q4. Okay, now tell me, on a scale of 0 to 100, how much you think that you know what you should do in this game to win as much money as possible (or, if you can't win, to avoid losing as much as possible). 0 means that you have no idea of what you should do and feel that you still need to explore the game more and 100 means that you know exactly what you should do and have no doubts that that would be the best strategy.

Q5. Now suppose I told you that you could only select cards from one of the decks until the end of the game, but that you were allowed to choose now the deck from which you would draw your cards. Which of the four decks would you pick?

REFERENCES

- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine. *Econometrica*, 21, 503-546.
- Amiez, C., Procyk, E., Honore, J., Sequeira, H., & Joseph, J. P. (2003). Reward anticipation, cognition, and electrodermal activity in the conditioned monkey. *Experimental Brain Research*, 149(3), 267-275.
- Annoni, J. M., Ptak, R., Caldara-Schnetzer, A. S., Khateb, A., & Pollermann, B. Z. (2003). Decoupling of autonomic and cognitive emotional reactions after cerebellar stroke. *Annals Of Neurology*, 53(5), 654-658.
- Barnes, J. (1998a). Psyche. In R. L. Gregory & O. L. Zangwill (Eds.), *The oxford companion to the mind* (2nd ed.). Oxford: Oxford Univeristy Press.
- Barnes, J. (1998b). Plato. In R. L. Gregory & O. L. Zangwill (Eds.), *The oxford companion to the mind* (2nd ed.). Oxford: Oxford Univeristy Press.
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal Of Behavioral Decision Making*, 16(3), 215-233.
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50(1-3), 7-15.
- Bechara, A., Tranel, D., Damasio, H., & Damasio, A. R. (1996). Failure to respond autonomically to anticipated future outcomes following damage to prefrontal cortex. *Cerebral Cortex*, 6(2), 215-225.

- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997a). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*, 1293-1295.
- Bechara, A., Tranel, D., Damasio, H., & Damasio, A. R. (1997b). An anatomical system subserving decision-making. *Society for Neuroscience Abstracts*, *23*, 495.
- Bechara, A., Damasio, H., Tranel, D., & Anderson, S. W. (1998). Dissociation of working memory from decision making within the human prefrontal cortex. *Journal Of Neuroscience*, *18*(1), 428-437.
- Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *Journal Of Neuroscience*, *19*(13), 5473-5481.
- Bechara, A., Tranel, D., & Damasio, H. (2000). Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions. *Brain*, *123*, 2189-2202.
- Bechara, A., Dolan, S., Denburg, N., Hinds, A., Anderson, S. W., & Nathan, P. E. (2001). Decision-making deficits, linked to a dysfunctional ventromedial prefrontal cortex, revealed in alcohol and stimulant abusers. *Neuropsychologia*, *39*(4), 376-389.
- Bechara, A., & Damasio, H. (2002). Decision-making and addiction (part i): Impaired activation of somatic states in substance dependent individuals when pondering decisions with negative future consequences. *Neuropsychologia*, *40*(10), 1675-1689.

- Bechara, A., Dolan, S., & Hinds, A. (2002). Decision-making and addiction (part ii): Myopia for the future or hypersensitivity to reward? *Neuropsychologia*, *40*(10), 1690-1705.
- Bechara, A. (2003). Decisions, uncertainty, and the brain: The science of neuroeconomics. *Journal Of Clinical And Experimental Neuropsychology*, *25*(7), 1035-1037.
- Bechara, A., & Martin, E. M. (2004). Impaired decision making related to working memory deficits in individuals with substance addictions. *Neuropsychology*, *18*(1), 152-162.
- Bechara, A., & Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games And Economic Behavior*, *52*(2), 336-372.
- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (2005). The iowa gambling task and the somatic marker hypothesis: Some questions and answers. *Trends In Cognitive Sciences*, *9*(4), 159-162.
- 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*, 266-286.
- Bernoulli, D. (1738/1954). Specimen theoriae novae de mensura sortis. Pp. 175-192 in commentarii academiae scientiarum imperialis petropolitanae, tomus v (papers of the imperial academy of science in petersburg, vol. V). *Econometrica*, *22*, 23-26.

- Bickel, W. K., & Marsch, L. A. (2001). Toward a behavioral economic understanding of drug dependence: Delay discounting processes. *Addiction*, *96*(1), 73-86.
- Blair, R. J. R., & Cipolotti, L. (2000). Impaired social response reversal - a case of 'acquired sociopathy'. *Brain*, *123*, 1122-1141.
- Blair, R. J. R., Colledge, E., & Mitchell, D. G. V. (2001). Somatic markers and response reversal: Is there orbitofrontal cortex dysfunction in boys with psychopathic tendencies? *Journal Of Abnormal Child Psychology*, *29*(6), 499-511.
- Bolla, K. I., Eldreth, D. A., London, E. D., Kiehl, K. A., Mouratidis, M., Contoreggi, C., et al. (2003). Orbitofrontal cortex dysfunction in abstinent cocaine abusers performing a decision-making task. *Neuroimage*, *19*(3), 1085-1094.
- Bowman, C. H., & Turnbull, O. H. (2003). Real versus facsimile reinforcers on the iowa gambling task. *Brain And Cognition*, *53*(2), 207-210.
- Bowman, C. H., Evans, C. E. Y., & Turnbull, O. H. (2005). Artificial time constraints on the iowa gambling task: The effects on behavioural performance and subjective experience. *Brain And Cognition*, *57*(1), 21-25.
- Brase, G. L., Fiddick, L., & Harries, C. (2006). Participant recruitment methods and statistical reasoning performance. *Quarterly Journal Of Experimental Psychology*, *59*(5), 965-976.
- Buehner, M. J., & May, J. (2004). Abolishing the effect of reinforcement delay on human causal learning. *The Quarterly Journal of Experimental Psychology*, *57B*(2), 179 - 191.

- Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the bechara gambling task. *Psychological Assessment, 14*(3), 253-262.
- Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal Of Risk And Uncertainty, 19*(1-3), 7-42.
- Campbell, M. C., Stout, J. C., & Finn, P. R. (2004). Reduced autonomic responsiveness to gambling task losses in huntington's disease. *Journal Of The International Neuropsychological Society, 10*(2), 239-245.
- Cannon, W. B. (1927). The james-lange theory of emotion: A critical examination and an alternative theory. *American Journal of Psychology, 39*, 106-124.
- Carter, S., & Pasqualini, M. C. S. (2004). Stronger autonomic response accompanies better learning: A test of damasio's somatic marker hypothesis. *Cognition & Emotion, 18*(7), 901-911.
- Cavedini, P., Riboldi, G., D'Annuncci, A., Belotti, P., Cisima, M., & Bellodi, L. (2002). Decision-making heterogeneity in obsessive-compulsive disorder: Ventromedial prefrontal cortex function predicts different treatment outcomes. *Neuropsychologia, 40*(2), 205-211.
- Cella, M., Dymond, S., Cooper, A., & Turnbull, O. H. (2007). Effects of decision-phase time constraints on emotion-based learning in the iowa gambling task. *Brain and Cognition, 64*(2), 164-169.

- Churchland, P. M. (1993). *Matter and consciousness: A contemporary introduction to the philosophy of mind (revised edition)*. Cambridge, Massachusetts: MIT Press.
- Clark, L., & Robbins, T. W. (2002). Decision-making deficits in drug addiction. *Trends In Cognitive Sciences*, 6(9), 361-363.
- Clark, L., & Manes, F. (2004). Social and emotional decision-making following frontal lobe injury. *Neurocase*, 10, 398-403.
- Crone, E. A., & van der Molen, M. W. (2004). Developmental changes in real life decision making: Performance on a gambling task previously shown to depend on the ventromedial prefrontal cortex. *Developmental Neuropsychology*, 25(3), 251-279.
- Crone, E. A., Somsen, R. J. M., Van Beek, B., & van der Molen, M. W. (2004). Heart rate and skin conductance analysis of antecedents and consequences of decision making. *Psychophysiology*, 41(4), 531-540.
- Crone, E. A., Bunge, S. A., Latenstein, H., & van der Molen, M. W. (2005). Characterization of children's decision making: Sensitivity to punishment frequency, not task complexity. *Child Neuropsychology*, 11(3), 245-263.
- Damasio, A. R., Tranel, D., & Damasio, H. (1991). Somatic markers and the guidance of behaviour: Theory and preliminary testing. In H. S. Levin, H. M. Eisenberg & A. L. Benton (Eds.), *Frontal lobe function and dysfunction*. (pp. 217-229). New York: Oxford University Press.

- Damasio, A. R. (1994). *Descartes error: Emotion, reason and the human brain*. New York: Avon.
- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions Of The Royal Society Of London Series B-Biological Sciences*, 351(1346), 1413-1420.
- Damasio, A. R. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. New York: Harcourt.
- Damasio, H., Bechara, A., & Damasio, A. R. (2002). Do somatic markers mediate decisions on the gambling task? Reply. *Nature Neuroscience*, 5(11), 1104-1104.
- Dawson, M. E., Schell, A. M., & Filion, D. L. (1990). The electrodermal system. In J. T. Cacioppo & L. G. Tassinary (Eds.), *Principles of psychophysiology: Physical, social, and inferential elements* (pp. 295-324). New York: Cambridge University Press.
- Dempster, F. N., &. (1996). Distributing and managing the conditions of encoding and practice. *In Memory*, 317 - 344, Bjork, E.L. and Bjork, R.A. (eds). *Academic Press: San Diego*.
- Descartes, R. (1965). Rene descartes (1596 - 1650) on mechanism in human action. In R. J. Herrnstein & E. G. Boring (Eds.), *A source book in the history of psychology* (pp. 266-272). Cambridge, Massachusetts: Harvard University Press.

- Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience And Biobehavioral Reviews*, 30(2), 239-271.
- Edwards, W. (1954). The theory of decision making. *Psychological Bulletin*, 51(4), 380-417.
- Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4), 848-881.
- Evans, C. E. Y., Kemish, K., & Turnbull, O. H. (2004). Paradoxical effects of education on the iowa gambling task. *Brain And Cognition*, 54(3), 240-244.
- Evans, C. E. Y., Bowman, C. H., & Turnbull, O. H. (2005). Subjective awareness on the iowa gambling task: The key role of emotional experience in schizophrenia. *Journal Of Clinical And Experimental Neuropsychology*, 27(6), 656-664.
- Fellows, L. K., & Farah, M. J. (2003). Ventromedial frontal cortex mediates affective shifting in humans: Evidence from a reversal learning paradigm. *Brain*, 126, 1830-1837.
- Fellows, L. K., & Farah, M. J. (2005). Different underlying impairments in decision-making following ventromedial and dorsolateral frontal lobe damage in humans. *Cerebral Cortex*, 15(1), 58-63.

- Fernie, G., & Tunney, R. J. (2006). Some decks are better than others: The effect of reinforcer type and task instructions on learning in the iowa gambling task. *Brain And Cognition*, 60(1), 94-102.
- Fernie, G. and Tunney, R. J. (2008). Decision-making in the Iowa Gambling Task. In F. Columbus (Ed.). *The Psychology of Decision-Making*. New York: Nova.
- Fiorina, M. P. (1971). Critique and comment. A note on probability matching and rational choice. *Behavioural Science*, 16, 158-166.
- Fischer, A. R. H., Blommaert, F. J. J., & Midden, C. J. H. (2005). Combining experimental observations and modelling in investigating feedback and emotions in repeated selection tasks. *User Modeling And User-Adapted Interaction*, 15(5), 389-424.
- Fishbein, D., Hyde, C., Eldreth, D., London, E. D., Matochik, J., Ernst, M., et al. (2005). Cognitive performance and autonomic reactivity in abstinent drug abusers and nonusers. *Experimental And Clinical Psychopharmacology*, 13(1), 25-40.
- Franken, I. H. A., Georgieva, I., Muris, P., & Dijksterhuis, A. (2006). The rich get richer and the poor get poorer: On risk aversion in behavioral decision-making. *Judgement and Decision Making*, 1(2), 153-158.
- Frith, C. (1995). *The cognitive neuropsychology of schizophrenia*. London: Psychology Press.
- Garon, N., & Moore, C. (2004). Complex decision-making in early childhood. *Brain And Cognition*, 55(1), 158-170.

- Goel, V., Grafman, J., Tajik, J., Gana, S., & Danto, D. (1997). A study of the performance of patients with frontal lobe lesions in a financial planning task. *Brain, 120*, 1805-1822.
- Grant, S., Contoreggi, C., & London, E. D. (2000). Drug abusers show impaired performance in a laboratory test of decision making. *Neuropsychologia, 38*(8), 1180-1187.
- Gray, J. A. (1970). Psychophysiological basis of introversion-extraversion. *Behaviour Research And Therapy, 8*(3), 249-266.
- Greenberg, & Weiner, B. (1966). Effects of reinforcement history upon risk-taking behavior. *Journal Of Experimental Psychology, 71*(4), 587-592.
- Gutbrod, K., Krouzel, C., Hofer, H., Muri, R., Perrig, W., & Ptak, R. (2006). Decision-making in amnesia: Do advantageous decisions require conscious knowledge of previous behavioural choices? *Neuropsychologia, 44*(8), 1315-1324.
- Harmsen, H., Bischof, G., Brooks, A., Hohagen, F., & Rumpf, H. J. (2006). The relationship between impaired decision-making, sensation seeking and readiness to change in cigarette smokers. *Addictive Behaviors, 31*(4), 581-592.
- Heims, H. C., Critchley, H. D., Dolan, R., Mathias, C. J., & Cipolotti, L. (2004). Social and motivational functioning is not critically dependent on feedback of autonomic responses: Neuropsychological evidence from patients with pure autonomic failure. *Neuropsychologia, 42*(14), 1979-1988.

- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal Of The Experimental Analysis Of Behavior*, 4(3), 267-272.
- Herrnstein, R. J. (1990). Rational choice theory - necessary but not sufficient. *American Psychologist*, 45(3), 356-367.
- Herrnstein, R. J. (1997). *The matching law: Paper in psychology and economics* (1st ed.). New York: Russell Sage Foundation.
- Herrnstein, R. J., & Vaughan, W. J. (1980). Melioration and behavioural allocation. In J. E. R. Sradon (Ed.), *Limits to action: The allocation of individual behaviour*. New York: Academic Press.
- Herrnstein, R. J., & Prelec, D. (1991). Melioration - a theory of distributed choice. *Journal Of Economic Perspectives*, 5(3), 137-156.
- Herrnstein, R. J., & Prelec, D. (1992). A theory of addiction. In H. Rachlin & D. I. Laibson (Eds.), *The matching law* (pp. 160-187). London, England: Harvard University Press.
- Herrnstein, R. J., Loewenstein, G. F., Prelec, D., & Vaughan, W. (1993). Utility maximization and melioration - internalities in individual choice. *Journal Of Behavioral Decision Making*, 6(3), 149-185.
- Hertwig, R., & Ortmann, A. (2001). Experimental practices in economics: A methodological challenge for psychologists? *Behavioral And Brain Sciences*, 24(3), 383-451.

- Hinson, J. M., Jameson, T. L., & Whitney, P. (2003). Impulsive decision making and working memory. *Journal Of Experimental Psychology-Learning Memory And Cognition*, 29(2), 298-306.
- Hornak, J., O'Doherty, J., Bramham, J., Rolls, E. T., Morris, R. G., Bullock, P. R., et al. (2004). Reward-related reversal learning after surgical excisions in orbito-frontal or dorsolateral prefrontal cortex in humans. *Journal Of Cognitive Neuroscience*, 16(3), 463-478.
- Izquierdo, A., Suda, R. K., & Murray, E. A. (2004). Bilateral orbital prefrontal cortex lesions in rhesus monkeys disrupt choices guided by both reward value and reward contingency. *Journal Of Neuroscience*, 24(34), 7540-7548.
- James, W. (1884). What is an emotion? *mind*, 19, 188-205.
- Jameson, T. L., Hinson, J. M., & Whitney, P. (2004). Components of working memory and somatic markers in decision making. *Psychonomic Bulletin & Review*, 11(3), 515-520.
- Jones, M. R., & Myers, J. L. (1966). A comparison of 2 methods of event randomization in probability learning. *Journal Of Experimental Psychology*, 72(6), 909-&.
- Kahneman, D., & Tversky, A. (1979). Prospect theory - analysis of decision under risk. *Econometrica*, 47(2), 263-291.
- Keppel, G. (1991). *Design and analysis: A researcher's handbook*. Englewood Cliffs, New Jersey: Prentice-Hall, Inc.

- Kerr, A., & Zelazo, P. D. (2004). Development of "hot" executive function: The children's gambling task. *Brain And Cognition*, 55(1), 148-157.
- Kleeberg, J., Bruggimann, L., Annoni, J. M., van Melle, G., Bogousslavsky, J., & Schlupe, M. (2004). Altered decision-making in multiple sclerosis: A sign of impaired emotional reactivity? *Annals Of Neurology*, 56(6), 787-795.
- Kudadjie-Gyamfi, E., & Rachlin, H. (2002). Rule-governed versus contingency-governed behaviour in a self-control task: Effects of changes in contingencies. *Behavioural Processes*, 57, 29-35.
- Lange, C. (1885). The mechanisms of emotion. In E. Dunlap (Ed.), *The emotions* (pp. 33–90). Baltimore, Maryland: Williams and Wilkins.
- Leland, D. S., Richardson, J. S., Vankov, A., Grant, S. J., & Pineda, J. A. (1999). P300 and the iowa gambling task: Neural basis of decision-making and addiction. *Journal Of Cognitive Neuroscience*, 13-13.
- Levy, R. and Goldman-Rakic, P.S., (2000). Segregation of working memory functions within the dorsolateral prefrontal cortex. *Experimental brain research*, 133 (1), 23 -32.
- Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal Of Experimental Psychology*, 89(1), 46-&.
- Lichtenstein, S., & Slovic, P. (1973). Response-induced reversals of preference in gambling - extended replication in las-vegas. *Journal Of Experimental Psychology*, 101(1), 16-20.

- Lorch, R. F., & Myers, J. L. (1990). Regression-analyses of repeated measures data in cognitive research. *Journal Of Experimental Psychology-Learning Memory And Cognition*, *16*(1), 149-157.
- MacPherson, S. E., Phillips, L. H., & Della Sala, S. (2002). Age, executive function, and social decision making: A dorsolateral prefrontal theory of cognitive aging. *Psychology And Aging*, *17*(4), 598-609.
- Maia, T. V., & McClelland, J. L. (2004). A reexamination of the evidence for the somatic marker hypothesis: What participants really know in the iowa gambling task. *Proceedings Of The National Academy Of Sciences Of The United States Of America*, *101*(45), 16075-16080.
- Mazas, C. A., Finn, P. R., & Steinmetz, J. E. (2000). Decision-making biases, antisocial personality, and early-onset alcoholism. *Alcoholism-Clinical And Experimental Research*, *24*(7), 1036-1040.
- McGinn, C. (2003). Fear factor: A review of looking for spinoza by antonio damasio. *New York Times Book Review*, 2011.
- Mintzer, M. Z., & Stitzer, M. L. (2002). Cognitive impairment in methadone maintenance patients. *Drug And Alcohol Dependence*, *67*(1), 41-51.
- Mitchell, D. G. V., Colledge, E., Leonard, A., & Blair, R. J. R. (2002). Risky decisions and response reversal: Is there evidence of orbitofrontal cortex dysfunction in psychopathic individuals? *Neuropsychologia*, *40*(12), 2013-2022.

- Monterosso, J., Ehrman, R., Napier, K. L., O'Brien, C. P., & Childress, A. R. (2001). Three decision-making tasks in cocaine-dependent patients: Do they measure the same construct? *Addiction, 96*(12), 1825-1837.
- Myers, J. L. (1976). Probability learning and sequence learning. In W. K. Estes (Ed.), *Handbook of learning and cognitive processes: Approaches to human learning and motivation*. (pp. 171-205). Hillsdale, NJ: Erlbaum.
- Nahm, F. K. D., Tranel, D., Damasio, H., & Damasio, A. R. (1993). Cross-modal associations and the human amygdala. *Neuropsychologia, 31*(8), 727-&.
- Neufeld, R. W. J., Vollick, D., Carter, J. R., Boksman, K., & Jette, J. (2002). Application of stochastic modeling to the assessment of group and individual differences in cognitive functioning. *Psychological Assessment, 14*(3), 279-298.
- Neumark, E. D., & Shuford, E. H. (1959). Comparison of predictions and estimates in a probability learning situation. *Journal Of Experimental Psychology, 57*, 294-298.
- Nielen, M. M. A., Veltman, D. J., de Jong, R., Mulder, G., & den Boer, J. A. (2002). Decision making performance in obsessive compulsive disorder. *Journal Of Affective Disorders, 69*(1-3), 257-260.
- North, N. T., & O'Carroll, R. E. (2001). Decision making in patients with spinal cord damage: Afferent feedback and the somatic marker hypothesis. *Neuropsychologia, 39*(5), 521-524.

- O'Doherty, J., Kringelbach, M. L., Rolls, E. T., Hornak, J., & Andrews, C. (2001). Abstract reward and punishment representations in the human orbitofrontal cortex. *Nature Neuroscience*, *4*(1), 95-102.
- Ongur, D., & Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral Cortex*, *10*(3), 206-219.
- Overman, W. H. (2004). Sex differences in early childhood, adolescence, and adulthood on cognitive tasks that rely on orbital prefrontal cortex. *Brain And Cognition*, *55*(1), 134-147.
- Oya, H., Adolphs, R., Kawasaki, H., Bechara, A., Damasio, A., & Howard, M. A. (2005). Electrophysiological correlates of reward prediction error recorded in the human prefrontal cortex. *Proceedings Of The National Academy Of Sciences Of The United States Of America*, *102*(23), 8351-8356.
- Patterson, J. C., Ungerleider, L. G., & Bandettini, P. A. (2002). Task-independent functional brain activity correlation with skin conductance changes: An fmri study. *Neuroimage*, *17*(4), 1797-1806.
- Peters, E., & Slovic, P. (2000). The springs of action: Affective and analytical information processing in choice. *Personality And Social Psychology Bulletin*, *26*(12), 1465-1475.
- Petry, N. M., Bickel, W. K., & Arnett, M. (1998). Shortened time horizons and insensitivity to future consequences in heroin addicts. *Addiction*, *93*(5), 729-738.

- Plato, (1961). The collected dialogues. Hamilton, E. and Cairnes, H. (Eds.).
Princeton, Hew Jersey. Princeton University Press.
- Rachlin, H. (1989). *Judgment, decision and choice. A Cognitive/behavioral synthesis.*
New York: W.H. Freeman and Company.
- Rachlin, H. (2004). *The Science of Self-Control.* London, England: Harvard
University Press.
- Reavis, R., & Overman, W. H. (2001). Adult sex differences on a decision-making
task previously shown to depend on the orbital prefrontal cortex. *Behavioral
Neuroscience, 115*(1), 196-206.
- Rescorla, R. A. (1997). Response-inhibition in extinction. *Quarterly Journal Of
Experimental Psychology Section B-Comparative And Physiological
Psychology, 50*(3), 238-252.
- Ritter, L. M., Meador-Woodruff, J. H., & Dalack, G. W. (2004). Neurocognitive
measures of prefrontal cortical dysfunction in schizophrenia. *Schizophrenia
Research, 68*(1), 65-73.
- Rogers, R. D., Everitt, B. J., Baldacchino, A., Blackshaw, A. J., Swainson, R.,
Wynne, K., et al. (1999). Dissociable deficits in the decision-making cognition
of chronic amphetamine abusers, opiate abusers, patients with focal damage to
prefrontal cortex, and tryptophan-depleted normal volunteers: Evidence for
monoaminergic mechanisms. *Neuropsychopharmacology, 20*(4), 322-339.
- Rolls, E. T. (1990). A theory of emotion, and its application to understanding the
neural basis of emotion. *Cognition & Emotion, 4*(3), 161-190.

- Rolls, E. T. (2000). The orbitofrontal cortex and reward. *Cerebral Cortex*, 10(3), 284-294.
- Rolls, E. T. (2004). The functions of the orbitofrontal cortex. *Brain And Cognition*, 55(1), 11-29.
- Rolls, E. T. (2005). *Emotion explained*. Oxford: Oxford University Press.
- Rolls, E. T., Hornak, J., Wade, D., & McGrath, J. (1994). Emotion-related learning in patients with social and emotional changes associated with frontal-lobe damage. *Journal Of Neurology Neurosurgery And Psychiatry*, 57(12), 1518-1524.
- Rotheram-Fuller, E., Shoptaw, S., Berman, S. M., & London, E. D. (2004). Impaired performance in a test of decision-making by opiate-dependent tobacco smokers. *Drug And Alcohol Dependence*, 73(1), 79-86.
- Sanfey, A. G., Hastie, R., Colvin, M. K., & Grafman, J. (2003). Phineas gauged: Decision-making and the human prefrontal cortex. *Neuropsychologia*, 41(9), 1218-1229.
- Savage, L. J. (1967). Historical and critical comments on utility. In W. Edwards & A. Tversky (Eds.), *Decision making. Selected readings* (pp. 96-110). Bungay, Suffolk: Penguin Books.
- Schmitt, W. A., Brinkley, C. A., & Newman, J. P. (1999). Testing damasio's somatic marker hypothesis with psychopathic individuals: Risk takers or risk averse? *Journal Of Abnormal Psychology*, 108(3), 538-543.

- Silberberg, A., Thomas, J. R., & Berendzen, N. (1991). Human choice on concurrent variable-interval variable-ratio schedules. *Journal Of The Experimental Analysis Of Behavior*, 56(3), 575-584.
- Silva, F. J., & Gross, T. F. (2004). The rich get richer: Students' discounting of hypothetical delayed rewards and real effortful extra credit. *Psychonomic Bulletin & Review*, 11(6), 1124-1128.
- Suzuki, A., Hirota, A., Takasawa, N., & Shigemasa, K. (2003). Application of the somatic marker hypothesis to individual differences in decision making. *Biological Psychology*, 65(1), 81-88.
- Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Review Monographs*.
- Tomb, I., Hauser, M., Deldin, P., & Caramazza, A. (2002). Do somatic markers mediate decisions on the gambling task? *Nature Neuroscience*, 5(11), 1103-1104.
- Tranel, D., Bechara, A., Damasio, H., & Damasio, A. R. (1996). Fear conditioning after ventromedial frontal lobe damage in humans. *Society for Neuroscience Abstracts*, 22, 1108.
- Tranel, D., Bechara, A., & Damasio, A. (1999). Decision making and the somatic marker hypothesis. In M. S. Gazzaniga (Ed.), (Second Edition ed., pp. 1047-1061). Cambridge, Massachusetts: MIT Press.

- Tranel, D., Bechara, A., & Denburg, N. L. (2002). Asymmetric functional roles of right and left ventromedial prefrontal cortices in social conduct, decision-making, and emotional processing. *Cortex*, 38(4), 589-612.
- Tranel, D., & Hyman, B. T. (1990). Neuropsychological correlates of bilateral amygdala damage. *Archives Of Neurology*, 47(3), 349-355.
- Tunney, R. J., & Shanks, D. R. (2002). A re-examination of melioration and rational choice. *Journal Of Behavioral Decision Making*, 15(4), 291-311.
- Turnbull, O. H., Berry, H., & Bowman, C. H. (2003). Direct versus indirect emotional consequences on the iowa gambling task. *Brain And Cognition*, 53(2), 389-392.
- Turnbull, O. H., Evans, C. E. Y., Bunce, A., Carzolio, B., & O'Connor, J. (2005). Emotion-based learning and central executive resources: An investigation of intuition and the iowa gambling task. *Brain And Cognition*, 57(3), 244-247.
- Turnbull, O. H., & Evans, C. E. Y. (2006). Preserved complex emotion-based learning in amnesia. *Neuropsychologia*, 44(2), 300-306.
- Turnbull, O. H., Evans, C. E. Y., Kemish, K., Park, S., & Bowman, C. H. (2006). A novel set-shifting modification of the iowa gambling task: Flexible emotion-based learning in schizophrenia. *Neuropsychology*, 20(3), 290-298.
- van't Wout, M., Kahn, R. S., Sanfey, A. G., & Aleman, A. (2006). Affective state and decision-making in the ultimatum game. *Experimental Brain Research*, 169(4), 564-568.

- van den Bos, R., Houx, B. B., & Spruijt, B. A. (2006). The effect of reward magnitude differences on choosing disadvantageous decks in the iowa gambling task. *Biological Psychology*, *71*(2), 155-161.
- Vaughan, W., Kardish, T. A., & Wilson, M. (1982). Correlation versus contiguity in choice. *Behaviour Analysis Letters*, *2*(3), 153-160.
- Vaughan, W., & Herrnstein, R. J. (1987). Choosing among natural stimuli. *Journal Of The Experimental Analysis Of Behavior*, *47*(1), 5-16.
- Vaughan, W., & Herrnstein, R. J. (1997). Stability, melioration, and natural selection. In H. Rachlin & D. I. Laibson (Eds.), *The matching law* (pp. 194-225). London, England: Harvard University Press.
- Volkow, N. D., & Fowler, J. S. (2000). Addiction, a disease of compulsion and drive: Involvement of the orbitofrontal cortex. *Cerebral Cortex*, *10*(3), 318-325.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behaviour* (1st ed.). Princeton: Princeton University Press.
- von Neumann, J., & Morgenstern, O. (1947). *Theory of games and economic behaviour* (2nd ed.). Princeton: Princeton University Press.
- Vuchinich, R. E., & Tucker, J. A. (1988). Contributions from behavioral theories of choice to an analysis of alcohol-abuse. *Journal Of Abnormal Psychology*, *97*(2), 181-195.
- Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, *14*, 101-180.

Wilder, K. E., Weinberger, D. R., & Goldberg, T. E. (1998). Operant conditioning and the orbitofrontal cortex in schizophrenic patients: Unexpected evidence for intact functioning. *Schizophrenia Research*, *30*(2), 169-174.

Yechiam, E., Stout, J. C., Busemeyer, J. R., Rock, S. L., & Finn, P. R. (2005a). Individual differences in the response to forgone payoffs: An examination of high functioning drug abusers. *Journal Of Behavioral Decision Making*, *18*(2), 97-110.

Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005b). Using cognitive models to map relations between neuropsychological disorders and human decision-making deficits. *Psychological Science*, *16*(12), 973-978.