

AN EXPERIMENTAL INVESTIGATION OF THE
SECRETARY PROBLEM:
FACTORS AFFECTING SEQUENTIAL SEARCH
BEHAVIOUR

A dissertation

submitted in partial fulfilment

of the requirements for the Degree

of

Doctor of Philosophy in Psychology

in the

University of Canterbury

by Y. C. Hsiao

University of Canterbury

2018

TABLE OF CONTENTS

TABLE OF CONTENTS	2
LIST OF TABLES	8
LIST OF TABLES (CONT.)	9
LIST OF FIGURES	10
LIST OF FIGURES (CONT.)	11
LIST OF APPENDICES	12
LIST OF APPENDICES (CONT.)	13
ACKNOWLEDGEMENTS	14
ABSTRACT	15
1. GENERAL INTRODUCTION	16
1.1 THE SECRETARY PROBLEM.....	16
1.2 THE OPTIMAL STOPPING RULE	17
1.3. THEORETICAL AND EXPERIMENTAL WORK ON THE EXTENSIONS OF THE CLASSICAL SECRETARY PROBLEM	22
1.4 MY EXPERIMENTS	28
2. THE EFFECT OF TIME COST ON SEARCH BEHAVIOURS: EXPERIMENT ONE.	33
2.1. INTRODUCTION	33
2.2 METHODS	37

2.2.1 <i>Experimental design and procedures</i>	37
2.3. SIMULATIONS AND HYPOTHESES	40
2.4 RESULTS	45
2.4.1. <i>Stopping Positions</i>	48
2.4.2 <i>Optimal decision-making evaluation</i>	49
Total chosen price (ECUs)	49
Net payoff	49
2.4.3 <i>Evidence of Learning</i>	53
2.4.4 <i>Regression analysis</i>	54
2.5. DISCUSSION	57
3. A HOUSE-SELLING CONTEXT IMPROVES DECISION MAKING FROM A SEQUENCE OF OFFERS. EXPERIMENT TWO.....	59
3.1 INTRODUCTION	59
3.1.1 <i>Different decisions under different context frames</i>	62
3.1.2 <i>Evidence of making better decisions with context</i>	64
3.1.3 <i>How much information do we need to activate existing schema?</i>	65
3.1.4 <i>Hypothesis</i>	65
3.2 METHOD	67
3.2.1 <i>Experimental Design</i>	67

3.2.2 <i>Procedures and Participants</i>	70
3.3 RESULTS	70
3.3.1 <i>Optimal decision-making evaluation</i>	72
Stopping position	72
Total chosen price (ECU).....	72
3.3.2 <i>Does context framing effect last?</i>	73
Stopping position	73
Total chosen price	75
3.3.3 <i>Pearson Correlation Analysis</i>	76
3.4 DISCUSSION	77
4. WHY DO PEOPLE SEARCH TOO LITTLE? NOT ENOUGH TIME, OR TRYING TO AVOID REGRET? EXPERIMENT THREE.....	79
4.1 INTRODUCTION	79
4.1.1. <i>Feeling of regret</i>	79
4.1.2. <i>Feeling of overconfidence</i>	82
4.2 METHOD	86
4.2.1 <i>Participants</i>	86
4.2.2 <i>Experimental Design and Procedures</i>	86
Part 1	86

Part 2	89
4.3 RESULTS	91
4.3.1 <i>Average Results</i>	96
4.3.2 <i>Does self-confidence correlate with search behaviours?</i>	97
4.3.3 <i>Do participants follow the “classical stopping rule”?</i>	98
4.3.4 <i>What were the participants doing?</i>	99
4.3.5 <i>What was the penalty for not following the classical stopping rule?</i>	102
4.4 DISCUSSION	103
5. DECISION MAKING IN SMALL GROUPS. EXPERIMENT FOUR.	106
5.1 INTRODUCTION	106
5.1.1 <i>Theories of group decision making</i>	107
5.1.2 <i>Literature review</i>	109
5.2 METHOD	111
5.3 RESULTS	112
5.3.1 <i>Stopping position and treatments</i>	112
5.3.2 <i>Evidence of learning</i>	113
5.3.3 <i>Lorge-Solomon Analysis</i>	114
5.4. DISCUSSION	115

6. ALL MONEY IS NOT CREATED EQUAL: HOW ACTUAL AND OPTIMAL BEHAVIOUR VARIES WITH THE INCENTIVE STRUCTURE. EXPERIMENT FIVE.....	116
6.1 INTRODUCTION	116
6.1.1. <i>Why the effect of monetary incentive matters</i>	118
6.1.2. <i>How do monetary incentives motivate behaviour?</i>	120
6.1.3. <i>How do incentive structures work?</i>	121
6.1.4. <i>Hypotheses</i>	123
6.2 METHOD	124
6.2.1 <i>Participants</i>	124
6.2.2 <i>Experimental Design and Procedures</i>	125
6.3 RESULTS	131
6.3.1 <i>Stopping position</i>	134
6.3.2 <i>Total chosen price</i>	135
6.3.3. <i>Optimal price count</i>	136
6.3.4. <i>Correlation Results</i>	137
6.3.5. <i>Actual earnings</i>	139
6.4 DISCUSSION	140
7. GENERAL DISCUSSION AND CONCLUSION	143

7.1 SUMMARY OF RESULTS	143
<i>Chapter 2. The effect of time cost on search behaviours. experiment one.</i>	143
<i>Chapter 3. A house-selling context improves decision making from a sequence of offers.</i> <i>experiment two.</i>	143
<i>Chapter 4 Why do people search too little? Not enough time, or trying to avoid regret?</i> <i>experiment three.</i>	144
<i>Chapter 5 Decision making in small groups. experiment four.</i>	144
<i>Chapter 6 All money is not created equal: How actual and optimal behaviour varies with the</i> <i>incentive structure. experiment five.</i>	144
7.2 IMPLICATIONS OF THE FINDINGS AND CONCLUSIONS	145
7.3 LIMITATIONS AND FUTURE RESEARCHES	155
REFERENCE	159
APPENDICES.....	184

LIST OF TABLES

Table 2.1 Simulation summary statistics and between treatments statistical tests	42
Table 2.2. Summary prediction results of net payoff after applying the stopping rules to the sequences in the experiment when there is 0 ECU and 20 ECUs	44
Table 2.3. Experiment summary statistics and between treatments statistical tests	46
Table 2.4. Summary of statistical findings for the comparison between the no, unannounced time, monetary, and announced time cost treatments.....	50
Table 2.5. Coefficients of multiple regression for the tested factors in the no cost, unannounced time cost, monetary cost and announced time cost treatments	56
Table 3.1. Data for the price offer sequences employed in the experiment	69
Table 3.2. Experiment summary statistics and between treatments statistical tests	71
Table 3.3. Summary of statistical findings for the session halves and the treatments	75

LIST OF TABLES (CONT.)

Table 4.1 Position in the sequence at which the offer was accepted as a function of round number for low and high delay for the house selling and secretary hiring frame	92
Table 4.2 Offer prices in the house selling frame for each round	94
Table 4.3 Candidate scores in the secretary frame for each round.....	95
Table 4.4 Summary statistics of overconfidence, self-efficacy and depression scores.....	98
Table 4.5 Percentages of accepted offers following increasing or decreasing sequences...	100
Table 6.1 Participant data for the treatment groups.....	125
Table 6.2 Summary data of the prices presented in each round	130
Table 6.3 The summary statistics of position in the sequence at which the offer was accepted (panel A), the sum of 10 chosen prices (panel B) and the number of rounds the optimal price is selected (panel C) for the treatments	132
Table 6.4 Pearson correlation coefficient with the position in the sequence at which the offer was accepted, the number of rounds the optimal price is selected, and the sum of 10 chosen prices in the treatments	138

LIST OF FIGURES

Figure 1. 1. The illustration of the secretary problem.	19
Figure 2.1. The average payoff for all stopping rules in the cost is equal to 0 ECU and 20 ECU (for both time and monetary) treatments. The filled bar represents 0 ECU and blank bar is 20 ECU.	43
Figure 2.2. The half session average stopping position in the first (round 1-5) and second half (round 6-10) of the session for the no cost, unannounced time cost, monetary cost and announced time cost treatments.....	54
Figure 3.1. The average stopping position averaged across participant in the first half of session (round 1 – 5) and the second half of session (round 6 – 10) in the no context, house frame, and house with info treatments.....	74
Figure 3.2. The average total chosen price averaged across participant in the first half (round 1 – 5) and second half (6 – 10) sessions in the no context, house frame and house with info treatments.	76

LIST OF FIGURES (CONT.)

Figure 4.1. Frequency distribution of the average stopping positions (averaged across 10 selling houses and 10 hiring secretary rounds) of the participants in the low and high delay treatments.....	102
Figure 5.1. The average stopping position in the group and individual treatments.	113
Figure 5.2. The half session average stopping position in the first (round 1–5) and second half (round 6–10) of the session for both group and individual treatments.	114
Figure 6.1. The average stopping position of incentive structures averaged across frames (house selling and no context frames) in each round.	135
Figure 6.2. Percentage distribution of the number of rounds the optimal price is selected (averaged across house selling and no context frames) of participants in the commission base, best only and no incentives structures.	137

LIST OF APPENDICES

Appendix 2.A Instructions for the no cost and unannounced time cost treatments	184
Appendix 2.B Instructions for the announced time cost treatment	186
Appendix 2.C Instruction for the monetary cost treatment	188
Appendix 2.D The price offer sequences used in Experiment One, Two and Five	190
Appendix 2.E Simulations and Hypotheses	191
Appendix 3.A Instruction for the house frame treatment.....	192
Appendix 3.B Instruction for the no context treatment.....	194
Appendix 3.C Summary prediction results of three dependent variable after applying the decision rules to the sequences in the experiment.....	196
Appendix 3.D Summary of statistical findings for the comparison between the no context and house frame, no context and house with info, and house frame and house with info treatments	198

LIST OF APPENDICES (CONT.)

Appendix 4.A General Knowledge and Probability Evaluation Task.....	200
Appendix 4.B The Self-efficacy scale	204
Appendix 4.C The depression scale	206
Appendix 4.D The demographic questionnaire	212
Appendix 6.A The commission base treatment instruction with the house selling frame ..	214
Appendix 6.B The best only treatment instruction with the house selling frame	216
Appendix 6.C The no incentive treatment instruction with the house selling frame	218
Appendix 6.D The commission base treatment instruction with the no context frame.....	219
Appendix 6.E The best only treatment instruction with the no context frame	221
Appendix 6.F The no incentive treatment instruction with the no context frame	223

ACKNOWLEDGEMENTS

I would like to thank Professor Simon Kemp and Professor Maroš Servátka for the supervision, the new perspectives and the encouragement they provided.

The Macquarie University Research Excellence Scholarship allowed me to complete this research, and the University of Canterbury PhD student conference scholarship funded the presentation of part of this research at the International Conference on Humanity and Social Sciences in Macau in 2017, and the Macquarie University Postgraduate Research Fund from Macquarie University funded the presentation at the 2017 European Meeting of the Economics Science Association Conference in Vienna.

I am grateful to everyone for their time and effort in participating in my experiments, and to the people who helped me along the way to make the experiments possible.

Finally, I would like to thank my family, parents and friends for their support and encouragement, which made this dissertation possible.

ABSTRACT

The underlying cause of people stopping their search prior to the theoretical optimal solution in commonly found sequential search problems, often known as the secretary problem, is as yet undetermined. In the secretary problem, an immediate decision must be made from offers that are presented one at a time in a random sequence without re-call. The factors that might cause the early stopping bias were examined in a series of experiments. In Experiment One, I experimentally investigated the effect of time search cost in a version of the secretary problem framed as a house-selling context and found that, if each additional search involved a pre-announced time delay, people shortened the search. In Experiment Two, I found that if there were a house-selling context, people searched more optimally than with no context, even without house information. Experiment Three extended the findings from Experiment Two, framed the secretary problem as either a house-selling or a secretary-hiring task, and found evidence suggesting that regret aversion was a cause of the early stopping bias. Experiment Four investigated group decision making and found that groups chose less optimal prices than the aggregate performance of individuals. In Experiment Five, I examined the effect of an incentive structure that rewarded only finding the optimal prices and found that people did indeed make more optimal decisions in this treatment. Overall, these findings suggest that when we make sequential search decisions in life, exhausting and exploring all offers is never optimal, time spent can never be regained, nor can the rejected offer. Taking note of the context and gathering information prior to making decisions does lead to better results. Settling for offers other than the best can save time and effort, and the earnings will not be much different than when only aiming for the best.

1. GENERAL INTRODUCTION

1.1 THE SECRETARY PROBLEM

When Alan Turing invented the very first computer, the machine was designed to go through all the possible scenarios, in the hope of solving the Germans' coded messages during World War II. Today, computers are confined by time and space, just like human beings are, and it is no longer feasible to solve problems using this approach, not even for a machine that can work tirelessly around the clock. As modern-day problems become more and more complex, for problems where the necessary information is lacking or has an excessive amount of possibility, algorithms have been programmed into computers to help solve these problems efficiently and effectively, one of which is the optimal stopping rule (Christian & Griffith, 2016). This optimal stopping rule is derived from a recreational mathematics problem called the secretary problem. The secretary problem has attracted a good deal of attention since it was first brought to our general attention by Martin Gardner. Called the game of googol, it was first published in the February 1960 Martin Gardner column of mathematics games in Scientific American (Gardner & Mathematical Association of America, 2009). According to Gardner, the game was devised in 1958 by John Fox of the Minneapolis-Honeywell Regulator Company and Gerald Marnie of the Massachusetts Institute of Technology. See Ferguson (1989) and Freeman (1983) for historical reviews. The secretary problem is also known as the game of googol (Gardner & Mathematical Association of America, 2009), the marriage problem (Lindley, 1961) and several other synonyms. The most basic formulation of the problem is often referred to as the classical secretary problem or the standard secretary problem. It is defined as follows: suppose a manager wants to hire a secretary, and the manager receives a bonus only when she hires the very best candidate. The manager is only allowed to hire from a given list of

people. She knows exactly how many people are on the list before the interviewing process, and the candidates she interviews are presented randomly one at a time. During the interviewing process, she has to decide right away whether to hire the candidate, and it is not possible to re-call candidates rejected earlier. If she has not hired any candidate before the end, she must hire the last candidate. How does she maximise her chances of receiving this bonus?

1.2 THE OPTIMAL STOPPING RULE

The secretary problem is a form of the sequential decision-making problem that we face quite often in everyday life. Most of the research in sequential decision-making problems involves the investigation of optimal search, and how and why people are unable to choose optimally. A statement about "insufficient search" or "over search" would only be meaningful with respect to a theoretical optimal. The optimal strategy in making decisions for the classical secretary problem is a stopping rule; this states that the manager should reject the first $r - 1$ candidates (let candidate M be the best applicant among these $r - 1$ candidates) to learn about the distribution of the candidates' quality, so the first $r - 1$ candidates serve as the *information set*, and she then selects the first subsequent candidate that is better than candidate M . The optimal stopping rule is therefore suggesting that the manager learn about distribution in the information set, does not accept anyone regardless of their quality and only consider accepting a candidate after having interviewed (rejected) everyone in the information set. This seemingly complicated problem turns out to have a simple, even elegant, solution. For an arbitrary r , the probability that the best candidate is selected can be shown as follows (see Lindley, 1961; Gilbert & Mosteller, 1966, for detailed proofs):

$$P(r) = \sum_{i=1}^n P(\text{candidate } i \text{ is selected} \cap \text{candidate } i \text{ is the best}). \quad (1)$$

The goal is to find the probability that candidate i is selected and is the best, otherwise no bonus is earned, as per the assumption of the classical secretary problem. The secretary problem is a case of a dependent event, where the probability of the first event occurrence influences the probability of the second event, and so on. The probability of selecting the best candidate depends in part on the probability of the best candidate being presented in the information set, which in turn depends on the size of the information set. If the information set is too large and contains a large proportion of the candidates, the probability of selecting the best is reduced, given that all candidates in the information set are rejected. If the information set is too small, the manager does not receive adequate knowledge about the distribution, and the probability of finding the best candidate is also reduced. The key issue then is: how large should the information set be?

The multiplication rule of probability states:

$$P(A \cap B) = P(B | A) \times P(A).$$

(The notation $P(B | A)$ means "the probability of B , given that A has happened."),

so,

$$P(A) = \text{Probability candidate } i \text{ is the best}$$

$$P(B) = \text{Probability candidate } i \text{ is selected}$$

$$P(B | A) = P(\text{candidate } i \text{ is selected} | \text{candidate } i \text{ is the best})$$

Thus, Equation 1 can be written as

$$P(r) = \sum_{i=1}^n P(\text{candidate } i \text{ is selected} \mid \text{candidate } i \text{ is the best}) \times P(\text{candidate } i \text{ is the best}) \quad (2)$$

Figure 1.1 will be helpful to visualize the problem.

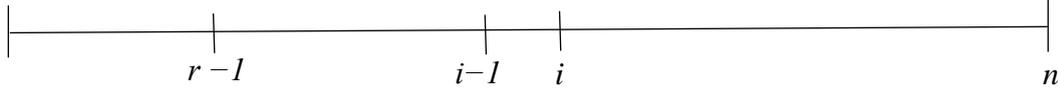


Figure 1.1. The illustration of the secretary problem.

The probability of candidate i being the best among the n candidates is $\frac{1}{n}$ because each candidate is equally likely to be the best. This leads to Equation 3:

$$P(r) = \sum_{i=1}^n P(\text{candidate } i \text{ is selected} \mid \text{candidate } i \text{ is the best}) \times \frac{1}{n}. \quad (3)$$

Then the probability that candidate i is selected, given that candidate i is the best, can be broken up into two different sums. First, if the best candidate is presented in the first $r - 1$ candidates (i is 1 through to $r - 1$), the selection process will always reject the best candidate, and therefore the probability of candidate i being selected and being the best is equal to zero. The first sum is therefore a zero, for each i from 1 through to $r - 1$ as expressed in Equation 4. The limits on the second sum in Equation 4 is because for each candidate, starting with our stopping point r , continuing to the last candidate n (which are the candidates in the selection set) the manager needs to identify the probability of not choosing a candidate before the best candidate i presents, but after the $r - 1$ candidates. (That is to say, there may be a candidate better than M but before we get to the best at candidate i .) The selection process stops after she chooses a candidate, so she needs to find

the probability of not choosing a candidate before the best presents. The probability can be calculated using the number of target candidates divided by the total number of candidates.

$$P(r) = \left[\sum_{i=1}^{r-1} 0 + \sum_{i=r}^n P(\text{the best of the first } i-1 \text{ candidates is in the first } r-1 \text{ candidates} \mid \text{candidate } i \text{ is the best}) \right] \times \frac{1}{n} \quad (4)$$

There are $i-1$ total candidates between the start and the best candidate i . The probability of selecting a candidate after the $r-1$ candidates and before the best candidate i is

$$\frac{(i-1) - (r-1)}{i-1}$$

Therefore, the probability of not choosing a candidate after rejecting $r-1$ candidates and before the best candidate i is

$$1 - \left(\frac{i-r}{i-1} \right) = \left(\frac{i-1}{i-1} \right) - \left(\frac{i-r}{i-1} \right) = \frac{r-1}{i-1}$$

So, for $i=r$ though to n

$$P(r) = \sum_{i=r}^n \frac{r-1}{i-1} \times \frac{1}{n} \quad (5)$$

Since r is just a constant number and i is a variable, $r-1$ can move outside the sum brackets,

$$P(r) = \frac{r-1}{n} \sum_{i=r}^n \frac{1}{i-1} \quad (6)$$

When taking the limit to infinity and making n arbitrarily large, the values of 1 remain the same and become comparatively too small to matter. If n is large, we can approximate it by the integral function below [i.e., (10)], with $x = r/n$ and $t = i/n$.

This gives us (writing $\ln x = \log_e x$):

$$= x \times [\ln(1) - \ln(x)] = x \times [0 - \ln(x)] = -x \ln x. \quad (7)$$

To maximize the probability of this function, take the derivative of the function with respect to x and set it to 0.

So,

$$\ln x = 1$$

$$x = e^{-1} \text{ or } \frac{1}{e}$$

Euler's number, e , is a mathematical constant equal (approximately) to 2.7182. This gives,

$$P(r) = \frac{1}{e} = \frac{1}{2.7184} = 0.3679 \approx 37\% \quad (8)$$

and $r = n/e$.

This optimal stopping rule may fail in two ways. The first way is that the best candidate may be in the information set. However, this may not be the most critical issue if one chooses a very small information set. The manager will never find anyone better than this and is thus finally forced to accept the n th candidate. If $n = 20$ and $r = 7$, the probability of having to select a twentieth candidate who is not optimal is a high 0.35. (The total probability of selecting the twentieth candidate with this strategy is about 0.38, as the twentieth candidate might actually be the best.) The second way in which the optimal

stopping rule might fail is that the first candidate encountered from position r onwards who is better than the best in the information set may not be the best candidate in the remainder of the set, i.e., the overall best. The probability of this occurring for $n = 20$ is 0.266 ($1 - 0.35 - 0.384$).

1.3. THEORETICAL AND EXPERIMENTAL WORK ON THE EXTENSIONS OF THE CLASSICAL SECRETARY PROBLEM

Simon (1956) suggested a bounded rationality approach to making sequential search decisions in everyday life even before the secretary problem was brought to researchers' attention. Instead of optimising the expected outcome as in traditional economic theory, people set an aspirational level and are satisfied once the aspirational level is achieved. Such an approach allows decision makers to effectively achieve a variety of needs in situations where the optimal strategy is unknown to the decision maker, that is, there is no need for a utility function to be postulated. For example, when a rat forages for food, it learns to choose time-conserving paths that lead to sufficient food for survival rather than paths which might attain the maximum amount of food but potentially risk survival in the process.

Theoretical analyses have been conducted with variations of the classical secretary problem. The assumptions of the classical secretary problem are limiting for many decision-making scenarios in life, and, especially, the assumption that only selecting the absolute best is a positive result. There are two reasons for this. Firstly, a manager may still gain utility from alternatives other than the best (Bearden, Rapoport, & Murphy, 2006). Hiring, say, the second-best secretary available would often be good enough and yield positive utility from the work the secretary provides. There are very few, if any, examples of real-life situations where only the best choice produces utility. Findings from incorporating such an assumption might not apply in the real-life settings. Secondly, in real life, it is not usually

clear what the best choice would have been. If you do not interview the best secretary, you will never know there really was one who was better than your choice.

In the attempt to address the limitation of the assumption that only the best yields a payoff, the assumption has been relaxed in various ways. The version of the secretary problem that relaxed the payoff structure is often referred to the generalised secretary problem (Bearden & Murphy, 2000). Mucci (1973) and Yeo and Yeo (1994) provide theoretical grounds for how to derive the optimal stopping rule, when monotone payoffs $w_1 \geq w_2 \geq \dots \geq w_n$ are assigned to the best n candidates; for example, when choosing the best candidate receives the highest payoff, choosing the second best receives the second highest, and so on. This stopping rule then becomes a multithreshold rule. The classical secretary problem is a special case of the generalised secretary problem when $w_1 = 1$ and $w_n = 0$ for all $n > 1$ (Bearden & Murphy, 2007). After the first threshold of the information set and prior to reaching the second threshold, the manager should accept a candidate who is the best seen so far. If no one fits the criteria, after the second threshold, she needs to lower her standard and accept either the best or second-best candidate that she has interviewed before reaching the third threshold, and so on. Bearden, Rapoport and Murphy (2006) provide a numerical example of this solution for the case of 50 candidates, with a payoff of 16 for choosing the best, 8 for the second-best, 4 for the third-best, 2 for the fourth-best and 1 for the fifth-best candidate. The optimal stopping rule suggests that the manager should reject the first 17 candidates; i.e., the first 35%, and only choose the best seen so far (every interviewed candidates) between candidates 18 to 35. If this fails, accept the candidate who is either the best or second-best candidate that interviewed so far between candidates 36 to 43. If this still fails, she should choose either the best, second- or third-best candidate interviewed so far between candidates 44 to 47.

Gilbert and Mosteller (1966) presented optimal solutions for the problem when a payoff of 1 can be earned once the manager chooses either the best or second-best candidate, and she receives nothing otherwise. They found the maximum probability of receiving the payoff when $n \rightarrow \infty$ is actually 0.574. This optimal stopping rule has two thresholds. The first is to reject the first $0.347 \times n$ candidates, then choose the next best seen so far until she reaches the $0.667 \times n$ th candidate. If this fails, after $0.667 \times n$ candidates, she should choose the best or second-best candidate interviewed so far. For example, in the case of 100 candidates, she should reject the first 35 candidates, and only accept the best candidate between candidates 36 to 67. From candidate 68 and onwards, she should choose either the best or second-best candidate among the interviewed candidates.

Gilbert and Mosteller (1966) also studied the problem when the manager is allowed to accept two candidates, and a payoff is earned when either candidate is the best. They found the first decision to accept should be made after rejecting the first n divided by e to the power of $2/3$ candidates. The second decision to accept should be made after $\frac{n}{e}$ candidates. As $n \rightarrow \infty$, the maximum probability of finding the best candidate with two decisions is approximately 59%. For example, if there are 50 candidates to choose from, the first decision should be made between candidate 13 and 19; that is, the first information set is 12. The second decision is made after candidate 19. Both candidates are accepted only if the candidate is the best seen so far. Tamaki (1979) also allows the acceptance of two candidates, but they need to be the best and second-best candidates. This version of the optimal stopping rule rejects the first $0.229 \times n$ candidates. After $0.607 \times n$ candidates, either the best or second-best candidate seen so far are chosen. With an example of 50 candidates, for the first decision, the best candidate between candidates 13 to 31 is accepted. As for the second decision, the candidate is accepted who is either the best or second best

among all interviewed after candidate 31. When $n \rightarrow \infty$, the maximum probability of earning the payoff is approximately 22.5%.

Lindley (1961) derived the optimal stopping rule for choosing the candidate with the minimum expected ranking of the candidate, where the best candidate has a ranking of 1; Moriguti (1993) constructed an algorithm for computing the optimal thresholds for this extension. In the case of 25 candidates, the manager should reject approximately the first one third of candidates, that is, 8 candidates. Then she should select a candidate only if the candidate is the best seen so far up to candidate 13. If no candidate has been selected, from candidates 14 to 16, the manager should accept a candidate only if that candidate happens to be either the best or second best of those interviewed so far. If that still fails, then the best, second- and third-best candidate between candidates 17 and 19 are accepted, and so on.

The cost of a search is another variation that is often considered, as search cost is inevitable in any sequential search decision. Introducing costs associated with the consideration of each choice have been examined in various studies (Yeo, 1998; Zwick, Rapoport, Lo & Muthukrishnan, 2003). Yeo (1998) theoretically set up a version of the secretary problem in which he introduced a fixed, per-candidate interview search cost with positive payoffs from selecting any of the k best candidates. The payoffs increase with the quality (and thereby rank) of the candidate. The problem is then to find the optimal stopping rule and use it to evaluate the probabilities of success and search length given the distribution of the quality of candidates interviewed. Yeo used numerical methods to arrive at the conclusion that, in such a setup, as the search cost increases, it is optimal to decrease search activity. Zwick, Rapoport, Lo and Muthukrishnan (2003) used a simulation to identify the size of the information set for three different per-unit search costs: \$0, \$0.1 and \$0.3, in which there was some probability of re-calling previously rejected candidates (not

just the most recently rejected, but any previously rejected). The manager would only earn a payoff when the best candidate is selected. In the case of 20 candidates, the information set was found to be 13, 10 and 6 candidates for the respective unit search costs given above; these information sets also yielded the highest predicted payoff.

In addition to the theoretical investigation of the secretary problem, there is also prior experimental work. Theoretical investigation produces an optimal stopping rule to variants of the secretary problem. The optimal stopping rule does not predict where the optimal candidate lies, only a strategy that maximises the probability of finding the optimal candidate. Experimental research examines how people conduct searches and compares that with the theoretical predicted optimal search, both in the classical problem and its extensions. Often in experimental research, the theoretical derived optimal stopping rule is applied to the experimental design to obtain the predicted optimal amount of search that people should conduct. A general empirical result for the secretary problem is that people tend to stop their search too soon, sometimes even stopping their search within the theoretical information set (e.g., in the experiment of Zwick et al.; 2003).

Seale and Rapoport (1997) studied the secretary problem by incorporating the classical assumptions: a known number of candidates were presented randomly one at a time without re-call, and the participant only earned a payoff from selecting the best candidate. Seale and Rapoport found that their participants stopped the search too early, and they suggested that time search cost might be the reason for early stopping. There was no monetary search cost (henceforth monetary cost) in Seale and Rapoport's 1997 experiment, but the participants still needed time to search each candidate, and there were 100 potential candidates presented per trial. The classical optimal strategy is no longer applicable to

compare to participants' behaviours with time search cost (henceforth time cost). Perhaps they maximised their payoff when the endogenous time cost was taken into account.

Other variants of the secretary problem have also found early stopping. For example, Seale and Rapoport (2000) presented a candidate population of unknown size, and the participants were asked to select the top-ranked candidate possible, without knowing the number of candidates available in the selection pool. They reported that the participants stop their search earlier than the theoretical optimum predicts. The theoretical information set was obtained using simulation and it predicted that, in the case of 40 candidates, it is optimal to reject 15% of the candidates, that is, six candidates, before considering whether to accept. Bearden, Rapoport and Murphy (2006) tested a variant of the secretary problem with a payoff scheme that depended on the ranking of the candidate selected. Six positive payoffs were available: participants received \$25 for selecting the best candidate, \$13 for the second-best candidate, \$6 for the third-best, \$3 for the fourth, \$2 for the fifth and \$1 for the sixth, with $n = 60$ candidates available in the selection pool. The results also showed that the participant tended to stop searching too soon compared with the theoretically predicted optimum. The theoretical optimum predicted for the first information set was 35%, as mentioned previously, which is slightly lower than the optimum the classical solution predicts.

Zwick, Rapoport, Lo, and Muthukrishnan (2003) investigated the secretary problem, allowing re-call of a previously rejected choice, in searching for apartments with a certain probability that the apartments were still available. In addition to allowing for re-calling a previously passed-over apartment, Zwick et al. studied the effect of monetary search cost using a 2×2 design with the two factors being the size of the apartment set and whether there was a search cost in searching for another apartment (0.3 cents for a 20-apartment

treatment and 0.1 cents for the 60-apartment treatment). The participants received a payoff of \$10 if they selected the overall best. The results show that participants stopped their search sooner with a search cost than with no search cost. However, the participants were considered as searching for too long compared with the predicted optimal result they obtained from simulation, where the effect of search cost was incorporated. As previously mentioned, their simulation results suggest a smaller information set with an increase in search cost. However, the participants in the cost treatments searched for longer than the simulation predicted under the search-cost conditions.

1.4 MY EXPERIMENTS

In total, this dissertation presents the results of five experiments that present different but related manipulations of variables in the secretary problem. Each is presented in its own chapter below, and each chapter begins with an introduction relevant to that chapter. What follows in this section is a quick overview of the themes of these chapters and does not present a complete rationale for the studies presented in them.

Previous experiments have consistently found that people behave suboptimally in making sequential search decisions. The underlying causes for early stopping behaviour are still unknown. It is possible that the endogenous time search cost, as suggested by Seale and Rapoport (1997), is the cause, and was not accounted for in the theoretical analysis and yet time cost is inevitable in any type of searching, and just like any other resource, time is scarce and limited. The cost of time is often measured using a monetary value, for example, the opportunity cost of time. The more time we spend on searching, the higher this cost. A higher cost suggests a lower payoff. If we aim to maximise our overall payoff, as economists propose, reducing cost is a sensible and appropriate way to behave. In this dissertation, I examine whether the time cost has a similar shortening effect to monetary

cost. If time cost also shortens a search, then stopping prior to the theoretical predicted optimal stopping position would be an optimal behaviour by itself in an attempt to reduce time cost. I have also conducted a simulation to derive the optimal stopping rule for my design. The experimental and simulation findings are reported in Experiment One (Chapter Two).

In addition to time search cost, the effect of context has also not been considered in the optimal theorem and yet it is a well-documented finding that decisions change systematically based on how the context is framed (e.g., Tversky & Kahneman, 1981; McNeil, Pauker, Sox, & Tversky, 1982). Moreover, there are pieces of experimental evidence which suggest that people make better decisions when the task is presented in a context rather than without (e.g., Sugiyama, Tooby, Cosmides, 2002; Cosmides, Barret & Tooby, 2010; Griggs & Cox, 1982). To derive the best solution for every context, the human brain has evolved to think through the context while making decisions. A context may assist the activation of existing schemas that contain effective strategies constructed from previous, similar experiences. This allows us to make judgements and decisions in an effective and efficient manner, without having to construct a solution from scratch every time we encounter the same or similar experiences. Thus, the effect of context can potentially be another cause of the suboptimal behaviour found in sequential decision making. Experiment Two in Chapter 3 was therefore designed to examine the effect of context.

Besides the time cost and context effects, cognitive biases – for example, overconfidence and regret aversion – may also cause people to choose suboptimally. Chapter Four reports a study (Experiment Three) of whether these biases affect stopping behaviour. Overconfidence is one of the most predominant biases found in human decision

making (e.g., Moore & Small, 2007; Swann & Gill, 1997). This cognitive bias describes the tendency of a person's subjective confidence in their judgements or abilities to be systematically and persistently higher than their objective accuracy. Overestimating one's own ability in making optimal decisions, or believing one has sufficient knowledge when one does not, can result in insufficient searching. As in the secretary problem, the decision maker does not know the distribution of the candidates' quality prior to making decisions. The optimal theorem requires an information set to gather information on the distribution prior to accepting a candidate. Overconfidence might, therefore, lead to the shortening of the information set and observation of fewer candidates overall. Similarly, avoidance of the feeling of regret could potentially be another cause for observing fewer than the optimal number of candidates. While the secretary problem does not allow re-call of the previously rejected candidates, if one turns down a candidate and nothing better is offered later, the manager knows that she has passed over a better candidate, and this is likely to produce the feeling of regret (see Bell, 1982; Loomes & Sugden, 1982 on regret theory). But, if she accepts a candidate early on, then she will not subsequently find out about the later, better candidate that is yet to come, and there will be less regret. In the attempt to minimise or eliminate regret, the manager is likely to shorten the amount of the search.

Many sequential search decisions are made in a group in our daily life; for example, buying houses are often decided by couples or family members. Experiment Four (Chapter Five) examined whether group decision making eliminates the early stopping behaviour found in individual decision making. Critical decisions are often made by groups, assuming that groups make better decisions than individuals. However, theoretical research into group decision making is inconclusive as regards its performance. The wholistic theory predicts that a group makes a better decision than an individual due to an interaction between group members, and results in the whole being greater than the sum of the individuals (see Stasser

& Birchmeier, 2003 for more details). Alternatively, the reductionist proposes that the group is at best only equal to the sum of its parts if group interaction goes smoothly, otherwise, a group's performance will be less than the sum of what each individual member can achieve, due to reasons such as the Ringelmann effect (Ringelmann, 1913). Take solving the secretary problem as an example: if one group member knows how to make the optimal decision, the group could solve the problem successfully and the group could outperform most of the individuals in it. Furthermore, the other group members could influence the decision and steer the decision away from the optimal. Group members normally discuss their preferences and exchange opinions about the problem at each stage, until they reach a consensus decision among the group members. It is possible that the group can make worse decisions due to the polarising effect or to social loafing. The polarising effect refers to the way that decisions tends to shift towards the extreme after group discussion (e.g., Moscovici & Zavalloni, 1969; Bray & Noble, 1978). For example, Myers and Bishop (1970) found that a group of prejudiced American high school students polarised their attitude to become more prejudiced after group discussion. Therefore, if two group members are regret-averse, they may become more regret-averse after discussion. People could search less with a regret-aversion tendency; they could shift towards a shorter search after discussion than they would have done if they had made individual decisions. Furthermore, no single member is fully responsible for the final decision; this makes social loafing in a group possible. Ingham, Levinger, Graves and Peckham (1974) showed that people were less motivated when working in a group. Therefore, some group members may not exert the same amount of effort in solving the problem as they would when working alone.

It is a common belief that exerting more effort can result in a higher performance. This is why monetary incentives are often offered in the workplace and laboratory. However, effort-motivating incentives are often found to be ineffective in tasks that require

more skills to perform. The way to best motivate performance in skill-dominant tasks remains unknown. Experimental research has found that people often use information from the situation and environment to work out how they are expected to behave. For example, when people are served food with a large bowl, they consume more food compared with when they are served with a small bowl (Wansink & Cheney 2005; Wansink, Van Ittersum & Painter, 2006). It is possible that the monetary incentive structure operates in a similar way. Experiment Five (Chapter Six) experimentally examines whether an incentive structure that might signal a higher expectation of one's ability to perform a task can lead to an increase in performance. Note that if one exerts more effort by conducting more searches – for example, 20 searches – one will not necessarily choose a better candidate than an individual who searches less – for example, with only 10 searches. As the optimal theorem of the secretary problem dictates, the key to maximise the probability of finding the optimal xxx is knowing the size of the information set. Therefore, exerting more effort, that is, more searches alone, without knowing the size of information set does not necessarily result in finding a better candidate.

For the purpose of consistency, the word “offer” refers to the choices available to the participants during the experiment, and it will be employed for the remainder of the dissertation, so as to avoid confusion from different experimental set-ups and designs. For example, in the house-selling frame treatment, “offer” refers to the price offer available. In the hiring-employee frame, it is referring to the candidates. In the “no context” frame, it refers to numbers presented to the participants.

2. THE EFFECT OF TIME COST ON SEARCH BEHAVIOURS: EXPERIMENT ONE.

2.1. INTRODUCTION

Consider the situation of selling a house, where the seller reviews one offer at a time to buy and must decide whether to accept or reject the standing offer, without knowing what the subsequent offers will be. The decision cannot be re-called, meaning that if an offer is rejected, the seller cannot change his mind and accept this offer later on. How can the seller find the highest offer? Since it is unlikely that the first offer is the highest, it may be reasonable to reject the first few offers to learn about their distribution. While a seller might have some knowledge about the real estate market in general, each house has its specificities; the aggregate price level may not convey enough information regarding buyers' values for a particular house. The search process of reviewing and rejecting offers is costly and time-consuming. Monetary search costs have been shown to shorten the search in various contexts, but not much is known about the effect of time search costs. Since time is also a scarce and valuable resource, it is crucial to understand the effect it has on sequential decision making. Search theory predicts that an increased search cost of any type will lead to a shorter search (McCall, 1970; see also Yeo, 1998). Empirical literature provides evidence that, as the (monetary) search cost (henceforth, monetary cost) increases, people do indeed search less (e.g., Zwick, Rapoport, Lo & Muthukrishnan, 2003). But what if, instead of incurring a monetary search cost (henceforth, monetary cost) people have to devote their time to search for better alternatives? Search models typically do not stipulate how different types of costs influence the length of search activity. For example, McCall (1970) stated that the length of search activity depends on the opportunity cost of time, which could be approximated by the perceived wage rate (see Lippman & McCall, 1976; Goldman & Johansson, 1978 for more details on theoretical framework), suggesting that

time and monetary costs are theoretically equivalent. Theoretically, an increase in the time search cost (henceforth, time cost) should decrease the search activity, just as an increase in the monetary cost would. However, apart from its different nature, the time cost is also specific in that it is usually not known *ex ante* and has to be inferred as the search progresses. In this chapter, I test whether an increase in the time cost does result in less searching and whether monetary cost shortens the search in the secretary problem, as predicted by search models and previous studies. The contribution can, therefore, be viewed as extending the empirical analysis of search behaviour to the domain of time.

Is there any existing evidence that would suggest that people behave differently when time and money are at stake? Behavioural literature does provide such evidence in various domains. For example, Gino and Mogilner (2014) found that people are less likely to cheat in a mathematical task when primed with a time-related construct than when primed with money. Okada and Hoch (2004) demonstrated that people are more risk-seeking when investing their time than when their investment is monetary. In a similar fashion, people are more prone to the sunk-cost fallacy with past monetary investment than with time (Soman, 2001). However, it has also been shown that a scarcity of time and a scarcity of money can impact people in similar ways (Mullainathan & Shafir, 2013). A scarcity of money makes poor farmers focus their mental effort on saving, neglecting the risk of adverse circumstances and not buying enough insurance (e.g., against drought or flooding), even when the insurance is subsidized so as to be affordable (Giné, Townsend, & Vickery, 2008). A similar tunnelling effect is displayed when the scarcity concerns people's time. For example, people who are experiencing scarcity of time drive and talk on the phone simultaneously and neglect the risk of hazardous driving (Strayer, Drews, & Crouch, 2006). The tunnelling effect refers to people focusing on things that they experience a shortage of

(e.g., money, time or food), while ignoring potential issues that might result from activities mitigating said shortage (Mullainathan & Shafir, 2013).

Given that some studies suggest time and money affect people's decisions differently, whereas in other domains their impact could be similar, the implications of time cost for sequential decision making are not obvious. In Experiment One, I examine whether the presence of time cost shortens sequential search. In particular, I embed the explorations in the house-selling frame described in the method section (also known as the secretary problem) in which the participants earn payoffs based on the offers they accept. In contrast to the classical secretary problem (see Ferguson, 1989, for a discussion), which assumes people derive utility only from the optimal offer, this experiment allows people to earn money also from suboptimal offers (Bearden, Rapoport, & Murphy, 2006); that is, people obtain utility (earn money in the experiment) based on the true value of the offer, instead of zero payoffs when anything other than the best offer is selected.

Job search is another prominent example of sequential decision making. McCall (1970) gives an analogous prediction to Yeo (1998) regarding the effect of search cost on job search. In McCall's model, the job seeker knows the distribution of wage offers, the search is unlimited, and there is no re-call. He receives one randomly selected wage offer at a time and decides whether to reject or accept it. If the offer is accepted, the search ends and the job seeker receives a wage equal to the offer. The job seeker thus considers the current wage offer against the prospects of being unemployed and receiving his outside option (representing unemployment compensation, welfare payment, and/or leisure benefits) and facing the search cost again in the next round(s) to obtain a new draw from the wage distribution. McCall shows that a relatively high cost compared to the wage offer

encourages the job seeker to limit his search, holding everything else constant. In the extreme, it might even mean that the job seeker prefers to remain unemployed.

Job search models share several assumptions with the secretary problem, including the random presentation of job offers and no re-call. Their other features, such as knowledge of the distribution of offers and an unbounded search horizon, do not coincide with the classical secretary problem. However, despite the difference in assumptions between the job search and secretary problems and the resulting difference in decision-making environments, laboratory experiments studying job search also find that participants stop searching too soon (e.g., Cox & Oaxaca, 1989; Cox & Oaxaca, 2000). Cox and Oaxaca (2000) experimentally tested the job search behaviour with known and unknown wage distribution, finding that people terminate their search early in both cases. In the treatment with known distribution, the participants knew exactly the distribution of the offers that were drawn. In the treatment with unknown distribution, the participants had information about the two possible distributions where the offers were drawn. The offers were sampled with replacement, there were 10 offers available in each round, but a maximum of 20 periods was allowed to accept an offer. This provides a possibility of re-calling a previously rejected offer by design. Their results showed that the extent of search was in fact similar in these two cases.

Both the job search model and the secretary problem give similar theoretical predictions regarding the effect of search cost. The result that monetary cost in the secretary problem shortens search is well documented in both theoretical and empirical literature. To the best of our knowledge, the effect of time cost, which is theoretically equivalent to monetary cost, has not been tested experimentally, although the time cost of searching is unavoidable in a sequential search; but, because of its different nature, one might question

whether its effect is also behaviourally equivalent to monetary cost. The main contribution of this paper is therefore in identifying how different search costs influence search behaviour.

2.2 METHODS

2.2.1 EXPERIMENTAL DESIGN AND PROCEDURES

The experimental task is framed as selling 10 houses; one house in each round. A description of a house, consisting of the floor area, the number of bedrooms, the suburb and the year the house was built in, is presented at the beginning of the round, prior to any price offer. To enhance the link between the experiment and the outside-the-lab world, the house descriptions presented to participants were taken from houses sold in 12 different suburbs in Christchurch, New Zealand during October 2014. All information was obtained from the Quotable Value Limited database (qv.co.nz). In each round, a participant can review up to 20 price offers for the given house. The price offers are presented one at a time. Once a price offer is presented, the participant decides whether to reject the offer or whether to accept it. Once the decision is made, there is no re-call. If the participant has not accepted an offer prior to the final (20th) offer, the participant is forced to accept the final offer, regardless of its value. The actual price offer sequences are presented in Appendix 2.C.

The experiment consisted of four treatments implemented in an across-subject design, varying the search cost. In the no search cost treatment (henceforth no cost), the price offers were presented without any search cost being imposed, neither time nor money. In the monetary search cost treatment (henceforth monetary cost), each participant had to pay 20 experimental currency units (ECUs) to obtain each new offer. The monetary cost was cumulative, that is, the more offers the participant obtained, the more monetary cost had to be paid in total. For example, if a participant accepts the 20th offer for a house, the total

monetary cost is 20 x 20 ECUs compared with no cost. There are two time cost treatments, announced time cost treatment (henceforth announced time cost) and unannounced time cost treatment (henceforth unannounced time cost). In the announced time cost treatment, the participants were instructed with a five second time delay prior to presenting each offer. There was no such instruction in the unannounced time cost treatment; the unannounced time cost treatment had identical instructions to the no cost treatment. To make a parallel comparison between monetary and time costs, the treatment of announced time delay was introduced, as the existence of monetary cost must be pre-announced in the experiment to avoid deceiving the participants.

However, in everyday life, time costs (waiting times) are usually not announced ex ante. To make a parallel comparison with an actual house-selling scenario, the treatment of unannounced time delay was included, where the participants were not informed about the time delay in the instructions but had to learn about it through experience as the experiment progressed. In the unannounced time delay treatment, the time cost had to be individually inferred. This design thus makes the experiment less likely to find a treatment effect due to the time delay incurred with each additional search, making it a conservative test of the conjecture that time cost operates in the same way as monetary cost.

In both time search cost treatments, each participant had to wait five seconds before a new offer was displayed on the screen. The time cost was also cumulative, the more offers the participant searched through, the longer was the wait in total. For example, if a participant accepted the 20th offer for a house, the additional time cost is 20×5 seconds = 100 seconds, compared with no cost. Therefore, a five second time delay can cumulatively yield a significant waiting time.¹ This was indeed the case as the unannounced time cost

¹ Initially, I ran a 10-second delay session that lasted three hours, but many participants dropped out of the

treatments session in this experiment lasted on average 15 minutes longer (approximately one hour and 10 minutes) than the no cost treatment (approximately 55 minutes). The announced time cost lasted on average one hour. The monetary cost treatment lasted 10 minutes less than the no cost treatment.

One hundred and eighty-eight undergraduate students from the University of Canterbury in Christchurch, New Zealand participated in the experiment. Forty-eight participants participated in the no cost treatment, forty-four in the monetary cost treatment forty-three in the unannounced time cost treatment and fifty-three in the announced time cost treatment. The participants were selected randomly from the database using the ORSEE recruitment system (Greiner, 2015). The participants earned 12.2 New Zealand dollars (NZD) on average. In the no cost and the two time cost treatments (unannounced and announced), all experimental earnings have been exchanged into NZD at the announced exchange rate (The details are in the Appendix 2.A–B). For the monetary cost treatment, the payoff is the house's accepted price and subtracting the cumulative monetary search cost; more details are in Appendix 2.C. The payoff protocol was single-blind, meaning that the experimenter was able to track participant decisions to their identity. The search task was implemented using z-Tree software (Fischbacher, 2007).

After arriving at the lab, the participants were randomly assigned to a cubicle and read the instructions (provided in Appendix 2.A–C) at their own pace. Any questions were answered in private. The experiment began with two practice rounds, followed by 10 paid rounds. To allow comparisons between participants, the same 10 random sequences, generated prior to the experiment, were employed for each participant in each session (see

experiment. Subsequently, I recalibrated the time cost (unannounced and announced) treatments to a five-second delay.

Appendix 2.D for details). Each sequence was generated from a normal distribution using the actual mean house price in each suburb, with an upper and lower limit of two actual standard deviations. No data are available for actual house price distributions in NZ suburbs, but the distribution is likely to tail off at both ends as the normal distribution does and a uniform distribution does not: very cheap and very expensive houses are rare. Actually the distribution is unlikely to be important: Kahan, Rapoport and Jones (1967) tested search behaviour with different underlying distributions and found no effect on participant decisions, the price offers were presented in ECUs. The participants were informed that their pay in NZD would be based on their cumulative payoffs from all 10 rounds, with an exchange rate of 1000 ECUs = 1 NZD.

One might think that having more time in the cost treatment with an additional five-second delay prior to presenting an offer could yield better decisions. This design addresses this issue by not imposing a time limit on participants' decisions in any of the treatments; that is, they could take as long as they needed to evaluate each offer presented to them. For an example of how time pressure can affect decisions see Kocher, Pahlke and Trautmann (2013).

2.3. SIMULATIONS AND HYPOTHESES

What is the optimal stopping rule in this variation of the secretary problem? To derive a prediction for the impact of search cost on search behaviour, I conducted a simulation to evaluate the performance of different stopping rules in each treatment. Each simulation compares the payoffs resulting from 20 different stopping rules (as there was a maximum of 20 offers), which contain all possible stopping positions (i.e., the n th offer in a given sequence that has been accepted; where $1 \leq n \leq 20$). Each simulation cycle generates a set of 20 random offers from a standardized normal distribution, using the mean and

standard deviation for each house from the values used in the experiment. Once a set of offers has been generated, the offers are (implicitly) ordered from the highest to lowest and assigned a rank within this particular order. These offer values and the rank for each offer are recorded to test the performance of each stopping rule. The simulation runs separately for each house with 2 million iterations (Zwick et al., 2003, ran 10,000 iterations in their simulation). The simulation shows results from testing the decision rule by generating 20 different sets of random numbers and running it for 2 million iterations (more details on simulation procedurals are presented in Appendix 2.D.) To compare the performance of stopping rules across treatments, I ran the simulation using search cost of 0 (representing the no cost treatment) and 20 ECUs (representing the monetary and time cost treatments). In general, the opportunity cost of time depends on the wage rate a participant can earn. The experiment is parameterized to yield a maximum payment of NZD 14 per hour. Using the implemented exchange rate (NZD 1 = 1000 ECUs), a 5-second delay thus equals approximately 2 cents (NZD 14/3600 seconds) or 20 ECUs.

I compared the performance of all stopping rules for all treatments by using the average payoff (in ECUs) they yield. The average payoff represents the mean value earned from accepting an offer according to a given stopping rule, net of the cumulative search cost. The average payoff statistic indicates which stopping rule yields the highest payoff.

It is almost trivial to see that due to the accumulation of search costs every stopping rule, except for “Accept the first offer” yields a higher payoff in no cost compared to the treatment with cost (Figure 2.1). The average payoff in treatments with cost is statistically significantly lower than in no cost ($p < .001$) according to both Mann-Whitney and the t-tests presented in Table 2.1. According to the simulation, the stopping rule that prescribes “Accept the next highest offer after seeing 4 offers” yields the highest payoff at no cost. In

the cost treatment, the optimal stopping rule is to always “accept the first offer.” The simulation thus suggests that in this environment, it is optimal to stop the search sooner (i.e., accept an earlier offer) when the search is costly as compared to a situation when it is not. This is due to the cumulative nature of the search cost which, given this parameterization, causes the overall payoff to decrease if one searches for more than one offer.

Table 2.1

Simulation summary statistics and between treatments statistical tests

	0 ECU*	20 ECU*	t-test		Cohen's d	Mann-Whitney Test	
	Average (SD)	Average (SD)	t(38)	<i>p</i>	<i>d</i>	U	<i>p</i>
Net payoff (ECUs)	520.2 (20.7)	147.9 (98.1)	17.1	< .001	5.4	2.0	< .001

Note: SD refers to standard deviation. *0 ECU refers to no cost treatment, 20 ECU refers to both monetary and announced time cost treatments.

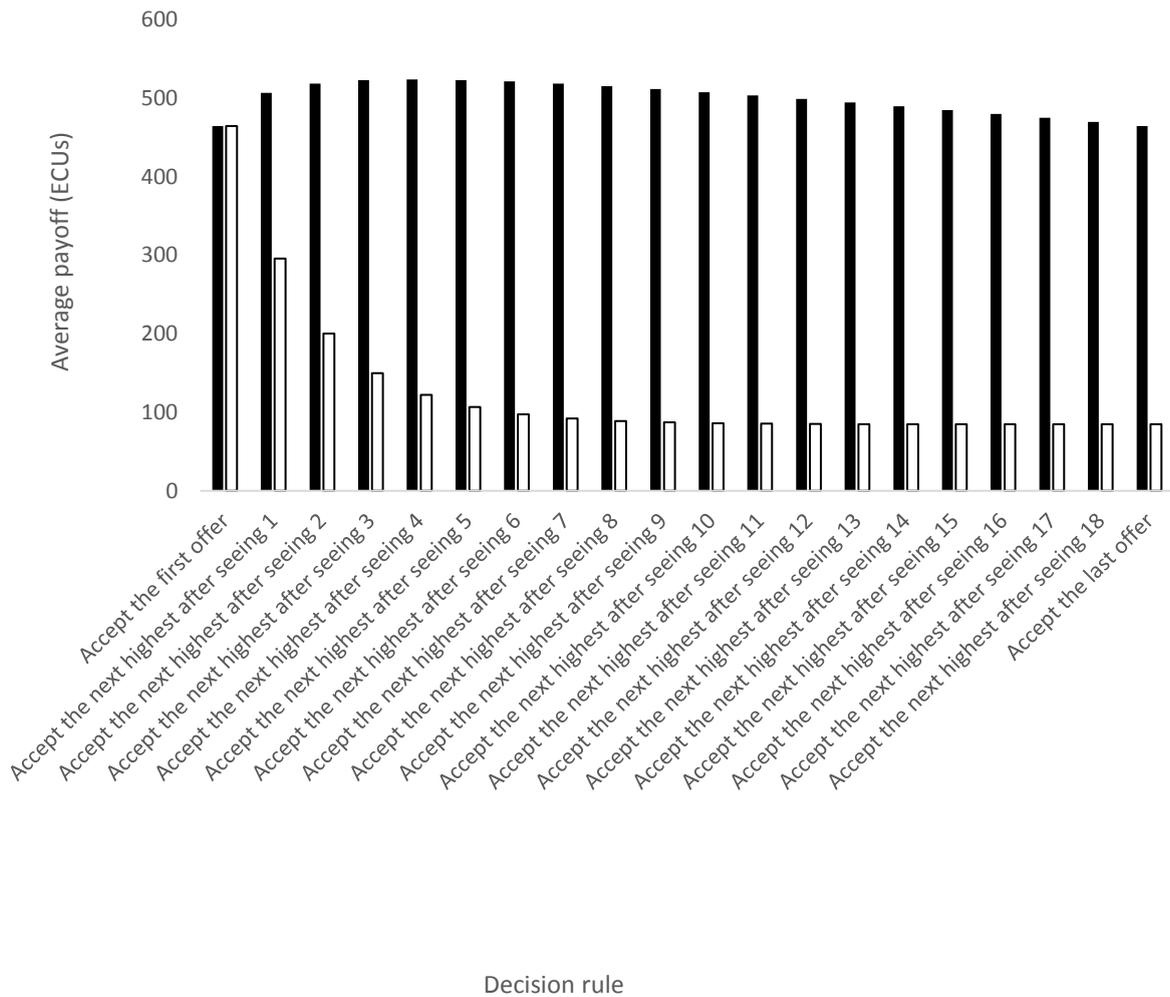


Figure 2.1. The average payoff for all stopping rules in the cost is equal to 0 ECU and 20 ECU (for both time and monetary) treatments. The filled bar represents 0 ECU and the blank bar is 20 ECU.

In addition to the simulation findings, Table 2.2 lists the net payoffs (resulting from all possible decision rules) for the 10 actual sequences of offers used in the experiment. As presented in Table 2.2, the optimal stopping rule that yields the highest net payoff when there is no cost is stopping rule 7. When it cost 20 ECUs per search, the stopping rule that

yields the highest net payoff is stopping rule 1. This displays very similar result to the simulation.

Table 2.2

Summary prediction results of net payoff (ECUs) after applying the stopping rules to the sequences in the experiment when there is 0 ECU and 20 ECUs (for both time and monetary). The average net payoff (in ECUs) is the mean value earned from accepting an offer according to a given stopping rule net of the cumulative search cost.

Decision rule	Variable	Average Net payoff (ECU)	
		0 ECU	20 ECU
Pick first offer		454.7	434.7
Choose the next biggest after seeing 1		662.9	525.7
Choose the next biggest after seeing 2		735.2	522.3
Choose the next biggest after seeing 3		722.8	378.9
Choose the next biggest after seeing 4		756.2	370.3
Choose the next biggest after seeing 5		761.1	373.2
Choose the next biggest after seeing 6		775.4	344.4
Choose the next biggest after seeing 7		793.0	352.0
Choose the next biggest after seeing 8		763.4	292.7
Choose the next biggest after seeing 9		780.8	308.1
Choose the next biggest after seeing 10		582.5	139.6
Choose the next biggest after seeing 11		582.5	139.6
Choose the next biggest after seeing 12		630.1	183.2
Choose the next biggest after seeing 13		609.6	183.2
Choose the next biggest after seeing 14		446.1	7.7
Choose the next biggest after seeing 15		446.1	7.7
Choose the next biggest after seeing 16		446.1	7.7
Choose the next biggest after seeing 17		446.1	7.7
Choose the next biggest after seeing 18		446.1	7.7
Pick last offer		407.7	7.7
Average		612.4	229.7

In line with previous research and the simulation results. I hypothesised that (1) participants will shorten their search when there is a monetary search cost, and people search longer when there is no search cost imposed; (2) assuming that the effects of monetary and time cost are qualitatively equivalent (given my calibration), the amount of search when there is a monetary cost should be the same as in the case when there is a time cost (presented as time delay), when time cost is also pre-announced; (3) in line with the theoretical assumption that monetary and time costs are qualitatively equivalent, I conjecture that imposing a time delay before a new offer is presented will cause the participants to search less than in the case where there is no such delay, just like monetary cost would, regardless of whether time cost is unannounced (3a) or announced (3b); (4) it is possible that without prior announcement of a time delay, it may require some learning before participants adjust their search behaviour with the presence of a time delay; therefore, the amount of search when there is an unannounced time delay may be more than when there is a prior announcement.

2.4 RESULTS

In light of the simulation results, I examined participants' behaviour in four ways. First, the stopping positions for all four treatments were computed to determine whether introducing time (both unannounced and announced) or monetary cost leads to shorter searches. Second, I examined the relationship between the observed stopping positions and payoffs to establish whether it is optimal to shorten a search if it involves time (either unannounced or announced) or monetary cost (I performed the analysis with and without subtracting the opportunity cost of time from the final payoffs). The participants are assumed to be aiming to maximise their expected payoffs after cost in this experiment (therefore, risk-neutral decision makers). Third, I examined the effect of learning in the

amount of search conducted in the first and second half of sessions. Four, I examined the factors that influenced participants' decision to accept an offer using regression analysis.

There are three dependent variables: first, the amount of search exerted in the task for different treatments, that is, the position in the sequence where the participant accepts the offer (henceforth stopping position) is evaluated; second, the prices they selected, that is, the cumulative sum of 10 chosen prices obtained by accepting the offers (henceforth total chosen price) is used; third, the amount of money they earned after cost, the sum of 10 chosen prices subtracting the cumulative search cost (henceforth net payoff). The offer and the net payoff are both presented in ECUs. The results relating to these three dependent variables are shown in Table 2.3 (Panels A, B and C respectively).

Table 2.3

Experiment summary statistics and between treatments statistical tests.

Panel A. Stopping position

	Average (SD)	Tukey Post Hoc test***			Analysis of Variance		
		<i>p</i>			<i>F</i> (<i>df</i>)	<i>p</i> (MS error)	Partial η^2
<i>No cost</i> ⁽¹⁾	11.6 (2.8)	.21 ⁽²⁾	<.001 ⁽³⁾	<.001 ⁽⁴⁾	28.0 (3, 184)	<.001 (8.05)	.31
<i>Unannounced time cost</i> ⁽²⁾	10.4 (2.2)	.21 ⁽¹⁾	<.001 ⁽³⁾	.27 ⁽⁴⁾			
<i>Monetary cost</i> ⁽³⁾	6.4 (3.5)	<.001 ⁽¹⁾	<.001 ⁽²⁾	<.001 ⁽⁴⁾			
<i>Announced time cost</i> ⁽⁴⁾	9.4 (2.8)	<.001 ⁽¹⁾	.27 ⁽²⁾	<.001 ⁽³⁾			
Sequence optimal before cost*					13.1		
Sequence optimal after cost**					5.6		

Panel B. Total chosen price (ECUs)

	Average (SD)	Tukey Post Hoc test***			Analysis of Variance		
			<i>p</i>		<i>F</i> (<i>df</i>)	<i>p</i> (MS error)	Partial η^2
<i>No cost</i> ⁽¹⁾	7338.8 (580.7)	.98 ⁽²⁾	<.001 ⁽³⁾	1.0 ⁽⁴⁾	11.90 (3,184)	<.001 (3.7E+5)	.16
<i>Unannounced time cost</i> ⁽²⁾	7387.9 (436.3)	.98 ⁽¹⁾	<.001 ⁽³⁾	.98 ⁽⁴⁾			
<i>Monetary cost</i> ⁽³⁾	6732.3 (899.2)	<.001 ⁽¹⁾	<.001 ⁽²⁾	<.001 ⁽⁴⁾			
<i>Announced time cost</i> ⁽⁴⁾	7342.9 (413.4)	1.0 ⁽¹⁾	.98 ⁽²⁾	<.001 ⁽³⁾			
Sequence optimal (before cost)*					9228		

Panel C. Net payoff (ECUs)

	Average (SD)	Tukey Post Hoc test***			Analysis of Variance		
			<i>p</i>		<i>F</i> (<i>df</i>)	<i>p</i> (MS error)	Partial η^2
<i>No cost</i> ⁽¹⁾	7338.9 (580.7)	<.001 ⁽²⁾	<.001 ⁽³⁾	<.001 ⁽⁴⁾	131.3 (3,184)	<.001 (3.3E+5)	.68
<i>Unannounced time cost</i> ⁽²⁾	5303.7 (462.0)	<.001 ⁽¹⁾	.60 ⁽³⁾	.51 ⁽⁴⁾			
<i>Monetary cost</i> ⁽³⁾	5459.2 (777.6)	<.001 ⁽¹⁾	.60 ⁽²⁾	1.0 ⁽⁴⁾			
<i>Announced time cost</i> ⁽⁴⁾	5469.0 (466.6)	<.001 ⁽¹⁾	.51 ⁽²⁾	1.0 ⁽³⁾			
Net sequence optimal**					7807.0		

*The average actual optimal price/position from the sequences used in the experiment prior to search cost; averaged across 10 sequences.

** The average optimal price/position from the sequences used in the experiment after search cost; averaged across 10 sequences.

*** This is the result of pairwise comparisons between two treatments, the small numbers in parentheses indicate the treatment compared to.

Note. SD shows the standard deviation. *df* refers to the degree of freedom.

2.4.1. STOPPING POSITIONS

To test the hypotheses, I computed the average stopping position for each participant across all 10 rounds and compared them between treatments. I observed that participants did indeed stop their search earlier in the cost treatments time (both unannounced and announced; hypothesis 3a & 3b) and monetary costs (hypothesis 1) than in the no cost treatment, with the average stopping positions being 10.4, 9.6, 6.4 and 11.6, respectively. Thus, on average (averaged across the unannounced and announced time cost treatments) the participants search through approximately two fewer offers if the time cost was involved and five fewer offers while paying the monetary cost – detailed results were reported in Table 2.3 Panel A. The participants in the no cost treatment stopped sooner than in the average sequence optimal before cost; the stopping position of the actual optimal offer from the sequence employed in the experiment. Both time and monetary cost treatments stopped later than the sequence optimal after cost; the stopping position of the highest price subtracting the cumulative search cost of each offer.

Analysis of variance found a significant main effect between the average stopping positions; results are reported in Table 2.3 panel A. Tukey HSD post hoc tests showed that participants in the no cost treatment chose to stop at a significantly later position than the monetary cost ($p < .001$; hypothesis 1) and announced time cost (with $p < .001$; hypothesis 3b) treatments. There was no significant difference found between the no cost and unannounced time cost treatments ($p = .21$; hypothesis 3a). However, the Kolmogorov–Smirnov test, reported in Table 2.4, detected the difference to be close to statistically significant ($p = .06$). The announced time cost treatment ($p < .001$, hypothesis 2) had a significantly longer search than the monetary cost treatment. Also, there was no significant difference found in the amount of search between unannounced and announced time cost

treatments ($p = .27$, hypothesis 4). However, the Kolmogorov–Smirnov test, reported in Table 2.4, detected the difference to be statistically significant ($p = .04$).

2.4.2 OPTIMAL DECISION-MAKING EVALUATION

To illustrate the impact of search cost on participant's payoffs, I compared the realized participants' payoffs between treatments. I did so by (a) comparing the total chosen prices (ignoring the monetary value of time and the monetary cost spent on searching), and (b) after subtracting search costs.

TOTAL CHOSEN PRICE (ECUs)

Participants in time cost (averaged across the unannounced and announced, $M = 7365.5$) selected the highest total chosen price (ECUs) on average compared to the no cost ($M = 7338.8$) and monetary cost ($M = 6732.3$), also see Table 2.3 Panel B. Tukey HSD post hoc tests confirmed that the monetary cost participants chose a significantly lower total price than no cost participants ($p < .001$). There was no significant difference found between no cost and the two time cost treatments (unannounced, $p = .98$; announced, $p = 1.00$). Similar differences were also found with the Kolmogorov–Smirnov test and reported in Table 2.4.

NET PAYOFF

Next, I compared the average net payoffs, which were calculated by taking the chosen price for each house and subtracting the monetary value of total time delay experienced to reach that offer in the time cost treatments (in both unannounced and announced). For example, if the 19th offer was accepted and its value is 500 ECU, the payoff equals $500 \text{ ECU} - (19 \times 20 \text{ ECU}) = 120 \text{ ECU}$, as the participant had to wait for 19×5 seconds (and the opportunity cost of 5 seconds is 20 ECU as explained earlier). In the monetary cost treatment, if the 19th offer is accepted and its value is 500 ECU, the net payoff equals $500 \text{ ECU} - (19 \times 20 \text{ ECU})$

= 120 ECU. So, for the same offer in the same position, both time and monetary cost treatments yielded the same net payoff.

Participants in the monetary cost ($M = 5459.2$) treatment earned a higher net payoff than the time cost treatments (averaged across the unannounced and announced time cost, $M = 5387.0$). Tukey HSD post hoc tests confirmed significant differences between the no cost and the three cost treatments (all with $p < .001$), see Table 2.3 Panel C for more details. Similar differences were also found with Kolmogorov–Smirnov test. These results are reported in Table 2.4.

Table 2.4

Summary of statistical findings for the comparison between the no, unannounced time, monetary, and announced time cost treatments.

No cost and monetary cost (hypothesis 1)

Variables	Treatments	Average (SD)	Kolmogorov–Smirnov test	
			Z	<i>p</i>
Average stopping position of 10 rounds	<i>No cost</i>	11.6 (2.8)	3.1	< .001
	<i>Monetary cost</i>	6.4 (3.5)		
Total chosen price (ECU)	<i>No cost</i>	7338.8 (580.7)	2.3	< .001
	<i>Monetary cost</i>	6732.3 (899.2)		
Net payoff (ECU)	<i>No cost</i>	7338.9 (580.7)	4.6	< .001
	<i>Monetary cost</i>	5459.2 (777.6)		

Monetary cost and announced time cost (hypothesis 2)

Variables	Treatments	Average (SD)	Kolmogorov–Smirnov test	
			Z	<i>p</i>
Average stopping position of 10 rounds	<i>Monetary cost</i>	6.4 (3.5)	2.20	< .001
	<i>Announced time cost</i>	9.4 (2.8)		
Total chosen price (ECU)	<i>Monetary cost</i>	6732.3 (899.2)	2.50	< .001
	<i>Announced time cost</i>	7342.9 (413.4)		
Net payoff (ECU)	<i>Monetary cost</i>	5459.2 (777.6)	.86	.44
	<i>Announced time cost</i>	5469.0 (466.6)		

No cost and unannounced time cost (hypothesis 3a)

Variables	Treatments	Average (SD)	Kolmogorov–Smirnov test	
			Z	<i>p</i>
Average stopping position of 10 rounds	<i>No cost</i>	11.6 (2.8)	1.34	.06
	<i>Unannounced time cost</i>	10.4 (2.2)		
Total chosen price (ECU)	<i>No cost</i>	7338.8 (580.7)	.45	.99
	<i>Unannounced time cost</i>	7387.9 (436.3)		
Net payoff (ECU)	<i>No cost</i>	7338.9 (580.7)	4.70	< .001
	<i>Unannounced time cost</i>	5303.7 (462.0)		

No cost and announced time cost (hypothesis 3b)

Variables	Treatments	Average (SD)	Kolmogorov–Smirnov test	
			Z	<i>p</i>
Average stopping position of 10 rounds	<i>No cost</i>	11.6 (2.8)	1.90	.002
	<i>Announced time cost</i>	9.4 (2.8)		
Total chosen price (ECU)	<i>No cost</i>	7338.8 (580.7)	.57	.91
	<i>Announced time cost</i>	7342.9 (413.4)		
Net payoff (ECU)	<i>No cost</i>	7338.9 (580.7)	4.70	< .001
	<i>Announced time cost</i>	5469.0 (466.6)		

Unannounced time cost and announced time cost (hypothesis 4)

Variables	Treatments	Average (SD)	Kolmogorov–Smirnov test	
			Z	<i>p</i>
Average stopping position of 10 rounds	<i>Unannounced time cost</i>	10.4 (2.2)	1.40	.04
	<i>Announced time cost</i>	9.4 (2.8)		
Total chosen price (ECU)	<i>Unannounced time cost</i>	7387.9 (436.3)	.71	.70
	<i>Announced time cost</i>	7342.9 (413.4)		
Net payoff (ECU)	<i>Unannounced time cost</i>	5303.7 (462.0)	1.0	.24
	<i>Announced time cost</i>	5469.0 (466.6)		

2.4.3 EVIDENCE OF LEARNING

A repeated-measures ANOVA with treatments (no, unannounced time, monetary and announced time cost) as a between-subjects factor and halves (*First* vs. *Second*) as a within-subjects factor, detected a significant effect of *halves*, $F(1,184) = 6.63$, $p = .01$, $MS_{error} = 3.95$, partial eta-squared = .035. The interaction effect between stopping position in halves and treatments was not significant, $F(3,184) = 2.07$, $p = .11$. Tukey HSD post hoc tests comparing the average stopping positions between the first and second half of the session showed no significant difference in no cost ($p = .32$), unannounced time cost ($p = .99$), monetary cost ($p = 1.0$) and announced time cost ($p = .10$). The amount of search remains consistent throughout the session in all treatments; there was no evidence of further shortening the amount of search after learning time delay in unannounced time cost. Figure 2.2 graphically present the results.

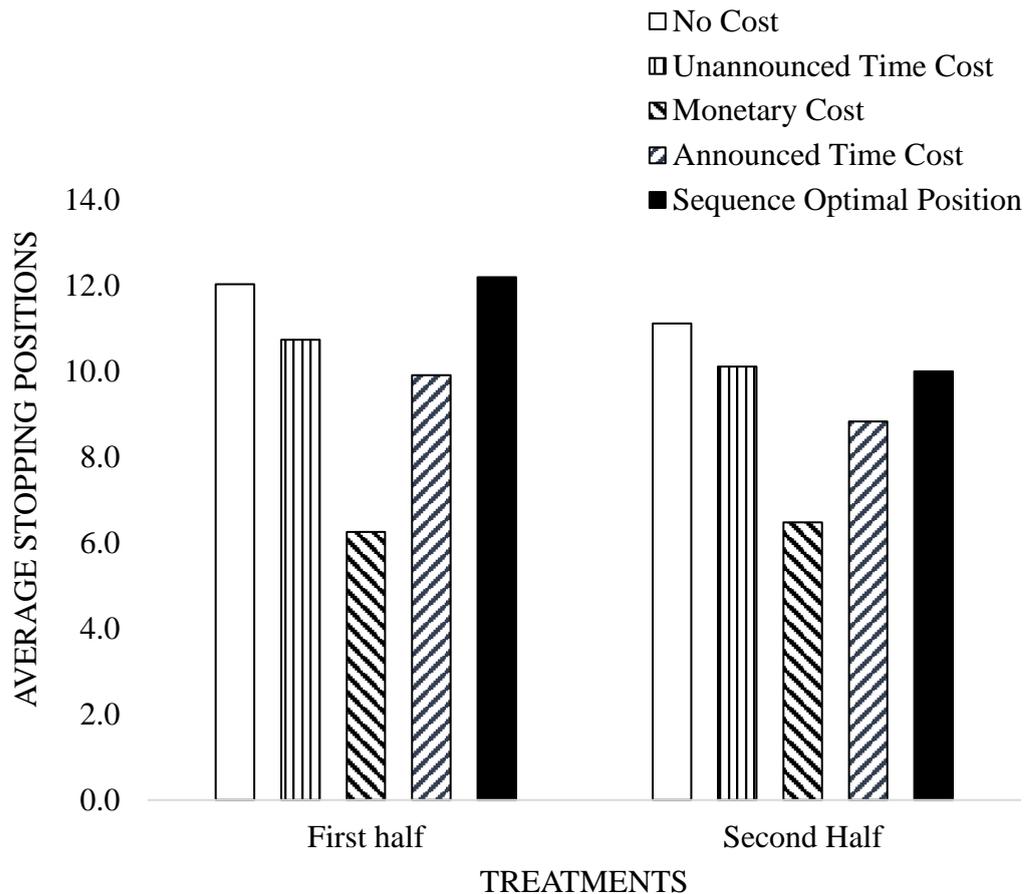


Figure 2.2. The half session average stopping position in the first (round 1–5) and second half (round 6–10) of the session for the no-cost, unannounced time cost, monetary cost and announced time cost treatments.

2.4.4 REGRESSION ANALYSIS

I next analysed potential factors influencing participants’ search behaviour using multiple regression. Table 2.5 Panel A represents the results in predicting stopping positions, and the results for predicting chosen prices are presented in Table 2.5 Panel B. In all four treatments, the last rejected offer prior to acceptance was significant in predicting the participants’ stopping position. In the no cost and time cost treatments (unannounced and announced), the higher the last rejected offer prior to acceptance, the longer the search.

But, the reverse was true in the monetary cost treatment. Also, in the no cost and the time cost (unannounced and announced) treatments, the higher the price they rejected in the sequence, the sooner the participants stopped. Yet, in the monetary cost treatment, the higher the price they have rejected, the longer they searched. For all four treatments, when participants have accepted higher prices in the previous round, they conducted more searches in the current round. For the unannounced time cost treatment, the longer the participants searched in the previous round, the less they searched in the current round. But, the reverse was true in the monetary cost treatment.

Interestingly, one highly significant factor was found in predicting the price that the participant chose in all four treatments. The highest seen offer was the best predicting factor in the no cost and time cost (unannounced and announced) treatments. The participants chose higher prices when they had seen a higher price from the previously reviewed sequence. Yet, the last rejected offer prior to acceptance was the only significant factor in predicting chosen price in the monetary cost treatment. When there was a monetary cost, people made decisions from a shorter term perspective (the most recent rejected price) than for the no cost treatment and the treatments with time cost.

Table 2.5

Coefficients of multiple regression for the tested factors in the no-cost, unannounced time cost, monetary cost and announced time cost treatments.

Panel A. Stopping position as the predicting variable

Factors	Treatments	<i>No cost</i>	<i>Unannounced time cost</i>	<i>Monetary cost</i>	<i>Announced time cost</i>
		Coef. (<i>p</i>)	Coef. (<i>p</i>)	Coef. (<i>p</i>)	Coef. (<i>p</i>)
Last rejected offer prior to acceptance		.23 (<i>< .001</i>)	.12 (.02)	-.70 (<i>< .001</i>)	.16 (.001)
Highest seen price		-.29 (<i>< .001</i>)	-.24 (<i>< .001</i>)	.58 (<i>< .001</i>)	-.32 (<i>< .001</i>)
Accepted offer in previous round		.26 (<i>< .001</i>)	.28 (<i>< .001</i>)	.32 (<i>< .001</i>)	.20 (<i>< .001</i>)
Stopping position in previous round		-.04 (.39)	-.09 (.05)	.22 (<i>< .001</i>)	.07 (.12)
R ²		.16 (<i>< .001</i>)	.17 (<i>< .001</i>)	.26 (<i>< .001</i>)	.14 (<i>< .001</i>)

Panel B. Chosen price as the predicting variable

Factors	Treatments	<i>No cost</i>	<i>Unannounced time cost</i>	<i>Monetary cost</i>	<i>Announced time cost</i>
		Coef. (<i>p</i>)	Coef. (<i>p</i>)	Coef. (<i>p</i>)	Coef. (<i>p</i>)
Last rejected offer prior to acceptance		-.006 (.86)	-.05 (.07)	.99 (<i>< .001</i>)	-.05 (.008)
Highest seen price		.77 (<i>< .001</i>)	.87 (<i>< .001</i>)	-.002 (.94)	.97 (<i>< .001</i>)
Accepted offer in previous round		-.002 (.93)	-.009 (.74)	.002 (.82)	.04 (.01)
Stopping position in previous round		.09 (.002)	.02 (.45)	.009 (.29)	.02 (.28)
R ²		.63 (<i>< .001</i>)	.73 (<i>< .001</i>)	.97 (<i>< .001</i>)	.88 (<i>< .001</i>)

Note. Coef. refers to the standardized coefficient of the tested factors.

2.5. DISCUSSION

I investigated the effect of time and monetary cost on search behaviour. The implemented sequential search task was framed as selling houses, which in everyday life is a time-consuming process. I first conducted simulations to show how (monetary) search cost should lead to a shorter search in my experimental setting, and which, if participants considered the opportunity cost of time, should have a similar effect to time cost. The experiment confirms the simulation results as the participants did indeed search less when the monetary search cost was introduced as an experimental manipulation (hypothesis 1). People also conducted less search when there was a monetary cost than a time cost (hypothesis 2). When incurring a time delay, people searched less than in the no cost treatment when there was a prior announcement of the delay (hypothesis 3b) but not when there was no such announcement (hypothesis 3a). However, when comparing the amount of search between announced and unannounced time delay, the statistical tests yielded both significant and non-significant findings (hypothesis 4). This suggests that the search behaviour may be influenced by the time search cost even without having announced it prior to the search.

I observed that participants in the time and monetary cost treatments stop their search later than the simulation suggests. This may not be surprising in the unannounced time cost treatment, given that participants had to infer the time delay through experience. Despite this conservative design feature, a result is observed with the Kolmogorov–Smirnov test that is close to statistically significant. (The simulation was based on being fully informed about the existence and magnitude of the announced time cost). However, there was no evidence of learning and further shortening the search in the unannounced time cost

treatment, as there was no significant difference in the stopping position between the first and second half of session in all treatments. The shortening effect was indeed stronger when the delay was pre-announced at the beginning of the experiment, and yet it was still smaller than the simulation predicts.

There are at least two potential underlying causes for the difference between the observed amount of search of time and monetary costs. First, the perceived value of participants' time might not equal their actual opportunity cost of time (the way I calculated it and used in the time cost treatments). Some previous research also reports that money and time are processed differently (Lee, Lee, Bertini, Zauberman & Ariely, 2015). A 10-second delay treatment was conducted prior to the 5-second delay treatment, and over half of the participants dropped out of the experiment in that session. People who stayed for the whole session had the average stopping position of 6.9. This potentially suggests that people treat time and monetary costs in different ways. When time cost is too much to afford, people refuse to pay it instead of trying to minimise it the way they would with monetary cost. Second, the implemented framing of selling houses might have interacted with the perception of time cost (and its monetary value as calculated in the monetary cost treatment). This second explanation is in line with the demonstrated finding that different frames can elicit systematically different decisions (see, e.g., Tversky & Kahneman, 1981), and in some cases, people make better decisions with frame than without (see Experiment Three for reference). Housing is usually associated with positive images (e.g., comfort, security, childhood memories) and when something is being perceived more positively, it may also be perceived as more important, something more worth doing; that is, the halo effect (e.g., Nisbett & Wilson, 1977; Kahneman, 2011), hence they may be more inclined to invest more time, therefore search longer than they would if only the experimental payoffs were taken into account, as in the simulation.

3. A HOUSE-SELLING CONTEXT IMPROVES DECISION MAKING FROM A SEQUENCE OF OFFERS. EXPERIMENT TWO.

3.1 INTRODUCTION

Information, experience and knowledge are meaningless without considering the context. For example, the decision to drive 50 kilometres per hour can be either hazardous or cautious. It may be hazardous to drive at a speed of 50 kilometres per hour in bad weather and it may be cautious to drive at the same speed on a sunny day. Previous research has shown that decisions are sensitive to how the context is worded or framed (e.g., Tversky & Kahneman, 1981; McNeil, Pauker, Sox, & Tversky, 1982; Meyerowitz & Chaiken, 1987; Hoffman, McCabe, Shachat, & Smith, 1994; Dufwenberg, Gächter & Hennig-Schmidt, 2011). Framing a context is known to have an effect that makes people's decisions or behaviour change systematically, based on the description given. The question is: Do people make better decisions in a context? For example, they might solve more problems. There is evidence that framing improves decisions in various tasks (e.g., Griggs & Cox, 1982; Sugiyama, Tooby, Cosmides, 2002; Cosmides, Barret & Tooby, 2010). This ability to make better decisions under context can be explained by the dual-processing theory. Dual-processing theory proposes that most daily decisions are made by associating a new situation with existing knowledge in similar experiences, rather than forming new knowledge and information for each new experience (Kahneman, 2003). Therefore, people can use existing schemas that contain effective strategies constructed from previous experiences to make decisions. However, it is not clear how much information is necessary to activate this process. In this experiment, I test whether providing a house-selling context in a variant of the secretary problem does result in better decisions than without the context.

The contribution can be viewed as extending the empirical analysis of the context-framing effect to the domain of schema activation.

A schema is a system of organising and perceiving new information, which is then encoded as the default assumptions of the world. Schemas forms mental structures that describe how the world works, and how we interact with the world (see Bower, & Cirilo, 1985; Dimaggio, 1997; Narvaez & Bock, 2002, for more details). For example, when someone holds a schema that maximising profit is the best approach for making decisions, he or she will consistently re-apply this schema in various economic decisions. Since every decision is made within a context, it is therefore, crucial to understand the link between the information needed in a context and the implications for decision making.

Is there any existing evidence that would suggest that people make better decisions in a context? Behavioural literature does provide such evidence in various domains. For example, Griggs and Cox (1982) found that people were more likely to solve a reasoning problem correctly when framed as a drinking-age problem than as an abstract, context-free problem. Similar results were obtained by Pollard and Evans (1987). The same reasoning problem is also more likely to be solved correctly when framed as a social exchange problem than as an abstract, context-free problem (e.g., Cosmides, 1989; Cosmides & Tooby, 1992; Cosmides & Tooby, 2005; Cosmides, Barret & Tooby, 2010). In a similar fashion, Eger and Dickhaut (1982) found that accounting students performed better in a probability judgement task that was framed in an accounting context than as an abstract, context-free problem.

Although different contexts can lead to different decisions, the issue of how much information in a context is necessary to change behaviour remains unknown. In this chapter, I experimentally examine whether the presence of a simple context without descriptions also

facilitates better performance. This experiment can therefore provide evidence of whether a context without description does indeed lead to better decisions by the participants, and empirically supports the result of Pollard and Evans (1987; albeit in a different type of task) that the description of a context does not facilitate better decisions. However, no previous studies have considered the effect of context on the secretary problem or any other sequential decision-making task.

The human brain is often thought to be the result of an evolutionary process to resolve problems and enhance survival back in the Stone Age as hunter-gatherers (Cosmides, Tooby, 1992). The brain enables a collection of cognitive mechanisms that guide our behaviour and decision making. For example, cooperation is found in both babies and chimpanzees (Warneken, Chen, & Tomasello, 2006), altruistic behaviours in babies at 14 months of age (Warneken & Tomasello, 2007), as is the concept of object permanence in infants (Baillargeon, & DeVos, 1991) and monkeys (Churchland, Chou & Lisberger, 2003). These mechanisms are likely acquired through evolution to better solve problems and survive. To bear in mind the context while making a decision is possibly one of these mechanisms. Consider this scenario: Approximately 500,000 years ago, a *Homo sapiens* is resting in the bush after a long day of work and sees a shadow of a rock with a long rope attached. What is the likely response? Options are (1) Keep resting, (2) Escape from the bush. The *Homo sapiens* who decides to escape has taken into account the context when making that decision; predatory animals are often found in the bush, and they tend to first hide to observe the target before attacking. Those who do not consider the context, would not make the decision to escape, and they would be less likely to survive and pass down their genes. Yet, even now, there is no definite theory explaining the effect of context framing; nor does the mathematically derived optimal decision rule account for the effect of context in making decisions. The gap between theory and actual behaviour, that is, stopping

prior to the theoretical optimal, may be because the optimal decision rule fails to account for the effect of context.

3.1.1 DIFFERENT DECISIONS UNDER DIFFERENT CONTEXT FRAMES

A long line of research has found that decisions change under different context frames (e.g., Levin & Gaeth, 1988; Duchon, Dunegan & Barton, 1989; Gamliel & Peer, 2010). An early study of the issue is that of Kahneman and Tversky (1984), who explored how different phrasing with the same outcome affected people's preference in hypothetical life-and-death decision-making scenarios. The decision was presented to participants either with positive framing, for example, 2 of 3 people would live or with negative framing, 1 of 3 people would die. They found that although the outcome was the same in both scenarios, most people preferred the positively framed scenario that presented as how many people would live instead of die.

Dual-processing theory has been proposed by numerous researchers to explain why different decisions result from how the problem is framed (e.g., Kahneman, 2003; Evans, 2008). These dual-processing theorists propose that the decision-making process relies on both intuitive/heuristic and analytic/executive processes. System one involves implicit (unconscious) processing that uses intuitive and heuristic forms of reasoning that operate in most of our everyday reasoning and decision making. It is a domain-specific and contextualised, fast, and automatic response requiring very little effort. System two involves explicit (conscious) processing; these analytic/executive operations tend to be slow, controlled, serial and effortful (De Neys, 2006). The context framing effect can be caused by different decision-making schemas belonging to system one.

Some schemas are activated chronically due to the regular contact of environmental context (Bargh, Lombardi, & Higgins, 1988). For example, when an individual learns from

repeatedly looking for a car-parking space that is closest to the destination, he or she learns that a certain way facilitates finding the best parking space, and other ways do not. These schemas are activated involuntarily. They are formed from previous experiences and are then used to organise or integrate new information (see Bower, & Cirilo, 1985; Dimaggio, 1997; Narvaez & Bock, 2002, for more detailed discussion on schemas). Once schemas are formed, they operate constantly in the brain and are activated by stimuli that resemble the stimuli that were present when the schema was first created (see Higgins & Chaires, 1980 Narvaez, & Bock, 2002 for more detail). To return to the previous example. If the experiences of finding car-parking spaces couple with experiences in searching for the best car to buy within a given price range, the brain may form a fuller mental model of how to make sequential search decisions generally. This may then be activated when a similar situation, for example, finding an apartment, arises. Heider and Simmel (1944) found evidence to support the hypothesis of re-applying existing schema to explain a new experience. In their experiment, the participants watched a short animated film. The film had a motionless large square with a door that opens and closes in one side. First, a big triangle appeared inside the square, then a small triangle and circle appeared. As the door flapped open and shut, the geometric figures slid around the screen. After ninety seconds or so, the small triangle and the small circle disappeared, and the big triangle breaks down the big square.².

After the participants watched the animated film, they were asked to describe what they saw. Only 3 of 114 participants reported seeing geometric shapes moving around a screen. The majority reported a narrative and attributed agency and intent to the shapes: for example, a romance story between the small triangle and the small circle, the big triangle is

² The complete film can be watched on YouTube <https://www.youtube.com/watch?v=VTNmLt7QX8E>

the angry parent who wants to separate them, and so on. This experiment demonstrated that people can explain a new situation using a similar, existing mental structure or schema. Although we may feel we are experiencing novelty every day, the novelty is perceived and interpreted by existing schemas without consciously processed in the brain (Wegner & Wheatley, 1999).

3.1.2 EVIDENCE OF MAKING BETTER DECISIONS WITH CONTEXT.

Just as a schema can be activated through the same or a similar stimulus encountered in previous experiences, making decisions in a context allows us to effectively resolve problems and make decisions without starting from scratch every time. Experimentation on the Wason selection task demonstrates how context framing enhances people's ability to solve problems. The Wason selection task is a logic puzzle to test deductive reasoning. The participant is presented a set of four cards placed on a table, where each card has a number on one side and a letter on the other side. The visible faces of the cards show A, D, 4 and 7. Which card(s) must you turn over in order to test the truth of the proposition that if a card has a vowel on one side, then it has an even number on the other side?

Only 5 % of the participants were able to solve the context-free problem correctly (Johnson-Laird & Wason, 1970). A drastic increase in accuracy is consistently reported in versions of the task that involve a social exchange to detect cheaters (e.g., Cosmides, 1989; Cosmides & Tooby, 1992; Cosmides & Tooby, 2005; Cosmides, Barret & Tooby, 2010), cross-cultural experiments with the SchiWiar of Ecuadorian Amazonia (Sugiyama, Tooby, Cosmides, 2002), and as a drinking-age problem (Griggs & Cox, 1982). Griggs and Cox reported that when people were asked to solve the problem as a drinking-age problem, 73% of them were able to solve it correctly. In this form of the problem, the task framed the problem as a police officer who is ensuring drinking-age regulations are being followed in a

bar. The participants were presented with 4 cards, each with information about one person in a bar. One side of a card tells a person's age and on the other side is what the person is drinking. The task is to identify the card(s) that violate this rule: If a person is drinking beer, then the person must be over 19 years of age.

3.1.3 HOW MUCH INFORMATION DO WE NEED TO ACTIVATE EXISTING SCHEMA?

Pollard and Evans (1987) argued that two different aspects that change between the context-free problem and the context-framed problem can potentially contribute to better performance. For example, the drinking-age problem provides content (drinks and age), as well as the context or scenario (policeman checking patrons' age in a bar), within which this content is relevant. Pollard and Evans (1987) demonstrated experimentally that both the relevant content and context of the drinking-age problem were necessary to facilitate an increase in accurate responses. However, they suggested that context may be a stronger contributor to people's performance than content.

3.1.4 HYPOTHESIS

The following experiment was set up to examine how much information is necessary to facilitate better decisions. To the best of my knowledge, whether information in a context is required to elicit better decision making has not been tested experimentally, although the context is presented in every decision. Nor, to my knowledge has the effect of context framing been explored in the secretary problem. For this experiment, the manipulated variable is the degree of context from an issue encountered in everyday life. The basic task provides a house-selling context in which the participants earn payoffs based on the offers they accept, and the offers are presented sequentially. Measures of how good the decision is can then be based on the stopping position and total earnings. I assume that the participants are risk-neutral decision makers and that they aim to maximise their expected payoffs. The

first hypothesis, formulated with reference to experiments on the Wason selection task, is that people make better decisions when presented with a context than without.

Theoretically, this would be because having a context allows one to better assess the situation and access existing schema constructed from a similar experience, for example, selling a car or even buying and selling a house in real life. When there is no context available, the person experiences difficulty in determining what schema to apply and the chance of applying an inappropriate schema is increased. Thus, treatments with a house-selling context are expected to produce higher total earnings and closer to the optimal amount of search than the no context treatment.

Furthermore, it is possible that having a fairly complete description of the house – for example, floor area, number of bedrooms, the year the house was built – is also critical in making better decisions, and yet such information could also be confusing and distracting, resulting in misuse of the information. Some of the house information could be perceived by the participants to play a bigger role in determining the house value than other information. In line with Pollard and Evans (1987), I hypothesise that having content does not facilitate better decisions. Thus, the total earnings and stopping position are expected to be the same in the two house-selling context treatments: those with and without the house description.

Overall, in line with previous research, I hypothesised that (1) people make better decisions with a house-selling context than in the treatment with no context; (2) quality of decisions when presented in a house-selling context with house information is no different to that without house information.

3.2 METHOD

3.2.1 EXPERIMENTAL DESIGN

The experiment aimed to analyse the effect of context on search activity in a sequential decision-making task. The experiment consisted of three treatments implemented in an across-subject design, varying the context of the task. The available price offers were identical throughout all the treatments. More details on the price offers are presented in Table 3.1. In the house with information treatment (henceforth house with info), the experiment is framed as a selling-house task (the house with info treatment is also reported in Experiment One as the treatment with no cost), the selling house instructions read as follows (also see Appendix 2.A):

“You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. You will be given a brief description of the house that will be followed by a series of price offers. The price offers are randomly generated by the computer and available one at a time. Once a price offer is presented, you can either accept or reject it. If you accept the price offer, the house will be sold at the price you accepted. All sales are final. If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. In total there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. *Therefore, make your decisions carefully.* There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.”

There were two practice rounds followed by 10 paid rounds. A description of a house, consisting of the floor area, the number of bedrooms, suburb, and year the house was built, was presented prior to the price offer in the house with info treatment. Each round

featured a different house description, and a different randomized sequence of price offers. The details of mean and standard deviation of the price offers are provided in Table 3.1. All the house descriptions used in the experiment were taken from houses sold in October 2014 in 12 different Christchurch suburbs. This information was obtained from the Quotable Value Limited database (qv. co. nz).

The house frame treatment (henceforth house frame) employs an identical instruction to the house with info treatment, except that no description of the houses was presented prior to the price offers (also see Appendix 3.A.)

In the no context treatment (henceforth no context), the experiment description instructions explained that the task was to choose a number, without any selling-house framing, description, or indeed mention, of a house. The instructions for the no context treatment were as follows (Also see Appendix 3.B):

“You will participate in 10 rounds. In each round, you will be asked to decide whether to accept or reject a number. The numbers are randomly generated by the computer and available one at a time. Once a number is presented, you can either accept or reject it. If you accept the number, you receive the amount represented by the number (in experimental currency units, as will be explained below). All decisions are final. If you reject the number, the number will disappear; you cannot go back to the previously rejected number. In total there are 20 numbers available; if you have not accepted a number prior to the 20th number, you will be forced to accept the 20th (i.e. the final) number. Therefore, make your decisions carefully.”

The no context treatment also had two practice rounds and 10 paid rounds.

All sequencing and numerical information was identical across the three treatments. Thus, only ten randomised sequences of numbers were employed in the entire experiment; Table 3.1 shows detailed information on the sequences employed.

Table 3.1

Data for the price offer sequences employed in the experiment.

Variable Round	Sequence optimal		Predicted optimal		Min.	Average price	SD
	Position	Price	Position	Price			
1	8	848	8	848	276	509.6	165.4
2	10	875	8	818	2	469.2	284.4
3	10	708	10	708	207	437.6	147.2
4	20	733	20	733	267	518.5	145.5
5	13	578	10	484	186	331.2	114.4
6	10	1574	9	1400	89	714.3	447.4
7	19	581	19	581	197	369.2	128.1
8	3	966	20	541	250	636.4	234.4
9	14	1740	12	1264	105	756.4	396.2
10	4	625	20	553	250	440.4	101.3
Average	11.1	922.8	13.6	793.0	183	518.3	216.4

Notes. Sequence optimal position = the position at which the highest price presents in the sequence employed; Sequence optimal price = the highest price value in each round, also the maximum; Predicted optimal position = the stopping position predicted by the optimal stopping rule – the stopping rule that yields the highest earning, see Appendix 3.C for more details; Predicted optimal price = the price at the position predicted by the optimal stopping rule, refer to Appendix 3.C for details. Min. = the lowest price in each round. Average price = the average price in the sequence employed for each round. SD = the standard deviation of 20 price offers in each round.

3.2.2 PROCEDURES AND PARTICIPANTS

One hundred and thirty-seven students from the University of Canterbury in Christchurch, New Zealand participated in the experiment, which was run at New Zealand Experimental Economics Laboratory (NZEEL). A total of 46 students participated in the no context treatment, 43 and 48 students participated in the house frame and house with info respectively. A session lasted on average forty-five minutes and the participants earned NZD12.1 on average. The average earning in the no context frame is NZD11.8, both house frame and house with info receive NZD12.3. The total price chosen by the participant in ECU is then converted into New Zealand dollars using the conversion rate of 1000 ECU to 1 New Zealand dollar and is explained in the written instructions. The payoff protocol was single-blind, meaning that the experimenter was able to match participant decisions to their identity. The remaining procedural details were identical to Experiment One.

3.3 RESULTS

First, I considered the average results obtained from the experiment. The average results examined the two dependent variables that assess the quality of the decision making. Then I examined whether people improved their decisions over time, by comparing the decisions in the first and second half of the session. Finally, I deduced correlation results between the search activities and the chosen prices.

There were two dependent variables. First, the amount of search exerted in the task for different treatments, that is, the position in the sequence where the participant accepts the offers (henceforth stopping position) was evaluated. Second, the amount of money they earned, the cumulative sum of 10 chosen prices obtained by accepting the offers (henceforth total chosen price) was used. The price offer/number is presented in ECU. The results relating to these two dependent variables are shown in Table 3.2 (Panels A, B respectively).

Table 3.2

Experiment summary statistics and between treatments statistical tests.

Panel A. Stopping position

	Average	SD	Tukey Post Hoc test***		Analysis of Variance			
			<i>p</i>	<i>F</i> (<i>df</i>)	MS error	<i>p</i>	Partial η^2	
<i>No context</i> ⁽¹⁾	9.8	3.6	.07 ⁽²⁾	.015 ⁽³⁾	4.3 (2, 134)	9.4	.015	.06
<i>House framing</i> ⁽²⁾	11.2	2.8	.07 ⁽¹⁾	.60 ⁽³⁾				
<i>House with info</i> ⁽³⁾	11.6	2.8	.015 ⁽¹⁾	.60 ⁽²⁾				
Sequence optimal*			11.1					
Predicted optimal**			13.6					

Panel B. Total chosen price (ECU)

	Average	SD	Tukey Post Hoc test***		Analysis of Variance			
			<i>p</i>	<i>F</i> (<i>df</i>)	MS error	<i>p</i>	Partial η^2	
<i>No context</i> ⁽¹⁾	6880.3	750.0	.005 ⁽²⁾	<.001 ⁽³⁾	8.1 (2, 134)	.00003	<.001	.11
<i>House framing</i> ⁽²⁾	7269.5	361.4	.005 ⁽¹⁾	.84 ⁽³⁾				
<i>House with info</i> ⁽³⁾	7338.9	580.7	<.001 ⁽¹⁾	.84 ⁽²⁾				
Sequence optimal*			9228.0					
Predicted optimal**			7930.0					

*Sequence optimal refers to the actual optimal position/price from the sequences used in the experiment, more details in Appendix 2.D. (Average column 2 in Table 3.1 shows the stopping position, average column 3 in Table 3.1 shows the average price for 10 houses.)

** This is the predicted result from applying the optimal decision rule presented in Appendix 3.C. (Average column 4 in Table 3.1 shows the stopping position, average column 5 in Table 3.1 shows the average price for 10 houses.)

*** This is the result of pairwise comparisons between two treatments; the small numbers in parentheses indicate the treatment compared to.

Note. SD shows the standard deviation, *df* refers to the degree of freedom.

3.3.1 OPTIMAL DECISION-MAKING EVALUATION

STOPPING POSITION

Participants in the no context treatment chose to stop at a significantly earlier average position ($M = 9.8$) than those in house frame ($M = 11.2$) and house with info ($M = 11.6$). Tukey HSD post hoc tests showed a significant difference between no context and house with info treatment ($p = .015$, hypothesis 1), and a finding that was close to significant between no context and house frame ($p = .07$, hypothesis 1). There was no significant difference between house frame and house with info ($p = .60$, hypothesis 2). (Similar results were obtained from non-parametric test, Kolmogorov–Smirnov test; see Appendix 3.D.) Thus, participants invested in fewer searches when no context was given. As a result, no context stopped their search approximately 1.3 positions prior to the sequence optimal position ($M = 11.1$), which is approximately 11.7% less. In contrast, the house frame and house with info treatments on average were only approximately 0.3 positions more than the sequence optimal position, which is only approximately 2.7% away from the optimal (Table 3.2 panel A).

TOTAL CHOSEN PRICE (ECU)

Participants in house with info selected the highest total chosen price (ECU) on average ($M = 7338.9$) compared with those in the house frame ($M = 7269.5$) and no context ($M = 6880.3$) frame. On average, the house-selling context treatments (averaged across house frame and house with info) yielded approximately 80% of the maximum total price (which is the sum of the highest price in 10 rounds), and no context earned approximately 75%. Analyses of variances showed a significant main effect of treatment on the average

total chosen price, as shown in Table 3.2 Panel B. Tukey HSD post hoc tests confirmed significant differences between the total chosen price in the no context and the other two treatments (the house frame: $p = .005$, and house with info: $p < .001$, hypothesis 1). There was no significant difference between house frame and house with info ($p = .84$, hypothesis 2). (Similar differences were also found with the Kolmogorov–Smirnov test; See Appendix 3.D.)

3.3.2 DOES CONTEXT FRAMING EFFECT LAST?

STOPPING POSITION

Next, to examine the duration of the context framing effect, decisions on the first half (rounds 1–5) of the session were compared to the second half (rounds 6–10) of the session. First, I computed the average stopping position for each participant across the first half and second half of the session and compared the results across the treatments (no context, house frame, and house with info). Results are presented in Table 3.3. Analysis of variance results showed that there was a significant main effect between the average stopping positions in the halves (first or second), and the frames (no context, house frame or house with info). Participants in the second half of the session stopped significantly earlier than in the first half. There is also a statistically significant ($p = .007$) interactive effect found between the frames and session halves. This result indicates that the inferior decisions of people without context increases rather than decreases in the long term. The results are graphically presented in Figure 3.1.

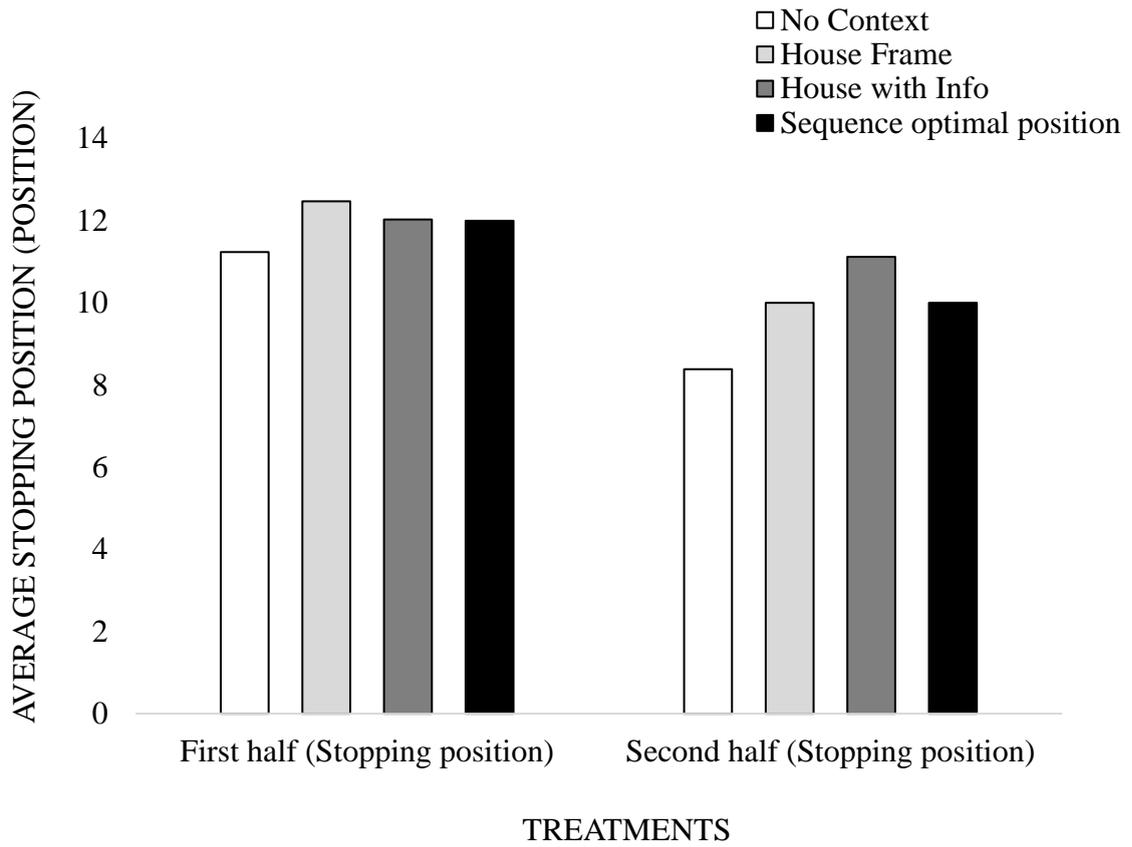


Figure 3.1. The average stopping position averaged across participants in the first half of session (rounds 1–5) and the second half of session (rounds 6–10) in the no context, house frame, and house with info treatments.

Table 3.3

Summary of statistical findings for the session halves and the treatments.

		Analysis of Variance				
		<i>F</i>	<i>df</i>	MS error	<i>p</i>	Partial η^2
Stopping position	Main effect	61.0	1, 134	4.84	< .001	.31
	Interaction effect with framing	5.1	2, 134		.007	.07
Total chosen price	Main effect	472.4	1, 134	53690	< .001	.78
	Interaction effect with framing	1.6	2, 134		.21	.02

Note. *df* refers to the degree of freedom

TOTAL CHOSEN PRICE

The total chosen prices for each participant across the first half (rounds 1–5) and second half (6–10) of the session were computed and compared across treatments (no context, house frame, and house with info). Results are graphically presented in Figure 3.2. Analysis of variance results showed that there is a significant main effect between the total chosen price in the session halves (first or second), and the frames (no context, house frame or house with info), but the interaction effect was not significant ($p = .21$). Results are presented in Table 3.3.

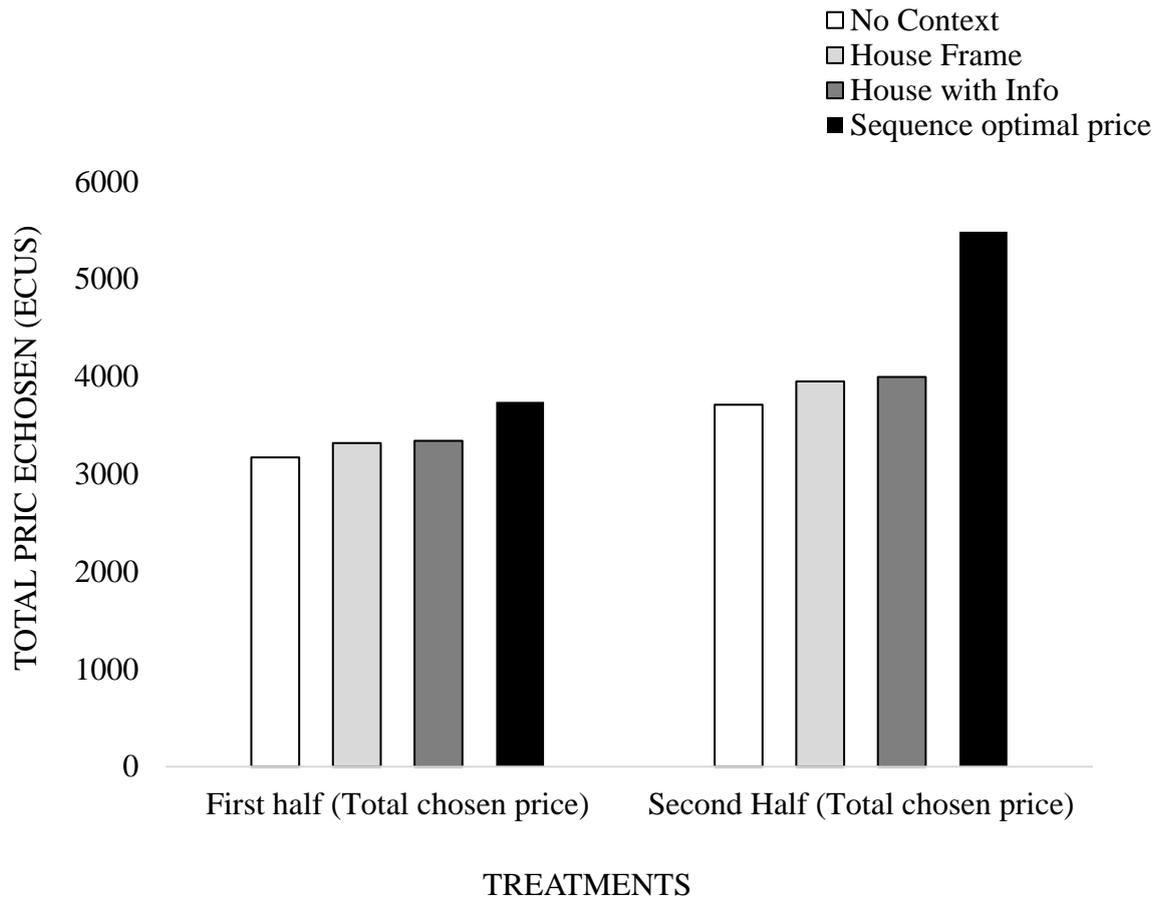


Figure 3.2. The average total chosen price averaged across participant in the first half (rounds 1–5) and second half (6–10) sessions in the no context, house frame and house with info treatments.

3.3.3 PEARSON CORRELATION ANALYSIS

Pearson correlation analysis was used to examine the relationship between the average stopping positions, the total chosen prices and the optimal prices count. The average stopping positions were obtained by averaging 10 stopping position across each participant within treatments. The total chosen prices were obtained from each participant by adding up the 10 prices they selected. The optimal prices count was obtained by adding the number of rounds when each participant had chosen the optimal price out of 10 rounds.

In general, there is a positive and significant correlation between the number of rounds in which the optimal price is selected and the total chosen price for all three treatments ($r = .57, p < .001$). This is expected because the optimal price is the highest price in each round, and selecting more rounds of optimal price should result in choosing higher total prices. There is a large positive and significant correlation between the length of the search and the chosen price ($r = .61, p < .001$), as well as the optimal price count ($r = .24, p = .005$). The longer the participants searched, the more optimal prices they selected. The longer the participants searched in the experiment, the higher their chosen prices. They make more optimal decisions with more searches. However, individual correlation analysis for each treatment shows this is not always true in the treatments with context frame (house frame, $r = .12, p = .44$, house with info, $r = .56, p < .001$ respectively). Only the house with info treatment shows a significant positive correlation.

3.4 DISCUSSION

This study investigated the effect of context on sequential search decisions. The implemented sequential search task was framed as selling houses, which is a significant life decision. The experiment confirms hypothesis 1, as the participants indeed chose higher prices and were closer to the optimal amount of search when the house-selling context was introduced as an experimental manipulation. This result was consistent with the conjectured explanation that having a context can potentially activate existing schemas that enhance the decision-making ability, which is commonly observed in previous reasoning task experiments.

One might think that, without a context, people would conduct a longer search, which allows them to collect sufficient information prior to making decisions. However, this was not found in this experiment. There are several possible underlying reasons. First,

perhaps, people are more emotionally attached to houses, therefore more willing to conduct more searches. Second, people may expect buying and selling houses to take longer. Third, people may behave more randomly (after applying inappropriate schemas) without a context, and hence do not explore different possibilities properly.

Furthermore, this experiment also confirmed hypothesis 2 that the descriptive information did not add value to the effect of context. Perhaps people already stored this information and knowledge in their mental structure or schema, and activating the schema allowed access to this information. In any case, simply stating the context of the decision seemed sufficient to facilitate performance.

There was a long-lasting effect of making inferior decisions when there was no context. People without a context moved still further away from the optimal amount of search in the second half of the session and consistently chose lower prices. There are at least two potential underlying causes for making inferior decisions when there is no context. First, without a context, system one is unable to effectively associate a new experience with existing knowledge and strategies that we have obtained from past experiences in a similar situation. This means participants would need to come up with the strategy in the experiment through trial and error; hence the larger variation found in their search behaviours (see the standard deviations in Table 3.2). Second, people do not know the quality of their decisions until they have reviewed all the offers, and yet they cannot go back to the previously rejected offers. This makes it difficult to identify an effective strategy during the experiment, particularly when there is no context.

4. WHY DO PEOPLE SEARCH TOO LITTLE? NOT ENOUGH TIME, OR TRYING TO AVOID REGRET? EXPERIMENT THREE.

4.1 INTRODUCTION

This study investigated why people make suboptimal decisions of early stopping in the sequential search problem. I also investigated whether individuals choose according to the “classical” optimal stopping rule in this sequential search problem when the decision-making situations are framed as selling houses or hiring secretaries, and how these decisions are affected by adding a time delay. I also investigated whether regret or overconfidence causes the early stopping behaviour.

Decisions regarding buying or selling a house are something that people have experienced in their lives or may do so. People do not have the opportunity to learn the best way of making these decisions. The consequence of mistakes in making decisions regarding houses can be long-lasting and severe. Investigating how people make decisions in selling houses and examining the underlying rationales for making suboptimal decisions can help improve the quality of decisions and of wellbeing. As mentioned previously, time search cost is a critical part of any sequential search problem, as it is unavoidable in any type of search activity. However, the underlying causes for early stopping might not always be the time search cost involved; there can be other, more psychological costs, such as regret or overconfidence.

4.1.1. FEELING OF REGRET

It was stated in the general introduction that employing the classical stopping rule, or indeed any stopping rule that separates an information set of candidates from a set of

available offers, can fail in two different ways. For example, in the secretary problem, failure may occur because one misses out on a good candidate who is encountered in the information set or because one chooses a suboptimal candidate early during the rejection phase. The cost of the two errors can be objectively determined, but there is also a subjective factor which suggests that the first type of failure might incur a higher psychological cost: the feeling of regret.

The feeling of regret as a psychological factor in decision making has been studied extensively (Connolly & Zeelenberg, 2002). Regret is a cognitive and emotional state of feeling sorrow and loss from what could have been achieved (Zeelenberg, Beattie, van der Pligt, & de Vries, 1996). Ritov and Baron (1995) investigated experimentally the feeling of regret after participants made decisions in different frames. Their results suggest that the participants experienced a higher level of regret over an action that resulted in negative outcomes compared with inaction with the same negative outcomes, and that was because taking an action can be perceived as a cause of such negative outcomes. This leads to the next question: whether the feeling of control influences the feeling of regret. The answer seems to be no. Researchers found that whether people have control over an outcome makes little difference to the feeling of regret they experienced (Connolly, Ordóñez, & Coughlan, 1997; Zeelenberg, van Dijk & Manstead, 1998). While the feeling of regret is viewed as a consequence of making a decision, anticipation of regret is found to also affect decisions (Simonson, 1992). According to regret theory (e.g., Bell, 1982; Loomes & Sugden, 1982), when decision makers are given the outcomes of the alternative offer that they have passed over, they are more likely to anticipate the feeling of regret, and this anticipation of regret influences their decisions prior to making them. People alter their future decisions based on a feeling of regret, but current decisions are based on whether they anticipate regret.

People are found to anticipate regret when they expect they will be able to compare the outcomes; for example, a good versus bad outcome, from passing over an option (Zeelenberg & Pieters, 2004). Zeelenberg and Pieters investigated regret in two kinds of lotteries: The National State Lottery and the Postcode Lottery in Netherlands. The National State Lottery is a traditional lottery where people purchase a lottery ticket that is assigned a random number. Alternatively, one's postcode becomes the ticket number in a Postcode lottery (where the country has very precise postcodes). Thus, people who did not purchase a ticket in the Postcode lottery may still find out that they would have won if they had played. The results show that people anticipated a higher level of regret in the Postcode Lottery where they can find out about the positive outcome when not playing, compared with the National State Lottery. As this research shows, the level of anticipated post-decisional regret that people can experience is influenced by the ability to compare the outcomes, and it further influences decisions on purchasing lottery tickets.

Ultimately, people want to minimize regret when making decisions, including both regret and anticipated regret. If one turns down a candidate and nothing better is offered later in the secretary problem, the manager will be aware that she has passed over a better candidate. If she accepts a candidate early on, then the manager will not subsequently find out about the later, better candidate that is yet to come. Thus, these two possible sub-optimal decisions may not have the same psychological effect. The first will produce more anticipated regret than the second. After rejecting a few candidates, a decision to pass over more candidates is much more likely to produce anticipated regret than a decision to accept (cf. Loomes & Sugden, 1982; Zeelenberg & Beattie, 1997).

People can anticipate regretting their decision and adjust their behaviour accordingly by choosing earlier and reducing the size of the information set. In addition to this general

prediction, consideration of regret leads us to make a more specific one. Under the classical stopping rule, if the best candidate is in the information set, the manager is forced to choose the final candidate. On average this final candidate will indeed be average – less good than half of the candidates already seen – and this could be expected to produce a high level of regret. I therefore predicted that there should be great reluctance to choose this candidate and that the probability of this happening experimentally should be considerably less than predicted by the classical stopping rule: approx. 0.37% if $n \rightarrow \infty$.

4.1.2. FEELING OF OVERCONFIDENCE

Another possible psychological cause for choosing sub-optimally is overconfidence. The overconfidence bias is a tendency in which individuals' subjective confidence in their judgments or abilities is systematically and persistently higher than their objective accuracy (Pallier et al. 2002). Keren (1997) attributes overconfidence to two possible causes: cognitive and motivational. Kahneman and Tversky (1996) claimed that overconfidence bias is a cognitive bias, because people often rely on judgmental heuristics to make decisions. The motivational source indicates that overconfidence serves as a self-motivating mechanism in that it fulfils one's needs to believe in one's efficacy to make progress.

Studies have long shown that people are frequently overconfident about their physical talent (Moore & Small, 2007), occupational expertise (Haun, Zeringue, Leach, & Foley, 2000), and social skill (Swann & Gill, 1997). Consequently, overconfidence can have detrimental effects on performance and decision making. Excess entry/trading in the financial market (Trinugroho & Sembel, 2011) and company failure (Camerer & Lovo, 1999) are some of the major issues that have been reported as being associated with overconfident decision making.

In regard to understanding the effects of overconfidence, it is vital to identify what constitutes overconfidence. A study by Hilton et al. (2011) suggested that overconfidence is not a single construct, but rather built up of multiple constructs. There are two major aspects underlying the concept of overconfidence: miscalibration and positive illusion.

Miscalibration is a type of cognitive bias that explains the tendency to overestimate the accuracy of people's knowledge (Michailova, 2011). Typically, in studies on calibration, participants are asked to answer a series of general knowledge questions by choosing from two or more possible answers and to state their confidence of being correct for each answer between 50% (for two options) and 100%. Calibration is calculated by taking the percentage of questions a participant has answered correctly minus the participant's average confidence in the answers to these questions. Individuals are believed to be accurately calibrated if, over time, the percentage confidence responses by the participant match with the actual accuracy (Adams, 1957). Still, people consistently demonstrate overconfidence with their miscalibration; the proportion of accurate answers is significantly lower than the proportion estimated by the participants. For example, Oskamp (1962) studied the diagnosis of clinical psychologists and found a major discrepancy between the level of confidence (53%) expressed by clinical psychologists and their diagnosis accuracy rate (28%). Furthermore, their confidence in their decisions grows with experience, and results in further deviation, while their accuracy remains the same. Miscalibration was measured by Deaves, Lüders and Luo, (2009) through a general knowledge task. For each general knowledge question, the participants estimated the upper and lower bounds of the 90% confidence interval in which the actual numerical answer would fall. Overconfidence was calculated by taking the percentage of times that the confidence intervals contained correct answers. A well-calibrated individual will have 90% of the correct answers falling inside this interval; the lower the percentage, the higher the overconfidence.

The second predominant dimension of overconfidence was studied in the context of positive illusion, which includes the better-than-average effect, illusion of control and unrealistic optimism (Skala, 2008). Hilton et al. (2011) found that positive illusion is correlated with unrealistic optimism, the illusion of control and the better-than-average effect, and also predicted miscalibration in probability evaluation tasks. The better-than-average effect refers to the belief that one is better than the average population, and that one overestimates the chance of experiencing positive outcomes (Alicke & Govorlin, 2005). The most commonly known study that demonstrated the better-than-average effect is the self-report driving skill study. Svenson (1981) found that 82 percent of undergraduate students ranked themselves as safe drivers and belonging to the top 30 percentile of the safest drivers. Similar findings were found by Sümer, Özkan and Lajunen, (2006) when the general public was tested. The illusion of control is the inclination of people to exaggerate their control over the outcomes of the event, even when one has no influence over them (Thompson, 1999). The illusion of control relates to unrealistic optimism and contributes to the psychological bias that is found in mentally healthy individuals, known as “positive illusion” (Johnson et al., 2006). For example, some might believe they have a lower probability of having cancer than other people, even if they have no control over whether they will have cancer or not. Studies showed that the probabilities of desirable events are overestimated by optimists, even in the cases when they have no control (Griffin & Brenner, 2004). Kahneman and Riepe (1998) established that undergraduates perceived themselves to be less likely to be diagnosed with cancer or cardiovascular disease before the age of 50 than their roommate. The positive illusion is then further strengthened in stressful and competitive situations, for example, financial trading. Fenton-O’Creevy et al. (2003) examined illusion of control in traders using computer-based trials. The participants were given an index for their trading performance and asked to rate their success, without

knowing that their index was randomly assigned. Their actual performance is inversely related to their propensity for the illusion of control. The traders with a higher level of the illusion of control performed worse than the traders with a lower level of the illusion of control.

The effect of psychological costs in a sequential search problem has not yet been examined. This study was structured around two of the common framings of the generalized secretary problem that people encounter in everyday life – housing and employment decision making – to examine whether regret and self-confidence influence decision making. The selling-houses experiment provided a frame for participants to make a decision as a seller. They were given a task to sell houses one at a time. They faced a series of random price offers on the house, and they were asked to either accept or reject the offer. The aim was to choose the highest price offer for the house. The hiring frame asked participants to behave as a staff in a recruitment company. Their task was to choose a secretary for several companies. The procedure was similar to the selling house experiment.

Experiment Three compared the decisions made with the predictions of the classical stopping rule, both in general and for particular sequences of price offers and candidates. I examined the effect of time search costs by varying the delay between the different offers (henceforth, the offer is referring to both the price offers for the selling-house frame and candidates for the hiring-secretary frame). Overall, I hypothesised (1) in line with the previous research, people would choose earlier offers than predicted by the classical stopping rule; (2) in line with the consequences of regret, contrary to the predictions of the classical stopping rule, there would be very few forced decisions of the last available offer; (3) in line with the finding from Experiment One, that the effect of adding an extra delay would make participants tend to accept earlier offers as simulation results predict; (4) that,

overall, the participants would perform sub-optimally on a range of different measures,(5) that the self-confidence of the participants would be related to their search behaviours. People with overconfidence should believe they have adequate knowledge in making the decision; that is, they should overestimate their ability to make the decision, and thus search less. The optimal stopping rule in the secretary problem contains an information-gathering stage, which allows decision makers to learn about the distribution prior to making the decision. A short information set produces less searching; i.e., they observe fewer overall candidates.

4.2 METHOD

4.2.1 PARTICIPANTS

A total of 71 people from the University of Canterbury, New Zealand, participated individually in the experiment, 33 people in the low-delay treatment and 38 in the high-delay treatment. They were recruited by advertisements both on a Facebook webpage and posters on campus. The participants either receive a \$10 supermarket voucher or two first-year psychology course credits as an incentive. They were within an age range from 18 to 50 years old, and the median age in the range was 21 years. The experiments lasted on average one hour.

4.2.2 EXPERIMENTAL DESIGN AND PROCEDURES

PART 1

The experiment consisted of two parts, the first of which was the computerized task, the second the self-confidence evaluation task. For the computerized task, the participants were given instructions about the first of the two frames, house selling, used in the experiment. For the house-selling experiment, the initial instructions were:

“You are a real estate agent in Christchurch city, and your job is to sell houses. You will be given a brief description of the house and a series of price offers. The price offers will be presented in a random sequential order. Once a price offer is presented, you can either reject or accept an offer. The house will be sold at that price if you accept. Once the decision to reject or accept the offer is made, it cannot be re-called. There are a total 20 offers available; if you have not accepted an offer prior to the 20th offer presented, you will be forced to accept the final offer.”

There were two practice rounds and then ten experimental rounds. Each experimental round (of a possible 20 offers) featured a different house description, and a different randomized sequence of prices; moreover, the mean and standard deviation of the house price offers differed between the rounds. The house descriptions showed the suburb where the house was located, its floor area, the number of bedrooms and year the house was built. The house description used in each round was taken from houses sold in October 2015 in 12 different Christchurch suburbs. The information was obtained from the Quotable Value Limited database (qv.co.nz). The random sequence used for the housing frame was generated using the median house price in each suburb, with an upper and lower limit of 2 standard deviations and using a normal distribution. Each participant was presented with the same randomized sequence. (See Table 4.2 and Table 4.3 for details of the sequence).

For the secretary-hiring frame, the instructions read:

“You are on the staff of a recruitment company. Your job is to hire executive assistants. You will be given a description of the company, and a series of candidates with an overall score will be provided to help with the decision making. The candidates will be presented in a random sequential order. Once a price offer is presented, you can either reject or accept the candidate. The position will be filled by the candidate if you accept. Once the

decision is made it cannot be recalled. There are a total of 20 candidates available. If you have not accepted a candidate before the 20th, you will be forced to accept the final candidate.”

Again, there were two practice and ten experimental rounds, and the means of the overall score varied between rounds and with slightly different positions. The job description shows the annual salary the vacancy is offering – for example \$48,000, \$28,000; the location of the job – for example regional clinics, an advertising agency; the position vacancy – for example, medical secretary, personal assistant. The job description was obtained from seek.co.nz in Oct 2014. The descriptions of the candidates emphasized different qualities. The randomized sequence of scores was different for the different rounds, but the same for all participants. The competency scores within each round were each based on the mean and they varied according to a normal distribution limited to two standard deviations on either side. (See Table 4.1 and Table 4.3 for more detail.)

All participants performed in both the house-selling and secretary-hiring rounds. Half of the participant performed the house-selling rounds first, followed by the secretary-hiring rounds. The other half of the participants performed the secretary-hiring rounds first, followed by the house-selling rounds. The participants were randomly assigned to either condition. Different randomized sequences were used for the house-selling and secretary-hiring frames. No feedback was given at any stage in the experiment.

Participants performed their entire experiment under one of two delay treatments. In the low-delay treatment, there was a constant 1-second delay after a participant rejected an offer and before the next one was presented. In the high-delay treatment, this delay was 1 second, 2 seconds or 5 seconds, with an equal chance of each delay. The randomized

sequences used were identical and presented in the same order in the two delay treatments. Thirty-three participants were in the low-delay and 38 in the high-delay treatment.

PART 2

The second part of the experiment consisted of four sections. The first was a general knowledge quiz, which consisted of 21 general knowledge questions regarding New Zealand (see Appendix 4.A). The questions were not related to psychology and economics to avoid biases on an individual's own expertise so as to avoid a relative advantage from participants' different expertise or knowledge (Deaves et al., 2009; Michailova, 2011). I counterbalanced for the hard–easy effect, and the questions consisted of an equal number of three levels (difficult, medium and easy). The hard–easy effect refers to the tendency to increase overconfidence as the difficulty of a task increases (Lichtenstein, & Fischhoff, 1977). People tend to overestimate their ability in solving a difficult task, and yet they tend to under-estimate their ability in solving an easier task. If the questionnaire consists of difficult questions only, people are more likely to be overconfident. If the questionnaire consists of easy questions only, more people will be showing under-confidence. The questions were also controlled for gender bias, and excluded gender-specific questions, for examples, sports, and fashion-related questions. Each question had three possible answers; the participants could choose only one. Participants were asked to state their level of confidence in the accuracy of their answer to each corresponding question between 33% (totally guessing) to 100% (absolute certainty). Work on probability evaluation procedure demonstrates reliable test–retest consistency and internal consistency when measuring miscalibration (Jonsson & Allwood, 2003). Also, Bruin, Parker and Fischhoff, (2007) found, when using probability evaluation technique to assess calibration, that the outcome correlates with various decision-making competency measures, and the ability to resist

reversal of preference under framing. If the average self-reported accuracy is higher than the average accuracy of their answer in the questions, the participant is considered to be overconfident.

Section 2 (Appendix 4.B) and 3 (Appendix 4.C) measured the positive illusion construct in overconfidence. Section 2 also aimed to measure self-efficacy of the participants. The General Self-Efficacy scale (Jerusalem & Schwarzer, 1992) is a 10-item scale developed to assess optimistic self-beliefs to cope with a variety of difficulties in life. It captures one's belief that one is able to complete a novel or a difficult task. Evidence suggests that the General Self-Efficacy scale has high reliability, stability, and construct validity (Schwarzer, Mueller & Greenglass, 1999). The scale has been found to be reliable across different ethnic groups and countries, and it measures only one global dimension (Leganger, Kraft, & Røysamb, 2000; Scholz, Doña, Sud, & Schwarzer, 2002). Cronbach alpha ranges from 0.75 to 0.94 across different language versions (Rimm & Jerusalem 1999; Luszczynska et al. 2005). The General Self-Efficacy Scale is positively associated with other social cognitive constructs, for example, intention and outcome expectations, which further demonstrated high validity on the scale (Luszczynska et al. 2005). Section 3 focuses on participants' depression tendencies. Empirical study shows that mentally healthy individuals tend to exhibit psychological biases that encourage optimism, collectively known as “positive illusions” (Peterson, 2000). People who score higher on the depression scale are not likely to be overconfident/overoptimistic about their ability and expertise. The Quick Inventory of Depressive Symptomatology (Rush et al., 2003) is a 16-item self-rating scale and provides each item with four possible answers describing symptoms of increasing severity associated with a score from 0 to 3. The Quick Inventory of Depressive Symptomatology scale showed high reliability with the Cronbach's alpha of 0.86 (Trivedi et al., 2004) and 0.72 correlations with the Hamilton Depression Rating Scale (Rush et al.,

2005). The Hamilton Depression Rating Scale (Hamilton, 1960) is perceived to be the most widely used scale for controlled clinical trials in depression (Cusin, Yang, Yeung, & Fava, 2009).

The fourth section (Appendix 4.D) elicited demographic information and previous histories regarding the experience of the participants in selling or purchasing houses and hiring employees.

4.3 RESULTS

In this section, I first consider the average results obtained from the experiment. Then I examine the correlation between self-confidence and the search behaviours. I also ask to what extent the participants followed the classical stopping rule suggested by the analysis of the secretary problem outlined above. I then consider whether deviation from the classical stopping rule produced markedly poorer performance in the experiment. Next, I present results related to individual differences. Finally, I present a regression analysis to examine the factors that influence decisions prior to making.

There were two important dependent variables for each frame. For the house selling frame, these were the position in the sequence where the participant accepts the offer and the price obtained by accepting the offer at this position in the sequence. Results relating to these two dependent variables are shown in Table 4.1 and Table 4.2 respectively. Similarly, for the secretary-hiring frame, statistics relating to the position in the sequence at which the secretary is hired are shown in the bottom half of Table 4.1 and the average candidate scores for the hire are shown in Table 4.3.

Table 4.1

Position in the sequence at which the offer was accepted as a function of round number for low and high delay for the house-selling and secretary-hiring frame.

Round	Best*	Low				High			
		Mean	%Best	% 20 th	%3/20	Mean	%Best	% 20 th	%3/20
House-selling:									
1	6	12.2	30	30	48	8.5	29	18	76
2	17	7.4	12	3	39	5.3	5	3	37
3*	14	11.4	18	6	27	9.7	24	3	24
4	14	9.8	42	6	79	4.8	8	3	84
5	9	6.3	24	3	27	6.5	3	0	3
6	12	6.6	6	3	33	5.4	3	0	39
7	4	9.5	21	12	55	3.2	34	3	68
8	13	11.9	36	3	42	10.4	5	3	11
9	6	11.2	42	18	73	7.5	45	5	84
10	13	8.3	55	0	88	7.2	42	0	82

Secretary-hiring:

1	6	7.3	36	9	73	4.7	50	3	87
2	15	11.1	6	12	58	9.6	3	5	66
3	18	5.6	3	0	67	5	0	0	71
4	10	5.3	18	0	39	4.3	8	3	13
5	1	13.7	0	6	52	13.3	0	3	45
6	6	15.1	15	42	24	9.4	34	13	42
7	4	11.3	21	3	70	9.3	34	0	76
8	6	3.8	15	0	61	2.4	5	0	39
9	18	6.5	0	0	55	4.7	0	0	61
10	8	8.2	27	15	64	5.7	39	0	68

Notes. Best = position at which the best offer appears; % Best = percentage of participants choosing the best offer; % 20th = percentage of participants accepting the 20th offer; % 3/20 = percentage of participants choosing one of the best three offers.

*For all but one round in the house-selling frame, the optimal stopping position suggested by theory is either on the best offer or on the twentieth offer (if the best offer is in the first seven), except for Round 3 where it is on offer 9. Fifty-five percent of the respondents in the low-delay treatment and 53 % of those in the high-delay treatment accepted this ninth offer. For the hiring-secretary frame, in Rounds 3 and 9, the candidates predicted by the classical stopping rule are candidate 13 and 16 respectively; 18% of participants chose candidate 13 for Round 3 in both delay treatments, and 12% chose candidate 16 in the low-delay, 13% in the high-delay treatment.

Table 4.2

Offer prices in the house-selling frame for each round. The table shows the actual average offer (averaged over the 20 offers), the actual best offer, the average and standard deviation used to generate the price offer sequence, the offer that would be accepted under the classical theory, and the average offers accepted for the two delay treatments. All results are in NZ\$(000).

Round	Avg. Actual	Best Actual	Avg. Exp	SD Exp	Classical Theory	Low	High
						Avg.	Avg.
1	426	522	383	52	365	458	477
2	632	687	638	41	687	649	639
3	382	437	383	32	415	414	413
4	604	631	601	19	631	623	623
5	362	547	383	97	547	420	341
6	440	538	445	85	538	485	493
7	573	646	567	32	523	602	610
8	500	538	502	23	538	525	517
9	451	564	445	64	525	527	534
10	474	633	502	93	633	618	618
Av.	484	574	484.9	53.8	540	532	527

Notes. Avg. = average, SD = standard deviation

Table 4.3

Candidate scores in the secretary frame for each round. The table shows the actual average candidate score (averaged over the 20 applications), the actual best candidate, the average and standard deviation used to generate the candidate score, the score of the candidate that would be hired under the classical theory, and the average score of the candidate accepted for the two delay treatments.

Round	Avg. Actual	Best Actual	Avg. Exp	SD Exp	Classical Theory	Low	High
						Avg.	Avg,
1	418	489	420	45	404	473	480
2	448	463	450	13	463	460	460
3	513	671	490	93	662	595	589
4	453	565	450	86	565	512	500
5	116	287	310	22	255	229	213
6	421	444	420	18	393	416	429
7	300	331	310	19	295	323	327
8	555	708	560	87	620	600	546
9	515	608	490	64.1	605	552	548
10	367	457	370	43	457	409	425
Av.	410.6	502.3	427	49.01	471.9	456.9	451.7

Notes. Avg. = average, SD = standard deviation

4.3.1 AVERAGE RESULTS

Analysis of variance was used to examine the effect of the frame (housing or secretary), the delay (low or high) and earlier (first five) or later (second five) rounds on the average positions in the sequence where the participant accepts the offer. Participants chose to stop at a significantly later position in the sequence ($M = 9.1$) under the low-delay than the high-delay ($M = 6.8$) treatment ($F(1, 69) = 18.9$, $MS_{error} = 36.5$, $p < .0001$, partial $\eta^2 = 0.215$). There was also a significant interaction ($F(1, 69) = 4.22$, $p = .044$, partial $\eta^2 = .038$) between the frame and the delay, by which the effect of increased delay was a little more pronounced in the secretary than the housing frame. This interaction might be due either to differences between the frames or differences in the sequences used in the frames. There were no other statistically significant ($p < .05$) main or interactive effects.

Separate analyses were used to investigate factors (delay and early or later round) affecting the average house selling price and an average score of the successful secretary candidate. For house selling, the average price obtained from averaging across low- and high-delay treatments rose between the early ($M = \$505,588$) and the later rounds ($M = \$552,723$; $MS_{error} = 78.4 \times 10^9$; $F(1,69) = 257.9$, $p < .0001$). There was also a significant interaction between rounds and the delay ($F(1, 69) = 8.90$, $MS_{error} = 2.71 \times 10^9$, $p < .004$), with the difference between early and later rounds greater for the longer delay. There was no main effect of delay. For secretary hiring, the corresponding analysis produced no significant main or interactive effects, although there was a slight suggestive tendency for candidates to have a higher score in the later ($M = 457$) than the earlier rounds ($M = 451$; $F(1,69) = 3.51$, $MS_{error} = 1482$, $p = .065$).

Perhaps the key result emerging from these analyses is that, consistent with previous research, participants reviewed fewer offers when the time cost of looking at each offer rose. However, the decrease in offers viewed was not reflected in the quality of the offer accepted.

4.3.2 DOES SELF-CONFIDENCE CORRELATE WITH SEARCH BEHAVIOURS?

The overconfidence rate of each participant was calculated by taking the average percentage of correctly answered questions minus the percentage of self-estimate in correctly answered questions. The percentage of correctly answered questions was obtained by taking the number of correctly answered question and dividing by the number of questions they needed to answer; that is, 21 questions for each participant. The self-confidence in correctly answering the questions is obtained by averaging across 21 questions for each participant. For example, if a participant was correct 71% of the time, that is, answered 15 out of 21 question correctly, and estimated himself or herself as being correct 77% of the time, then the participant is overestimating with 6%. If a participant answered the questions 71% of the time, yet they estimated themselves as correct 60% of the time, they were under-estimating their ability and their scored was expressed as negative 11%. Therefore, a positive score in percentage means the person is overconfidence and negative score is under-confidence. Participants were found, on average, to be more overconfident under the low ($M = 7.96\%$, $t(63.34) = 12.67$, $p < .001$, Cohen's $d = 2.96$) than the high ($M = -40.12\%$) delay treatment. Summary statistics are reported in Table 4.4. In fact, 95% of the participants in the high-delay treatment reported being under-confident about correctly answering the general knowledge quiz. Yet, only 24% of people reported being under-confident in the low-delay treatment. There is no significant difference found between the low ($M = 31.8$) and high ($M = 30.8$) delay treatments in the self-efficacy scores according to t-test ($p = .24$). T-test also shows there is also no significant difference found in the depression scores between the low ($M = .4$) and high ($M = .3$, $p = .49$) delay treatments.

The correlation analysis shows that the overconfidence rate of both low-delay ($r = -.189, p = .292$) and high-delay treatment ($r = -.119, p = .475$) did not significantly correlate with their average stopping positions. In the low-delay condition, both the self-efficacy score ($r = -.040, p = .827$) and the depression score ($r = .131, p = .468$) were uncorrelated with the average stopping positions. However, in the high-delay treatment, the self-efficacy score was found to positively and significantly correlated with their average stopping positions ($r = .355, p = .029$) but not the depression scores ($r = -.167, p = .317$).

Table 4.4

Summary statistics of overconfidence, self-efficacy and depression scores.

	Overconfidence		Self-Efficacy		Depression	
	Average (SD)	Range (Min–Max%)	Average (SD)	Range (Min – Max)	Average (SD)	Range (Min–Max)
Low delay	8.0% (12.3%)	(–15.7 to 30.8)	31.8 (3.6)	(23–39)	0.4 (0.6)	(0–2)
High delay	–40.1% (19.4%)	(–64.0 to 5.4)	30.8 (3.6)	(23–39)	0.3 (0.6)	(0–2)

Note: SD refers to the standard deviation. The total score of self-efficacy scale is range from 10 to 40. The higher the score, the higher the self-confidence. The total score of depression scale is in the range from 0 to 27. The higher the score on the depression scale, the higher the risk of being depressed.

4.3.3 DO PARTICIPANTS FOLLOW THE “CLASSICAL STOPPING RULE”?

If the classical stopping rule for finding and selecting the best offer were employed then participants should have invariably used the first seven offers to establish a reserve price or candidate score, and then accepted the next offer that was higher than the highest of the initial seven. If, as happened in three of the house-selling and five of the secretary-hiring rounds, the best offer or score was actually one of the first seven, they should then be forced to accept the

twentieth offer. As the sequences turned out, for all the other rounds but three (Round 3 of house selling and Round 3 and 9 of secretary hiring), the best offer was either in the first seven or the best offer was the first offer higher than any offer in the first seven. Thus, if the classical stopping rule was followed by participants, they should have accepted the twentieth offer on eight of the total of 20 rounds and the best offer on nine of the remaining 12. But in fact, examination of Table 4.1 shows quite clearly that very few participants did this. Indeed, the only round for which the majority of participants pick the offer predicted by the classical stopping rule is Round 3 of the housing frame.

Given that the participants were not following the classical stopping rule, one might ask whether they followed some variant of it. For example, they might follow a strategy of attempting to find one of the best three offers. However, as Table 4.1 shows, relaxing the criterion to one of the best three shows that the participants fulfil this criterion only about half the time (even though on only one round – Round 1 of the secretary frame – were all of the best three offers inside the first seven).

4.3.4 WHAT WERE THE PARTICIPANTS DOING?

Table 4.5 shows the percentages of accepted offers, taken over all participants and both frames and delay treatments, after increasing or decreasing runs in price or score of 1, 2, 3 or 4. The remainder accepted the very first offer. Most acceptances followed increasing runs, but many participants accepted offers that were less than the previous offer. The large majority of offers were accepted in an increasing run, suggesting a tendency to loss aversion.

Table 4.5

Percentages of accepted offers following increasing or decreasing sequences of 0, 1, 2 or 3.

Accepted first offer	11.2%
After increasing run of 1	46.3%
After increasing run of 2	30.1%
After increasing run of 3	5.4%
After increasing run of 4	0.2%
After decreasing run of 1	3.6%
After decreasing run of 2	1.4%
After decreasing run of 3	0%
After decreasing run of 4	0.4%

Brief statements about the strategy they had used in the house-selling frame were collected from the participants. (Most said they used similar strategies in the secretary-hiring frame.) Thirty-six of the 71 participants claimed to have used a strategy in which they refused to accept the initial offers and only considered offers after that. Fifteen of these participants described the size of the initially rejected set as a “few” or “some” offers; the rest gave numbers. These ranged from 1 (6 participants) to 10 (2 participants), with the others in the range 3–5 (13 participants). Thus, the basic strategy of separating an information-gathering set

from a candidate population was frequently followed although the size of the information-gathering set was generally smaller than seven. Six participants stated that as the number of offers increased they lowered their threshold for acceptance; for two this was lowered after 10 offers, for one after 15. Some participants actually made use of their knowledge of the housing market in deciding on initially acceptable offers, but no evidence suggested that this resulted in finding more rounds of optimal offer.

If the participants used different strategies, then different participants should, on average, differ in the number of offers they reviewed before choosing. Figure 4.1 shows that this is what happened. There are consistent differences between participants in the same treatment and these differences are substantial. Note incidentally, that if all participants were to follow the same strategy, there would be no variation at all. For example, if participants were all using the classical stopping rule, they would accept exactly the same offers and candidates. The most frequent average stopping positions were 9 and 10 in the low-delay, and 7 and 8 in the high-delay conditions.

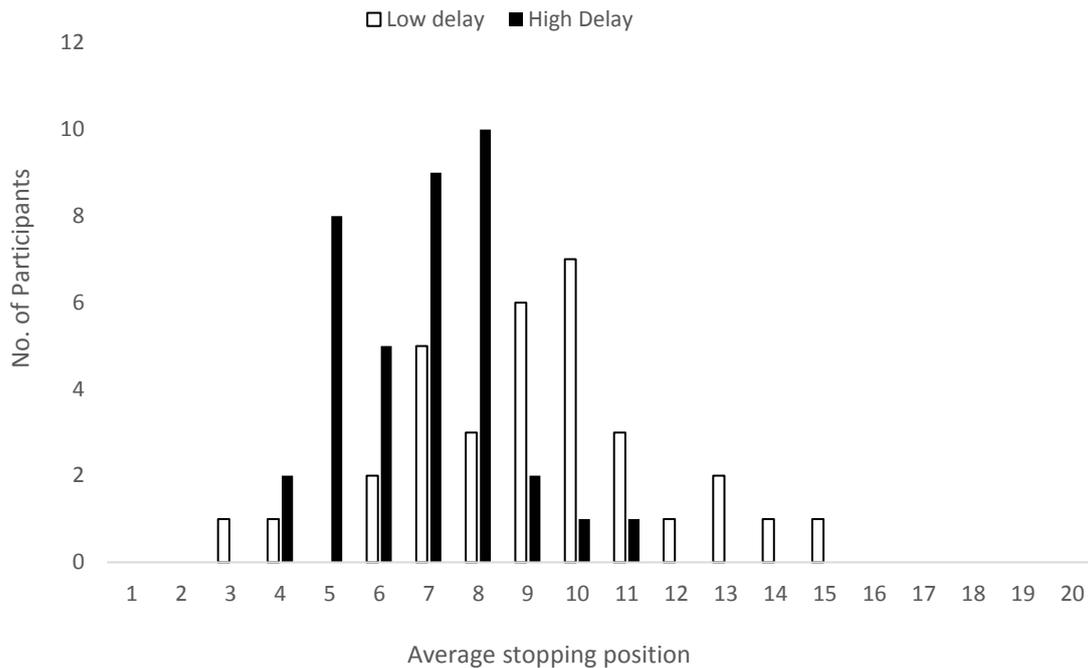


Figure 4.1. Frequency distribution of the average stopping positions (averaged across 10 selling houses and 10 hiring secretary rounds) of the participants in the low- and high-delay treatments.

4.3.5 WHAT WAS THE PENALTY FOR NOT FOLLOWING THE CLASSICAL STOPPING RULE?

On average, the participants accepted offers earlier in the sequence than the classical stopping rule indicate they should, particularly in the high-delay treatment. For both the house selling and secretary-hiring frames, participants who viewed more offers ended up with, on average, higher house prices, or secretaries with higher scores. For the low-delay treatment, the Pearson correlation between the mean (of 10) stopping positions (for the housing-frame rounds) and mean house price accepted was 0.39, and that between the average stopping position and mean secretary score was 0.63. The respective

correlations for the high-delay treatment were 0.54 and 0.66. (All four correlations significantly differ from zero, $p < .05$.)

Column 3 in Table 4.3 and Table 4.4 shows the expected outcome, in terms of house selling price or secretary candidate score respectively, of following the classical stopping rule in the different rounds of the experiment, that is, to reject the first 7 offers and accept the next highest offer (37% of 20 offers is 7.4 offers in the information set). These results illustrate the point made in the introduction: Following the classical stopping rule can leave you stuck with a regrettable final offer. That this consideration was actually taken into account by the participants is evident from the strategies used by a few of them and, more importantly, by the low frequency of acceptance of the twentieth offer.

Columns 4 and 5 in Table 4.3 and 4.4 show the average results actually obtained by the participants. The results illustrate that while, on average, these results are slightly worse than those that would have been obtained by following the classical stopping rule, there are rounds on which the average participant does better than this strategy. Note, too, that the average results are almost always considerably better than those obtained by choosing randomly, and that, as pointed out above, although participants in the higher delay treatment chose earlier in the sequence, they paid no significant penalty for doing so.

Thus, in general there was a penalty for not following the classical stopping rule in this experiment, but it was not high.

4.4 DISCUSSION

The participants generally stopped their search earlier than would be predicted by the classical stopping rule and this search was shortened by adding an extra variable time search cost between offers. The data obtained for the specific stopping positions on

different rounds indicate that only a minority of participants accepted the particular offer that was predicted by the classical stopping rule. The departure from the behaviour predicted by the classical stopping rule was most marked for those rounds in which the best offer occurred early on, within the information set predicted by the classical stopping rule. The classical stopping rule should force the offer of the last option. In practice, in all but the very first round, selecting the twentieth offer was rare. In the first round, people were probably learning and attempting to develop the best strategy, and thus here it may have made more sense for them to review all 20 offers.

A large amount of evidence indicates that participants shortened the initial information set compared with the classical stopping rule. However, over and above that, some participants stated that they revised their threshold for acceptance downwards as they neared the end of the offer sequence, and examination of the data indicates a number of instances where many participants did not simply use a short initial information set and then refuse to accept any offer that did not exceed one in that shorter set (e.g., round 7 of house selling; rounds 5 and 8 of secretary hiring).

A number of criteria could be used to evaluate the performance of the participants. One such is whether they chose the best offer, and on average they did not, and they chose the best offer less often than they would have done by using the classical stopping rule. Another criterion is whether the average price chosen is close to the offers they might have obtained by following the classical stopping rule. On average, the offers accepted were earlier than those that might have been obtained under the classical stopping rule, but the payoffs in terms of house cost or secretary excellence were not very much lower. Thus, as has been found previously in other research (e.g., Gigerenzer & Goldstein, 1996), the strategies used by the participants were reasonably well adapted to the problem and

produced quite reasonable outcomes in a rather shorter time than the classical stopping rule. Moreover, unless the participants had either been versed in the mathematics of the secretary problem or had been forced to observe a large number of sequences, it is not clear how they might have arrived at a better solution.

Finally, the experiment produced results indicating that participants tried to reduce their regret. Participants avoided being left with a potentially poor forced offer on the last price and acted so as to minimize the possibility for regret by missing out on reasonable early prices. Evidence for the influence of regret on the decisions made is also provided by the demonstration that overall price is earlier, both in this and previous research, than is predicted by the classical stopping rule.

5. DECISION MAKING IN SMALL GROUPS.

EXPERIMENT FOUR.

5.1 INTRODUCTION

This experiment investigated whether groups make a better decision than individuals in the sequential search problem. Decisions regarding selling houses and hiring employees are often made in a group. For houses, it is most likely a joint decision of the couple who own the house or the family. As for hiring employees, a committee is often involved in deciding which candidate to hire. It is therefore critical to investigate whether groups are more effective than individuals in making these kinds of decisions. Group decision making is more costly and time-consuming than individual decision making. It costs more to hire expert boards or think tanks to make a decision than to have one person do it. Group decision making also requires spending time in discussion and exchanging information prior to making decisions. Is this extra cost worthwhile?

Most critical decisions are now made by groups, on the assumption that groups can make better decisions than individuals. In the decision-making process of a sequential search problem, group members discuss their preferences and exchange opinions about the problem at each stage and discuss whether to accept or reject the current candidate or offer, until they reach a consensus decision among the group members. This decision-making process of exchanging opinions and information is perceived to overcome some of the shortcomings in individuals' decision making; for example, insufficient knowledge or information to make a good decision. The issue of whether groups make a better decision than individuals in these commonly found sequential decision-making scenarios (selling houses or hiring employees) is as yet unexplored. In this experiment, I investigate the difference between group and individual search behaviour experimentally, by employing a

sequential search task that presents offers one by one in a random order, with an immediate decision and no re-call.

5.1.1 THEORIES OF GROUP DECISION MAKING

Levin (2009) defines group dynamics as two or more people within a social group or between groups, for longer than a few moments, behaviourally and psychologically interact with and influence one another and perceive one another. This definition via group dynamics argues that all groups have one thing in common, which is that their members interact. (Consider, for example, a group of people jogging together every morning.) However, a few students working individually in a computer room is not a group, because they are only a collection of unrelated individuals in a common place rather than an interacting group.

There are some known advantages of group decision making over individual decision making. For example, group decisions increase the perception of fairness and the acceptance of the decision made. Also, higher identification with the group decision is found, resulting in a higher commitment to the decision made by the group. (see Moscovici & Doise, 1994; Vroom & Jago, 1988 for reviews). Group decisions are able to aggregate information and preferences compared with individual decisions. Group decision making, in theory, can have a broader scope and synergy for solving a problem. Every individual group member may contribute unique information and expertise through discussion. Sharing information facilitates understanding, clarifies issues and avoids mistakes (e.g., Hollenbeck et al., 1995).

There are two main theories in psychology and sociology that attempt to explain the outcomes of group decision making; the wholistic theory and reductionism theory. The

wholistic theory predicts that a group makes a better decision than an individual, due to an interaction between group members, and further, results in the whole being greater than the sum of the individuals. Researchers with a wholistic viewpoint proposed that interaction and discussion with others leads people to perform tasks better than by themselves (see Pavitt, 1998 for a detailed review), make a higher quality decision (Postmes, Spears & Cihangir, 2001) and solve more problems (Stasser & Stewart, 1992). Shaw (1932) found in a series of accuracy tasks of solving a puzzle that four-person groups had a higher proportion of accurate solutions than the individuals. The wholistic theory implies that interacting with people is the underlying rationale for why groups' performances are better than individuals'.

In contrast, reductionism theory suggests that the group is at best only equal to the sum of its parts if group interaction goes smoothly, otherwise, a group's performance will be less than the sum of what each individual member can achieve. The larger the group, the less well it can perform compared with the aggregate performance of individuals with the same number. The Ringelmann effect describes the inclination for individuals in a group to become less and less productive with more people. Ringelmann (1913) found that when group members work jointly on a task, for example, pulling a rope, they exert less effort than when individuals combine separately. The group becomes increasingly inefficient with more group members. This violates the assumption of wholistic theory: that joint effort can lead to a better outcome than that produced by individuals.

The two theories propose contradicting predicted outcomes of group performance. A model has been proposed to evaluate and compare groups' performances and individual aggregate performance on problem solving (Lorge-Solomon, 1955). The simple form of the model is given as follows:

$$P_G = 1 - (1 - P_i)^n$$

where P_G denotes the probability that a group of size n will be able to solve the problem, and P_i denotes the probability that a given individual by themselves can solve the problem. $(1 - P_i)$ is the complement of P_i and denotes the probability that an individual is unable to solve the problem. $(1 - P_i)^n$ denotes the probability that no one in a group of size n can solve the problem, assuming that, if one person in the group can solve the problem, the problem is considered solved for the group. Say you has a 37% chance of solving a problem, the Lorge-Solomon model suggests that there is a 60% chance that you and another person together will solve the same problem. If the group of two has solved the problem more than 60% of the time, then the group performance has supported the wholistic theory that interaction enhances group performance compared with the aggregation of individuals' performances. By contrast, if groups solved the problem less than 60% of the time, reductionism theory is supported. If the groups have solved the problem exactly 60% of the time, group interaction has no effect in enhancing or reducing the performance of group, and therefore, neither wholistic or reductionism theory is supported.

5.1.2 LITERATURE REVIEW

Group interaction has been found to have a polarising effect on a participant's individual choices (e.g., Stoner, 1968; Moscovici & Zavalloni, 1969; Bray & Noble, 1978). The decision shifts towards the extreme after-group discussion. People may choose a riskier choice or a more cautious choice after discussion. The group polarising effect is not only found in risk-related choices. Myers and Bishop (1970) found that a group of prejudiced American high school students polarised their attitude to become more

prejudiced. The unprejudiced students shifted towards less prejudiced attitudes. Value theories are believed to best explain the polarising effect with group discussion (e.g., Pruitt, 1971; Moscovici, & Zavalloni, 1969). These theories describe the assumed decision-making process by which individuals tend to polarise their decisions towards the value held by the group. The predictions of value theories are supported by a positive significant correlation found between the individual participants' original risk attitudes and the extent of a shift towards the risk attitude of the group (Clark & Willems, 1969; Teger & Pruitt, 1967), because the original risk attitude can be perceived as the baseline of the attitude scale. Value theories emphasise the most dominant human value the individual holds. The stronger the attitude, the stronger the shift. Value theories assume that groups shift in the direction that the individual members already favour. It is an intensification of attitude and preference against other alternatives. Apart from polarising attitudes, groups also make decisions differently than individuals in experimental tasks. People working in a group need to coordinate their activities with each group member. The greater the number of group members, the harder it is to coordinate. Even for a simple task of pulling on a rope, it can be difficult to coordinate to achieve the best result. Ingham et al. (1974) explain the Ringelmann effect as due to coordination and motivation. They showed that people were less motivated when working in a group. Participants put in more effort when they believe they work alone, and they refer to this reduced effort as social loafing.

Social loafing occurs when the individuals' efforts cannot be clearly distinguished from each other in a group. If the members of the group are not punished or called out for insufficient effort, the individual members are likely to reduce their effort and the group produces a poor performance. Social loafing was found to be influenced by the size of group. People exert less effort with larger groups (Chidambaram, & Tung, 2005).

However, for tasks like an accuracy task, social loafing may be reduced, where performance depends on the most competent group member. If one member of the team knows the answer, the group has success as a whole. This leads to the conclusion that the group's performance depends upon the groups' most competent member. The group's performance is therefore enhanced, and higher than the average individual's performance.

Previous research is inconclusive about whether groups perform better than individuals, but one common finding is that groups behave differently to individuals and there are diverse outcomes when comparing group and individual behaviours. To the best of my knowledge, group performance in a sequential search task has not been tested experimentally, although groups are often used to make decisions in sequential search tasks. Previous research mostly supports the reductionism theory that the individuals perform more poorly in a group setting than they would if they were working alone. In line with this previous research, I hypothesised that (1) groups do not conduct more searches than individuals, and (2) groups will find less optimal prices than individuals.

5.2 METHOD

Seventy-nine undergraduate students from the University of Canterbury, participated in the decision-making experiment. Participants were from a third-year psychology course on judgment and decision making. The participants were rewarded with course credits as an incentive to participate.

This experiment used a within-subject design, and the participants were asked to participate in both individual and group decision making. Participants were randomly selected to be in one of the three experimental sessions. In the first session, the participants started with the selling-houses task individually then performed the hiring task in a group

of two or three. In the second session, the participants started the hiring task individually and then the selling-houses task in a group. The third session began with the group experiment on the hiring task and was followed by the individual selling-houses task. In the group task, participants needed to discuss the decision with their team member prior to decision making. The grouped participants also needed to rotate controlling the mouse between the group members. There were a total of 45 people in the individual selling-houses task and 25 people in the individual secretary task; 42 people participated in the group hiring task (21 groups) and 25 people (12 groups) in the group selling-houses task. Other details of the experimental procedural were as in Experiment Three.

5.3 RESULTS

5.3.1 STOPPING POSITION AND TREATMENTS

The average stopping position for each participant was computed across all 10 rounds in both tasks (selling houses and hiring secretaries), and compared between the group and individual treatments. Figure 5.1 depicts the frequencies in the percentage of the average stopping position for both treatments. There was no obvious difference in the search pattern between the group and individual treatments. Most of the people in the group (45%) and individual (38%) treatments stopped their search in position 9 and 10. Only 21% of people in the group and 22% of people in the individual treatment searched beyond the 10th offer.

An independent t test showed no significant difference ($t(110) = -.06, p = .96$), in stopping positions between the group ($M = 8.78, SD = 2.53$), and individual treatments ($M = 8.30, SD = 2.37$). A Mann-Whitney U test also confirmed the no significant difference result ($U = 1289.5, p = .93$).

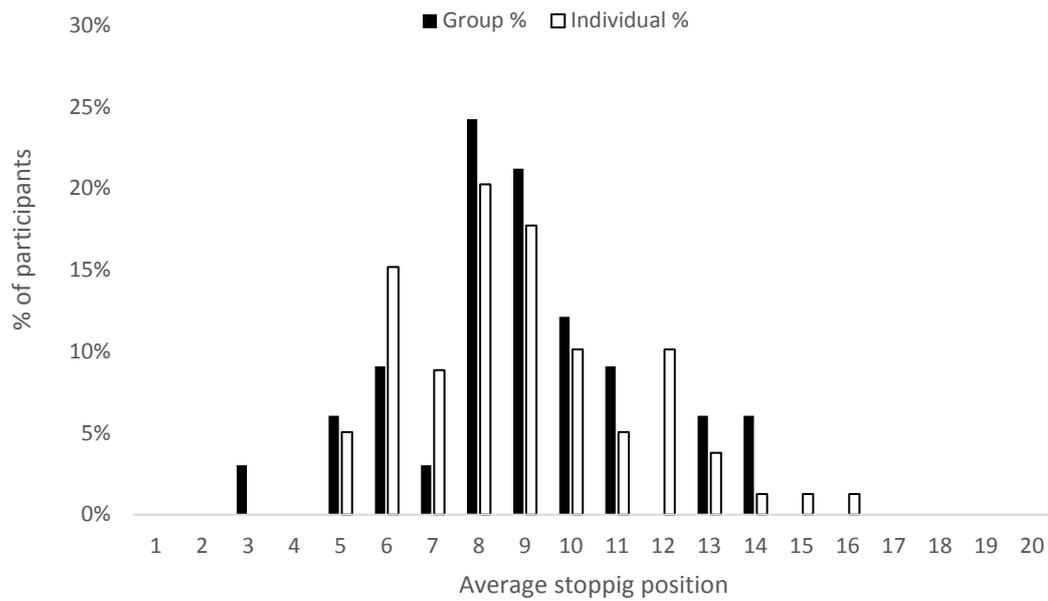


Figure 5.1. The average stopping position in the group and individual treatments.

5.3.2 EVIDENCE OF LEARNING

A repeated-measures ANOVA with treatments (group and individual) as a between-subjects factor and halves (first vs. second five offers) as a within-subjects factor, detected no significant effect of treatments ($F(1,110) = .003, p = .96$), but there was a significant main effect of halves ($F(1,110) = 11.6, p < .001, \text{partial } \eta^2 = .095$). The interaction effect between stopping position in halves and treatments was also not significant ($F(1,110) = .79, p = .37$, see Figure 5.2 for more details). So there was no relative difference in group versus individual performance as the experiment wore on.

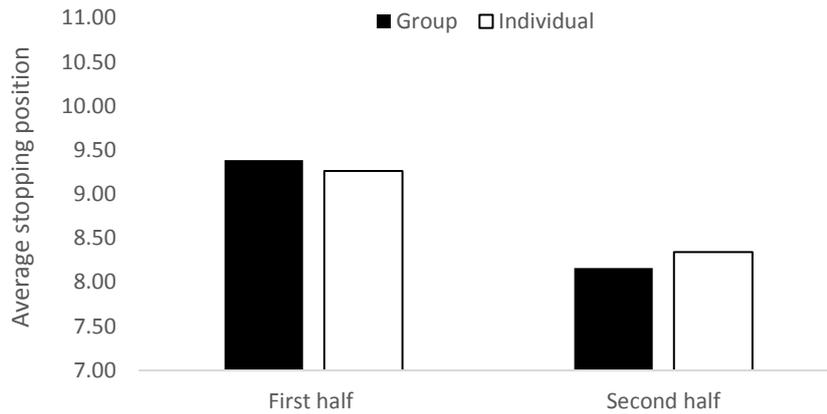


Figure 5.2. The half session average stopping position in the first (rounds 1–5) and second half (rounds 6–10) of the session for both group and individual treatments.

5.3.3 LORGE-SOLOMON ANALYSIS

The average probability of successfully solving the task for each participant was computed across all 10 rounds in both tasks (selling houses and hiring secretaries), and compared between the group and individual treatments. The average probability of successfully solving the task in the group treatment is 18.7%; that is, 1.87 rounds out of 10 rounds, and 18.0% in the individual treatment. The t test showed no significant difference between the treatments ($t(110) = -.64, p = .52$). The Lorge-Solomon model predicts a 32.8% chance that an aggregate of two people should solve the same problem if individuals solve the problem at 18.0% of the time, and yet the groups solved fewer tasks than an aggregate of two individuals. Hence, the reductionist view is supported.

5.4. DISCUSSION

On average, there was no difference found in the length of search between the group and individual treatments, nor was there a great difference in the decision pattern between groups and individuals. In both treatments, the participants tended to accept an offer with a higher value than the previous one. However, there was some support for the reductionist view that the performance of a group of two is actually worse than two individuals combined, if they solve the task individually.

In this experiment, participants were asked to make decisions both in a group of two and as an individual. However, if neither person knows how to solve the task, a group will not perform better than an individual solving the task. Due to the complexity of this particular task, group decision making may only perform better than individuals if at least one person in the group knows how to solve the problem. Perhaps, due to the nature of the secretary problem, it requires a larger group to increase the odds of any group member proposing an accurate or close-to-accurate strategy to solve the problem. Some participants reported that they had to alter their strategy in making a decision during the group experiment so that the other group member would agree to employ it. Because there are only two members in a group, they cannot use voting or another method in disagreement, and it may be that sometimes they ultimately employ a less efficient strategy after discussion and produce poorer performance than they would have done individually. It is possible that a different result may be found with a larger group of three or more people.

6. ALL MONEY IS NOT CREATED EQUAL: HOW ACTUAL AND OPTIMAL BEHAVIOUR VARIES WITH THE INCENTIVE STRUCTURE.

EXPERIMENT FIVE.

6.1 INTRODUCTION

Do all monetary incentive structures enhance performance? Money has the ability to fulfil basic needs such as buying food, shelter, and clothing, but also it signals one's worth and competence to other people (DeVoe, Pfeffer, & Lee, 2013). Monetary incentives are commonly used in the workplace and laboratory experiments to motivate performance. The rationale for providing an incentive is that people will exert more effort and an increase in effort can result in a higher performance, for example, a better outcome for a given task. However, empirical evidence has found that the effect of monetary incentives on performance is inconclusive (e.g., Jenkins, Mitra, Gupta, & Shaw, 1998; Camerer & Hogarth, 1999; Bonner, Hastie, Sprinkle & Young, 2000; Hsieh, Li, & Tsai, 2010; Mir, Trender-Gerhard, Edwards, Schneider, Bhatia, & Jahanshahi, 2011). Factors like skill, knowledge and experience, the complexity of the task and incentive structure are all potential reasons that an incentive-induced increased effort does not necessarily lead to a higher performance (Bonner & Sprinkle, 2002). For example, Fryer (2011) conducted randomized incentive experiments in over 250 schools and found that incentives offered for academic performance (better grades) were not effective, but incentives offered for effort, such as attendance, good behaviour and so on, were in fact effective. Note that the students can control how much effort they put into school work with increased attendance and good behaviour, but may not know how (have the skill) or have little success in

transferring effort into better performance, that is, the grades. The way to effectively motivate performance in such tasks is as yet unknown.

Neuroscientists and biologists have found that expectations of an event or outcome, imposed externally by the other people, environment or situation, can influence one's physiological response. For example, participants have a higher level of endorphins and experience less pain in a placebo treatment than without it (Lipman, Miller, Mays, Miller, North, & Byrne, 1990), and neuron activity can be shaped. For example, an fMRI study found that expectation enhances performance on a task, through reducing the activity (sharpening) of potentially distracting ambiguous neuronal activation in the brain region associated with performing the task (Kok, Jehee, & de Lange, 2012). It is possible that expectations of one's performance can be signalled by the incentive structure.

In this experiment, I tested whether an incentive structure that signals an expectation of performance to the participants might improve performance in a task that requires more skill to perform. The secretary problem is employed to test this conjecture, as this is a task for which exerting effort without skill does not necessarily result in a better performance in making optimal decisions. Most simply, as we have seen, an individual who searches through 20 offers during the task, and thus exerts more effort, will not normally choose more optimal prices than an individual who only searches through 10 offers. As the optimal theorem of the secretary problem dictates, the key to maximising the probability of finding the optimal price is estimating the size of the information set, which is the number of offers one needs in order to learn about the price distribution. Conducting more searches without knowing the size of the information set does not necessarily result in finding a higher price.

This experiment randomly assigned participants to one of three monetary incentive structures: the commission base, best only, and no incentive. There were two different frames for each: house-selling and no context. In the commission base structure, the participants receive an earning for their decisions according to the exchange rate announced in the instructions. For the best only structure, a fixed amount of positive earnings is available, but only when the optimal price is selected. The no incentive structure offers a fixed amount of earnings regardless of the selected price. The main contribution of this experiment is identifying how the different incentive structures influence behaviour and extend the empirical analysis of effort-enhancing monetary incentive structures to the domain of performance enhancement. Since monetary payments are the most commonly used tool in incentivising behaviours and performance, it is crucial to understand the link between how the incentive is set up and its effect on performance in tasks that require more skill to perform.

6.1.1. WHY THE EFFECT OF MONETARY INCENTIVE MATTERS

The rationale for providing monetary incentives is to motivate individuals into exerting more effort by associating effort with earning: the higher the effort, the higher the earning. A higher level of effort is assumed to result in a better performance of the task (Ariely, Gneezy, Loewenstein & Mazar, 2009). There are two general beliefs about incentive on performance. First, people who receive earnings based on their performance are expected to perform better than people who receive a fixed earnings regardless of their ability, e.g., people who receive \$0.5 when correctly recognizing an item in an item recognition task will perform better than participants who receive a \$10 flat fee for the same task. Second, people who get paid more for completing a task are expected to perform better than people who get paid less, e.g., when solving a puzzle, people who

receive \$2 when correctly solved a puzzle, will solve more (or even harder) puzzles than people who only receive \$1 per puzzle. However, Camerer and Hogarth (1999) reviewed 79 laboratory experiments and found that the effect of a monetary incentive on performance is, in fact, inconclusive. This is because factors such as the skills and cognitive abilities required to perform well differ with the tasks. Arguably, every task requires both skill and effort to perform. However, some tasks have a steeper learning curve that requires fewer skills, and people can learn to perform it quickly, for example, an item-recognition task; but some tasks have a flatter learning curve, requiring more skill to perform, for example, finding the optimal prices in the secretary problem. In such tasks it is often difficult for people to learn the required skills in the duration of an experiment. The type of task that has a steeper learning curve is, therefore, an effort-dominant task. The type of task requiring more skill to perform is referred to as a skill-dominant task. Monetary incentives are generally found to have an effect on performance with mundane tasks that are effort-dominant, and which people lack intrinsic motivation to perform, for example, clerical tasks (Riedel, Nebeker & Cooper, 1988), and item recognition and re-call tasks (e.g., Kahneman & Peavler, 1969, Libby & Lipe, 1992). As they derive no satisfaction or enjoyment from doing the task, the monetary incentive provides extrinsic motivation to perform. There are also tasks that people are intrinsically motivated to perform; in particular, tasks that are skill dominant and require more cognitive abilities due to their complexity (see Bonner and Sprinkle, 2002, for the framework of monetary incentives on effort and performance). In these cases, the monetary incentive is mostly found to have no effect on performance; for example, probability judgements using numerical or verbal expressions (Wallsten, Budescu & Zwick, 1993), a variant of the secretary problem task (Hey, 1987), trading and bidding tasks (Camerer, 1987) and ultimatum games (e.g., Bolle, 1990; Forsythe, Horowitz, Savin, & Sefton, 1994). In some

cases, monetary incentives have been found to impede performance; for example, in tasks predicting outcomes (e.g., Hogarth, Gibbs, McKenzie, & Marquis, 1991; Ashton, 1990), the Monty Hall problem (Friedman, 1998) and the Luchins water jar problem-solving task (McGraw & McCullers, 1979). Ariely, Gneezy, Loewenstein, & Mazar (2009) also found that, when offering a performance-contingent payment that varied in amount from small to very large (one month's salary) over different tasks, higher payments surprisingly produced worse performance than lower ones in tasks that required basic cognitive skills. This finding is compatible with the idea that high incentives can lead people to choke under pressure. Excessive incentives can lead to a decrement in performance (Yerkes & Dodson, 1908).

6.1.2. HOW DO MONETARY INCENTIVES MOTIVATE BEHAVIOUR?

Across human experiences, previous research commonly found that positive or negative expectation of the upcoming event suggested by the current situation helps to guide human behaviour to obtain reward or to avoid harm (Scott et al., 2007). Incentives like food, money and drugs motivate behaviours and activate the reward system through the release of dopamine to several brain regions, mainly the pleasure centre of the brain (see Schultz, 2000, for a review on how reward works in the brain and influences behaviour). Scott et al. (2007) found that expecting a monetary reward increased synaptic activity similar to placebo-induced dopamine release. An expectation of an event or outcome can also shape the perception and experience of events. For example, expectations were found to be modulators of pain (Montgomery & Kirsch, 1997; see Colloca & Benedetti, 2005, for a detailed review) and emotion (Petrovic et al., 2005) in placebo studies, and can also alter visual awareness (Sterzer, Frith & Predrag, 2010). Research on coping stress has found that participants who were instructed to expect

stressful events as functional and adaptive when watching a video, exhibited lower vascular resistance, increased cardiac efficiency and decreased attentional bias compared with those who did not hold such expectations during a stressful task. Their physiological responses, measured using electrocardiography (ECG), impedance cardiography (ICG), and blood pressure monitor, returned to baseline faster after the stressful task when they had such an expectation (Jamieson, Nock, & Mendes, 2013). Furthermore, the expectation improves task performance. Kok, Jehee, and de Lange (2012), using fMRI, found that prior expectations in visual perception facilitated performance in a visual orientation task by reducing the neural response in the primary visual cortex (V1). The reduced activity of distracting ambiguous neuron activation and sharpening neuron activity improves stimulus representation in the area. The improvement of the stimulus representation is then correlated with improved performance in the visual task when the grating orientation was expected, but not when it was unexpected.

6.1.3. HOW DO INCENTIVE STRUCTURES WORK?

People often use information from the situation and environment to work out how they are expected to behave. Consider, for example, the framing effect (Tversky, & Kahneman, 1981) and anchoring bias (Schkade & Kahneman, 1998). The environment or situation often signals an explanation of an event or behaviour. For example, when someone says that they are in a movie theatre, we infer that they are watching a movie. When people are served food with a large bowl, they consume more food compared with when they are served with a small bowl (Wansink & Cheney, 2005; Wansink, Van Ittersum & Painter, 2006). It is possible that people infer from the size of the bowl how many calories are appropriate to consume. This behavioural response is likely formed from past experiences. When we were younger, parents often served small portions of food on a

smaller plate. The size of the plate gets bigger when we are served with a larger portion of food as we grow older. Over time, the brain links the size of the plate to the portion of food that we are supposed to consume. The bowl used to serve the participants' food may be perceived as an indicator of how much food is appropriate for the individuals to consume. It may be similar with monetary incentives, as people can infer information about the task from the incentives. When people in a community were offered a large monetary reward for allowing the building of a nuclear waste site nearby, the willingness to support the nuclear site was lower than without such a reward (Frey & Oberholzer-Gee, 1997). The authors suggested that the presence and size of the monetary incentive signal that the risks involved for having such a waste site are high. Ariely, Brach and Meier (2009) found, in both laboratory and field experiments, that incentives are effective when people participate in charity events in private, but not in public. Receiving a monetary incentive for attending a charity event signals to other people that they are doing it for the money rather than trying to help others. Furthermore, people may also infer information about the task from the incentive structure rather than just the presence of a monetary incentive. Cole, Kanz, and Klapper (2015) conducted a field experiment with loan officers who were in charge of assessing risk and issuing loans in commercial banks. The loan officers were randomly assigned to one of three performance compensation structures: (a) the officers earned an origination bonus for successfully getting a loan approved; (b) the officers received a small bonus for loans that performed well, and a small penalty for loans that did not; (c) the loan officers received a large bonus for loans that they approved and which performed well, and a large penalty for loans they approved and which performed poorly. The loan officers in the large bonus/penalty group rejected more bad loan applications, exerted more effort and thus increased profit more for each loan they originated. The loan officers who received an origination bonus exerted a similar amount of effort and made a similar profit to the small

reward/penalty group. It is possible that the large bonus/penalty loan officers performed better after inferring a high expectation of their ability to perform the task from the incentive structure: “If I am getting paid a lot to make such decision, the decision must be a difficult one. I must be very good in making decisions to be trusted in making such a decision.” Furthermore, Cole et al. (2015) found that the loan officers who received an origination bonus rated the application significantly less risky in all loans, both those approved and those not, than those loan officers whose bonus was tied to a loan's success.

6.1.4. HYPOTHESES

The following experiment was set up to examine how incentive structures motivate performance in variants of the secretary problem. Measures of performance were the number of the optimal (best possible) prices found, total earnings, and stopping position. In the experiment by Kok et al. (2012), participants were hypothesised to perform better when the expectation that the best was possible was signalled. When monetary incentives are implemented in the laboratory setting, it is assumed that the participants aim to maximise their utilities. However, maximising utility is not always the same as maximising earnings. If a skill-dominant task such as this task, requires too great a cognitive load, one may perceive exerting cognitive load as an increased cost and therefore be more willing to settle for non-optimal earnings. The commission base structure signals that the participants can maximise their earnings through finding the optimal prices, yet it is also acceptable to not find the optimal decision, as positive earnings are still available for decisions other than the optimal. Hence, they are more likely to maximise their utilities in various ways, not necessary through earnings, for example, by choosing a good enough price that requires less cognitive load. The no incentive structure may also elicit similar behaviour, due to a lack of expectation to perform.

In contrast, the best only structure may signal that finding the optimal price is expected and possible, at least on occasion, and this is why earnings are only being offered when the optimal price is chosen. Furthermore, people in the best only structure may also be more willing to search longer, as a monetary incentive is only available on finding the optimal price. They are unlikely to stop searching until they have (or believe they have) found the optimal price.

In line with the previous research, I hypothesised (1) that the best only incentive structure would yield more rounds of optimal prices than the other incentive structures; (2) in line with the previous hypothesis, participants in the best only incentive structure would have higher overall earnings because they find more optimal prices; (3) in line with the information signalled by the incentive structures, the best only incentive structure will encourage participants to search longer than the other incentive structures; (4) in line with the findings from Experiment Two, the house-selling frame will encourage participants to search longer and choose higher prices; (5) in line with general incentive theory, the commission base structure would find more optimal prices and choose higher total prices than the no incentive structure.

6.2 METHOD

6.2.1 PARTICIPANTS

A total of 178 undergraduate students from the University of Canterbury participated individually in the experiment. They were recruited by advertisements on the Department of Psychology research pool website (ucpsyc.sona-systems.com). All the participants received two 100-level course credits for participating in the experiment (a show-up payment). In addition, they could receive different cash

incentives: The participants earned 10.50 NZD on average. The search task was implemented using E-Prime software (pstnet.com/products/e-prime). There was an age range from 18 to 40 years old, and median in the range 18 – 21 years old. The experiments lasted on average one hour.

6.2.2 EXPERIMENTAL DESIGN AND PROCEDURES

The experiment consisted of two parts, the first of which was the search task, the second a survey to elicit demographic information. The experimental procedures were similar to Experiment One. The experiment was a 2×3 across-subject design, with two frames, house selling and no context, and three monetary incentive structures, commission base, the best only, and no incentive. Participants performed their entire experiment in one treatment only. The number of participants that participated in each treatment is summarised in Table 6.1. In the house selling frame treatments, the participants were asked to accept price offers for 10 houses, a total of 20 price offers for each house (Appendix 6.A – 6.C present instructions of the three incentive structures in the house selling frame). In the no context frame, the participants were asked to accept numbers for 10 rounds (Appendix 6.D – 6.F present the instructions under no context frame). Table 6.2 presents summary information on the prices in each round. All treatments had two practice rounds prior to the 10 paid rounds.

In the commission base structure, the instructions explained their earnings are calculated by adding the 10 prices (presented in experimental currency units, henceforth ECUs) that the participants have selected. Then the total chosen price was converted into New Zealand Dollars using the conversion rate of 735 ECU to 1 New Zealand Dollar. The instructions for how the commission base payment was calculated in the house selling frame were as follows (Appendix 6.A presents more detail):

“The payoffs will be denoted in experimental currency units (ECUs).

$$735 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

Total ECUs you earn

=

Accepted price offer for House 1 + Accepted price offer for House 2 + +

Accepted price offer for House 10

Example: Suppose you accepted the price offer 450 for House 1, 260 for House 2, 380 for House 3..... 658 for House 10. The total amount of ECUs you earn is 450+260+380+.... +658.”

The wording in the no context frame of the commission base structure was exactly the same, except there was no mention of the house selling frame, Appendix 6.D presents detailed instructions. The participants in the best only treatments received NZD 4.60 for each round in which they selected the optimal price (the offer with the highest value in each round). The overall payment they received on completing the experiment was NZD 4.60 multiplied by the number of rounds they selected the optimal price. The instructions for how the best only payment calculated in the house selling frame are as follows (Appendix 6.B shows more details):

“You will earn NZD 4.60 if you have chosen the highest price offer for each house.

Your payoffs are determined as follows:

Total NZD you earn

= number of houses that you have selected the highest price offer * NZD 4.60

Example: Suppose you accepted the highest price offer for House 1, House 3 and House 10. The total amount of NZD you earn is $4.60 * 3 = 13.80$.”

The best only structure under no context frame was identical to the house selling frame, without the house selling context; full instructions are in Appendix 6.E. The participants in the no incentive treatments simply received \$9.50 for their participation; Appendix 6.C and Appendix 6.F for house selling and no context frames respectively.

For this experiment, the cash reward for the participants was set to be approximately NZD 10. The tuition fee for a Psychology 100-level course at the University of Canterbury was NZD 719. The two course credits thus might be worth approximately $719 \times 2\% = \text{NZD } 14.4$. Thus, the overall incentive was approximately NZD 24.4. This is 64% higher than the hourly rate of minimum wage in New Zealand³. In short, the participants were reasonably well incentivised in this experiment. Considering the participants were mostly 18-year-old, first-year college students, they were unlikely to earn this much money from other sources.

The conversion rate in the commission base treatments, the cash earning for selecting an optimal price in the best only structure and the cash payment in the no incentive structure were based on the results of previous findings. So, all treatments were

³ Obtained from <https://www.employment.govt.nz/hours-and-wages/pay/minimum-wage/minimum-wage-rates/>

intended to yield roughly the same earning on average, regardless of the incentive structures, which was set to be NZD 10. In the commission base structure, the conversion rate was calculated by taking the selected total price per participant from the no cost treatment in Experiment One and dividing by \$10. As for the best only structure, the payment amount for finding an optimal price was obtained by taking \$10 and dividing by the average number of rounds the optimal price was chosen in the no cost treatment in Experiment One. However, the average earnings in the best only structure in this experiment yielded higher earnings than predicted (due to the treatment effect of the incentive structure). The actual average earnings in the different treatments are presented in Table 6.2. More analysis regarding the actual earnings is available in the results, Section 6.3.5. The remaining procedural details were identical to those described in Experiment One.

Table 6.1

Participant data for the treatment groups.

	Number of Participants	Average cash earnings (NZD)
House-selling frame		
Commission base	26	9.9
Best only	30	13.2
No incentive	31	9.5
No Context		
Commission base	31	9.4
Best only	29	11.4
No incentive	31	9.5

Table 6.2

Summary data of the prices presented in each round.

	Optimal price*	Optimal position**	Average price***	Minimum price****	Standard deviation
Round 1	848	8	509.6	276	165.4
Round 2	875	10	469.2	2	284.4
Round 3	708	10	437.6	207	147.2
Round 4	733	20	518.5	267	145.5
Round 5	578	13	331.2	186	114.4
Round 6	1574	10	714.3	89	447.4
Round 7	581	19	369.2	197	128.1
Round 8	966	3	636.4	250	234.4
Round 9	1740	14	756.4	105	396.2
Round 10	625	4	440.4	250	101.3

Note: All the prices are presented in ECU, which is in the same format the participants see in the experiment.

*Optimal price refers to the highest price in each round, also known as the maximum price in each round.

**Optimal position is the position at which the optimal price was presented in that round.

***Average price was obtained by averaging across the 20 prices in each round.

****Minimum price was the lowest price in each round.

6.3 RESULTS

In this section, I first consider the average results obtained from the experiment. The average results examine dependent variables that are measures of overall subject decision making. Then I examine the correlation between the incentive structures and the search behaviours. Finally, I present results related to the actual earnings made by participants with respect to incentive structures.

There were three dependent variables. To examine the number of search activities exerted in the task for different incentive structures, the position in the sequence where the participant accepted the offers (henceforth stopping position) was evaluated. To examine the decision-making performance, the cumulative sum of the 10 chosen prices obtained by accepting the offers (henceforth chosen price) was used. Finally, the number of rounds the optimal price was selected (henceforth optimal count) was calculated. The results relating to these three dependent variables are shown in Table 6.3 (panels A, B, C respectively).

Table 6.3

The summary statistics of position in the sequence at which the offer was accepted obtained after averaging across 10 rounds (panel A), the sum of 10 chosen prices in ECUs (panel B) and the number of rounds the optimal price is selected (the offer with the highest value in the sequence; panel C) for the treatments.

Panel A. Stopping position (position)

Treatments	Average	Median	Standard Deviation	Range (Min–max)
House selling				
Commission base	8.3	6	6.2	1–20
Best only	9.4	8	6.9	1–20
No incentive	7.6	5	6.0	1–20
No context				
Commission base	8.6	6	7.1	1–20
Best only	10.1	8	7.2	1–20
No incentive	9.1	8	6.4	1–20
Sequence optimal*	11.1			3–20
Classical optimal**	13.6			

Panel B. Total chosen price (ECU)

Treatments	Average	Median	Standard Deviation	Range (Min–max)
House selling				
Commission base	7252.5	7416.5	563.8	4905–7951
Best only	7244.4	7265.5	495.6	5515–8320
No incentive	6983.5	7007	687.1	4833–8476
No context				
Commission base	6929.7	7038	605.8	5254–7709
Best only	6947.2	7103	665.2	4960–7729
No incentive	7081.5	7094	260.4	6638–7620
Sequence optimal*	9228			
Classical optimal**	7930			

Panel C. Optimal price count (round)

Treatments	Average	Median	Standard Deviation	Range (Min-Max)
House selling				
Commission base	2.0	2	1.2	0 – 5
Best only	2.9	3	1.5	0 – 6
No incentive	1.6	2	1.2	0 – 4
No context				
Commission base	1.9	2	1.2	0 – 4
Best only	2.5	3	1.5	0 – 5
No incentive	1.6	2	0.8	0 – 3
Sequence optimal*	10			
Classical optimal**	4			

*Sequence optimal refers to the actual optimal price from the sequences used in the experiment. ** This is the result predicted by the optimal theory of the classical secretary problem.

6.3.1 STOPPING POSITION

Table 6.3. Panel A shows that participants in the best only structure searched the longest in both house-selling ($M = 9.4$) and no context frames ($M = 10.1$) compared to the commission base structure ($M = 8.3, 8.6$ respectively) and no incentive ($M = 7.6, 9.1$ respectively). Because there are two frames for each incentive structure, I refer to structures instead of treatments when referring to the average result of two treatments (frame and no context) in the same structure. Figure 6.1 graphically presents the average stopping position of incentive structures averaged across frames (house-selling and no context frames) in each round. The best only structure consistently produced a longer search than the commission base and no incentive structures, except in rounds 5 and 7. Analysis of variance results showed a significant main effect of the incentive structures on the average stopping positions averaged across 10 rounds for each participant, ($F(2, 172) = 4.55, MS_{error} = 8.08, p = .012$, partial $\eta^2 = .50$). Participants in the best only structure chose to stop at a significantly later position ($M = 9.7$) than the commission base ($M = 8.4$) and no incentive ($M = 8.3$) structures. Tukey HSD post hoc tests showed that there was a significant ($p < .05$) difference between the commission base and best only structures ($p = .038$), as well as between the best only and no incentive structures ($p = .019$). But there was no significant difference between the commission base and no incentive ($p = .978$) structures. This suggests participants searched longer when only the optimal price yields a positive payoff. Note, too, that the amount of searching activities in the best only structure is closer to the sequence optimal ($M = 11.1$) in the experiment than the other two incentive structures.

Analysis of variance results also showed that participants chose to stop at a significantly later position in the sequence ($M = 9.3$) under the no context frame than the house-selling ($M = 8.4$) frame ($F(1, 172) = 4.08, p = .045, \text{partial } \eta^2 = .023$). However, there was no statistically significant ($p = .524$) interactive effect of frame and incentive structure.

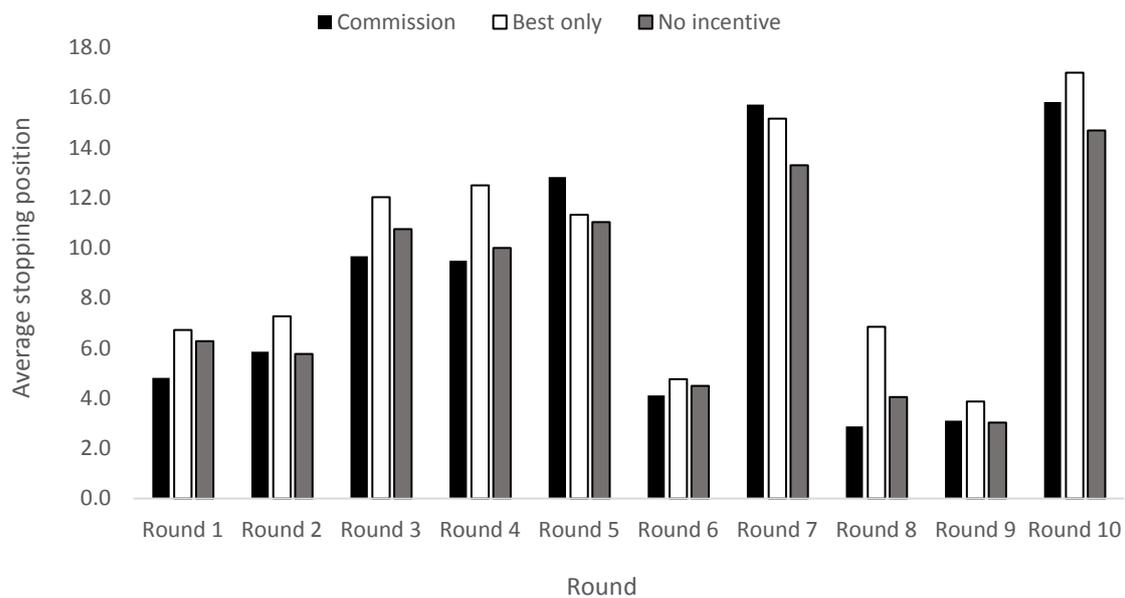


Figure 6.1. The average stopping position of incentive structures averaged across frames (house-selling and no context frames) in each round.

6.3.2 TOTAL CHOSEN PRICE

Analyses of variances found a significant main effect of the frames ($F(1, 172) = 4.227, MS_{error} = 3.18 \times 10^5, p = .012, \text{partial } \eta^2 = 4.23$). There were no other statistically significant ($p < .05$) main or interactive effects. However, the interaction effect was close to being significant ($F(2, 172) = 2.64, p = .074, \text{partial } \eta^2 = .03$). Thus, the incentive structure did not affect the prices chosen, despite the payoff being

immediately dependent on the price selected in the commission base structure. This further suggests that in the commission base structure, people maximise their utilities, not all through earnings. Table 6.3 Panel B shows more descriptive results on the total chosen price.

6.3.3. OPTIMAL PRICE COUNT

Participants in the best only structure selected the highest number of optimal prices on average ($M = 2.7$) compared with the commission base structure ($M = 1.95$) and no incentive ($M = 1.6$). Analyses of variance were again used to investigate the effect of incentive structure (the commission base, best only or no incentive), the frame (house-selling or no context) on the average number of rounds the optimal price was selected. The average number the optimal price was selected was calculated by averaging across each participant within the treatment. There was a significant main effect of the incentive structures on the number of optimal prices chosen, ($F(2, 172) = 12.45, MS_{error} = 1.51, p < .001, \text{partial } \eta^2 = .126$). The Tukey HSD post hoc test confirmed significant differences between the optimal count in the best only and commission base ($p = .002$), and the best only and no incentive ($p < .001$) structures. There was no significant difference between commission base and no incentive ($p = .306$) structures. There were no other statistically significant ($p < .05$) main or interactive effects. Detailed results are shown in Table 6.3 Panel C. Taken over both frames, 30.6% of the participants in the best only structure performed better or as well as the classical theorem in finding the optimal price (by finding 4 or more optimal prices), and 10.2% of them outperformed the classical theorem (by finding 5 or more optimal prices); 5.2% of the commission based structure participants performed as predicted by the classical theorem and 1.9% performed better. Only

3.2% of the no incentive participants performed as well as the classical theorem and 0% outperformed. Figure 6.2 graphically presents the frequency distribution of the number of rounds the optimal price is selected in the incentive structures.

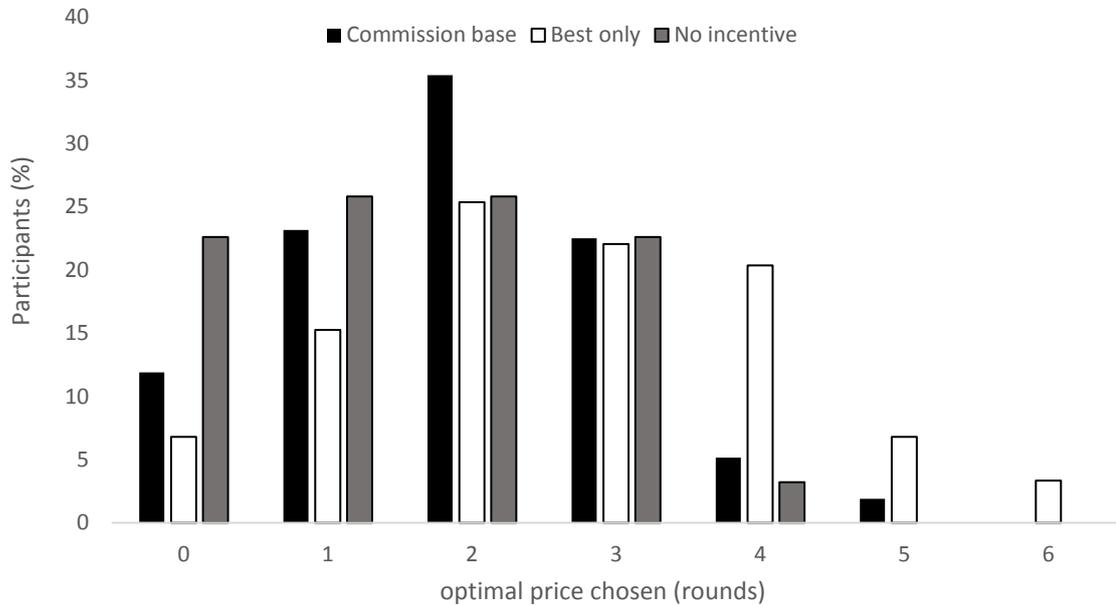


Figure 6.2. Percentage distribution of the number of rounds the optimal price is selected (averaged across house-selling and no context frames) of participants in the commission base, best only and no incentives structures.

6.3.4. CORRELATION RESULTS

Pearson correlation analysis was used to examine the relationship between the three dependent variables: the average stopping positions, the sum of 10 chosen prices and the number of rounds the optimal price is selected. The average stopping positions were obtained by averaging 10 stopping position across each participant. The sum of 10 chosen prices was obtained from each participant by adding up the 10 prices they selected. The optimal price count was obtained by adding the total number (out of 10 rounds) of the

optimal prices that each participant has selected. The correlation results are summarised in Table 6.4.

Table 6.4

Pearson correlation coefficient with the position in the sequence at which the offer was accepted, the number of rounds the optimal price is selected, and the sum of 10 chosen prices in the treatments.

	Stopping position		Chosen price	
	Coefficient* (Range)**	<i>p</i>	Coefficient* (Range)**	<i>p</i>
Optimal count	0.55 (0.27–0.61)	< 0.001	0.58 (0.50–0.73)	< 0.001
Stopping position			0.56 (0.27–0.83)	< 0.001

*Coefficient shows the correlation analysis results of three dependent variables in the six treatments.

*Range show the lowest correlation coefficient and the highest coefficient in the six treatments.

In general, there is a large positive and significant correlation between the length of the search and the number of rounds in which the optimal price is selected ($r = .547, p < .001$). The longer the participants searched the higher their chosen price ($r = .560, p < .001$). The more optimal prices the participant selected, the higher the overall price they selected ($r = .583, p < .001$).

Another correlation analysis was conducted within each incentive structure by averaging across frames (house selling and no context) with respect to the three dependent variables. The commission base structure had a strong correlation between the stopping position and chosen price ($r = .729, p < .001$) compared with the best

only ($r = .441, p < .001$) and no incentive ($r = .541, p < .001$). The strong correlation in the commission base structure shows that some people searched longer, so earned a higher payoff, and some people searched less and earned a lower payoff. The weaker correlation in the other two treatments suggests the amount of search was similar but earnings varied between individuals within the structure. This suggests that the amount of search activities in the commission base structure is more inconsistent than in the other structures. The average stopping position was approximately 8 in the commission base structure and the likelihood of finding a better price by continuing their search was quite high (60% chance of finding a higher price). The longer their search, the higher the price they will find and vice versa.

6.3.5. ACTUAL EARNINGS

The actual earnings in the commission base structure were calculated by taking the total chosen price (ECU) in 10 rounds divided by the conversion rate (735 ECU to 1 NZD). Also, see the actual average earning for each treatment in Table 6.1. In the best only structure, it is found by multiplying \$4.60 by the number of rounds they selected the optimal price. The participants in the no incentive structure received \$9.50 regardless of the price they selected. The average actual earnings of the incentive structures were obtained by averaging across the house-selling and no context frames. The average actual earning in the best only structure ($M = \$12.3$) is the highest compared with the commission base ($M = \$9.70$) and no incentive structure ($M = \$9.50$). Analyses of variances found a significant main effect of the incentive structures ($F(2, 172) = 9.65, MS_{\text{error}} = 15.56, p < .001, \text{partial } \eta^2 = .099$). The Tukey HSD post hoc test showed there was a significant difference in the actual earnings between the best only and commission base structures ($p < .001$), also

between the best only and no incentive structures ($p < .001$). But no difference is found between the commission base and no incentive structures ($p = .983$). There were no other statistically significant ($p < .05$) main or interactive effects. Thus, the best only structure had a higher earning than predicted from the previous experiment. The actual earnings in the commission base structure were the same as in the previous experiment.

6.4 DISCUSSION

This study investigated the effect of different incentive structures in a sequential search task which required some skill to perform well. The experiment found different performance with the different monetary incentive structures. Performance was measured by three dependent variables: the number of rounds that the optimal price was selected, the cumulative sum of the chosen price, and the search length. The best only structure found the most rounds of optimal prices and produced longer search, whereas there was no difference between the commission base and no incentive structures. However, there was no significant difference found in the total price chosen between the incentive structures. The house selling frame was found to have an effect in choosing a higher total price.

People employed different search strategies under the different incentive structures. In the best only structure, their searches were worthless if they did not select the optimal price; thus they persisted longer until they reached the optimal price. When they had (or believed they had) missed the optimal, they stopped their search even if the price was poor, as it was pointless to continue searching. A decision to choose poor prices in these circumstances could potentially contribute to the finding that the total chosen price in the best only structure was similar to the other incentive structures, even when they found

more optimal prices. One might argue that the cost of continued searching would be minimal given that the best only structure only receives a payoff when they found the optimal price, so they might be likely to search through all the offers when they are uncertain if they have passed the optimal or not. This was not found in the experiment, as the average stopping positions were earlier than the sequence optimal. In fact, the participants stopped their search prior to the optimal prices in all treatments including the best only. Also, even if the participants in the best only structure continue searching after the optimal price, the probability of choosing a high price decreased with more searches, and hence they would not have a higher total price.

There are at least two possible limitations in this experiment. First, the no incentive structure received NZD 9.5 instead of NZD 10. The original design was to avoid participants in the no incentive structure receiving a higher average earning than the other two structures. The average payoff for the other two structures was set to be approximately NZD10. Considering that the participants were all first-year students with no experience in buying or selling houses, they might be expected to perform more badly than predicted (from results found in the previous experiment) but this was not found in this experiment. Second, the earnings of the best only structure were presented in NZD directly, but in the commission base structure they were presented in ECUs first, then converted to NZD. The use of ECUs may lead to a money illusion effect and increase incentives (Fehr & Tyran, 2001), and yet the best only structure still outperformed the commission base structure. Furthermore, Drichoutis, Lusk, and Nayga (2015) reported that the use of ECUs does not have an effect on influencing behaviours when compared with using a monetary term directly.

To conclude, the experiment showed, firstly, that having a commission base, proportional incentive did not produce a better performance than simply having a flat payment for any of the different dependent variables considered. However, paying only for the best did lead to longer (although still not optimal) searches and (reasonably enough) to more frequently obtaining the best price. It is plausible that the effect of the best only condition arose because the condition itself signaled to participants that obtaining the best was possible (as indeed it often was). A more general implication is simply that the structure of the incentive matters in a complex task such as the secretary problem.

7. GENERAL DISCUSSION AND CONCLUSION

The results of the five experimental studies presented here indicate the broad range of factors that humans consider in their decision-making process. The general discussion first presents a summary of the findings from Experiments One to Five. Second, I discuss the implications of these experiments. Third, I discuss the limitations of the studies and present future directions of research for the secretary problem and sequential decision making in general.

7.1 SUMMARY OF RESULTS

CHAPTER 2. THE EFFECT OF TIME COST ON SEARCH BEHAVIOURS. EXPERIMENT ONE.

Result 1: People searched less when paying a monetary or time cost for each new offer.

Time cost had a smaller shortening effect than monetary cost

Result 2: People chose lower prices when incurring monetary cost

Result 3: People adjust their search behaviour with time delay even when the delay was not pre-announced.

Result 4: Search behaviour changes drastically when paying monetary cost.

CHAPTER 3. A HOUSE-SELLING CONTEXT IMPROVES DECISION MAKING FROM A SEQUENCE OF OFFERS. EXPERIMENT TWO.

Result 1: People conducted less search without a context, but the amount of search in the housing context was similar with or without house information.

Result 2: Having a context helped people achieve higher prices. But people had the same earnings with or without information.

Result 3: People without a context searched even less in later rounds

Result 4: People without a context chose lower prices from the start and this persisted over time.

CHAPTER 4 WHY DO PEOPLE SEARCH TOO LITTLE? NOT ENOUGH TIME, OR TRYING TO AVOID REGRET? EXPERIMENT THREE.

Result 1: The search behaviour suggested regret aversion.

Result 2: There was similar search shortening with the time delay, as in Experiment One.

Result 3: Self-confidence has little influence on search behaviour.

CHAPTER 5 DECISION MAKING IN SMALL GROUPS. EXPERIMENT FOUR.

Result 1: The group did not search longer than individuals.

Result 2: Two-person groups found fewer optimal prices than the aggregate performance of individuals.

CHAPTER 6 ALL MONEY IS NOT CREATED EQUAL: HOW ACTUAL AND OPTIMAL BEHAVIOUR VARIES WITH THE INCENTIVE STRUCTURE. EXPERIMENT FIVE.

Result 1: People searched for longer when they were rewarded only for finding the optimal price.

Result 2: People chose higher prices when the task has a house-selling frame, whereas an incentive structure had little effect on price choice.

Result 3: People found a greater number of optimal prices when they were being rewarded when finding the optimal price only than with other types of reward.

Results 4: When people were being rewarded with a commission on the price, they had a similar amount of search and found similar number of optimal prices as the people not being incentivised for their performance.

Result 5: People earned more money when they were rewarded only for finding the best.

Result 6: When people were being rewarded with a commission on the price, search behaviour varies between individuals.

7.2 IMPLICATIONS OF THE FINDINGS AND CONCLUSIONS

Throughout this dissertation, a number of factors that contribute significantly to one's sequential search decisions have been identified. These include the search costs, contexts, regret aversion and incentives. Although these underlying causes have been found to influence decisions, to date they have not been accounted for in the theoretical optimal solution. Experiment One examined the effect of time and monetary search costs, as search costs play a critical role in every search decision. The time people spend to evaluate price offers can be otherwise spent in getting one's own work done. When the time search cost equalled one percent of the optimal price offer, people's search behaviours aligned with the theoretical optimal solution (Seal & Rapoport, 1997). Even when time cost was not present in Seal and Rapoport's (1997) experiment, people would always incur time cost, whether decisions were being made in the laboratory or outside world. Experiment One found that time and monetary costs shortened the amount of search, but monetary cost shortened people's search by approximately 45%; the time cost shortened by approximately 14%. Despite a shortening effect with different magnitude being found between time and monetary costs, the effects of money and time on decisions do differ in various ways. For example, people settled for lower prices with monetary cost compared to no cost and time cost conditions. Also, people who incurred monetary cost had a narrow and shorter term perspective than those who incurred time costs. People were more concerned about the most recent rejected price when there was a monetary cost,

rather than being concerned about the highest price seen in the time cost or no cost conditions. Monetary cost appears to affect behaviours and decisions in more ways than time cost, although theoretically the effects of monetary and time cost are qualitatively equivalent.

Several other studies also suggest that people do treat time and money differently (e.g., Lee, Lee, Bertini, Zauberger & Ariely, 2015; Gino, & Mogilner, 2014). Very often, a monetary cost has a more adverse effect on decisions, as found in Experiment One, than time. For example, it appears to make people more prone to bias (Soman, 2001), less moral (Gino & Mogilner, 2014), and more risk-seeking (Okada & Hoch, 2004). One way to mitigate the adverse effect is perhaps to remind oneself of the underlying purpose of the task and shift the focus away from the incurred monetary costs. Ariely, Kamenica and Prelec (2008) asked their participants to build Lego Bionicles in their experiment. During the experiment, the participants were paid \$2 for building the first one, \$1.89 for the second, \$1.67 for the next and so on. The participants decided when to stop building. In the *meaningful* treatment, the completed Bionicles were placed on the desk in front of the participant and accumulated on the desk. In the *Sisyphus* treatment, once the participant finish the first Bionicle, and start working on the second one, the experimenter would immediately undo the first Bionicle into its pieces and return the pieces back into the box. In the *meaningful* treatment, people built a total of 10.6 Bionicles and received an average of \$14.4, whereas, in the *Sisyphus* treatment, people only built 7.2 Bionicles and earned an average of \$11.52. The authors claim that when people find the task meaningful and obtain fulfilment or satisfaction from doing the task, they are more productive than when only doing it for the money. Therefore, when selling a house in reality, for example, people may see it as an opportunity to spend time with loved ones, instead of focusing on the monetary

cost. This may mitigate the adverse effect of monetary costs and enhance the probability of selling the house to the best buyer.

However, when we are selling houses, we may have more information than the assumptions in the secretary problem suggested. We generally have an idea of how much the house is worth now from official government valuation, and we also have access to information about the housing market in a suburb and an entire city. In this case, we likely know the price range of the offer prior to the search, and perhaps something about the likely distribution of the offers within that range. Christian and Griffith (2016) proposed that, for a uniform distribution and a known price range, the optimal solution for the acceptable price that maximises earnings is $p \geq 1 - \sqrt{2c}$, where p is the acceptable offer and c the cost of the search. Let us say, you know the offers are between \$400,000 and \$500,000. The search cost for an additional offer is \$1. Then you should search until you receive an offer of \$499552.79 or more. However, if the search cost is \$10,000, the offer should be accepted if the price is \$455,279 or more. In a case where the search is very expensive and costs \$50,000 per search, an offer over \$400,000 should be accepted. Clearly, even in the case where information about offers is available to the decision maker when selling houses, search cost remains a significant factor, if not the most critical factor, in determining the amount of search.

Experiment Two expanded on the findings of Experiment One and examined whether contexts would influence search behaviour. Experiment Two showed that not only search costs but also the contexts presented in decision making are linked with the amount of search and the prices people choose in the task. It is well known that we make different decisions in different contexts, for example, the framing effect (e.g., Tversky & Kahneman, 1981; McNeil, Pauker, Sox, & Tversky, 1982; Hoffman, McCabe, Shachat, &

Smith, 1994; Dufwenberg, Gächter & Hennig-Schmidt, 2011). The effect of context is not present in the optimal theorem previously mentioned. Experiment Two found that people performed better in a house-selling context than without. The amount of information presented in the context did not influence the ability to make superior decisions. In fact, simply stating the context of a decision without providing detailed house information seemed sufficient to facilitate better performance. People without a context performed worse from the start and this effect persisted, even worsening over time. The results may arise because an existing schema that enhances our decision-making ability may only be activated when a context is presented. Furthermore, experiments without a context may not necessarily represent the actual decision-making behaviours in life, where no decision is made without a context.

The secretary problem can also be found in many contexts in our lives apart from selling houses, one of which is finding a partner. As people mostly likely would date one person at a time, and we cannot evaluate all potential candidates in our lifetime simultaneously, the quality of the future candidates cannot be known ahead of time. The question then becomes: How can one choose the best person to settle down with, without having the information on every candidate that one might want to marry? Re-calling a previously rejected candidate is often unlikely in most circumstances. If we apply the classical optimal theorem to this context, we should reject the first 37% of the people that we date, and choose the next best one. However, the total number of candidates is never known ahead of time, a modified version of the optimal stopping rule is found, which can be applied to time rather than people (Fry, 2015). Say you start dating at sixteen years old and would expect to get married by the age of forty. In the first 37% ($[40 - 16] \times 37\% = 8.9$) years, every candidate should be rejected before the age of 25 ($16 + 8.9$) regardless of

his/or her quality; after 25, marry the first person who is better than everyone that you have dated. But how would you compare the candidates? Unlike in the house-selling context, where information about the price offer range, search cost, and so on, are often readily available and can be objectively determined, a better candidate for marriage often cannot.

Christian and Griffith (2016) suggest that, for the best chance of success in the quest of finding the best partner, one should evaluate the candidates based on objective criteria, for example, their wealth percentile, rather than on subjective criteria, for example, feelings, that cannot be easily and objectively evaluated. Suppose that the decision maker has information about the distribution of the criteria being evaluated when finding a partner (as when selling houses), and the candidates' qualities can be evaluated with an objective measure. Suppose, too, that one knows the approximate wealth distribution of the candidates, such as 500,000 is above average, 700,000 is in 75th percentile, and so on. In this case, the decision maker would know immediately when a good candidate is presented, for example, above the 98th percentile, without needing to first gather information as the classical optimal theorem suggested. The optimal strategy, therefore, becomes a threshold rule, where a candidate will be accepted if he/she is above a certain threshold, such as the 95th percentile. This implies that even if the candidate is presented early on, he/she can still be accepted. The chance of choosing the best candidate is drastically increased when full information is available. When there is no information, the probability of finding the best is approximately 37% (see Lindley, 1961; Gilbert & Mosteller, 1966 for more details) and with full information, it is approximately 58% as $n \rightarrow \infty$ (Gilbert & Mosteller, 1966). Therefore, when making sequential search decisions in life, access to information matters, as it significantly enhances the possibility of success.

It is possible that the commonly found early stopping bias in the sequential search task is at least in part caused by the regret aversion. As there is a higher possibility of feeling regret after a long search, it might be optimal to shorten the search to avoid regret and reduce the psychological cost in one's decisions. Participants shortened the information set compared with the classical theorem, and this potentially suggests regret aversion. Strong evidence for regret aversion comes from the finding that only rather rarely did participants choose – or were forced to choose – the final offer. The overconfidence bias was another psychological factor considered. The results show that search behaviour was similar regardless of the level of self-confidence. However, the participants tended to stop their search in an increasing sequence rather than decreasing sequence, suggesting a tendency to loss aversion. While the non-incentivised procedure made the experiment a more theoretical exercise, it provided an opportunity for the participants to set their own goals (as they probably would in real life), and, in the process, provided data that indicate what other forces might be at work besides the logic of the experiment itself. This is an important consideration because there are a vast number of different ways in which an experiment of this kind could be incentivised (some are investigated in Experiment Five).

A good situation to apply findings from Experiment Three in life is another example of the secretary problem: finding a car parking space. Assume that you are driving on a long road with the parking spaces (both free and occupied) evenly distributed, approaching your destination. Your goal is to minimise the walking distance to your destination, and the occupancy rate – the probability of how many car parking spaces – is known. Once the decision is made to pass an empty space, it is not likely to be available later on, as there are other people behind you who are also looking for spaces, especially in

a highly occupied parking street. DeGroot (1970) found the optimal solution to this problem is to take the first empty spot less than $-\log_2/\log(1-p)$ spots from the destination, where p is the probability of any given space being available, i.e., the occupancy rate. So, if the occupancy rate of the parking space is 50%, you should wait until 1 parking space away from your destination then take the next free spot; basically, drive up to the last space before the destination. If the occupancy rate is 90%, then you should wait until 7 spaces away from the destination then consider parking the car. However, if the occupancy is 99%, then wait until 69 parking spaces away. If the occupancy rate is 99.9%, then you should start looking for parking space at 693 spaces away from the destination. From this example, we can easily see that when the occupancy rate is high, a small increase in the occupancy rate changes the result drastically. Also, finding a space requires time, attention and fuel, all of which are scarce resources; taking a spot early on, and avoiding regret at having to spend more time and effort to find one later on, may not be a bad strategy to employ.

Very often, when selling or buying houses, the decision is made jointly by the couple or the family who own the house. When hiring employees, a committee is often involved in deciding which candidate to hire. We often believe that two heads are better than one; groups can make better decisions than individuals. I explored whether group decision making can outperform individual decision making in solving the secretary problem. It is often assumed that group decision making may allow the group to avoid more cognitive biases than individuals, through discussion, and thus groups come up with better solutions. This was not found in solving the secretary problem. Although the problem can be solved when one member of the group knows how to solve the problem, the other group member might not be willing to adopt the solution due to the lack of

feedback in the experiment. Furthermore, group discussion can have the effect of shifting decisions away from the optimal or towards the optimal strategy. Experiment Four found that the group does not choose more optimal offers than individuals, and hence it provides evidence that supports the reductionist view in group decision making, where the group of two solves fewer problems than the aggregate performance of individuals. Group decision making in life can be found more often in some contexts, such as buying or selling houses, than others, such as finding partners or parking spaces. As Experiment Four demonstrated, having more people will not contribute to higher success, but, as previously mentioned, having more information would. “Two heads are better than one” is only true in solving sequential search problems when the other head knows more.

In addition to information and knowledge, there is also evidence that the incentive structure influences search behaviour. We often overlook how the expectation of an event or outcome helps to prepare us for the acquisition of a reward. In Experiment Five, the incentive structure that rewarded only finding the best resulted in finding more optimal prices than the other two structures. Both Experiments Two and Five included treatments with both a house-selling and a no context frame. I did not find people conducted less search without a context in Experiment Five. There can be at least three potential underlying causes for this. First, it may be due to a different population of students. In Experiment Two, the participants were mostly commerce majors whereas, in Experiment Five, the participants were mostly psychology majors. The difference in finding could be driven by self-selection into majors. Second, the incentive structure was different. The participants in Experiment Two received a cash-only reward in New Zealand Dollars, whereas people received a mixture of cash (NZD) plus course credits in Experiment Five. People may be inclined to work harder to earn the course credits than gain a cash reward.

Third, the overall value of the incentive was higher in Experiment Five. As previously mentioned, the cash value of NZD 10 plus two course credits is approximately NZD 25.4. This was higher than the maximum earning of NZD 15 in the Experiment Two design.

Experiment Five found that the commission base structure had very similar results to the no incentive structure. These two treatments gave rise to the same amount of search and a similar number of optimal offers. This was unexpected, as often an incentivised procedure can better motivate behaviour than a non-incentivised procedure. There are at least three potential underlying reasons for what I found. First, people in the no incentive structure may be intrinsically motivated to perform even when the extrinsic motivation, a monetary reward, was lacking. For example, Gneezy and Rustichini (2000) found that in a field experiment where high school students were collecting charity donations door-to-door, students collected more donations when there was no monetary incentive for them than with a monetary reward. Second, Gneezy, Meier and Rey-Biel (2011) suggested that when an extrinsic incentive is available, it signals to the participants that the experimenter perceives them to lack intrinsic incentive to perform, and people in the commission base structure may have had some performance decrement due to lack of trust. Third, as previously mentioned, participants may have already been inclined to work hard for the course credits, and both the commission base and no incentive structures offered the same course credits for completing the experiment.

The best only treatment was able to outperform the commission base and no incentive treatments in Experiment Five. They conducted longer searches and found more optimal prices than the other treatments, despite the expected earnings having been designed to be the same as the other two treatments. However, participants in the best only treatment were also three times more likely to accept the last offer than in the other

treatments. People continue searching even when they pass the optimal until they reach the last offer. The last offer tends to have an average payoff or a lower payoff when search cost is taken into account. As a result, people in the best only treatment had the same total chosen price as the commission base treatment, even when they found more optimal offers.

Not knowing when to stop can sometimes turn out to be problematic in life, where the number of offers is often infinite if there is no time constraint. Theoretically, if one is willing to wait, the search can last for a long time, as an offer will eventually arise, regardless its quality. There is a sequential decision-making problem for which there is no optimal solution, for example, the game of triple or nothing (Christian & Griffiths, 2016). Suppose you are given an endowment of \$10 to play a game, and you can play as many times as you want. But the rule of the game is that you need to bet all the money each time and you would have a 50% chance of tripling the money and 50% chance of losing your entire stake. How many times should you play this game to maximise the earnings? Theoretically, infinitely many times; you should never stop playing. But, following this strategy will only result in you eventually losing the entire stake. Knowing when to stop, when solving sequential search problems, can be as tricky as knowing how to choose optimally.

Overall, this dissertation examined specific cases of people confronting a sequential search problem and it is clear that these kinds of problems present themselves often in one form or another in our lives. Whether it involves buying or selling houses, hiring employees, finding partners or car parking spaces, the most critical question that we wish to answer is how one should make optimal decisions. To answer this question, we must determine what factors contribute to the optimal solution. As the experiments have demonstrated, various factors are often being considered simultaneously in the decision-

making process. The combination of these has not yet been captured by any mathematically derived optimal solution. The claim that we make suboptimal decisions is thus, in fact, often inaccurate and superficial. Instead, the decisions we make may be optimal when the effect of these factors that we consider to be critical and crucial to our desired outcome are taken into account, allowing us to maximise utilities, as the economists would say. In general, these findings suggest that when we buy or sell real houses we face the choice between spending time or monetary cost, and opting for time cost may avoid the adverse effect of monetary cost. These experiments would suggest that we should always take heed of the context while making decisions, as decisions do depend on the context. If possible, when comparing between offers, use measures that can be objectively determined to enhance the chance of success: More information allows higher quality decisions, as suggested by the optimal stopping rule. When all offers yield a payoff, aim for the second or third best to save time and effort, as the earnings will not be much different than when aiming for the best. Most importantly, as Christian and Griffith (2016) have stated, always be ready to stop, as the time and effort already spent are sunk costs that need to be disregarded. Exhausting and exploring all offers is never the optimal choice, as time spent can never be regained, nor can the rejected offers. The basic assumptions of the secretary problem have suggested that the secretary problem is, in fact, the study of the time itself, which represents the journey of human life. We only pass this journey of life once; we should make decisions wisely and timely and live without regrets.

7.3 LIMITATIONS AND FUTURE RESEARCHES

The experiments in this dissertation provided important findings regarding human decision making. Nevertheless, there are limitations to the experimental designs that might have affected results. First, in Experiment One, the experiment did not separate out the

motives behind different perceptions of time and money. The underlying cause of the different behaviours between the time and monetary search costs remains unclear. This opens the possibility of future extensions and robustness checks, as identification of the reason for this difference would require a design different from the current one.

Second, in everyday life situations, re-calling a previously rejected offer can sometimes be possible. There is some evidence (Zwick et al., 2003) that people are willing to search longer when a re-call is allowed, both with and without monetary search costs. Experiments can, therefore, be conducted to examine whether the different effects of time and monetary costs on search behaviour persist after allowing for re-call. In a similar fashion, the impact of time search cost could vary with the opportunity cost of time, which presumably differed for different participants in my experiments (Knowles & Servátka 2015). I leave these explorations for future theoretical and experimental studies.

Third, Experiment Two adapted a house-selling context to explore the change in decision-making behaviours, and yet, to date, there is no definite theory explaining the effect of context framing, or what type of framing one should use to achieve an intended change, for example, higher earnings. From that perspective, the results of a house-selling frame may not generalise to other contexts, such as the context of buying houses, selling or buying cars, choosing life partners, and so on. One should also test the robustness with respect to other types of context framing, to gain insights on the issues above. Furthermore, although Experiment Two found that people with no context had fewer searches than people with context, this result did not generalise to Experiment Five. A theory about the context of framing would be useful for this future research.

Fourth, a clear limitation in exploring the psychological insights on the early stopping bias in Experiment Three was the non-incentivised procedure, although in

Experiment Five people found similar performance in the non-incentivised procedure as one of the incentivised procedure (the commission base incentive structure). In Experiment Three, participants were asked to choose the best offer (highest price offer and the candidate with the highest score in each round), but there were no earnings associated with the quality of their decisions. Goette and Stutzer (2010) showed that noncash incentives such as T-shirts, lottery tickets, or gift cards could have a positive effect on blood donation, whereas a cash incentive often has a negative effect. Some participants only receive course credits in Experiment Three, which is also a type of noncash incentive; previous literature has shown that a noncash incentive has a different effect to cash incentives. Also, incentivised procedures may motivate decision-making behaviours in ways that non-incentivised procedures do not. Future research on how the decisions are made under different types of incentive in the secretary problem should also be considered.

Fifth, one of the limitations of Experiment Four was the size of the group. The groups consisted of just two group members, where, given a disagreement, the more dominant or assertive member may be more likely to make the decision, regardless of the accuracy of his or her decision. This is less likely to occur with a larger group, for example, three or more people. Some common approaches for making decisions in a group, such as voting, cannot really be carried out in a group of two. Varying group size is recommended for a future experiment.

Sixth, the total chosen price was similar regardless of its incentive structures in Experiment Five. It appears that the best only incentive structure motivated behaviours that were specific to the incentive structure when the reward was given but did not affect another aspect of performance, the overall price value, and yet choosing a higher price may also be an important performance measure in different tasks. A future experiment is

needed to explore the possibility of incentive structure in performance regarding price value and optimal prices, for example, a commission base incentive structure that also provides a bonus when people found the optimal prices.

The findings in this dissertation have important implications for research in psychology and economics decision making. The fact that most of the tested factors were not included in prior theories provides evidence for gaps between existing theories and actual decision-making behaviours. In addition to exploring behaviours experimentally, formulating a theory to better explain and predict behaviours should be carried out in the future.

REFERENCE

- Adams, J. K. (1957). A confidence scale defined in terms of expected percentages. *The American Journal of Psychology*, 70(3), 432–436. 10.2307/1419580
- Alicke, M., & Govorlin, O. (2005). *The better-than-average effect*. In M. D. Alicke, D. A. Dunning, & J. Krueger (Eds.), *The self in social judgment* (pp. 85–106). New York: The Psychology Press.
- Ariely, D., Bracha, A., & Meier, S. (2009). Doing good or doing well? Image motivation and monetary incentives in behaving prosocially. *The American Economic Review*, 99(1), 544–555. 10.1257/aer.99.1.544
- Ariely, D., Gneezy, U., Loewenstein, G., & Mazar, N. (2009). Large stakes and big mistakes. *Review of Economic Studies*, 76(2), 451–469. 10.1111/j.1467–937X.2009.00534.x
- Ariely, D., Kamenica, E. & Prelec, D. (2008). Man's search for meaning: The case of Legos. *Journal of Economic Behaviour & Organization*, 67(3–4), 671–677. 10.1016/j.jebo.2008.01.004
- Ashton, R. H. (1990). Pressure and performance in accounting decision settings: Paradoxical effects of incentives, feedback and justification. *Journal of Accounting Research*, 28, 148–180. 10.2307/2491253
- Baillargeon, R; DeVos, J. (1991). Object permanence in young infants: Further evidence. *Child Development*, 62(6), 1227–46. 10.2307/1130803
- Bargh, J. A., Lombardi, W. J., & Higgins, E. T. (1988). Automaticity of chronically accessible constructs in person×situation effects on person perception: It's just a

- matter of time. *Journal of Personality and Social Psychology*, 55(4), 599-605.
10.1037/0022-3514.55.4.599
- Bearden, J. N., & Murphy, R.O. (2007). *On Generalized Secretary Problems*. In:
Abdellaoui M., Luce R.D., Machina M.J., Munier B. (eds) *Uncertainty and Risk.*
Theory and Decision Library C (Series C: Game Theory, Mathematical
Programming and Operations Research), vol 41. Springer, Berlin, Heidelberg.
10.1007/978-3-540-48935-1_11
- Bearden, J. N., Rapoport, A. & Murphy, R. O. (2006). Sequential Observation and
Selection with Rank-Dependent Payoffs: An Experimental Study. *Management*
Science, 52(9), 1437–1449. 10.1287/mnsc.1060.0535
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*,
30(5), 961–981. 10.1287/opre.30.5.961
- Bolle, F. (1990). High reward experiments with high expenditure for the experimenter?
Journal of Economic Psychology, 11(2), 157–167. 10.1016/0167-4870(90)90001-
P
- Bonner, S. E., & Sprinkle, G. B. (2002). The effects of monetary incentives on effort and
task performance: Theories, evidence, and a framework for research. *Accounting,*
Organizations and Society, 27(4), 303–345. 10.1016/S0361-3682(01)00052-6.
- Bonner, S. E., Hastie, R., Sprinkle, G. B., & Young, S. M. (2000). A review of the effects
of financial incentives on performance in laboratory tasks: Implications for
management accounting. *Journal of Management Accounting Research*, 12(1), 19–
64. 10.2308/jmar.2000.12.1.19

- Bower, G. H., & Cirilo, R. K. (1985). Cognitive psychology and text processing. *Handbook of Discourse Analysis*, (pp. I: 71–I:105). Retrieved from https://stanford.edu/~gbower/1985/Cog_Psy_&_Text_Proc-Cirilo.pdf
- Bray, R. M., & Noble, A. M. (1978). Authoritarianism and decisions of mock juries: Evidence of jury bias and group polarization. *Journal of Personality and Social Psychology*, 36(12), 1424–1430. 10.1037/0022-3514.36.12.1424
- Bruin, W. B. d., Parker, A. M., & Fischhoff, B. (2007). Individual differences in adult decision-making competence. *Journal of Personality and Social Psychology*, 92(5), 938–956. 10.1037/0022-3514.92.5.938
- Camerer, C. F. (1987). Do biases in probability judgement matter in markets? Experimental evidence. *American Economic Review*, 77(5), 981–997. Retrieved from <http://www.jstor.org/stable/1810222>
- 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), 7–42. Retrieved from <http://www.jstor.org/stable/41760945>
- Camerer, C., & Lovo, D. (1999). Overconfidence and excess entry: An experimental approach. *The American Economic Review*, 89(1), 306–318. 10.1257/aer.89.1.306
- Chidambaram, L., & Tung, L. L. (2005). Is out of sight, out of mind? An empirical study of social loafing in technology-supported groups. *Information Systems Research*, 16(2), 149–168. 10.1287/isre.1050.0051

- Christian, B. & Griffith, T. (2016). *Algorithms to Live By: The Computer Science of Human Decisions*. Picador: New York, NY Reviewed.
- Churchland, M. M., Chou, I. H., & Lisberger, S. (2003). Evidence for object permanence in the smooth-pursuit eye movements of monkeys. *Journal of Neurophysiology*, *90*(4), 2205–18. 10.1152/jn.01056.2002
- Clark, R. D., & Willems, E. P. (1969). Risk preferences as related to judged consequences of failure. *Psychological Reports*, *25*(3), 827–830. 10.2466/pr0.1969.25.3.827
- Cole, S., Kanz, M., & Klapper, L. (2015). Incentivizing calculated risk-taking: evidence from an experiment with commercial bank loan officers *Journal of Finance*, *70*(2), 537–575. 10.1111/jofi.12233
- Colloca, L., & Benedetti, F. (2005). Placebos and painkillers: is mind as real as matter? *Nature Reviews Neuroscience*, *6*(7), 545–52. 10.1038/nrn1705
- Connolly, T., & Zeelenberg, M. (2002). Regret in decision making. *Current Directions in Psychological Science*, *11*(6), 212–216. 10.1111/1467-8721.00203
- Connolly, T., Ordóñez, L. D., & Coughlan, R. (1997). Regret and responsibility in the evaluation of decision outcomes. *Organizational Behaviour and Human Decision Processes*, *70*(1), 73–85. 10.1006/obhd.1997.2695
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, *31*(3), 187–276. 10.1016/0010-0277(89)90023-1

- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture* (pp. 163–228). New York, NY, US: Oxford University Press. Retrieved from [http://citeseerx.ist.psu.edu/viewdoc/download?doi = 10.1.1.140.7186&rep = rep1&type = pdf](http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.140.7186&rep=rep1&type=pdf)
- Cosmides, L., & Tooby, J. (2005). Social exchange: The evolutionary design of a neurocognitive system. In M. S. Gazzaniga (Ed.), *The new cognitive neurosciences, III*. Cambridge, MA: MIT Press
- Cosmides, L., Barrett, H. C., & Tooby, J. (2010). Adaptive specializations, social exchange, and the evolution of human intelligence. *Proceedings of the National Academy of Sciences, 107*, 9007–9014. 10.1073/pnas.0914623107
- Cox, J., & Oaxaca, R. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty, 2*(3), 301–329. 10.1111/j.1465–7295.1982.tb01146.x
- Cox, J. C., & Oaxaca, R. L. (2000). Good news and bad news: Search from unknown wage offer distributions. *Experimental Economics, 2*(3), 197–225. 10.1007/BF01669196
- Cusin, C., Yang, H., Yeung, A., & Fava, M. (2009). *Handbook of Clinical Rating Scales and Assessment in Psychiatry and Mental Health Series*. Current Clinical Psychiatry (Eds.), Chapter 2, Humana Press.
- DeGroot, M. H. (1970). *Optimal Statistical Decisions*. New York: McGraw–Hill.

- De Neys, W. (2006). Dual processing in reasoning two systems but one reasoner. *Psychological Science*, *17*(5), 428–433. 10.1111/j.1467–9280.2006.01723.x
- Deaves, R., Lüders, E., & Luo, G. Y. (2009). An experimental test of the impact of overconfidence and gender on trading activity. *Review of Finance*, *13*(3), 555–575. 10.1093/rof/rfn023
- DeVoe, S. E., Pfeffer, J., & Lee, B. Y. (2013). When does money make money more important? Survey and experimental evidence. *ILRREVIEW*, *66*(5), 1078–1096. Retrieved from <http://www.jstor.org/stable/24369572>
- DiMaggio, P (1997). Culture and cognition. *Annual Review of Sociology*, *23*(1), 263–287. 10.1146/annurev.soc.23.1.263.
- Drichoutis, A.C., Lusk, J. L., & Nayga, R. M. (2015). The veil of experimental currency units in second price auctions. *Journal of the Economic Science Association*, *1*(2), 182–196. 10.1007/s40881–015–0014–2
- Duchon, D., Dunegan, K. J., & Barton, S. L. (1989). Framing the problem and making decisions: the facts are not enough. *IEEE Transactions on Engineering Management*, *36*(1), 1, 25–7. 10.1109/17.19979
- Dufwenberg, M., Gächter, S., & Hennig-Schmidt, H. (2011). The framing of games and the psychology of play. *Games and Economic Behaviour*, *73*(2), 459–478. 10.1016/j.geb.2011.02.003 10.1016/j.geb.2011.02.003
- Eger, C. & Dickhaut, J. (1982). An examination of the conservative information processing bias in an Accounting framework. *Journal of Accounting Research*, *20*(2), 711–723. 10.2307/2490894

- Evans, J. S. B. T. (2008). Dual-processing accounts of reasoning, judgement, and social cognition. *Annual Review of Psychology*, 59(1), 255–278.
10.1146/annurev.psych.59.103006.093629
- Fehr, E., & Tyran, J. R. (2001). Does money illusion matter? *The American Economic Review*, 91(5), 1239–1262. 10.1257/aer.91.5.1239
- Fenton–O’Creevy, M., Nicholson, N., Soane, E., & Willman, P. (2003). Trading on illusions: Unrealistic perceptions of control and trading performance. *Journal of Occupational and Organizational Psychology*, 76(1), 53–68.
10.1348/096317903321208880
- Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical Science*, 4(3), 282–296. 10.1214/ss/1177012493.
- Fischbacher, U. (2007). z–Tree: Zurich toolbox for ready–made economic experiments. *Experimental Economics*, 10(2), 171–178. 10.1007/s10683–006–9159–421
- Freeman, P. R. (1983). The secretary problem and its extensions: A review. *International Statistical Review*, 51(2), 189–206. 10.2307/1402748
- Forsythe, R., Horowitz, J. L., Savin, N. E., & Sefton, M. (1994). Fairness in simple bargaining experiments. *Games and Economic behaviour*, 6(3), 347–369.
10.1006/game.1994.10
- Frey, B. S., & Oberholzer-Gee, F. (1997). The cost of price incentives: an empirical analysis of motivation crowding-out. *American Economic Review*, 87(4), 746–55.
Retrieved from <http://www.jstor.org/stable/2951373>

- Friedman, D. (1998). Monty hall's three doors: Construction and deconstruction of a choice anomaly. *American Economic Review*, 88(4), 933–946. Retrieved from <http://www.jstor.org/stable/117012>
- Fry, H. (2015). *The mathematics of love: Patterns, proofs and the search for the ultimate equation* (1st T Books hardcover ed.). New York: Ted Books.
- Fryer, Jr., R. G. (2011). Financial incentives and student achievement: evidence from randomized trials. *The Quarterly Journal of Economics*, 126(4), 1755–1798. 10.1093/qje/qjr045
- Gamliel, E. & Peer, E. (2010). Attribute framing affects the perceived fairness of health care. *Judgment and Decision Making*, 5(1), 11–20. Retrieved from <https://doaj.org/article/934a99ad75bf4281be37bb9310886ee0>
- Gardner, F. (1960). Mathematical games. *Scientific American*, 202(1). 6.10.1038/scientificamerican0160–150
- Gardner, M., & Mathematical Association of America. (2009). *Sphere packing, lewis carroll, and reversi: Martin gardner's new mathematical diversions*. Cambridge;New York: Cambridge University Press.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669. 10.1037/0033–295X.103.4.650
- Gilbert, J. P., & Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313), 35–73. 10.1080/01621459.1966.10502008

- Giné, X., Townsend, R., & Vickery, J. (2008). Patterns of rainfall insurance participation in rural India. *The World Bank Economic Review*, 22(3), 539–566.
10.1093/wber/lhn015
- Gino, F., & Mogilner, C. (2014). Time, money, and morality. *Psychological Science*, 25(2), 414–421. 10.1177/0956797613506438
- Gneezy, U. & Rustichini, A. (2000). Pay enough or don't pay at all. *Quarterly Journal of Economics*, 115(30), 791–810. 10.1162/003355300554917
- Gneezy, U., Meier, S. & Rey-Biel, P. (2011). When and why incentives (don't) work to modify behaviour. *Journal of Economic Perspectives*, 25(4), 191–210.
10.1257/jep.25.4.191
- Goette, L & Stutzer, A. (2010). Blood donations and incentives: evidence from a field experiment. Working paper. 10.1111/j.0042-7092.2007.00700.x
- Goldman, A., & Johansson, J. K. (1978). Determinants for search of lower prices: An empirical assessment of the economics of information theory. *Journal of Consumer Research*, 5(3), 176–186. Retrieved from <http://www.jstor.org/stable/2488667>
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1), 114–125.
10.1007/s40881-015-0004-4
- Griffin, D., & Brenner, L. (2004). Perspectives on probability judgment calibration, in Koehler, D. and Harvey, N. (Eds.), *Blackwell Handbook of Judgment and Decision Making*, Blackwell, Malden Mass, 177–199. 10.1002/9780470752937.ch9

- Griggs, R. A., & Cox, J. R. (1982). The elusive thematic–materials effect in Wason’s selection task. *British Journal of Psychology*, 73(3), 407–420. 10.1111/j.2044–8295.1982.tb01823.x
- Hamilton, M. (1960). A rating scale for depression. *Journal of Neurology, Neurosurgery & Psychiatry*, 23(1), 56–62. 10.1136/jnnp.23.1.56
- Haun, D. E., Zeringue, A., Leach, A., & Foley, A. (2000). Assessing the competence of specimen-processing personnel. *Laboratory Medicine*, 31(11), 633–637. 10.1309/8Y66–NCN2–J8NH–U66R
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behaviour. *The American Journal of Psychology*, 57(2), 243–259. 10.2307/1416950
- Hey, J. D. (1987). Still searching. *Journal of Economic Behaviour and Organization*, 8(1), 137–144. 10.1016/0167–2681(87)90026–6
- Higgins, E. T., & Chaires, W. M. (1980). Accessibility of interrelational constructs: Implications for stimulus encoding and creativity. *Journal of Experimental Social Psychology*, 16(4), 348–361. 10.1016/0022–1031(80)90027–X
- Hilton, D., Régner, I., Cabantous, L., Charalambides, L., & Vautier, S. (2011). Do positive illusions predict overconfidence in judgment? A test using interval production and probability evaluation measures of miscalibration. *Journal of Behavioural Decision Making*, 24(2), 117–139. 10.1002/bdm.678
- Hoffman, E., McCabe, K., Shachat, K., & Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, 7(3), 346–380. 10.1006/game.1994.1056

- Hogarth, R. M., Gibbs, B. J., McKenzie, C. R. M., & Marquis, M. A. (1991). Learning from feedback: Exactingness and incentives. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *17*(4), 734–752. Retrieved from <http://psy2.ucsd.edu/~mckenzie/Hogarthetal1991JEPLMC.pdf>
- Hollenbeck, J. R., Ilgen, D. R., Segoe, D. J., Hedlund, J., Major, D. A., & Phillips, J. (1995). Multilevel theory of team decision making: Decision performance in teams incorporating distributed expertise. *Journal of Applied Psychology*, *80*(2), 292–316. 10.1037/0021-9010.80.2.292
- Hsieh, S., Li, T., & Tsai, L. (2010). Impact of monetary incentives on cognitive performance and error monitoring following sleep deprivation. *Sleep*, *33*(4), 499–507. Retrieved from <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2849789/>
- Ingham, A. G., Levinger, G., Graves, J., & Peckham, V. (1974). The Ringelmann effect: Studies of group size and group performance. *Journal of Experimental Social Psychology*, *10*(4), 371–384. 10.1016/0022-1031(74)90033-X
- Jamieson, J. P., Nock, M. K., & Mendes, W. B. (2013). Changing the conceptualization of stress improves affective and physiological outcomes in social anxiety disorder. *Clinical Psychological Science*, *1*(4), 363–374. 10.1177/2167702613482119
- Jenkins, G. D. Jr., Mitra, A., Gupta, N., & Shaw, J. D. (1998). Are financial incentives related to performance? A Meta analytic review of empirical research. *Journal of Applied Psychology*, *83*(5), 777–787. 10.1037/0021-9010.83.5.777

- Jerusalem, M., & Schwarzer, R. (1992). *Self-efficacy as a resource factor in stress appraisal processes*. In R. Schwarzer (Ed.), *Self-efficacy: Thought control of action* (pp. 195–213). Washington, DC, US: Hemisphere Publishing Corp.
- Johnson, D. D. P., McDermott, R., Barrett, E. S., Cowden, J., Wrangham, R., McIntyre, M. H., & Rosen, S. P. (2006). Overconfidence in wargames: Experimental evidence on expectations, aggression, gender and testosterone. *Proceedings of the Royal Society B: Biological Sciences*, 273(1600), 2513–2520. 10.1098/rspb.2006.3606
- Johnson-Laird, P. N., & Wason, P. C. (1970). A theoretical analysis of insight into a reasoning task. *Cognitive Psychology*, 1(2), 134–148. 10.1016/0010-0285(70)90009-5
- Jonsson, A., & Allwood, C. M. (2003). Stability and variability in the realism of confidence judgments over time, content domain, and gender. *Personality and Individual Differences*, 34(4), 559–574. 10.1016/S0191-8869(02)00028-4
- Kahan, J., Rapoport, A., & Jones, L. (1967). Decision making in a sequential search task. *Perception & Psychophysics*, 2(8), 374–376. 10.3758/BF03210074
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58(9), 697–720. 10.1037/0003-066X.58.9.697
- Kahneman, D. (2011). *Thinking, fast and slow* (1st ed.). New York: Farrar, Straus and Giroux.
- Kahneman, D., & Peavler, W. S. (1969). Incentive effects and pupillary changes in association learning. *Journal of Experimental Psychology*, 79(2 pt1), 312–318. 10.1037/h0026912

- Kahneman, D., & Riepe, M. W. (1998). Aspects of investor psychology. *Journal of Portfolio Management*, 24(4), 52–65. 10.3905/jpm.1998.409643
- Kahneman, D. & Tversky, A. (1984). Choices, Values, and Frames. *American Psychologist*, 39(4), 341–350. 10.1037/0003-066X.39.4.341
- Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, 103(3), 582–591. 10.1037/0033-295X.103.3.582
- Keren, G. (1997). On the calibration of probability judgments: Some critical comments and alternative perspectives. *Journal of Behavioural Decision Making*, 10(3), 269–278. 10.1002/(SICI)1099-0771(199709)10:3 < 269::AID-BDM281>3.0.CO;2-L
- Knowles, S., & Servátka, M. (2015). Transaction costs, the opportunity cost of time and procrastination in charitable giving. *Journal of Public Economics*, 125, 54–63. 10.1016/j.jpubeco.2015.03.001
- Kocher, M. G., Pahlke, J., & Trautmann, S. T., (2013). Tempus fugit: Time pressure in risky decisions. *Management Science*, 59(10), 2380–2391. 10.1287/mnsc.2013.1711
- Kok, P., Jehee, J. F., & de Lange, F. P. (2012). Less is more: expectation sharpens representations in the primary visual cortex. *Neuron*, 75(2), 265–270. 10.1016/j.neuron.2012.04.034
- Lee, L., Lee, M. P., Bertini, M., Zauberman, G., & Ariely, D. (2015). Money, time, and the stability of consumer preferences. *Journal of Marketing Research*, 52(2), 184–199. 10.2139/ssrn.1542283

- Leganger, A., Kraft, P., & Røysamb, E. (2000). Perceived self-efficacy in health behaviour research: Conceptualisation, measurement and correlates. *Psychology & Health, 15*(1), 51–69. 10.1080/08870440008400288
- Levin, J. (2009). Cohesiveness in groups. *Encyclopaedia of Human Relationships, 1*, 231–233.
- Levin, I. P., & Gaeth, G. J. (1988). How consumers are affected by the framing of attribute information before and after consuming the product. *Journal of Consumer Research, 15*(3), 374–378. Retrieved from <http://www.jstor.org/stable/2489471>
- Libby, R., & Lipe, M. G. (1992). Incentives, efforts and the cognitive processes involved in accounting-related judgements. *Journal of Accounting Research, 30*(2), 249–273. 10.2307/2491126
- Lichtenstein, S., & Fischhoff, B. (1977). Do those who know more also know more about how much they know? *Organizational Behaviour and Human Performance, 20*(2), 159–183. 10.1016/0030-5073(77)90001-0
- Lindley, D. (1961). Dynamic Programming and Decision Theory. *Journal of the Royal Statistical Society. Series C (Applied Statistics), 10*(1), 39–51. 10.2307/2985407
- Lipman, J. J., Miller B. E., Mays K. S., Miller M. N., North W. C., & Byrne W. L. (1990). Peak B endorphin concentration in cerebrospinal fluid: reduced in chronic pain patients and increased during the placebo response. *Psychopharmacology, 102*(1), 112–116. 10.1007/BF02245754
- Lippman, S. A., & McCall, J. J. (1976). The Economics of job search: A survey. *Economic inquiry, 14*(2), 155–189. 10.1111/j.1465-7295.1976.tb00386.x

- Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92(368), 805–824. 10.2307/2232669
- Lorge, I., & Solomon, H. (1955). Two models of group behaviour in the solution of Eureka-Type problems. *Psychometrika*, 20(2), 139–148. 10.1007/BF02288986
- Luszczynska, A., Scholz, U., & Schwarzer, R. (2005). The general self-efficacy scale: Multicultural validation studies. *The Journal of Psychology*, 139(5), 439–457. 10.3200/JRLP.139.5.439–457
- McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 84(1), 113–126. 10.2307/1879403
- McGraw, K. O., & McCullers, J. C. (1979). Evidence of a detrimental effect of extrinsic incentives on breaking a mental set. *Journal of Experimental Social Psychology*, 15(3), 285–294. 10.1016/0022–1031(79)90039–8
- McNeil, B. J., Pauker, S. G., Sox, H. C., & Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *The New England Journal of Medicine*, 306(21), 1259–1262. 10.1056/NEJM198205273062103
- Meyerowitz, B. E., & Chaiken, S. (1987). The effect of message framing on breast self-examination attitudes, intentions, and behaviour. *Journal of Personality and Social Psychology*, 52(3), 500–510. 10.1037/0022–3514.52.3.500
- Michailova, J. (2011). Development of the overconfidence measurement instrument for the economic experiment. *MPRA Paper 34799*, University Library of Munich, Germany. Retrieved from <https://mpra.ub.uni-muenchen.de/id/eprint/26384>

- Mir, P., Trender-Gerhard, I., Edwards, M. J., Schneider, S. A., Bhatia, K. P., & Jahanshahi, M. (2011). Motivation and movement: The effect of monetary incentive on performance speed. *Experimental Brain Research*, 209(4), 551–559. 10.1007/s00221-011-2583-5
- Montgomery, G. H., & Kirsch, I. (1997). Classical conditioning and the placebo effect. *Pain*, 72(1–2), 107–113. 10.1016/S0304-3959(97)00016-X
- Moore, D. A., & Small, D. A. (2007). Error and bias in comparative judgment: On being both better and worse than we think we are. *Journal of Personality and Social Psychology*, 92(6), 972–989. 10.1037/0022-3514.92.6.972
- Moriguti, S. (1993). Basic theory of selection by relative rank with cost. *Journal of the Operations Research Society of Japan*, 36(1), 46–61.
- Moscovici, S., & Doise, W. (1994). *Conflict and consensus: A general theory of collective decisions*. Thousand Oaks, Calif; London; Sage Publications.
- Moscovici, S., & Zavalloni, M. (1969). The group as a polarizer of attitudes. *Journal of Personality and Social Psychology*, 12(2), 125–135. 10.1037/h0027568
- Mucci, A. G. (1973). Differential equations and optimal choice problems. *The Annals of Statistics*, 1(1), 104–113. 10.1214/aos/1193342386
- Mullainathan, S., & Shafir, E. (2013). *Scarcity: Why having too little means so much* (1st ed.). New York: Times Books, Henry Holt and Company.
- Myers, D. G., & Bishop, G. D. (1970). Discussion effects on racial attitudes. *Science*, 169(3947), 778–779. 10.1126/science.169.3947.778

- Narvaez, D., & Bock, T. (2002). Moral schemas and tacit judgement or how the defining issues test is supported by cognitive science. *Journal of Moral Education, 31*(3), 297–314. 10.1080/0305724022000008124
- Nisbett, R. E., & Wilson, T. D. (1977). The halo effect: Evidence for unconscious alteration of judgments. *Journal of Personality and Social Psychology, 35*(4), 250–256. 10.1037//0022–3514.35.4.250
- Okada, E., & Hoch, S. (2004). Spending time versus spending money. *Journal of Consumer Research, 31*(2), 313–323. 10.1086/422110
- Oskamp, S. (1962). The relationship of clinical experience and training methods to several criteria of clinical prediction. *Psychological Monographs, 76*(28), 1–27. 10.1037/h0093849.
- Pallier, G., Wilkinson, R., Danthiir, V., Kleitman, S., Knezevic, G., Stankov, L., & Roberts, R. D. (2002). The role of individual differences in the accuracy of confidence judgments. *The Journal of General Psychology, 129*(3), 257–299. 10.1080/00221300209602099
- Pavitt, C. (1998). Small group discussion: a theoretical approach. Retrieve from <http://www.uky.edu/~drlane/teams/pavitt/>
- Peterson, C. (2000). The future of optimism. *American Psychologist, 55*(1), 44–55. 10.1037/0003–066X.55.1.44
- Petrovic, P., Dietrich, T., Fransson, P., Andersson, J., Carlsson, K., & Ingvar, M. (2005). Placebo in emotional processing—induced expectations of anxiety relief activate a

- generalized modulatory network. *Neuron*, 46(6), 957–969.
10.1016/j.neuron.2005.05.023
- Pollard, P., & Evans, J. S. B. T. (1987). Content and context effects in reasoning. *The American Journal of Psychology*, 100(1), 41–60. 10.2307/1422641
- Postmes, T., Spears, R., & Cihangir, S. (2001). Quality of decision making and group norms. *Journal of Personality and Social Psychology*, 80(6), 918–930.
10.1037//0022–3514.80.6.918
- Pruitt, D. G. (1971). Choice shifts in group discussion: An introductory review. *Journal of Personality and Social Psychology*, 20(3), 339–360. 10.1037/h0031922
- Riedel, J. A., Nebeker, D. M., & Cooper, B. L. (1988). The influence of monetary incentive on goal choice, goal commitment and task performance. *Organizational Behaviour and Human Decision Processes*, 42(2), 155–180. 10.1016/0749–5978(88)90010–6
- Rimm, H., & Jerusalem, M. (1999). Adaptation and validation of an estonian version of the general self–efficacy scale (ESES). *Anxiety, Stress & Coping*, 12(3), 329–345.
10.1080/10615809908250481
- Ringelmann, M. (1913). "Recherches sur les moteurs animés: Travail de l'homme"
[Research on animate sources of power: The work of man]. *Annales de l'Institut National Agronomique*, 12(2), 1–40.
<http://gallica.bnf.fr/ark:/12148/bpt6k54409695.image.f14.langEN>

- Ritov, I., & Baron, J. (1995). Outcome knowledge, regret, and omission bias. *Organizational Behaviour and Human Decision Processes*, 64(2), 119–127. 10.1006/obhd.1995.1094
- Rush, A. J., Trivedi, M. H., Carmody, T. J., Ibrahim, H. M., Markowitz, J. C., Keitner, G. I., Kornstein, S. G., Arnow, B., Klein, D. N., Manber, R., Dunner, D. L., Gelenberg, A. J., Kocsis, J. H., Nemeroff, C. B., Fawcett, J., Thase, M. E., Russell, J. M., Jody, D. N., Borian, F. E. & Keller, M. B. (2005). Self-reported depressive symptom measures: sensitivity to detecting change in a randomized, controlled trial of chronically depressed, nonpsychotic outpatients. *Neuropsychopharmacology*, 30(2), 405–16. 10.1038/sj.npp.1300614
- Rush, A. J., Trivedi, M. H., Ibrahim, H. M., Carmody, T. J., Arnow, B., Klein, D. N., Markowitz, J. C., Ninan, P. T., Kornstein, S., Manber, R., Thase, M. E., Kocsis, J. H., Keller, M. B. (2003). The 16-item Quick Inventory of Depressive Symptomatology, clinician rating and self-report a psychometric evaluation in patients with chronic major depression. *Biological Psychiatry*, 54(5), 573–583. 10.1016/S0006-3223(02)01866-8
- Schkade, D. A., & Kahneman, D. (1998). Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science*, 9(5), 340–346. 10.1111/1467-9280.00066
- Scholz, U., Doña, B. G., Sud, S., & Schwarzer, R. (2002). Is general self-efficacy a universal construct? Psychometric findings from 25 countries. *European Journal of Psychological Assessment*, 18(3), 242–251. 10.1027//1015-5759.18.3.242

- Schultz, W. (2000). Multiple reward signals in the brain. *Nature Reviews Neuroscience*, 1(3), 199–207. 10.1038/35044563
- Schwarzer, R., Mueller, J., & Greenglass, E. (1999). Assessment of perceived general self-efficacy on the Internet: Data collection in cyberspace. *Anxiety, Stress, and Coping*, 12(2), 145–161. 10.1080/10615809908248327
- Scott, D. J., Stohler, C. S., Egnatuk, C. M., Wang, H., Koeppel, R. A., & Zubieta, J. K. (2007). Individual differences in reward responding explain placebo-induced expectations and effects. *Neuron*, 55(2) 325–336. 10.1016/j.neuron.2007.06.028
- Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the "secretary problem". *Organizational Behaviour and Human Decision Processes*, 69(3), 221–236. 10.1006/obhd.1997.2683
- Seale, D. A., & Rapoport, A. (2000). Optimal stopping behaviour with relative ranks: the secretary problem with unknown population size. *Journal of Behavioural Decision Making*, 13(4), 391–411. 10.1002/1099-0771(200010/12)13:4 < 391::AID-BDM359>3.0.CO;2-I
- Shaw, M. E. (1932). A comparison of individuals and small groups in the rational solution of complex problems. *American Journal of Psychology*, 44(3), 491–504. 10.2307/1415351
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*. 63(2), 129–138. 10.1037/h0042769
- Simonson, I. (1992). The influence of anticipating regret and responsibility on purchase decisions. *Journal of Consumer Research*, 19(1), 105–118. 10.1086/209290

- Skala, D. (2008). Overconfidence in psychology and finance-an interdisciplinary literature review. *Bank i Kredyt*, 4, 33–5. Retrieved from <https://ssrn.com/abstract=1261907>
- Soman, D. (2001). The mental accounting of sunk time costs: Why time is not like money. *Journal of Behavioural Decision Making*, 14(3), 169–185. 10.1002/bdm.370
- Stasser, G., & Birchmeier, Z. (2003). Group creativity and collective choice. In P. B. Paulus & B. A. Nijstad (Eds.), *Group creativity: Innovation through collaboration* (pp. 85–109). New York, NY, US: Oxford University Press. 10.1093/acprof:oso/9780195147308.003.0005
- Stasser, G., & Stewart, D. (1992). Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of Personality and Social Psychology*, 63(3), 426–434. 10.1037/0022-3514.63.3.426
- Sterzer, P., Frith, C., & Petrovic, P. (2008). Believing is seeing: expectations alter visual awareness. *Current Biology*, 18(16), 697–698. 10.1016/j.cub.2008.06.021
- Stoner, J. A. F. (1968). Risky and cautious shifts in group decisions: The influence of widely held values. *Journal of Experimental Social Psychology*, 4(4), 442–459. 10.1016/0022-1031(68)90069-3
- Strayer, D. L., Drews, F. A., & Crouch, D. J. (2006). A comparison of the cell phone driver and the drunk driver. *Human Factors*, 48(2), 11. 10.2139/ssrn.570222
- Sugiyama, L., Tooby, J., & Cosmides, L. (2002). Cross-Cultural Evidence of Cognitive Adaptations for Social Exchange among the Shiwiari of Ecuadorian Amazonia. *Proceedings of the National Academy of Sciences of the United States of America*, 99, 11537–42. 10.1073/pnas.122352999.

- Sümer, N., Özkan, T., & Lajunen, T. (2006). Asymmetric relationship between driving and safety skills. *Accident Analysis and Prevention*, 38(4), 703–711.
10.1016/j.aap.2005.12.016
- Svenson, O. (1981). Are we all less risky and more skilful than our fellow drivers? *Acta Psychologica*, 47(2), 143–148. 10.1016/0001–6918(81)90005–6
- Swann, W. B., & Gill, M. J. (1997). Confidence and accuracy in person perception: Do we know what we think we know about our relationship partners? *Journal of Personality and Social Psychology*, 73(4), 747–757. 10.1037/0022–3514.73.4.747
- Tamaki, M. (1979). A secretary problem with double choices. *Journal of the Operations Research Society of Japan*, 22(4), 257–264.
<https://pdfs.semanticscholar.org/6ab6/4613367117868f0e78a6af058a789b708284.pdf>
- Teger, A. I., & Pruitt, D. G. (1967). Components of group risk taking. *Journal of Experimental Social Psychology*, 3(2), 189–205. 10.1016/0022–1031(67)90022–4
- Thompson, S. C. (1999). Illusions of control: How we overestimate our personal influence. *Current Directions in Psychological Science*, 8(6), 187–190. 10.1111/1467–8721.00044
- Trinugroho, I., & Sembel, R. (2011). Overconfidence and excessive trading behaviour: An experimental study. *Management*, 6(7), 147–152. 10.5539/ijbm.v6n7p147
- Trivedi, M. H., Rush, A. J., Ibrahim, H. M., Carmody, T. J., Biggs, M. M., Suppes, T., Crismon, M. L., Shores–Wilson, K., Toprac, M. G., Dennehy, E. B., Whitte, B. & Kashner, T.M. (2004). The Inventory of Depressive Symptomatology, Clinician

- Rating (IDS–C) and Self–Report (IDS–SR), the Quick Inventory of Depressive Symptomatology, Clinician Rating (QIDS–C) and Self–Report (QIDS–SR) in public sector patients with mood disorders: a psychometric evaluation. *Psychological Medicine*, 34(1), 73–82. 10.1017/S0033291703001107
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453–458. 10.1126/science.7455683
- Vroom, V. H., & Jago, A. G. (1988). *The new leadership: Managing participation in organizations*. Englewood Cliffs, N.J: Prentice–Hall.
- Wallsten, T. S., Budescu, D. V., & Zwick, R. (1993). Comparing the calibration and coherence of numerical and verbal probability judgments. *Management Science*, 39(2), 234–248. 10.1287/mnsc.39.2.176
- Wansink, B., & Cheney, M. M. (2005). Super Bowls: serving bowl size and food consumption. *Journal of the American Medical Association*, 293(14), 1727–1728. 10.1001/jama.293.14.1727
- Wansink, B., van Ittersum, K., & Painter, J. E. (2006). Ice cream illusions: bowls, spoons, and self–served portion sizes. *American Journal of Preventive Medicine*, 31(3), 240–243. 10.1016/j.amepre.2006.04.003
- Warneken, F., Chen, F. & Tomasello, M. (2006). Cooperative activities in young children and chimpanzees. *Child Development*, 77(3), 640–663. 10.1111/j.1467–8624.2006.00895.x
- Warneken, F., & Tomasello, M. (2007). Helping and cooperation at 14 months of age. *Infancy*, 11(3), 271–294. 10.1111/j.1532–7078.2007.tb00227.x

- Wegner, D. M., & Wheathey, I. (1999). Apparent mental causation: sources of the experience of will. *American Psychologist*, *54*(7), 480–492. 10.1037/0003-066X.54.7.480
- Yeo, A. J., & Yeo, G. F. (1994). Selecting satisfactory secretaries. *Australian Journal of Statistics*, *36*(2), 185–198. 10.1111/j.1467-842X.1994.tb00861.x
- Yeo, G. F. (1998). Interview costs in the secretary problem. *Australian & New Zealand Journal of Statistics*, *40*(2), 215–219. 10.1111/1467-842X.00024
- Yerkes, R. M., & Dodson, J. D. (1908). The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology*, *18*, 459–482. 10.1002/cne.920180503
- Zeelenberg, M., & Pieters, R. (2004). Consequences of regret aversion in real life: The case of the Dutch postcode lottery. *Organizational Behavior and Human Decision Processes*, *93*(2), 155–168. 10.1016/j.obhdp.2003.10.001
- Zeelenberg, M., Beattie, J., van der Pligt, J., & de Vries, N. K. (1996). Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behaviour and Human Decision Processes*, *65*(2), 148–158. 10.1006/obhd.1996.0013
- Zeelenberg, M., van Dijk, W. W., & Manstead, A. S. R. (1998). Reconsidering the relation between regret and responsibility. *Organizational Behaviour and Human Decision Processes*, *74*(3), 254–272. 10.1006/obhd.1998.2780

Zwick, R., Rapoport, A., Lo, A. K. C., & Muthukrishnan, A. V. (2003). Consumer Sequential Search: Not Enough or Too Much? *Marketing Science*, 22(4), 503–519.
Retrieved from <http://www.jstor.org/stable/4129735>

APPENDICES

Appendix 2.A

Instructions for the no cost and unannounced time cost treatments (the instructions were identical in both cases; the treatments differed only in the 5-second time delay imposed in the unannounced time cost treatment, which, however, was not announced in the instructions).

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask that you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. You will be given a brief description of the house that will be followed by a series of price offers. The price offers are randomly generated by the computer and available one at a time. Once a price offer is presented, you can either accept or reject it. If you accept the price offer, the house will be sold at the price you accepted. All sales are final. If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. In total there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. *Therefore, make your decisions carefully.*

There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.

Practice scenarios

There will be two practice scenarios. These practice scenarios are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice scenarios.

How payoffs are determined

The payoffs will be denoted in experimental currency units (ECUs).

$$1000 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

$$\begin{aligned} & \underline{\text{Total ECUs you earn}} \\ & = \\ & \underline{\text{Accepted price offer for House 1} + \text{Accepted price offer for House 2} + \dots + \text{Accepted} \\ & \quad \underline{\text{price offer for House 10}}} \end{aligned}$$

Example: Suppose you accepted the price offer 450 for House 1, 260 for House 2, 380 for House 3, ..., 658 for House 10. The total amount of ECUs you earn is 450+260+380+....+658.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice scenarios, with no money payoffs. Then, you will make decisions in 10 scenarios with money payoffs.

Appendix 2.B

Instructions for the announced time cost treatment.

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask that you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. You will be given a brief description of the house that will be followed by a series of price offers. The price offers are randomly generated by the computer and available one at a time. There is a five-second waiting time before presenting each offer. Once a price offer is presented, you can either accept or reject it. If you accept the price offer, the house will be sold at the price you accepted. All sales are final. If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. In total there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. *Therefore, make your decisions carefully.*

There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.

Practice scenarios

There will be two practice scenarios. These practice scenarios are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice scenarios.

How payoffs are determined

The payoffs will be denoted in experimental currency units (ECUs).

$$1000 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

Total ECUs you earn

=

Accepted price offer for House 1 + Accepted price offer for House 2 + ... + Accepted price offer for House 10

Example: Suppose you accepted the price offer 450 for House 1, 260 for House 2, 380 for House 3, ..., 658 for House 10. The total amount of ECUs you earn is 450+260+380+....+658.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice scenarios, with no money payoffs. Then, you will make decisions in 10 scenarios with money payoffs.

Appendix 2.C

Instruction for the monetary cost treatment.

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask that you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. You will be given a brief description of the house that will be followed by a series of price offers, denoted in experimental currency units (ECUs). The price offers are randomly generated by the computer and available one at a time. Once a price offer is presented, you can either accept or reject it. If you accept the price offer, the house will be sold at the price you accepted. All sales are final. If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. In total, there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. For inspecting *each* price offer you will incur a cost of 20 ECUs. *Therefore, make your decisions carefully.*

There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.

Practice scenarios

There will be two practice scenarios. These practice scenarios are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice scenarios.

How payoffs are determined

The payoffs will be denoted in experimental currency units (ECUs).

$$1000 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

Total ECUs you earn

=

(Accepted price offer for House 1 - 20 ECUs * the number of inspected offers for House 1)
+ (Accepted price offer for House 2 - 20 ECUs * the number of inspected offers for House
2) + + (Accepted price offer for House 10 - 20 ECUs * the number of inspected offers
for House 10)

Example: Suppose you accepted the 3rd price offer for House 1 and the price offer has a value of 450, the 4th price offer for House 2 with a value of 260, the 1st offer for House 3 with a value of 500,....., the 5th price offer for House 10 with a value of 658.

The total amount of ECUs you earn is $(450 - 20*3) + (260 - 20*4) + (500 - 20*1) + \dots + (658 - 20*5)$.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice scenarios, with no money payoffs. Then, you will make decisions in 10 scenarios with money payoffs.

Appendix 2.D

The price offer sequences used in Experiment One, Two and Five.

Round \ Offer	1	2	3	4	5	6	7	8	9	10
1	388	739	310	420	292	494	522	252	789	341
2	488	803	290	637	264	225	252	709	829	459
3	683	221	637	727	344	272	562	966	996	453
4	321	729	372	561	266	994	255	885	241	625
5	625	159	619	643	396	602	370	737	799	504
6	744	150	207	663	445	987	292	449	722	387
7	279	299	455	568	266	523	533	910	1088	250
8	848	818	400	636	241	683	237	250	876	308
9	276	585	251	422	370	1400	262	933	503	492
10	678	875	708	336	484	1574	343	491	650	455
11	408	130	452	414	264	1413	220	450	890	353
12	435	795	516	479	186	184	460	394	1264	588
13	679	481	420	332	578	1081	294	899	645	438
14	465	2	607	494	244	558	535	372	1740	408
15	393	525	410	546	189	273	297	505	1179	481
16	397	429	324	724	565	1182	452	608	250	467
17	588	62	214	411	271	305	284	827	840	418
18	358	459	480	267	235	661	436	712	272	273
19	644	748	463	357	350	785	581	838	449	554
20	495	374	617	733	373	89	197	541	105	553

Appendix 2.E

Simulations and Hypotheses

The simulations are run using Processing 3 (version 3.02) software. Processing was initiated by Ben Fry and Casey Reas. It is developed by a small team of volunteers. It is free to download at <https://processing.org/>.

Each simulation cycle generates a set of 20 random offers from a standardized normal distribution, using the mean and standard deviation for each house. A total of 20 stopping rules containing all possible stopping positions (0 to 19) is evaluated. Once a set of the offer has been generated, the rank associated with each offer is calculated. These offers and ranks for each offer are recorded to test each stopping rule. Stopping rule 0 is to accept the first offer regardless of the value; stopping rule 1 means to reject the first offer and choose the next highest offer, stopping rule 2 refers to reject the first two offers and choose the next highest offer and so on. Stopping rule 19 is to accept the 20th/final offer regardless of the value. The simulation runs separately for each house (using the mean and the standard deviation specific for each house) with 2 million iterations.

The simulation yields the average payoffs in ECUs of each stopping rule when there is 0 ECUs and 20 ECUs. The average payoff gives an indication of which stopping rule is the optimal, i.e., yields the highest payoff, when all offers yield a positive payoff. The simulation results of the average payoff show that in 0 ECU, the optimal stopping rule is stopping rule 4, which is to reject the first 20% of offers. The overall payoff starts to decrease if one chooses a decision rule after stopping rule 4. The optimal stopping rule is to accept the first offer when there is 20 ECUs, accept the first offer yields the highest average payoff than all the other stopping rules tested.

Appendix 3.A

Instruction for the house frame treatment.

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask that you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. You will be given a series of price offers for each scenario. The price offers are randomly generated by the computer and available one at a time. Once a price offer is presented, you can either accept or reject it. If you accept the price offer, the house will be sold at the price you accepted. All sales are final. If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. In total there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. *Therefore, make your decisions carefully.*

There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.

Practice scenarios

There will be two practice scenarios. These practice scenarios are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice scenarios.

How payoffs are determined

The payoffs will be denoted in experimental currency units (ECUs).

$$1000 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

Total ECUs you earn

=

Accepted price offer for House 1 + Accepted price offer for House 2 + ... + Accepted price offer for House 10

Example: Suppose you accepted the price offer 450 for House 1, 260 for House 2, 380 for House 3....., 658 for House 10. The total amount of ECUs you earn is 450+260+380+....+658.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice scenarios, with no money payoffs. Then, you will make decisions in 10 scenarios with money payoffs.

Appendix 3.B

Instruction for the no context treatment.

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask that you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 rounds. In each round, you will be asked to decide whether to accept or reject a number. The numbers are randomly generated by the computer and available one at a time. Once a number is presented, you can either accept or reject it. If you accept the number, you receive the amount represented by the number (in experimental currency units, as will be explained below). All decisions are final. If you reject the number, the number will disappear; you cannot go back to the previously rejected number. In total there are 20 numbers available; if you have not accepted a number prior to the 20th number, you will be *forced to accept* the 20th (i.e. the final) number. *Therefore, make your decisions carefully.*

There is no time limit on how long the numbers will be available for, so take as long as you need to evaluate each number.

Practice rounds

There will be two practice rounds. These practice rounds are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice rounds.

How payoffs are determined

The payoffs will be denoted in experimental currency units (ECUs).

$$1000 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

Total ECUs you earn

=

Accepted number for Round 1 + Accepted number for Round 2 + ... + Accepted number for Round 10

Example: Suppose you accepted the number 450 for Round 1, 260 for Round 2, 380 for Round 3....., 658 for Round 10. The total amount of ECUs you earn is 450+260+380+...+658.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice rounds, with no money payoffs. Then, you will make decisions in 10 rounds with money payoffs.

Appendix 3.C

Summary prediction results of three dependent variable after applying the decision rules to the sequences in the experiment. The chosen price (in ECUs) is the sum of 10 price selected by applying each decision rule, the optimal count is the number of rounds selected the optimal price after applying each decision rule and the average stopping position is the nth position obtained by averaging across 10 rounds after applying the decision rules.

Decision rule	Variable	Chosen price (ECUs)	Optimal count (Rounds)	Average stopping position
Pick First Offer		4547	0	1
Choose the next biggest after seeing 1		6629	0	2.6
Choose the next biggest after seeing 2		7352	2	3.7
Choose the next biggest after seeing 3		7228	4	10.3
Choose the next biggest after seeing 4		7562	3	12.4
Choose the next biggest after seeing 5		7611	3	12.5
Choose the next biggest after seeing 6		7754	4	13.1
Choose the next biggest after seeing 7		7930	4	13.6
Choose the next biggest after seeing 8		7634	4	15
Choose the next biggest after seeing 9		7808	5	15.1
Choose the next biggest after seeing 10		5825	3	18.4
Choose the next biggest after seeing 11		5825	3	18.4
Choose the next biggest after seeing 12		6301	4	18.6
Choose the next biggest after seeing 13		6096	3	19.3
Choose the next biggest after seeing 14		4461	2	19.9
Choose the next biggest after seeing 15		4461	2	19.9
Choose the next biggest after seeing 16		4461	2	19.9

Choose the next biggest after seeing 17	4461	2	19.9
Choose the next biggest after seeing 18	4461	2	19.9
Pick Last Offer	4077	1	20

Appendix 3.D

Summary of statistical findings for the comparison between the no context and house frame, no context and house with info, and house frame and house with info treatments.

No context and house frame

Variables	Treatments	Average (Standard Deviation)	Kolmogorov–Smirnov test	
			<i>Z</i>	<i>p – value</i>
Average stopping position of 10 rounds	<i>No context</i>	9.81 (3.55)	1.05	.22
	<i>House Frame</i>	11.24 (2.80)		
Sum of 10 chosen price (round 1-10)	<i>No context</i>	6880.28 (749.99)	1.38	.045
	<i>House Frame</i>	7269.51 (361.37)		

No context and house with info

Variables	Treatments	Average (Standard Deviation)	Kolmogorov–Smirnov test	
			<i>Z</i>	<i>p – value</i>
Average stopping position of 10 rounds	<i>No context</i>	9.81 (3.55)	1.39	.04
	<i>House with info</i>	11.58 (2.77)		
Sum of 10 chosen price (round 1-10)	<i>No context</i>	6880.28 (749.99)	1.73	.005
	<i>House with info</i>	7338.85 (580.71)		

HouseFrame and house with info

Variables	Treatments	Average (Standard Deviation)	Kolmogorov–Smirnov test	
			<i>Z</i>	<i>p – value</i>
Average stopping position of 10 rounds	<i>House Frame</i>	11.24 (2.80)	.64	.80
	<i>House with info</i>	11.58 (2.77)		
Sum of 10 chosen price (round 1-10)	<i>House Frame</i>	7269.51 (361.37)	.95	.33
	<i>House with info</i>	7338.85 (580.71)		

Appendix 4.A

General Knowledge and Probability Evaluation Task

Please choose one answer for each question, and move the pointer to indicate the percentage of self-assessed accuracy

*Correct answer is in bold letter

Easy Questions

1 The capital of New Zealand is...?

1. **Wellington** 2. Auckland 3. Christchurch

What is the probability that you answered question 1 accurately?

33% 100%

2 Who is the current Prime Minister of New Zealand?

1. Helen Clark 2. **John Key** 3. David Cunliffe

What is the probability that you answered question 2 accurately?

33% 100%

3 Which city in New Zealand recently suffered a series of devastating earthquakes?

1. Wellington 2. Rotorua 3. **Christchurch**

What is the probability that you answered question 3 accurately?

33% 100%

4 The highest point in New Zealand is ...?

1. **Mount Cook** 2. Mount Tasman 3. Mount Hamilton

What is the probability that you answered question 4 accurately?

33% 100%

- 5 What is the third official language in New Zealand?
 1. English **2. New Zealand Sign Language** 3. Maori
 What is the probability that you answered question 5 accurately?
 33% 100%
- 6 What is the name of the national cricket team?
 1. Black Stars 2. All Blacks **3. The Black Caps**
 What is the probability that you answered question 6 accurately?
 33% 100%
- 7 Who is the director of the Lord of the Rings trilogy
 1. Tim Burton **2. Peter Jackson** 3. Oliver Stone
 What is the probability that you answered question 7 accurately?
 33% 100%
-

Medium Questions

- 8 The largest lake in NZ is ...?
1. Lake Taupo 2. Lake Victoria 3. Lake Wakatipu
 What is the probability that you answered question 8 accurately?
 33% 100%
- 9 The first capital of New Zealand is ?
 1. Auckland **2. Russell** 3. Wellington
 What is the probability that you answered question 9 accurately?
 33% 100%
- 10 The first European to discover New Zealand is...?
 1. Christopher Columbus 2. James Cook **3. Abel Tasman**
 What is the probability that you answered question 10 accurately?
 33% 100%

Appendix 4.B

The Self-efficacy scale

Please read each statement carefully, and circle just one answer per question that best describes you.

- 1) I can always manage to solve difficult problems if I try hard enough.
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 2) If someone opposes me, I can find the means and ways to get what I want.
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 3) It is easy for me to stick to my aims and accomplish my goals.
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 4) I am confident that I could deal efficiently with unexpected events.
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 5) Thanks to my resourcefulness, I know how to handle unforeseen situations.
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 6) I can solve most problems if I invest the necessary effort.
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 7) I can remain calm when facing difficulties because I can rely on my coping abilities
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 8) When I am confronted with a problem, I can usually find several solutions.
Not at all true Hardly true Moderately true Exactly true
1 **2** **3** **4**
- 9) If I am in trouble, I can usually think of a solution.

Not at all true	Hardly true	Moderately true	Exactly true
1	2	3	4

10) I can usually handle whatever comes my way.

Not at all true	Hardly true	Moderately true	Exactly true
1	2	3	4

Appendix 4.C

The depression scale

Please read each statement carefully, and circle one response to each question that best describes you for the past 7 days. During the Past 7 Days...

1. Falling asleep

A - I never take longer than 30 min to fall asleep.

B - I take at least 30 min to fall asleep, less than half the time.

C - I take at least 30 min to fall asleep, more than half the time.

D - I take more than 60 min to fall asleep, more than half the time.

2. Sleep during the night

A - I do not wake up at night.

B - I have a restless, light sleep with a few brief awakenings each night.

C - I wake up at least once a night, but I go back to sleep easily.

D - I awaken more than once a night and stay awake for 20 min or more, more than half the time.

3. Waking up too early

A - Most of the time, I awaken no more than 30 min before I need to get up.

B - More than half the time, I awaken more than 30 min before I need to get up.

C - I almost always awaken at least 1 hour or so before I need to, but I go back to sleep eventually.

D - I awaken at least 1 hour before I need to, and can't go back to sleep.

4. Sleeping too much

A - I sleep no longer than 7 – 8 hours/night, without napping during the day.

B - I sleep no longer than 10 hours in a 24-hour period including naps.

C - I sleep no longer than 12 hours in a 24-hour period including naps.

D - I sleep longer than 12 hours in a 24-hour period including naps.

5. Feeling sad

A - I do not feel sad.

B - I feel sad less than half the time.

C - I feel sad more than half the time.

D - I feel sad nearly all of the time.

Please complete either 6 or 7 (**Not Both**)

6. Decreased appetite

A - There is no change in my usual appetite.

B - I eat somewhat less often or lesser amounts of food than usual.

C - I eat much less than usual and only with personal effort.

D - I rarely eat within a 24-hour period, and only with extreme personal effort or when others persuade me to eat.

-Or-

7. Increased appetite

A - There is no change from my usual appetite.

B - I feel a need to eat more frequently than usual.

C - I regularly eat more often and/or greater amounts of food than usual.

D - I feel driven to overeat both at mealtime and between meals.

Please complete either 8 or 9 (**Not both**)

8. Decreased weight (Within the Last 2 Weeks)

A - I have not had a change in my weight.

B - I feel as if I've had a slight weight loss.

C - I have lost 2 pounds or more.

D - I have lost 5 pounds or more.

-Or-

9. Increased weight (Within the last 2 weeks)

A - I have not had a change in my weight.

B - I feel as if I've had a slight weight gain.

C - I have gained 2 pounds or more.

D - I have gained 5 pounds or more.

10. Concentration/decision making

A - There is no change in my usual capacity to concentrate or make decisions.

B - I occasionally feel indecisive or find that my attention wanders.

C - Most of the time, I struggle to focus my attention or to make decisions.

D - I cannot concentrate well enough to read or cannot make even minor decisions.

11. View of myself

A - I see myself as equally worthwhile and deserving as other people.

B - I am more self-blaming than usual.

C - I largely believe that I cause problems for others.

D - I think almost constantly about major and minor defects in myself.

12. Thoughts of death or suicide

A - I do not think of suicide or death.

B - I feel that life is empty or wonder if it's worth living.

C - I think of suicide or death several times a week for several minutes.

D - I think of suicide or death several times a day in some detail, or I have made specific plans for suicide or have actually tried to take my life.

13. General interest

A - There is no change from usual in how interested I am in other people or activities.

B - I notice that I am less interested in people or activities.

C - I find I have interest in only one or two of my formerly pursued activities.

D - I have virtually no interest in formerly pursued activities.

14. Energy Level

A - There is no change in my usual level of energy.

B - I get tired more easily than usual.

C - I have to make a big effort to start or finish my usual daily activities (for example, shopping, homework, cooking or going to work).

D - I really cannot carry out most of my usual daily activities because I just don't have the energy.

15. Feeling Slowed Down

A - I think, speak, and move at my usual rate of speed.

B - I find that my thinking is slowed down or my voice sounds dull or flat.

C - It takes me several seconds to respond to most questions and I'm sure my thinking is slowed.

D - I am often unable to respond to questions without extreme effort.

16. Feeling Restless

A - I do not feel restless.

B - I'm often fidgety, wring my hands, or need to shift how I am sitting.

C - I have impulses to move about and am quite restless.

D - At times, I am unable to stay seated and need to pace around.

Appendix 4.D

The demographic questionnaire

Please read each statement carefully and circle one most appropriate answer per question for you.

1) Have you ever purchased a house before?

Yes No

2) Have you ever hired an employee before?

Yes No

3) Have you ever sold a house before?

Yes No

4) Please indicate the price range of the house you expected to own ultimately (ultimately refers to after some post-graduation period of career development)

Under \$350,000	\$350,001- \$500,000	\$500,001 - \$650,000
\$650,001 - \$800,000	\$800,001- \$950,000	Above \$950,001

5) Please indicate the annual salary range of the job you expected to get ultimately (ultimately refers to after some post-graduation period of career development)

Under \$20,000	\$20,001 – \$40,000	\$40,001- \$60,000
\$ 60,001 - \$80,000	\$80,001- \$100,000	Over \$100,001

6) What is your age?

Under 20	21 - 30	31 - 40
41 - 50	51 - 60	Above 61

7) What is your gender?

Male Female

8) What is your annual income?

A) Under \$20,000

B) \$20,001- \$40,000

C) \$40,001- \$60,000

D) \$60,001- \$80,000

E) \$80,001- \$100,000

F) Over \$100,001

Appendix 6.A

The commission base treatment instruction with the house selling frame.

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash, and 2 course credits at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. You will be given a brief description of the house that will be followed by a series of price offers. The price offers are randomly generated by the computer and available one at a time. Once a price offer is presented, you can either accept or reject it. If you accept the price offer, you receive the amount represented by the offer (in experimental currency units, as will be explained below). If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. All decisions are final. In total there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. *Therefore, make your decisions carefully.*

There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.

Practice scenarios

There will be two practice scenarios. These practice scenarios are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice scenarios.

How payoffs are determined

The payoffs will be denoted in experimental currency units (ECUs).

$$735 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

Total ECUs you earn

=

Accepted price offer for House 1 + Accepted price offer for House 2 + ... + Accepted price offer for House 10

Example: Suppose you accepted the price offer 450 for House 1, 260 for House 2, 380 for House 3....., 658 for House 10. The total amount of ECUs you earn is 450+260+380+.... +658.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice scenarios, with no money payoffs. Then, you will make decisions in 10 scenarios with money payoffs.

Appendix 6.B

The best only treatment instruction with the house selling frame.

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash, and 2 course credits at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. You will be given a brief description of the house that will be followed by a series of price offers. The price offers are randomly generated by the computer and available one at a time. Once a price offer is presented, you can either accept or reject it. If you accept a price offer and that offer is the highest of the 20 price offers; you will receive NZD 4.60. Otherwise you receive nothing from that scenario. If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. All decisions are final. In total there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. *Therefore, make your decisions carefully.*

There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.

Practice scenarios

There will be two practice scenarios. These practice scenarios are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice scenarios.

How payoffs are determined

You will earn NZD 4.60 if you have chosen the highest price offer for each house.

Your payoffs are determined as follows:

Total NZD you earn

= number of houses that you have selected the highest price offer * NZD 4.60

Example: Suppose you accepted the highest price offer for House 1, House 3 and House 10. The total amount of NZD you earn is $\$4.60 * 3 = 13.80$.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice scenarios, with no money payoffs. Then, you will make decisions in 10 scenarios with money payoffs.

The no incentive treatment instruction with the house selling frame

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you will be paid \$9.50 in cash and 2 course credits at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 scenarios, in which you will be selling houses. In each scenario, you will be asked to decide whether to accept or reject a price offer for a particular house. Your aim is to try to accept the highest price you can. You will be given a brief description of the house that will be followed by a series of price offers. The price offers are randomly generated by the computer and available one at a time. Once a price offer is presented, you can either accept or reject it. If you accept the price offer, you will move on to the next scenario. If you reject the price offer, the offer will disappear; you cannot go back to the previously rejected offer. In total there are 20 price offers available for each house; if you have not accepted an offer prior to the 20th offer, you will be *forced to accept* the 20th (i.e. the final) offer. *Therefore, make your decisions carefully.*

There is no time limit on how long the price offers will be available for, so take as long as you need to evaluate each offer.

Practice scenarios

There will be two practice scenarios. These practice scenarios are there to help you become familiar with the software.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice scenarios. Then, you will make decisions in 10 scenarios.

Appendix 6.D

The commission base treatment instruction with the no context frame.

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash, and 2 course credits at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 rounds. In each round, you will be asked to decide whether to accept or reject a number. The numbers are randomly generated by the computer and available one at a time. Once a number is presented, you can either accept or reject it. If you accept the number, you receive the amount represented by the number (in experimental currency units, as will be explained below). If you reject the number, the number will disappear; you cannot go back to the previously rejected number. All decisions are final. In total there are 20 numbers available; if you have not accepted a number prior to the 20th number, you will be *forced to accept* the 20th (i.e. the final) number. *Therefore, make your decisions carefully.*

There is no time limit on how long the numbers will be available for, so take as long as you need to evaluate each number.

Practice rounds

There will be two practice rounds. These practice rounds are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice rounds.

How payoffs are determined

The payoffs will be denoted in experimental currency units (ECUs).

$$735 \text{ ECUs} = 1 \text{ NZD}$$

Your ECUs will be converted into NZD at this rate, and you will be paid in NZD when you leave the lab. The more ECUs you earn, the more NZD you earn.

Your payoffs are determined as follows:

Total ECUs you earn

=

Accepted number for Round 1 + Accepted number for Round 2 + ... + Accepted number for Round 10

Example: Suppose you accepted the number 450 for Round 1, 260 for Round 2, 380 for Round 3....., 658 for Round 10. The total amount of ECUs you earn is 450+260+380+...+658.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice rounds, with no money payoffs. Then, you will make decisions in 10 rounds with money payoffs.

The best only treatment instruction with the no context frame

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you may earn a considerable amount of money which will be paid to you in cash, and 2 course credits at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 rounds. In each round, you will be asked to decide whether to accept or reject a number. The numbers are randomly generated by the computer and available one at a time. Once a number is presented, you can either accept or reject it. If you accept the number and that number is the highest of the 20 numbers, you will receive NZD4.60. Otherwise you receive nothing from that round. If you reject the number, the number will disappear; you cannot go back to the previously rejected number. All decisions are final. In total there are 20 numbers available for each round; if you have not accepted a number prior to the 20th number, you will be *forced to accept* the 20th (i.e. the final) number. *Therefore, make your decisions carefully.*

There is no time limit on how long the numbers will be available for, so take as long as you need to evaluate each number.

Practice rounds

There will be two practice rounds. These practice rounds are there to help you become familiar with the software. You will not be paid for the decisions you make in these two practice rounds.

How payoffs are determined

You will earn NZD 4.60 if you have chosen the highest number for each round

Your payoffs are determined as follows:

Total NZD you earn

= number of rounds that you have selected the highest number * NZD 4.60

Example: Suppose you accepted the highest number for Round 1, Round 3 and Round 10. The total amount of NZD you earn is $4.60 * 3 = 13.80$.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice rounds, with no money payoffs. Then, you will make decisions in 10 rounds with money payoffs.

The no incentive treatment instruction with the no context frame

GENERAL INSTRUCTIONS

Overview

You are about to participate in a decision-making experiment. If you follow these instructions carefully you will be paid \$9.50 in cash and 2 course credits at the end of the experiment. If you have a question at any time, please raise your hand and the experimenter will approach you and answer your question in private. We ask you not to talk otherwise during the experiment. Also, please turn off your cell phone and do not use the computer for any other purpose than your participation in the experiment requires. If you break these rules, we will have to exclude you from the experiment and from all payments.

INSTRUCTIONS

You will participate in 10 rounds. In each round, you will be asked to decide whether to accept or reject a number. Your aim is to try to accept the highest number you can. The numbers are randomly generated by the computer and available one at a time. Once a number is presented, you can either accept or reject it. If you accept the number (in experimental currency units, as will be explained below), you will move on to the next round. If you reject the number, the number will disappear; you cannot go back to the previously rejected number. In total there are 20 numbers available for each round; if you have not accepted a number prior to the 20th number, you will be *forced to accept* the 20th (i.e. the final) number. *Therefore, make your decisions carefully.*

There is no time limit on how long the numbers will be available for, so take as long as you need to evaluate each number.

Practice rounds

There will be two practice rounds. These practice rounds are there to help you become familiar with the software.

Do you have any questions?

You are now ready to begin the experiment. First, we will conduct two practice rounds. Then, you will make decisions in 10 rounds.