Essays on search behaviour and procedural fairness

A dissertation
presented to the School of Economics of
the University of East Anglia
in candidacy for the degree of
Doctor of Philosophy
by

Mengjie Wang

University of East Anglia
School of Economics

Date of submission: September 2017

Abstract

This thesis consists of three studies that relate to search behaviour and procedural fairness. Chapter 2 investigates experimentally whether the search-deterring effect of time-limited offers is intensified by behavioural factors – specifically, feedback-conditional regret, reduced decision quality due to time pressure, and aversion to small-scale risk. The conclusion is that the search-deterring effect is intensified, particularly (and surprisingly) when consumers are not subject to high time pressure. There is no evidence of regret effects. Overall, individuals show aversion to small-scale risk. Chapter 3 proposes a new concept of fairness: strategy fairness. The conjecture is that inequalities will tend to be seen as acceptable if they come about through the workings of fair rules, even though they are the result of self-interested intentions. A model of strategy fairness is provided to show how the concept of strategy fairness can be incorporated into a more complete model. The concept of strategy fairness is tested using an experiment. It turns out that subjects are more willing to accept inequalities that are the result of fair procedures. The surprising result emerges that procedural unfairness makes both disadvantaged and advantaged players more likely to take. Chapter 4 introduces a new search competition game. The search competition takes the form of parallel searches without recall. This is related to two theoretical and experimental literatures: contest and search. However, no work has been done on analysing this type of game. A theoretical analysis of the search competition game is provided. Subjects' actual play in the experiment is compared with both the subgame perfect Nash equilibrium solution and the empirical best response. It shows that, relative to the implications of a rational-choice analysis, subjects tend to search too little.

Acknowledgements

Working as a PhD student at the University of East Anglia was a magnificent as well as challenging experience to me. In all these years, many people have helped me in shaping up my academic career. Here is a small tribute to all those people.

First and foremost, I would like to express my sincere gratitude to my supervisors Prof. Robert Sugden and Prof. Daniel John Zizzo for introducing me to the amazing world of research, and for the patient guidance, encouragement and advice that they have provided throughout my time as their student. I have been extremely lucky to have two outstanding supervisors who cared so much about my work, who responded to my questions and enquires so promptly, and who listened to my ideas so patiently.

I would like to give thanks to Prof. Peter Moffatt for providing me with valuable help and suggestions on econometrics for Chapter 2 and Chapter 4.

I would like to thank the staff of the School of Economics at the UEA, the Centre for Competition Policy (CCP), and the Centre for Behavioural and Experimental Social Sciences (CBESS) for their valuable comments and support throughout the course of conducting research at the UEA.

I also wish to thank my numerous fellow PhD students of the School of Economics at the UEA for accompanying me during all those years and helping me to enjoy my life besides work, as well as for the peer pressure from them which kept me working everyday.

Finally, I want to say thanks to my family, especially my parents. Without their love and support over the years, none of this would have been possible. They have always been there for me and I am thankful for everything they have helped me achieve.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

During my training to be a behavioural and experimental economist, I became interested in topics on search behaviour and procedural fairness in competition. My thesis is developed to make contributions to these two research areas both theoretically and experimentally.

The second chapter[1] of the thesis investigates whether the search-deterring effects of time-limited offers are intensified or mitigated by behavioural factors. Time-limited offers have become one of the most frequently applied sales tactics worldwide. Once presented a time-limited offer, a buyer needs to make a quick decision on whether to take it or leave it, as it will be withdrawn unless accepted immediately. Time-limited offers discourage consumer search by making it more costly to search for alternatives. Previous studies have found that the presence of search costs makes markets less competitive and induces higher prices (Diamond, 1971; Salop and Stiglitz, 1977). Armstrong and Zhou (2016) present a range of theoretical models in which profit-maximising monopolistic or oligopolistic firms use time-limited offers to deter search. They find that time-limited offers have similarly anti-competitive effects, i.e. increasing prices and/or reducing the degree of match between consumers and products. In Armstrong and Zhou's models, consumers are assumed to be rational and risk-neutral who choose strategies that maximise their expected utility. It is natural to consider whether the search-deterring effects that their models describe are intensified or mitigated by properties of consumers' choice behaviour that have been identified by behavioural research but which do not feature in traditional decision theory.

Chapter 2 is concerned with three behavioural mechanisms that might be expected to make real consumers more likely than their counterparts in models of rational risk-neutral search to choose time-limited offers – anticipated regret, time pressure, and extreme risk aversion. First, a consumer who rejects a time-limited offer and continues to search may find that the rejected offer was in fact the best available, and this may induce painful feelings of regret. A

[1] Chapter 2 is based on the working paper Sugden, Wang, and Zizzo (2015). I have made significant contributions to this working paper, which includes having the original research idea, working on the experimental design, experimental implementation, data analysis and writing up the original draft.

decision to accept the offer and stop searching can be a method of avoiding feedback that could cause regret. Second, consumers' choices between accepting and rejecting time-limited offers may be affected by the pressure of having to make a decision in a short period of time. The psychological literature suggests that people may respond to this pressure by using simplifying decision heuristics, such as relying on the 'affect heuristic' (Finucane et al. 2000), attending more to negative dimensions of gambles (Ben Zur and Breznitz 1981), focusing on less information (Payne et al 1988), or searching less in search tasks (Ibanez et al 2009). This body of evidence suggests that time pressure might make consumers more likely to accept time-limited offers. Third, experimental evidence often reveals high degrees of risk aversion in decisions that involve very small stakes. If individuals are significantly averse to small risks, time-limited offers may be more attractive to consumers than is implied by models of risk-neutral search.

The experiment contains two parts. In the first part of the experiment, subjects face simple and intuitive problems of price search, in which time-limited offers appear occasionally and without warning. In the second part, subjects face what appear to be completely different binary choice problems between lotteries and certainties, but which in fact are reframings of choices about time-limited offers that faced by subjects in price search tasks.

I find no evidence that the tendency for consumers to choose time-limited offers is intensified by desires to avoid regret. Surprisingly, time-limited offers are less likely to be chosen at high levels of time pressure. Overall, individuals show aversion to small-scale risk; this is stronger in price search than lottery choice. Allowing for this, at the individual level, and particularly at high levels of time pressure, responses to the time-limited offer and lottery tasks are very consistent.

Chapter 3 introduces a new concept of fairness: strategy fairness. Conventional economic theories assume that people are self-interested and make choices that maximize their own monetary payoffs. However, many studies have shown that theories with this self-interest assumption are not always sufficient to explain people's decisions in the real world. In many real world situations, social preferences have a significant impact on individual decisions. Both theoretical literature and experimental evidence suggest that people not only care about their own income, but are also concerned about the distribution of payoffs, intentions signalled by other people's actions, entitlements, desert, randomisation of the procedure which determines the payoff distribution, and social welfare.

Although the existing theories of social preferences can solve many seemingly puzzling behaviours for which conventional economic theories cannot give an explanation, there are still some behaviours that cannot be fully explained. For instance, Isoni et al. (2014) find that in their bargaining games, subjects are more willing to settle on efficient but unequal allocations rather than equal but inefficient allocations when they have equal opportunity to compete. In fact, there are many real world competitions in which participants are willing to take the unequal outcomes even if all participants reveal self-interest intentions of trying to win. The literature of social preferences has neglected a type of procedural fairness that may be particularly important in market environments and in public choice: fairness of a framework of rules within which individuals pursue self-interest. To fill in this gap of literature on social preferences, I propose a new concept of fairness: strategy fairness. This is tested using an experiment.

I provide a model of strategy fairness. In the model, strategy fairness is defined as equalities of opportunities between players in the competition. The theoretical analysis and its implication give suggestions about how to incorporate the concept of strategy fairness into a more complete model.

The design of the experiment was partially inspired by the price search task used in Chapter 2. In the first part of the experiment, pairs of subjects play a simple card game that has some similarities with the time-limited offer task in the first experiment. Each player is dealt a numbered card and is allowed a fixed number of opportunities to reject it and be dealt another card. The player who ends up with the highest-numbered card wins. The number of replacement opportunities may be the same for both players (procedural fairness) or different (unfairness). In the second stage, each subject is endowed with a number of lottery tickets; the winner receives three times as many as the loser. The same pair of subjects then play a 'vendetta game' in which they have alternating opportunities to take tickets from one another; in each taking move, the taker gains one ticket but the other subject loses three. This game is based on a design used by Bolle et al. (2014), but changes were applied to make the game more intuitive for subjects and the interfaces more engaging. The idea of the experiment is to measure players' attitudes to inequality by the propensity for them to 'take'.

I find that subjects are more willing to accept inequalities that are the result of fair procedures; the fair version of the card game induces roughly the same degree of acceptance of inequality as a control treatment in which inequality is generated by a real effort task. A surprising result emerges when the propensity to take is measured separately for the two players. It turns out that procedural unfairness makes both advantaged and disadvantaged players more willing to take.

In Chapter 3, the card game is used as a simple game to create equality in the first stage of the experiment. In fact, the card game is a search competition game which has not been investigated. Chapter 4 provides a detailed analysis of play in the game in Chapter 3. The search competition game is similar to games in contest literature, as in both games players compete for a fixed prize. The competition takes the form of parallel searches without recall. There are many real world examples of this kind of situation, such as competing firms searching for managers. Surprisingly, although search without recall as a problem for an individual agent is often investigated in the search literature, no previously study has been done on this type of game.

Existing research on contests and search behaviour suggests opposite intuitions about actual behaviour in the search competition, relative to Nash equilibrium predictions. Overbidding is one of the main phenomena observed in contest experiments, i.e. the average effort level is significantly higher than the risk-neutral Nash equilibrium prediction (Dechenaux et al. 2015. It suggests that contestants in contests may try too hard to win. Experimental evidence on search problems shows that people tend to search too little compared to the risk neutral prediction of the optimal strategy, which suggests that people may not try hard enough to find the best offer. The experiment in Chapter 4 is designed to find out which of these intuitions applies to the search competition.

A theoretical analysis of the search competition game is provided. The subgame perfect Nash equilibrium solution to the card game is a cut-off value for each card, which depends on the replacement opportunity (e.g. first, second, third for a player with three opportunities), the total number of replacement opportunities, and the number of replacement opportunities for the co-participant. I identify the subgame perfect Nash equilibrium solution to card games for a range of parameters. The actual play is analysed econometrically and compared with both the subgame perfect Nash equilibrium prediction and the empirical best response. The estimation of actual play is econometrically difficult as I cannot observe the distribution of cut-offs directly from the sample.[2] The data only tells whether the subject accepted or rejected a given card. The decisions on second and third cards can be observed only if the subject rejected the previous cards.

Relative to the implications of a rational-choice analysis, the results show that subjects have a strong bias against replacing cards. Because payoffs are tickets in binary lotteries, this effect is not simple risk aversion. It is interestingly analogous with the bias that I found in choices about time-limited offers.

In the process of developing my thesis, I learned how to design experiments to which subject can relate intuitively. The design principle for experiments in all my three chapters was not only to mimic situations described in theories, but also to create problems in which the subject has a real experience. My experiments are simplified representations of real world situations. At the same time, they provide subjects with all the information that, according to the theory, they need for the problem, such as numerical probabilities. I aimed to make my experimental tasks simple, easy to understand, engaging, and with vivid interfaces. For instance, the price search task in chapter 2 creates a price search problem which is similar to the situation faced by consumer. In the card tasks which are used in chapter 3 and chapter 4, subjects can feel vividly the influence of having different numbers of replacement opportunities on their probability of winning. In the vendetta game, subjects take lottery tickets from one another rather than (as in the original version of the game) changing one another's probabilities of winning. This gave subjects a vivid sense of the cost of stealing behaviour, which was represented by the waste of lottery tickets.

My thesis opens new avenues for exploration. First, one of the surprising findings in chapter 2

[2]Prof. Peter Moffatt provided suggestions for the creation of the econometric model.

is the degree of consistency between consumers' decisions about time-limited offers, made even under intense time pressure, and their responses to problems of binary choice between lotteries. The context-dependence of preferences is a recurring theme in behavioural economics, but this is a case in which the context-independence of preferences require a psychological explanation, and its implications for consumer decision making need thinking through. Second, chapter 3 reveals that procedural unfairness tends to make both advantaged players and disadvantaged players more willing to take. As the experiment is designed to mainly investigate people's willingness to accept inequalities that result from fair/unfair competition, I focus more on the behaviour of disadvantaged players. Therefore, the experiment does not provide sufficient data to conduct more in-depth analysis of advantaged players' behaviour. In future studies, it would be interesting to explore the behaviour of advantaged players in the competitions. Third, in chapter 4, I compare subjects' decisions not only with the subgame perfect Nash equilibrium solutions, but also with the empirical best response. The result shows that subjects' choices are more close to the prediction of the empirical best response than to the subgame perfect Nash equilibrium. Many studies show that Perfect Bayesian Equilibrium solutions are often inaccurate descriptions of behaviour in games, as they are derived under the assumption that all other players obey a given model solution. Compared to Perfect Bayesian Equilibrium, empirical best response provides a better measurement of success of social learning, as it reflects the true behaviour of other players in the same situation. In future studies that measure the success of individual reasoning in games, it would be useful to compare actual play with empirical best response.

Chapter 2

Take it or leave it: experimental evidence on the effect of time-limited offers on consumer behaviour

2.1 Introduction

A common tactic by salespeople is to make offers that (it is claimed) will be withdrawn unless accepted immediately. Familiar examples include the doorstep seller who claims that he is currently 'in the area' but will not be returning, the telephone seller who makes a 'special offer' that can be accepted only during that phone call, the internet site which offers a buy-now discount, and the seller of a used car who claims that another buyer has shown great interest in it and will be returning shortly. Writers who have infiltrated businesses report that sales staff are routinely instructed to use such tactics (Cialdini, 2003, p. 208; Bone, 2006, p. 71-73). There is growing concern that, by using cookies to track the identities of website visitors, internet sellers may remove low-price offers between a potential buyer's first and second visit. Viewed in the perspective of competition regulation, such *time-limited* (or *exploding*) *offers* create barriers to the search processes by which consumers compare prices.[1]

It has long been known that the presence of search costs makes markets less competitive and induces higher prices (Diamond, 1971; Salop and Stiglitz, 1977). Armstrong and Zhou (2016) present a range of theoretical models in which opportunities to make time-limited

[1]Evidence about these sales strategies is reviewed by Office of Fair Trading (2010) and Armstrong and Zhou (2016). Armstrong and Zhou refer to current controversies about the alleged misuse of cookies by websites selling airline tickets, but take no view about the truth of these allegations - perhaps because it is not clear that sellers would benefit by imposing time limits that are not announced in advance.

offers have similarly anti-competitive effects. In these models, profit-maximising monopolistic or oligopolistic firms use time-limited offers to deter search, with the result that there is an increase in prices and/or a reduction in the degree of match between consumers' preferences and the products they buy. For a time-limited offer to have these *search-deterring* effects, consumers' opportunities to investigate alternatives during the offer period must be significantly constrained, which will typically be the case only if that period is quite short and is determined separately for each consumer. In this paper, we will be concerned with time-limited offers that have these properties.[2]

In Armstrong and Zhou's models, consumers are risk-neutral agents who choose search strategies that maximise expected utility, given knowledge of the distributions from which prices are drawn. As Armstrong and Zhou (2016, p. 50) point out, it would be useful to investigate whether the search-deterring effects that their models describe are intensified or mitigated by properties of consumers' choice behaviour that have been identified by behavioural research but which do not feature in traditional decision theory. Given the multiplicity of potentially relevant behavioural mechanisms that need to be disentangled in such an investigation, the methodology of controlled experiments has obvious advantages. We report an experiment that was designed with that objective.

We know of two previous experimental investigations of the effects of time-limited offers on buyers' search behaviour. Huck *et al.* (2010) report an experiment in which subjects played the role of consumers buying a homogenous product sold by two shops at prices that were independently drawn from a common distribution. Each visit to a shop incurred a cost. In the baseline condition, each shop quoted a simple per-unit price, which was revealed only when the shop was visited. These were *free-recall* offers – that is, a consumer who returned to a previously-visited shop would find the original price (and no other) still available. Five other conditions represented different and more complex 'price frames'. In the 'time-limited offer' frame, the offer at the first shop visited had to be accepted or rejected at that visit; the second shop's offer was made with free recall. If the consumer returned to the first shop, a different price would be generated (also with free recall). The experiment found oversearch in the baseline condition (i.e., the frequency with which the first offer was accepted was lower than would have been the case for rational risk-neutral agents) but undersearch in the time-limited offer frame. This observation suggests that time-limited offers may be more attractive to real consumers than to their rational counterparts. However, Huck and Wallace (p. 79) note that many subjects behaved as if they did not realise that the first shop's second offer could be lower than its initial time-limited offer. This misunderstanding may have induced a bias towards the choice of time-limited offers.

[2] A contrasting type of time-limited offer is exemplified by special promotions in supermarkets and by 'sales' periods in department stores. Such offers, which are made available to all consumers for periods of several days or even weeks, may be components of a strategy of varying prices over time in order to discriminate between consumers with different propensities to search (Varian, 1980), to manipulate consumers' reference points (Heidhues and Kőszegi, 2014), or to counteract consumers' tendencies to procrastination (O'Donoghue and Rabin, 2001).

Brown *et al.* (2014) report an experiment that implemented a variant of the Armstrong and Zhou (2016) duopoly model, but using incentivised human subjects to represent firms. Firms could choose whether or not to make their offers time-limited or free-recall. In one treatment, consumers were represented by human subjects; in another, they were computer programs that implemented optimal search strategies. The experiment found oversearching: at the first firm visited, time-limited offers were chosen *less* frequently by human subjects than by optimal search programs. Brown et al.'s interpretation of this result is that that buyers recognise and dislike search-deterring tactics when used by sellers, and express this dislike by rejecting even relatively good time-limited offers (p. 24).

Our investigation is concerned with three behavioural mechanisms that do not depend on consumers' interpretations of sellers' intentions, and that might be expected to favour the choice of time-limited offers. Each of these mechanisms makes use of the fact that choosing between accepting and rejecting a time-limited offer is a choice between a certainty (paying the known offer price) and an uncertain prospect (continuing to search, without the option of recall). In the Armstrong and Zhou models, rational buyers are represented as risk-neutral. Under plausible assumptions about utility functions and if the value of the relevant purchase is low relative to a consumer's wealth, risk-neutral models closely approximate the implications of expected utility theory. However, there are a number of behavioural reasons for expecting consumers' responses to time-limited offers to be significantly risk-averse, favouring the choice of such offers. First, a consumer who rejects a time-limited offer and continues to search may find that the rejected offer was in fact the best available, and this may induce painful feelings of regret. Choosing the offer and terminating the search process can be a method of avoiding regret. Second, a decision to accept or reject a time-limited offer may be made under time pressure, and consumers may respond to this pressure by using simplifying decision heuristics that favour certainties. Third, experimental evidence often reveals high degrees of risk aversion in decisions that involve very small stakes.

We begin by explaining these mechanisms, outlining the basic principles of our experimental design and stating the main hypotheses that we will test (Section 2.2). We then describe the implementation of our experiment in more detail (Section 2.3), report our results (Section 2.4), and discuss their implications (Sections 2.5 and 2.6). To anticipate our conclusions, we find no evidence that the tendency for consumers to choose time-limited offers is intensified by desires to avoid regret. Surprisingly, time pressure made time-limited offers *less* likely to be chosen. However, our subjects' behaviour was predominantly risk-averse. In the absence of time pressure, choices between accepting or rejecting time-limited offers were systematically more risk-averse than when the same decision problems were framed as choices between lotteries. The implication is that rationality-based models understate the attractive force exerted by time-limited offers.

2.2 Basic principles of experimental design and hypotheses to be tested

Our experiment had two parts. In Part 1, each subject faced thirty *price search* tasks, of which fifteen (randomly positioned in the series of tasks, and not announced in advance) involved time-limited offers. In Part 2, each subject faced fifteen *lottery* tasks requiring binary choices between lotteries. In the price search tasks, the price offers between which subjects had to choose were random draws, made independently for each subject, for each task and for each offer, from pre-specified distributions which remained constant across the thirty tasks. Lottery tasks were constructed separately for each subject, to match the specific distributions of offers that that subject had faced in the fifteen price search tasks with time-limited offers. Viewed in the framework of expected utility theory, each lottery task was equivalent to the corresponding price search task. Subjects were not told about this correspondence and, because the frames were so different, it would have been very difficult to detect.[3] To allow between-subject tests of the effects of regret, each experimental session was randomly assigned to one of two treatments, which differed in terms of the feedback provided after a time-limited offer had been accepted in a price search task. To allow within-subject tests of the effects of time pressure, both treatments incorporated variation in the time allowed for making a decision about a time-limited offer.

The price search tasks were designed to provide a controlled and stylised representation of environments in which consumers engage in price search and in which time-limited offers appear relatively infrequently and without prior notice. In each task, the subject was given a 'budget' and instructed to buy a 'good' by spending from this budget; any unspent surplus constituted her earnings from the task. The subject was able to see six price offers, which appeared sequentially on a computer screen with short time gaps between them. In most cases, offers were free-recall. Although free-recall offers could be accepted at any time, the subject was free to wait until all such offers had appeared and then choose whichever of these she judged best (presumably the lowest). In fifteen of the tasks, however, one of the first three offers to appear would be flagged as time-limited. In this case, the subject was able to accept the offer only in the interval before the next offer appeared.

We gave subjects full information about the random processes that generated the six price offers in each task. These processes were given an intuitive interpretation in terms of 'deals' of 'cards' (see later) and were held constant throughout the experiment. Thus, the actual values of the offers in any task provided no information about offers in later tasks. In this sense, repetition did not provide any opportunities for learning. However, although we expected

[3]The random draws of offer prices were represented to subjects as 'deals' of cards, made at the start of each task. That representation reflected the actual sequence of operations by the computer program: offer prices for each task were determined only as that task appeared on a subject's screen. Given this feature of the experiment, price search tasks necessarily preceded lottery tasks. As we explain later, it seems unlikely that our main results are affected by the order of the two types of task.

that subjects would understand how the offers were generated, we did not expect that they would be able to calculate optimal solutions to the problems they faced: such calculations would be mathematically challenging even in the absence of the time constraints our design imposed. Our design strategy was to give each Part 1 task the 'feel' of a natural price search problem rather than to mimic the properties of theoretical models of search. By framing the structure of the task in simple terms and repeating it many times, we tried to ensure that subjects would gain an intuitive understanding of it and converge on patterns of response similar to those that they would use in natural price search problems. We take it that the theories of search that we test are intended to predict the behaviour of ordinary consumers in everyday settings.[4]

The experiment was designed to investigate three behavioural mechanisms that might be expected to make real consumers more likely than their counterparts in models of rational risk-neutral search to choose time-limited offers.

2.2.1 Feedback-conditional regret

The hypothesis that decision-making behaviour is influenced by anticipated regret was proposed by Bell (1982) and Loomes and Sugden (1982). Cues which prompt individuals to anticipate possible future regret have been found to make individuals less likely to take risks in consumer choice (Simonson, 1992) and in sexual behaviour (Richard *et al.*, 1996). In the context of time-limited offers, the dependence of regret on feedback about the outcomes of non-chosen options, theorised by Humphrey (2004), is particularly significant.

Consider a choice between two lotteries with monetary outcomes, defined on the mutually exclusive and exhaustive events $E_1, ..., E_n$. The 'safe' lottery S pays s in every event; the 'risky' lottery R pays $r_1, ..., r_n$, depending on which event obtains. First suppose that, whichever lottery is chosen, the true state will be revealed. Then, according to regret theory, when the decision-making agent considers choosing S, she will anticipate feelings of regret in those events E_i where $s_i < r_i$; these anticipations have a negative impact on the expected utility of S. Conversely, when she considers choosing R, she will anticipate regret in those events E_j where $r_j < s_j$; these anticipations have a negative impact on the expected utility of R.[5] Now suppose instead that the true state will be revealed only if R is chosen. In this case, as before, choosing R exposes the agent to the possibility of regret-inducing feedback; but now choosing S cuts off such feedback. Thus, a regret-averse agent might choose S if there will be feedback but R if there will not be. Zeelenberg *et al.* (1996) find evidence of this pattern of behaviour.

The Dutch Postcode Lottery provides an extreme example of how this tendency can be exploited by sellers. By drawing a winning postcode rather than a winning ticket number,

[4]This general methodological strategy is explained and defended by Sitzia and Sugden (2011).

[5]For simplicity, we ignore anticipated feelings of 'rejoicing' (the opposite of regret). In the regret literature it is common to assume that regret is a much stronger emotion than rejoicing.

this lottery design exposes individuals who do not bet to the possibility of extreme regret if their neighbours bet and win. Zeelenberg and Pieters (2004) find that anticipated regret is positively correlated with preferences for participating in a postcode lottery rather than a conventional alternative. Time-limited offers may exploit a similar tendency. Rejecting a time-limited offer exposes a consumer to the possibility of regret, but if offers are discovered only by active search, a decision to stop searching cuts off feedback that could cause regret about having accepting such an offer.

To investigate this possibility, we allocated each subject to one of two different feedback treatments. In the *No Feedback* treatment, if an offer was accepted before all six had been shown, the remaining offers were not revealed. Thus, in accepting a time-limited offer, a subject cut herself off from feedback that could reveal that she would have done better if she had rejected it. In the *Regret Feedback* treatment, whenever an offer was accepted, any remaining offers were immediately displayed. We tested the following hypothesis:

> *Hypothesis 1:* Other things being equal, the probability with which time-limited offers are chosen is greater when feedback is absent than when it is present.

2.2.2 Time pressure

Consumers' choices between accepting and rejecting time-limited offers may be affected by the pressure of having to make a decision in a short period of time, rather than being able to think about the decision more carefully. The psychological literature suggests that people respond to time pressure by using relatively simple heuristics that are adapted to maintaining decision accuracy with reduced cognitive effort, rather than by using truncated forms of reasoning processes that would generate correct decisions in the absence of time constraints. Payne *et al.* (1988) find that, under severe time pressure, experimental subjects focus on a subset of the available information and change their information-processing strategies. Finucane *et al.* (2000) find that time pressure increases subjects' reliance on the 'affect heuristic', which induces a negative correlation between judgements of risk and benefit (for example, a tendency to underestimate the radiation risk from the use of x-rays in hospitals). In an investigation of choices between pairs of gambles under different levels of time pressure, Ben Zur and Breznitz (1981) find that subjects spend proportionately more time attending to negative dimensions of gambles under high time pressure and so make more risk-averse decisions. Kocher *et al.* (2013) find that time pressure has no effect on risk attitudes revealed in choices between prospects offering gains, but increases risk aversion in choices between prospects which may impose losses. In a costly search task with free recall, Ibanez *et al.* (2009) find that participants tend to search less when each decision about whether to terminate search is subject to time pressure, but that this effect dissipates as participants gain experience of the task. Taken together, this body of evidence suggests a psychological mechanism by which time pressure would make consumers more likely to accept time-limited offers. We investigated the effect of time pressure by randomising the time gap between the appearance of offers, and hence the length of time for which time-limited offers were available. The time gap (which was constant within a task) could be either 4s or 12s. Given the simplicity and repetitive nature of the tasks, we expected that 12s would be sufficient for most subjects to make a considered (although not necessarily optimal) decision, but that 4s would be perceived as creating significant time pressure. We tested the following hypothesis:

> *Hypothesis 2:* Other things being equal, the probability with which time-limited offers are chosen is greater when time pressure is greater.

2.2.3 Risk aversion

In experiments involving choice among lotteries with low-value consequences, subjects often reveal high degrees of risk aversion. Rabin (2000) shows that, if conventional expected-utility models are fitted to the patterns of behaviour typically observed in such experiments, these models imply manifestly unrealistic degrees of risk aversion for decisions involving higher (but still modest) stakes. This inconsistency can be eliminated if, as in prospect theory, utility is defined as a function of changes in wealth relative to a reference point, rather than as a function of wealth itself, since this allows there to be particularly high degrees of risk aversion in the neighbourhood of any reference point (Kahneman and Tversky, 1979). If individuals are significantly averse to small risks, time-limited offers may be more attractive to consumers than is implied by models of risk-neutral search. Hence the following hypothesis:

> *Hypothesis 3:* In choosing whether or not to accept time-limited offers, experimental subjects tend to reveal risk aversion.

Given the degree of risk aversion typically observed in experiments, confirmation of Hypothesis 3 might be thought unsurprising, particularly as there is a good deal of existing evidence that in search tasks, both with and without recall, experimental subjects tend to search less than would be optimal for risk-neutral agents (e.g. Rapoport and Tversky, 1970; Sonnemans, 1998; Seale and Rapoport, 1997; Schunk and Winter, 2009a).[6] It is perhaps more interesting to ask whether individuals display the same attitudes to risk when deciding whether to reject a time-limited offer as they do when choosing whether to enter a risky lottery. The equivalence between these two choice problems might be less salient to most people than it is to decision theorists. Thus, if individuals use simplifying heuristics, the heuristics primed by price search problems might not be the same as those primed when decision problems are explicitly framed as choices between lotteries. The lottery tasks allow us to investigate this issue.

[6] Schunk and Winter find no significant relationship between search behaviour and a measure of risk attitude. They suggest that observed under-searching may be due to loss aversion rather than risk aversion.

Since behavioural theory does not provide unambiguous predictions about the direction that differences between the two tasks might take, we tested the following directionless hypothesis:

> *Hypothesis 4:* Risk attitudes revealed in decisions about accepting or rejecting time-limited offers are systematically different from those revealed in binary choices between certainties and lotteries.

2.3 Design details and implementation

2.3.1 Overall structure of experiment

As explained above, each session of the experiment was randomly assigned to one of the two treatments (No Feedback or Regret Feedback). Part 1 of each session consisted of thirty price search tasks; Part 2 consisted of fifteen lottery tasks. At the beginning of each part, each subject received a copy of the instructions for that part; these instructions were read aloud by the experimenter. These instructions are reproduced in Appendix 1. Each subject then completed a computerised questionnaire which tested her understanding of the tasks. If a subject made a mistake, the computer would show her the correct answer and the relevant part of the instructions. Subjects were invited to ask the experimenter for clarification.

In the price search tasks, subjects' budgets and offer prices were expressed in an experimental currency unit, *experimental points* (EP). Each subject's earnings from each task (conditional on the task being selected for payment: see later) were equal to the budget (always 10EP) minus the price of the accepted offer (never greater than 10EP). Thus, the subject had an incentive to choose the offer with the lowest price. Each lottery task was a binary choice between two lotteries with outcomes expressed in experimental points. 'Lottery 1' gave a stated outcome for sure. 'Lottery 2' had five possible outcomes, each with a probability of 0.2.

At the end of each session, the computer randomly selected two price search tasks and one lottery task for each subject. Each subject's earnings were the sum of her earnings from the two selected price search tasks and the outcome of the lottery she had chosen in the selected lottery task. If she had chosen Lottery 2, its outcome was determined by a random process, using the stated probabilities. Subjects were paid at the exchange rate of £1 for every 2.5EP.

2.3.2 Price search tasks

The basic structure of these tasks was described in Section 1 above; here we fill in the details. Offers were represented on subjects' screens by pictures of 'cards'. As with a pack of playing cards, all cards looked identical when viewed face down. When viewed face up, each card showed a price offer. At the start of a task, a subject saw a row of six cards, all face down. In terms of the prices they might offer, there were two types of card, *red* and *blue*. In all tasks, one of the cards was red and five were blue. The red card had been randomly assigned to one of the first three positions in the row, but the subject was not told which position this was. The offer on each blue card was a price drawn independently and randomly from the set {0.00EP, 0.01EP, ..., 10.00EP}. The offer on the red card was a price drawn at random from the set {0.00EP, 0.01EP, ..., 4.00EP}. Thus, the red card was particularly likely to show a low-price offer. Subjects were fully informed about the distributions of offers on the two types of card. Our reasons for using these particular parameters will be explained later.

During each task, the cards were turned over by the computer, one by one and working from left to right. The time interval between cards being turned over was fixed within each task at either 4s or 12s. Which interval was used in each task was randomised, independently for each participant, subject to the constraint that each participant faced each interval in fifteen tasks. There was no time limit for completing the task as a whole. Subjects could not move on to Part 2 of the experiment until everyone in the session had completed Part 1. Thus, an individual subject who chose not to see all six offers was unlikely to save time by doing this.

The offers on blue cards were always free-recall. Whenever a free-recall card (of either colour) was turned over to display its offer, a button appeared below it on the screen, labelled 'Click to choose this offer', together with the message 'This offer will stay available throughout the task'. The card remained turned over until the end of the task. The subject was free to choose that offer at any time from then on, by clicking on the button. Thus, it was possible for a subject to wait until all cards had been turned over and then to choose from the complete set of free-recall offers, with no time constraint. The offer on the red card could be either free-recall (described to subjects as a 'standard red card') or time-limited (a 'time-limited red card'). We will classify tasks as 'free-recall' or 'time-limited' according to the properties of the red card. Whether the red card offer was free-recall or time-limited was randomised for each task, independently for each participant, subject to two constraints: that each subject faced fifteen free-recall tasks and fifteen time-limited tasks; and that in each of these sets of fifteen tasks, the split between 4s and 12s time intervals was either 7:8 or 8:7. Subjects were not told which type of offer the red card would make until it was turned over. If the offer was time-limited, a 'Click to choose this offer' button appeared below the card, together with the message 'This offer will stay available for 4 seconds' or 'This offer will stay available for 12 seconds'. A countdown clock showed the number of seconds remaining. At the end of this time interval, if the offer had not been accepted, the card was turned back, making the offer price no longer visible, and its choice button, availability message and countdown clock disappeared. Simultaneously, the next card in the sequence was turned over.

In price search tasks, the only difference between the two treatments occurred after an offer was accepted. In the No Feedback treatment, the offer price that the subject had chosen was shown on the screen, and the task ended at that point. Thus, if the subject chose an offer before all the cards had been turned over, she could not get any information about the offers on the remaining cards. In the Regret Feedback treatment, any remaining cards were turned over, revealing their prices. The subject was not allowed to change her decision at this stage.

Figure 2.1 shows a typical screen shot for a time-limited task. In this example, the third card has just been turned over. It is a time-limited offer with a price of 1.75EP. The subject has 3s remaining in which to choose whether to accept this offer. The two preceding free-recall prices are 2.27EP and 4.09EP.

Figure 2.1: Screen shot of a typical time-limited task

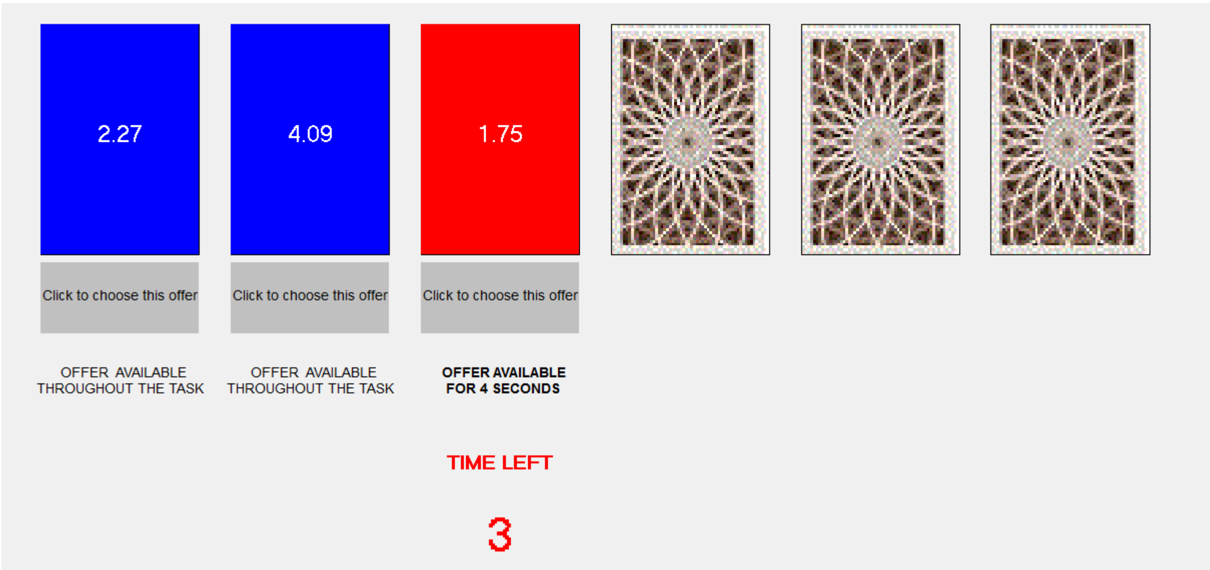

Viewed in the framework of standard decision theory, free-recall price search tasks are trivial. In such a task, a rational subject would wait until all prices had been revealed and then choose the lowest. A time-limited task is also trivial if the time-limited price is greater than or equal to any preceding free-recall offer: in such a case, the time-limited offer would always be rejected. A significant decision problem occurs only when a time-limited offer is better than any preceding offer and (as was always the case in our design) other offers remain to be revealed. In this case, there are only two relevant options for a rational subject: either to accept the time-limited offer; or to reject it, wait until all offers have been revealed, and then choose the best free-recall offer. Time-limited tasks which present such problems will be called *consequential*.

For reasons of statistical power, we needed to use parameters that would ensure that, in the experiment as a whole, approximately equal numbers of subjects would accept and reject time-limited offers in consequential tasks. We therefore chose the parameters so that, conditional on facing a consequential task, there was a 0.5 probability that a risk-neutral expected-utility-maximising subject would accept the time-limited offer. For similar reasons, we needed that, in a significant proportion of consequential tasks, the expected values of 'accept' and 'reject' would be relatively close to one another, so that subjects could be expected to find the decision problem relatively difficult. In order to achieve these objectives, it was necessary that time-limited offers appeared early in the offer sequence and, on average, had lower prices than free-recall offers. The fact that the time-limited offers in our design were relatively attractive mirrors the logic behind the use of such offers in retail markets. If, after rejecting a time-limited offer, a consumer can continue to search in a market in which free-recall offers exist, she is unlikely to be induced to accept a time-limited offer unless she perceives it as attractive in relation to the distribution of offers in the market as a whole. We used equal numbers of free-recall and time-limited tasks because we wanted the appearance of time-limited offers to be relatively infrequent and unpredictable. The presence of red cards in the free-recall tasks ensured that free-recall and time-limited tasks had exactly the same ex ante distribution of offers, and made salient to subjects that favourable distributions of offers (i.e. offers on red cards) were not necessarily time-limited. As we will explain later, the free-recall tasks were also useful for testing subjects' understanding of the structure of the experiment.

2.3.3 Lottery tasks

As explained in Section 2.3.1, in Part 2 each subject faced fifteen lottery tasks, each requiring a choice between two lotteries. The parameters of these tasks were set separately for each subject, so that for each subject there was a one-to-one correspondence between the fifteen lottery tasks and the fifteen time-limited tasks from Part 1.[7] In defining this correspondence, we consider two alternative ways in which the subject might have responded to the relevant price search task. (These are not the only possible responses, but all other responses would be non-rational according to the criteria presented in Section 2.3.2.) *Response 1* is to accept the time-limited offer. *Response 2* is to reject the time-limited offer, not to accept any offer until all have been revealed, and then to choose the lowest free-recall offer. Lotteries 1 and 2 represent the distributions of earnings implied by Responses 1 and 2 respectively. That is, the certainty offered by Lottery 1 represents the earnings that the subject would have made in the matched price search task, had she accepted the time-limited offer. The risky prospect offered by Lottery 2 represents the distribution of possible earnings for the subject, conditional on the information that would have been available to her when the time-limited offer was visible, had she rejected that offer, waited until all free-recall offers had been revealed, and then chosen the lowest of these. Notice that the specification of each lottery task varied only according to the randomly-determined *parameters* of the corresponding time-limited task; it was independent of the subject's *behaviour* in that task. Notice also that if a price search task is non-consequential, the corresponding lottery task is one in which Lottery 1 is stochastically dominated by Lottery 2. Conversely, if there is stochastic dominance in the lottery task, the

[7]Each subject faced the lottery tasks in the same order as she had faced the corresponding time-limited price search tasks. Recall, however, that the latter tasks were randomly distributed among the thirty price search tasks.

price search task is non-consequential.[8]

For example, consider the price search task that generates the screen shot shown in Figure 2.1. In this task, the time-limited price is 1.75EP. Since each task had a budget of 10EP, Response 1 gives earnings of 8.25EP with certainty. Thus, the outcome of the corresponding Lottery 1 is 8.25EP. The earnings from Response 2 are given by 10EP minus the lowest of the five free-recall prices. Two of these prices (2.27EP and 4.09EP) are already known; each of the others will be a random draw from the set {0.00EP, 0.01EP, ..., 10.00EP}. Thus, earnings from Response 2 are described by a well-defined probability distribution over {0.00EP, 0.01EP, ..., 10.00EP}. The corresponding Lottery 2 is a discrete approximation of that distribution. For each quintile of the actual distribution of Response 2 earnings, this approximation preserves the expected value of earnings conditional on earnings falling into that quintile. It therefore also preserves the unconditional expected value of earnings. Figure 2.2 shows a screen shot of the lottery task that corresponds to the price search task of Figure 2.1. It should be apparent from a comparison of the two figures that, although the two tasks are conceptually connected, there is no superficial resemblance between them that might facilitate memory connection or between-tasks learning.[9]

Figure 2.2: Screen shot of a typical lottery task

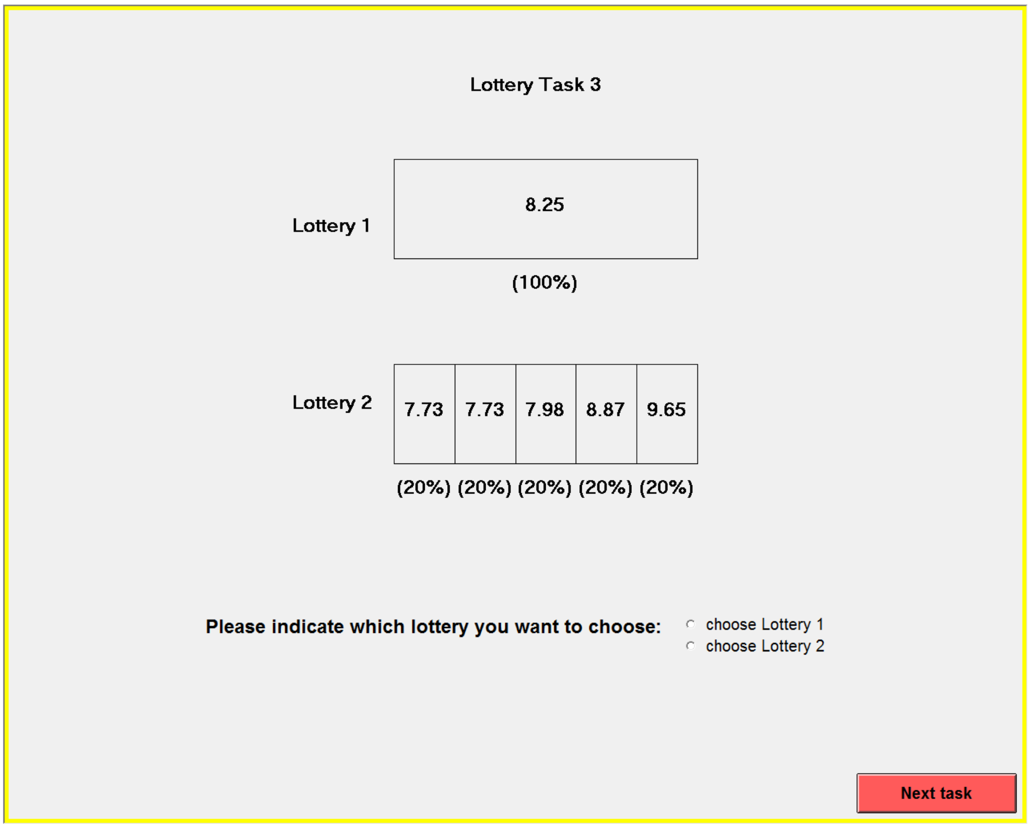

[8]This statement is subject to a technical qualification. Because each Lottery 2 is a discrete approximation to the true distribution of possible earnings in the relevant card task rather than that distribution itself, there is not an exact correspondence between non-consequential problems in card tasks and stochastic dominance in lottery tasks.

[9]Zizzo (2005) contains a summary of empirical evidence demonstrating how transfer of knowledge is easy when superficial features match but hard when they do not.

Each subject completed all fifteen lottery tasks before receiving any information about the actual outcome of any Lottery 2. At the end of the experiment, one lottery task was selected at random for payment. If the subject had chosen Lottery 2 in this task, the computer selected one of the five outcomes of that lottery at random, and this outcome determined the subject's earnings from Part 2. As far as Part 2 is concerned, the only difference between the two treatments occurred at this stage, and then only if the subject had chosen Lottery 1. In the No Feedback treatment, Lottery 2 was not resolved in this case. In the Regret Feedback treatment, the computer determined an outcome for that lottery and this was displayed on the subject's screen. In each treatment, the Part 2 instructions made clear what information would be provided at the end of the experiment. This difference between the two treatments maintained the correspondence between price search tasks and lottery tasks.

2.3.4 Implementation

The experiment was conducted in fourteen sessions at the Centre for Behavioural and Experimental Social Science laboratory at the University of East Anglia during the summer of 2014. The experiment was programmed and conducted with the experimental software z-Tree (Fischbacher, 2007). Participants were recruited from the general university population by email, using the Hroot online recruitment system (Bock *et al.*, 2014). Altogether there were 209 participants (101 male and 108 female), of whom 105 took part in the No Feedback treatment and 104 in the Regret Feedback treatment. Most of the participants were students from a wide range of academic disciplines and with an age range from 18 to 55. The experiment lasted about 50 minutes. Payments to participants ranged from £6.31 to £11.82, with an average of £10.43.

2.4 Results

2.4.1 Subjects' understanding of tasks

We begin by checking that subjects' behaviour showed a basic understanding of the tasks. In this section, we give an overview of the relevant findings; full details are given in Appendix 2.

One simple test is whether subjects chose dominated options. We will say that a subject chose a dominated option in a price search task if, at the moment at which she accepted an offer, another offer with a strictly lower price was visible and available. In the 6270 ($= 209 \times 30$) responses to price search tasks, there were only 107 cases (1.7 per cent) in which dominated offers were chosen. This percentage was exactly the same for 4s and 12s tasks. As explained in Section 2.3.3, the lottery tasks included cases in which Lottery 1 was stochastically dominated by Lottery 2. Of the 3135 ($= 209 \times 15$) cases of lottery tasks, 595 were choices between dominating and dominated lotteries; the dominated lottery was chosen in only 22 of these cases (3.7 per cent). We interpret these data as evidence that subjects had a basic understanding of both types of task.

In the price search tasks, as explained in Section 2.3.2, a subject incurred no money cost and almost no time cost by delaying a choice of a free-recall offer. One might reasonably claim that a fully rational subject would never accept a free-recall offer until she had seen all six offers. However, behaviour contrary to this principle does not necessarily reveal misunderstanding of the tasks. For example, it could result from a short-sighted desire to avoid immediate delays even though the total time spent on the experiment would be unaffected. Or it could result from subjects using satisficing heuristics that accept offers that are 'good enough'.

We classify a response to a price search task as *impatient* if the subject accepted a free-recall offer (whether red or blue) before all offers had been revealed, and if that offer had a non-zero price and was not dominated (as defined above). Because choices of time-limited offers are never classified as impatient, observations of impatience should be expected to be more frequent in free-recall tasks. In fact, 9.5 per cent of all responses in time-limited tasks and 23.3 per cent of all responses in free-recall tasks were impatient.

Because impatience is a property of responses to free-recall offers, it can be analysed more cleanly by using the data from free-recall tasks. (An analysis of the data from time-limited tasks leads to similar conclusions, but is more convoluted: see Appendix 2 for details of this analysis.) Impatient responses were more frequent in 12s free-recall tasks (where they made up 26.7 per cent of all responses) than in 4s free-recall tasks (19.8 per cent). This difference is highly significant ($z = -4.442$, $p < 0.001$).[10] This effect is not surprising: the longer the delay between each offer, the more patience is required if a subject is to see all of them.

It is useful to subdivide impatient responses according to the number of offers (from one to five) that were visible when the subject accepted an offer. Impatient responses were distributed fairly evenly over these five *stages* of the tasks. For any given subject who made an impatient response at any given stage of a free-recall task, we define the *loss due to impatience* as the difference between (i) the price the subject actually paid and (ii) the expected value of the price she would have paid, had she waited until all six offers had been revealed and then chosen the lowest price.[11] These values are remarkably low. Averaging over all five stages, the average loss per instance of impatience was only 0.13EP in 4s free-recall tasks and 0.18EP in 12s free-recall tasks; even at the first stage (i.e. when the first offer was accepted without any other offer having been seen), average losses were only 0.22EP and 0.29EP respectively. To provide a scale of reference: the ex ante expected value of the price to be paid by a fully rational subject was 1.11EP; random choice would imply an expected price of 5.00EP. The implication is that, even when making impatient responses, most subjects were using decision-making

[10] For each subject, we calculated the total number of impatient responses, separately for 4s and 12s tasks, and then ran a two-tailed Wilcoxon signed-rank test.

[11] This expected value is conditional on the subject's information at the relevant stage. The derivation of such expectations is explained in Appendix 3.

heuristics that were approximately rational. This in turn implies that impatient subjects typically terminated the search process only when they were able to accept very favourable offers.

2.4.2 The data to be analysed

From now on, we will be concerned only with subjects' decisions between accepting and rejecting time-limited offers in consequential tasks, and with subjects' choices in the corresponding lottery tasks. Thus, the price search choice data that we analyse exclude the 19 per cent of cases in which the time-limited price was greater than or equal to the lowest preceding free-recall price. These data also exclude the very few cases in which a subject accepted a free-recall offer before seeing the time-limited offer.[12] To maintain a one-to-one relationship between observations from price search tasks and lottery tasks, we also exclude the data from the corresponding lottery tasks. (By virtue of the principles used to construct lottery tasks, almost all of the excluded lottery tasks were ones in which Lottery 1 was stochastically dominated by Lottery 2.)

In testing Hypotheses 1 to 4, we compare the frequencies with which time-limited offers were chosen in different treatments or at different levels of time pressure. Recall that, in each time-limited task faced by each subject, the price on each of the six cards was determined randomly, as was the position (first, second or third) of the time-limited offer in the sequence of cards. Therefore, the consequential tasks faced by a subject often included tasks in which the time-limited offer was worth much more than the expected value of the best free-recall price, and/or tasks in which it was worth much less. In comparing responses across treatments, it is important to control for this variation.

We do this by calculating, separately for each subject and for each relevant task, the expected value of the lowest free-recall price, conditional on the subject's information at the stage at which the time-limited offer appears – that is, conditional on the actual values of offers that had already been revealed but using ex ante expectations of the others. From this expected value we subtract the actual time-limited price, to arrive at the *expected value difference* for that subject/task combination. (The derivation of expected value difference is explained in more detail in Appendix 3). The zero point on the scale of expected value difference corresponds with cases in which a fully rational, risk-neutral expected-utility-maximising individual would be indifferent between accepting and rejecting a time-limited offer. In our analysis, however, this variable serves primarily as a one-dimensional summary statistic of the net advantage of accepting rather than rejecting a time-limited offer. A subject's response to any given consequential time-limited task can be summarised by the expected value difference in that task and her decision to accept or reject.

For any given category of time-limited tasks and aggregating across all subjects, we can

[12] 14 of the 3135 responses to time-limited tasks (0.04 per cent) were of this kind.

investigate the relationship between the expected value difference and the relative frequency with which the time-limited offer is accepted. It is natural to expect that different subjects will have different attitudes to risk, and that the attitudes to risk revealed by any given subject will be subject to stochastic variation. Thus, one would expect this relationship – *the choice frequency function* – to be upward-sloping, with the S-shape that is characteristic of logistic functions. Throughout the paper, we show Lowess-smoothed plots of choice frequency functions. To provide a convenient summary statistic of the attitudes to risk revealed in any given category of time-limited task, we estimate the choice frequency function for that category using a logit model with subject-level clustering and then report the estimated probability with which the time-limited offer is chosen when the expected value difference is zero. We will call this the *standardised probability* of the choice of the time-limited offer in the relevant category of tasks.

2.4.3 Feedback effects

Figure 2.3 shows Lowess-smoothed choice frequency functions for the No Feedback and Regret Feedback treatments. It is clear from this diagram that, contrary to Hypothesis 1, time-limited offers were slightly *less* likely to be chosen in the No Feedback treatment. The standardised probability of choosing the time-limited offer is 0.590 for the No Feedback treatment and 0.645 for the Regret Feedback treatment. To test whether behaviour in the two treatments is significantly different, we estimate a logit model with a dummy variable to represent the differential effect of the Regret Feedback treatment. The coefficient of this variable is positive but not quite significant at the 10 per cent level ($p = 0.104$) (see Appendix 4, section 1). As far the hypothesis that we set out to test is concerned, our first main result is clear:

> *Result 1*: There is no support for Hypothesis 1, i.e. the hypothesis that time-limited offers are more likely to be chosen when feedback is absent than when it is present.

Figure 2.3: Comparison between responses in No Feedback and Regret Feedback Treatments

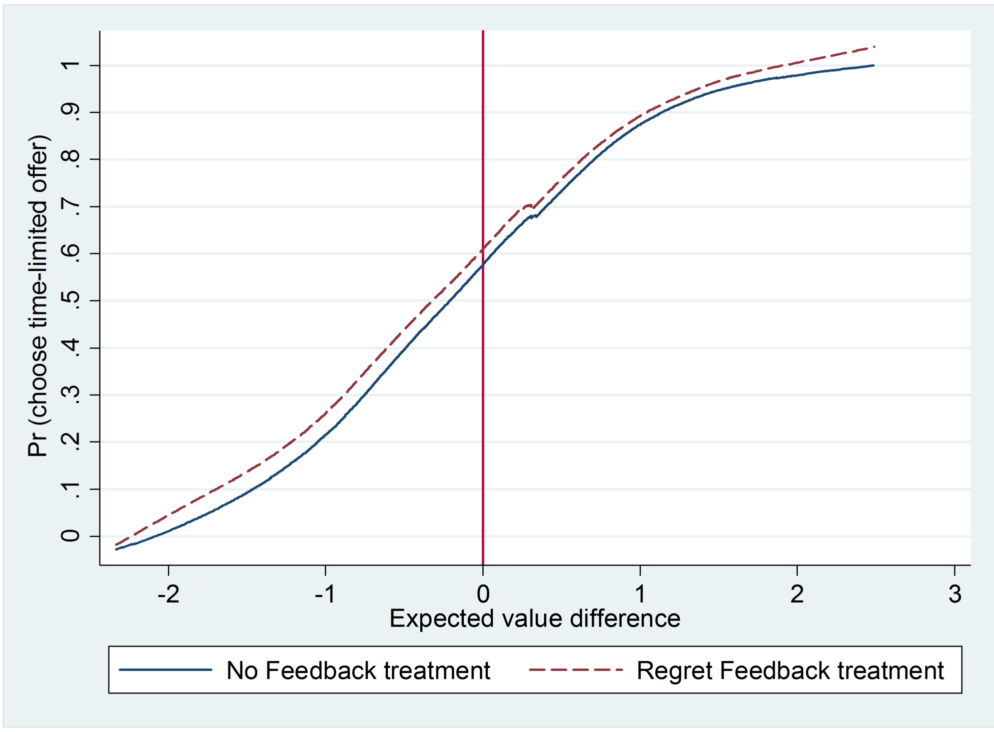

In Sections 2.4.3 and 2.4.4, we pool the data from the two treatments. Because the treatments are identical in all respects other than the provision of feedback, this pooling cannot induce systematic biases in our results; it merely introduces a small amount of additional noise.

2.4.4 Time pressure

We now use the same method as in Section 2.4.2 to test Hypothesis 2 about the effect of time pressure. Figure 2.4 plots Lowess-smoothed choice frequency functions for the 4s and 12s time-limited tasks.

Figure 2.4: Comparison between responses to 4s and 12s tasks

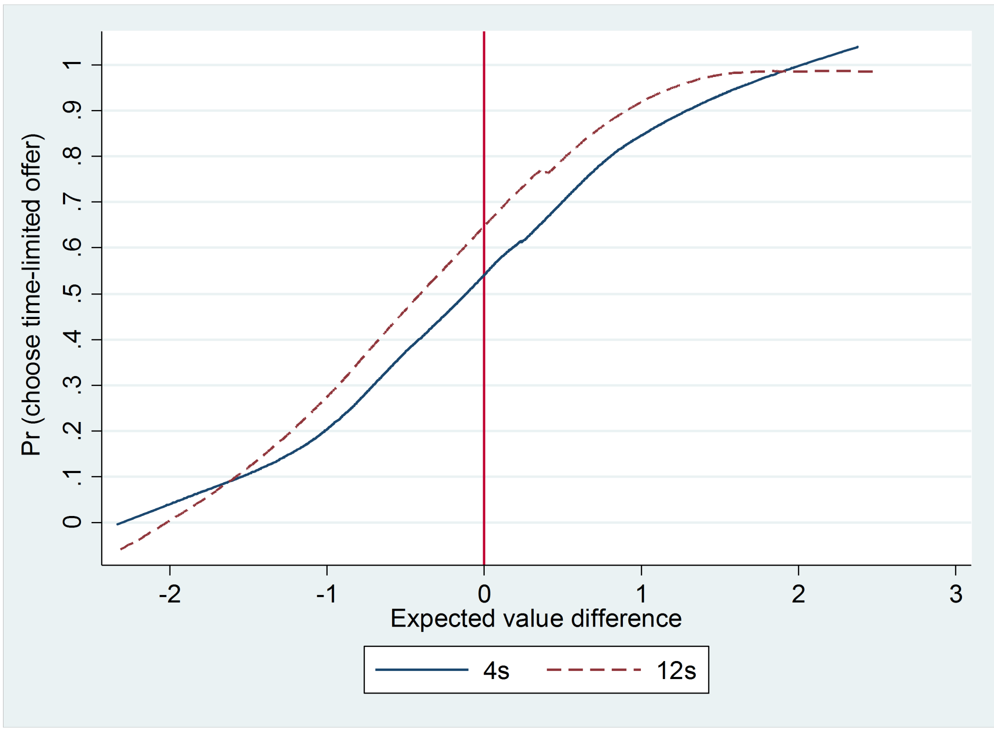

It is immediately obvious that, in direct opposition to Hypothesis 2, time-limited offers are *less* likely to be chosen in 4s tasks than in 12s tasks. The standardised probability of choosing the time-limited offer is 0.551 for 4s tasks and 0.691 for 12s tasks. To test whether behaviour is significantly different at the two levels of time pressure, we estimate a logit model with subject-level clustering. A dummy variable takes the value 0 for 4s tasks and 1 for 12s tasks. The coefficient of this variable is positive and strongly significant ($p < 0.001$) (see Appendix 4, section 2).

Choice frequency functions can also be used to show the degree of dispersion in responses, controlling for variation in expected value difference. It is clear from Figure 2.4 that the degree of dispersion is similar at the two levels of time pressure. Thus, there is no evidence to suggest that individual decision-making was subject to more stochastic variation when the time available was shorter, as would be implied by many psychological theories of decision processes (e.g. Busemeyer and Townsend, 1993; Stewart *et al.*, 2006). It seems that the tendency for time-limited offers to be chosen less frequently in 4s tasks was not an artefact of decision error or imprecision.

Another possibility is that this unexpected effect was an artefact of impatience. As shown in Section 2.4.1 and Appendix 2, subjects were more impatient in the 12s tasks, possibly because of boredom with the relatively long delays between the appearance of successive free-recall offers. Accepting a time-limited offer rather than waiting to see more free-recall offers could be an expression of impatience. To test for such an effect, we calculated, for each subject, the number of impatient choices in free-recall tasks and the number of time-limited offers accepted in time-limited tasks. Figure 2.5 shows a jittered scatterplot of these data. The correlation coefficient is only 0.065, and is not significantly different from zero ($p = 0.345$). The implication is that impatience and the propensity to choose time-limited offers are essentially orthogonal.

Figure 2.5: Jittered scatterplot of relationship between impatience and choice of time-limited offers

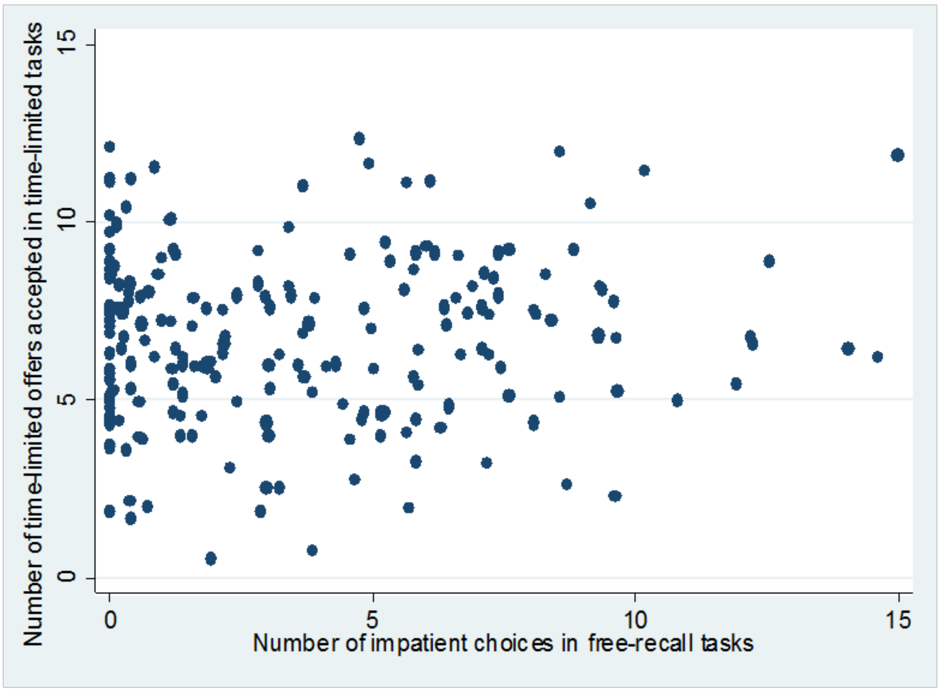

Postponing to Section 2.5 the discussion of other possible explanations, we state our second main result:

> *Result 2*: Contrary to Hypothesis 2, time-limited offers are less likely to be chosen when time pressure is high than when it is low.

For both 4s and 12s tasks, the standardised probability of choosing the time-limited offer was significantly greater than 0.5 ($p = 0.028$ for 4s tasks, $p < 0.001$ for 12s tasks). Hence:

> *Result 3*: Consistently with Hypothesis 3, subjects' choices between accepting and rejecting time-limited offers reveal an overall tendency towards risk aversion. This tendency is much stronger when time pressure is low than when it is high.

2.4.5 Comparisons between price search tasks and lottery tasks

We now consider subjects' responses to the lottery tasks. Analogously with our treatment of price search tasks, we define expected value difference for any given lottery task as the (certain) value of Lottery 1 *minus* the expected value of Lottery 2. Recall that, for each <subject, task> pair that is included in our analysis, there is a corresponding lottery task. Payoffs in this lottery task correspond with potential earnings in the price search task. Our definition ensures that every lottery task has the same expected value difference as the price search task to which it corresponds. Aggregating over all subjects and all (non-excluded) lottery tasks, we can plot the relative frequency of Lottery 1 choices as a function of the expected value difference.[13] Figure 2.6 shows the Lowess-smoothed choice frequency function for lottery tasks, superimposed on the 4s and 12s time-limited offer functions from Figure 2.4.

Figure 2.6: Comparison between choices in lottery tasks and price search tasks

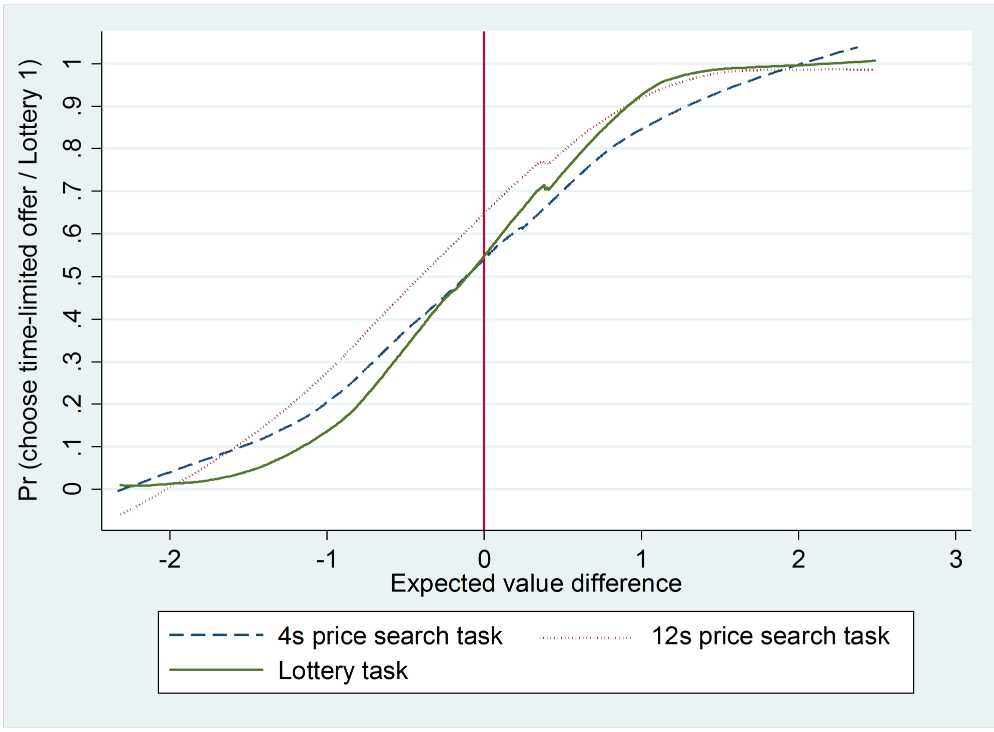

The standardised probability of choosing Lottery 1 is 0.577. This is significantly different from 0.5 ($p < 0.001$), indicating an overall tendency towards risk aversion. To test whether attitudes to risk revealed in the lottery tasks are significantly different from those revealed in the price search tasks, we estimate a logit model with subject-level clustering. We use two dummy variables. One takes the value 1 in 4s price search tasks, 0 otherwise; the other takes the value 1 in 12s price search tasks, 0 otherwise. The 4s dummy variable is negative but not significant ($p = 0.959$); the 12s dummy variable is positive and highly significant ($p < 0.001$) (see Appendix 4, section 4). Since lottery tasks were faced later in the experiment than price search tasks, differences in risk aversion between the two types of tasks might in principle be an artefact of gradual changes in risk aversion over the course of the experiment. But we found no trend in risk aversion over the thirty price search tasks (see Section 2.4.6).

[13]If the lottery choice data are disaggregated between the No Feedback and Regret Feedback treatments, the choice frequency functions for the two treatments are almost identical. Plots of these functions are shown in Appendix 4, section 3.

We conclude:

> *Result 4*: When time pressure is low, decisions about accepting or rejecting time limited offers are more risk-averse than are binary choices between certainties and lotteries. At high levels of time pressure, there is no significant difference between attitudes to risk in the two types of decision problem.

Notice that the tests that support Result 4 do not take account of differences in dispersion of attitudes to risk. It is clear from inspection of Figure 2.4 that there is less such dispersion in lottery tasks than in price search tasks. This is perhaps not surprising. In a lottery task, the probability distributions of payoffs for the two options are described explicitly to subjects. In a price search task, in contrast, it is a difficult problem to work out from first principles the distribution of payoffs implied by a rejection of the time-limited offer. If subjects' beliefs about this distribution are wholly or partly derived from experience, or if they use experience-based decision heuristics which bypass the formation of such beliefs, the random elements and time lags of the learning process will introduce additional noise into their responses.

Figure 2.6 shows a broad pattern of similarity between aggregate responses to time-limited tasks (particularly at 4s) and aggregate responses to lottery tasks. But it is also useful to investigate consistency between responses to the two types of task at the level of the individual subject. For each subject and for each consequential time-limited price search task faced by that subject, we can compare her response to that task with her response to the matched lottery task. We will say that these responses show *acceptance consistency* if the subject chose the time-limited offer in the price search task and Lottery 1 in the lottery task, and *rejection consistency* if the time-limited offer is rejected and Lottery 2 is chosen. Aggregating across subjects and task pairs, the relative frequency of the two types of consistency can be plotted as functions of expected value difference.

Figure 2.7 shows the Lowess-smoothed frequency functions for the two types of consistency for 4s tasks. The sum of these relative frequencies (the *consistency rate*) is also plotted. Figure 2.8 shows the corresponding plots for 12s tasks.

Figure 2.7: Consistency of individual choices between 4s price search tasks and lottery tasks

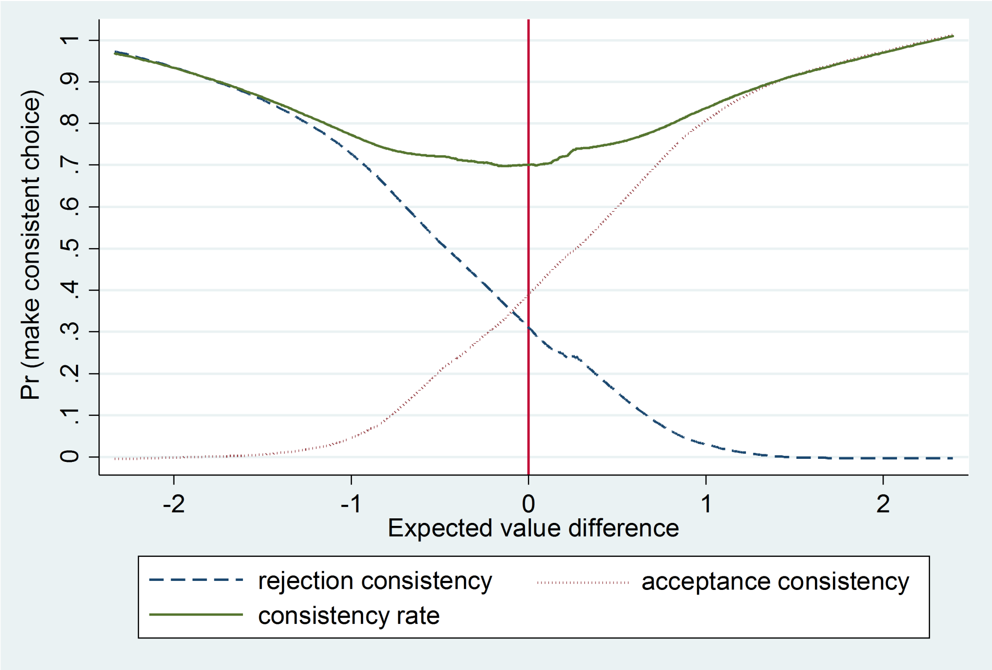

Figure 2.8: Consistency of individual choices between 12s price search tasks and lottery tasks

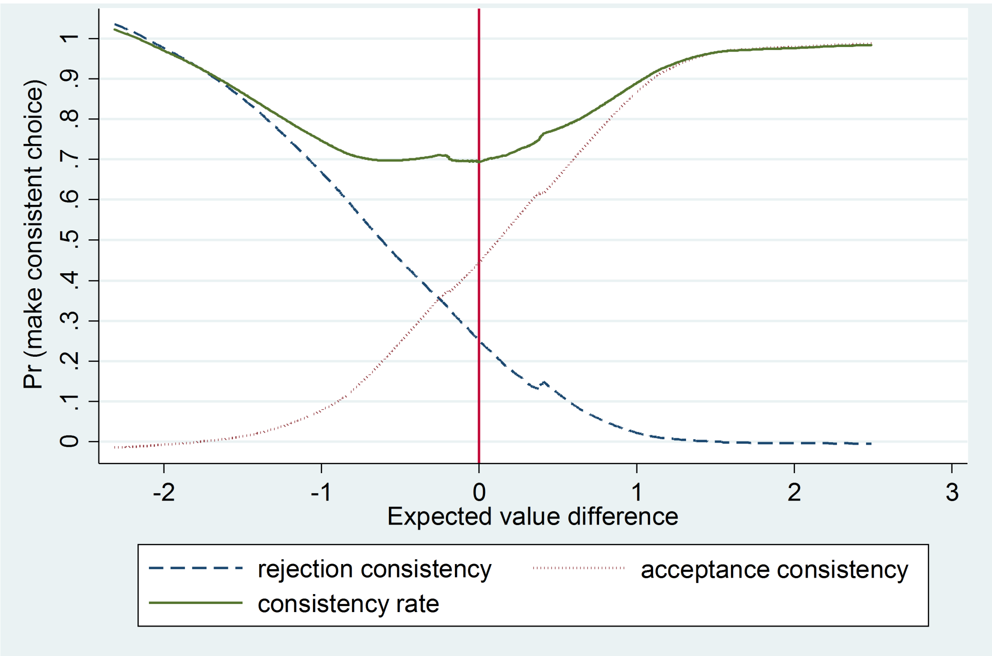

In interpreting consistency rates, it is useful to have a benchmark. Previous experimental research has generated data about how subjects respond when *exactly the same* problem of choosing between two lotteries is faced twice in the same experiment. These choice problems have normally been designed so that neither option is obviously better than the other. Consistency rates in such problems have typically been found to lie in a range from

70 to 82 per cent.[14] Intuitively, one might expect much less consistency when, as in our experiments, the two decision problems are framed very differently. A further reason for expecting inconsistencies arises from the distinction between *decisions from description* (that is, problems of choice under uncertainty in which decision-makers are given explicit information about probabilities) and *decisions from experience* (problems that are repeated several times and for which probabilities have to be inferred from experience). Our lottery tasks are clearly decisions from description, while (for reasons explained in Section 2.2) subjects might treat our price search tasks as decisions from experience. It is known that behaviour is systematically different in the two environments. In particular, when individuals are making decisions by description, they tend to over-weight small probabilities (Kahneman and Tversky, 1979), but when decisions are made by experience, small probabilities tend to be *under*-weighted (Hertwig *et al.*, 2004).

Viewed against the benchmark of repeated identical decision problems, consistency rates in our experiment are remarkably high. In both 4s and 12s tasks, the consistency rate reaches a minimum in the region of expected value difference at which each option is chosen with approximately the same frequency; in both cases, the minimum value of this rate is just below 70 per cent. Hence:

> *Result 5*: At the level of individual subjects, the rate of consistency between time-limited offer choices and corresponding lottery choices is similar in magnitude to previously-observed consistency rates between identical lottery choice problems.

This degree of consistency between price search tasks and lottery tasks suggests that, for a typical subject, both tasks tapped into some common body of attitudes to risk: preferences were not simply constructed in response to specific decision problems. Further, it suggests that subjects' capacity to make considered decisions about such offers was not greatly impaired by time pressure, even in the 4s tasks. Nevertheless, they made systematically different decisions when time pressure was relaxed.

2.4.6 Learning

Finally, we consider how subjects' responses evolved over the course of the experiment. To investigate this, we regress subjects' decisions in time-limited tasks on a set of explanatory variables, some of which capture possible learning mechanisms. All the regressions we report are estimations of logit models in which subject-level random effects are controlled.

Results for our three main models are presented in Table 2.1. In these models, the dependent variable is the probability that a subject chooses the time-limited offer in a consequential time-limited task. Because subjects in the No Feedback and Regret Feedback treatments received

[14] Loomes *et al.* (2002) report an experiment with a consistency rate of 82 per cent and cite four other studies which found consistency rates between 70 and 80 per cent.

different information, the distinction between these treatments is potentially significant in explaining learning. Model 1 uses data from all time-limited tasks; Models 2 and 3 respectively use only data from the No Feedback and Regret Feedback treatments.

Table 2.1: Regression results using expected value difference

	Overall		NF		RF	
	(1)		(2)		(3)	
	β	ME	β	ME	β	ME
Treatment	0.329	0.078				
	(0.200)	(0.047)				
Expected value difference	2.399***	0.572***	2.692***	0.662***	2.139***	0.496***
	(0.104)	(0.026)	(0.168)	(0.043)	(0.128)	(0.030)
Available time	0.692***	0.164***	0.634**	0.154**	0.715***	0.164***
	(0.130)	(0.030)	(0.195)	(0.047)	(0.175)	(0.039)
Good offers rejected	2.105***	0.502***	2.960***	0.728***		
	(0.272)	(0.065)	(0.450)	(0.111)		
Good offers seen					1.430***	0.331***
					(0.351)	(0.082)
Period	0.003	0.001	0.005	0.001	0.004	0.001
	(0.008)	(0.002)	(0.012)	(0.003)	(0.011)	(0.002)
Constant	−0.548**		−0.728**		−0.386	
	(0.201)		(0.276)		(0.245)	
Observations	2456		1246		1210	
LR chi2	536.210		256.673		280.878	
Prob > chi2	0.000		0.000		0.000	
Baseline predicted probability	−0.794		−1.061		−0.643	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. We used panel data to estimate all these models. The data used to estimate model 2 contains 1246 observations from 105 subjects. The data used to estimate model 3 contains 1210 observations from 104 subjects. For each model, the left column contains coefficients, and the right column report marginal effects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

Treatment in Model 1 is a dummy that takes the value 1 if subjects are in the Regret Feedback treatment. *Expected value difference* is as defined in Section 2.4.2. *Available time* is a dummy which takes the value 1 in 12s tasks. *Period* represents the position of the task in the sequence of price search tasks faced by the subject; 1 is the first task and 30 is the last. This variable picks up overall trends in learning.

We define two additional variables to pick up the effects of experience of the realisations of offer values. For each consequential time-limited task faced by each subject, we define the time-limited offer as *good* if its price is strictly less than the lowest price in the set of five free-recall offers and *bad* otherwise.[15] Notice that the good/bad distinction is based on ex post realisations of offer values, not ex ante expectations, and is unaffected by whether or not the time-limited offer was chosen. At the end of each time-limited task in the Regret Feedback treatment, subjects always knew whether the time-limited offer was good or bad. For subjects in this treatment, experiential learning might be mediated by memories of good and bad offers observed in previous tasks. Such memories might be interpreted as experiences of regret (if a good offer was rejected or a bad offer accepted) or rejoicing (in the opposite cases). Alternatively, they might be interpreted as encoding information about the distribution of the difference in value between the time-limited offer and the best free-recall offer. The variable *Good offers seen* is defined (for Model 3 only) as the proportion of previous time-limited tasks in which the time-limited offer was good. (In a subject's first time-limited task, this variable takes the value zero.)

For Models 1 and 2, we are forced to use a less clean measure of a subject's experience of good and bad offers. To maintain as close a parallel as possible with *Good offers seen*, we define *Good offers rejected* as the number of previous tasks in which good time-limited offers were not chosen as a proportion of the total number of previous tasks in which time-limited offers were not chosen. (In the first task in which a subject did not choose the time-limited offer, this variable takes the value zero.) Although subjects who made impatient decisions in the No Feedback treatment did not see all five free-recall offers, the value of the best of these offers is a good approximation to the value of the best free-recall offer actually seen by the subject (see Section 2.4.1). A further confound needs to be taken into account. For subjects who (independently of experience) have a relatively strong propensity to accept time-limited offers, the proportion of rejected offers that are found to be good will be relatively low. This mechanism induces a tendency for the choice of time-limited offers to vary *negatively* with *Good offers rejected.* Thus, finding a positive effect would be particularly strong evidence of reinforcement or regret-based learning.

Table 2.1 shows the results of these regression models, including the coefficients and the marginal effects. In all three regressions, unsurprisingly, the coefficients for *Expected value difference* are positive and highly significant ($p < 0.001$ in all cases). *Available time* has a positive and highly significant effect in all three regressions, in line with the findings reported in Section 2.4.4. In Model 1, the coefficient for *Treatment* is positive but not quite significant at the 10 per cent level ($p = 0.100$), paralleling the findings reported in Section 2.4.3. *Period* is insignificant in all three regressions, indicating the absence of any trend in the probability with which time-limited offers are chosen.[16]

[15] The random realisation of the value of each card in each task was determined before the subject made any decisions about turning over cards. Thus, 'good' and 'bad' cards are well-defined even if the subject never saw them turned over.

[16] As an additional check, we compared the Lowess-smoothed choice frequency functions for the choice of the time-limited offer (i) using data only from the first fifteen card tasks faced by each subject and (ii) using only the last fifteen card tasks. We made these comparisons separately for 4s and 12s tasks. In each case, there was no significant difference between the functions. See Appendix 4, section

In Model 3, *Good offers seen* has a positive and highly significant effect, indicating reinforcement or regret-based learning. Similarly, *Good offers rejected* has a positive and highly significant effect in Models 1 and 2. The parallel between the effects of these two experience variables strongly suggests that, despite the confounds discussed in the previous paragraph, reinforcement or regret-based learning is at work in both treatments. To investigate whether this learning mechanism was affected by time pressure, we tried adding variables to pick up interactions between *Available time* and *Good offers seen* (in Model 3) or *Good offers rejected* (in Models 1 and 2). We found no significant interaction effects (see Appendix 4, section 6).

In our analysis so far, we have used *Expected value difference* as a summary measure of the net advantage to be expected from choosing rather than rejecting the time-limited offer. However, it is also useful to investigate how subjects used specific items of information about the offers in a task. We therefore estimated three additional models in which *Expected value difference* was replaced by three *offer variables* that were directly observed by subjects: the price of the time-limited offer (*Red card*), the position of that offer in the sequence of cards (*Position*), and the lowest free-recall price revealed before the appearance of the time-limited offer (*Best blue card*). Because *Best blue card* is undefined if the time-limited offer is the first in the sequence, we restrict the analysis to cases in which that offer appeared second ($Position = 0$) or third ($Position = 1$). Notice that the distribution of outcomes resulting from the rejection of a time-limited offer is unambiguously better, the lower the values of *Best blue card* and *Position* (i.e., the more free-recall offers remain to be revealed). These models are reported in Table 2.2.

5.

Table 2.2: Regression results using offer variables

	(1)		(2)		(3)	
	β	ME	β	ME	β	ME
Treatment	0.307	0.064				
	(0.224)	(0.047)				
Red card	-2.146^{***}	-0.452^{***}	-2.461^{***}	-0.514^{***}	-1.895^{***}	-0.403^{***}
	(0.128)	(0.028)	(0.215)	(0.050)	(0.155)	(0.033)
Position	0.829^{***}	0.175^{***}	0.762^{***}	0.159^{***}	0.860^{***}	0.183^{***}
	(0.178)	(0.038)	(0.264)	(0.056)	(0.237)	(0.050)
Best blue card	0.097^{**}	0.020^{**}	0.074	0.015	0.129^{**}	0.028^{**}
	(0.040)	(0.008)	(0.061)	(0.013)	(0.052)	(0.011)
Available time	0.673^{***}	0.143^{***}	0.681^{***}	0.143^{***}	0.647^{***}	0.139^{***}
	(0.169)	(0.036)	(0.255)	(0.054)	(0.226)	(0.049)
Good offers rejected	1.866^{***}	0.393^{***}	2.780^{***}	0.580^{***}		
	(0.338)	(0.071)	(0.566)	(0.117)		
Good offers seen					1.232^{***}	0.262^{***}
					(0.461)	(0.098)
Period	-0.001	-0.000	-0.013	-0.003	0.011	0.002
	(0.010)	(0.002)	(0.016)	(0.003)	(0.014)	(0.003)
Constant	1.190^{**}		2.089^{**}		0.469	
	(0.595)		(0.900)		(0.770)	
Observations	1444		738		706	
LR chi2	292.027		134.373		155.455	
Prob > chi2	0.000		0.000		0.000	
Baseline predicted probability	-0.197		-0.224		-0.212	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. We used panel data to estimate all these models. The data used to estimate model 2 contains 1444 observations from 105 subjects. The data used to estimate model 3 contains 706 observations from 104 subjects. For each model, the left column contains coefficients, and the right column report marginal effects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

In all three models, *Red card*, *Position* and *Best blue card* have their expected signs. The first two variables are highly significant in all three models. *Best blue card* is significant at the 5 per cent level in Model 1 (using data from both treatments) and Model 3 (using only data from the Regret Feedback treatment). As in the models reported in Table 2.1, *Good offers seen* and *Good offers rejected* are highly significant; there are no significant interactions between these variables and *Available time* (see Appendix 4, section 7). Nor are there any significant interactions between the offer variables and *Available time* (see Appendix 4, section 8).

2.5 Discussion

In this section, we discuss three unexpected features of our results.

The first of these is the absence of a regret feedback effect in the predicted direction. Previous experiments and surveys have found a tendency for individuals to choose options that reduce their exposure to regret (Zeelenberg *et al.*, 1996; Zeelenberg and Pieters, 2004). Our No Feedback treatment was set up so that choosing a time-limited offer would eliminate the feedback that could otherwise induce regret, but we found no evidence that this contributed to the attractiveness of such offers.

One possible explanation of this difference in findings is that the regrets that our time-limited tasks were likely to induce were relatively mild. Consider any consequential task in which the time-limited offer was rejected and in which this offer was lower than the lowest free-recall offer. We define the associated 'regret' as the lowest free-recall offer minus the time-limited offer. Using data only from cases that gave rise to regret, the mean (median) regret from rejecting time-limited offers was 1.13EP (0.70EP) in the Regret Feedback treatment and 1.46EP (1.07EP) in the No Feedback Treatment. In the Regret Feedback treatment, if the time-limited offer was accepted and if this offer was higher than the lowest free-recall offer, 'regret' is defined as the time-limited offer minus the lowest free-recall offer. In this case, the mean (median) regret was 0.93EP (0.75EP). In contrast, the experiments reported by Zeelenberg et al. involve hypothetical choices between binary lotteries in which the worse outcome is always winning nothing. The case studied by Zeelenberg and Pieters is particularly extreme: a person who chooses not to buy a ticket in the Dutch Postcode Lottery is exposed to the possibility of massive regret. Such extreme cases are obviously relevant in some natural-world cases. That said, if exposure to regret is measured in relative terms (i.e. based on the comparison in value between 'what is' and 'what might have been'), the modest exposure to regret in our experiment is typical of a number of natural-world price search problems.

Although our subjects seem not to have been influenced by *anticipations* of regret, their willingness to accept time-limited offers was affected by their previous experience of offer realisations (as explained in Section 2.4.5). Since this experience took different forms in the No Feedback and Regret Feedback treatments, small differences in choice frequencies between the two treatments (such as those evident in Figure 2.3) are not particularly surprising. However, the absence of a regret feedback effect suggests that the mechanism by which experience was encoded and recalled was not driven by the aversive effects of regret.

The second unexpected feature of our results is the high degree of consistency, at both the aggregate and individual levels, between time-limited offer choices made even under high time pressure and the corresponding lottery choices. This consistency is particularly noteworthy given the very different presentation of the two decision problem. It is also noteworthy as subjects' decisions about time-limited offers were strongly influenced by realisations of offer prices in previous tasks. That influence suggests that price search tasks were treated as decisions from experience. In contrast, since no lottery risk was resolved until the end of the experiment, lottery tasks were necessarily decisions from description.[17]

It is clear from the results reported in Section 2.4.6 that, in choosing whether or not to accept a time-limited offer, subjects took account of the value of that offer and of its position in the sequence of offers. The evidence also suggests that they took account of the value of the best previous free-recall offer. The implication is that subjects were using quite sophisticated heuristics that were capable of identifying the main determinants of the distributions of earnings implied by the two options they faced. Even so, if our findings are viewed in the perspective of behavioural economics, the individual-level consistency of attitudes to risk between even 4s price search tasks and lottery tasks is a striking regularity that needs to be explained. Part of the explanation may be that the risky options in our experiment (i.e., rejecting the time-limited offer or choosing Lottery 2) had relatively low variance. Thus, these tasks did not give much scope for the under- and over-weighting of small probabilities that is a major cause of the description-experience gap (see Section 2.4.5).

The final unexpected feature of our results is that time-limited offers were more likely to be chosen when the time available for decision-making was 12s than when it was 4s. Contrary to our prior expectations, we found no direct evidence that subjects used simpler decision heuristics when time pressure was greater. At both time intervals, subjects took account of the same three items of task-specific information (the position and price of the time-limited offer and the value of the best preceding free-recall offer) and ignored the same fourth item (the nature of the feedback they would receive). At both intervals, they adapted their decisions in the light of realisations of offer values in previous tasks. At both intervals, the degree of dispersion in responses and the (very low) frequency of choices of dominated offers were almost the same. Impatient choices were more common in 12s tasks, but impatience was not correlated with the acceptance of time-limited offers. The only systematic difference we have been able to find is that decisions were more risk-averse at 12s than at 4s. This difference was highly significant.

The most natural interpretation of this finding is that the 12s time interval allowed subjects to take account of some additional factor, not present in lottery tasks, that favoured the choice of the time-limited offer. To put this another way, if one models the choice between accepting and rejecting a time-limited offer in terms of the corresponding probability distributions of outcomes, one is neglecting some force of attraction that is exerted by time-limited offers but which decision-makers tend to neglect under time pressure. We conjecture that this attraction is associated with the certainty of the time-limited offer price and with the salience of this certainty in a price-search problem. The 12s interval may have allowed subjects time to think

[17] To check whether subjects' lottery choices might have been influenced by realisations of offer values in the preceding card tasks, we estimated logit models in which the dependent variable was the choice of Lottery 1 and the independent variables included *Good offers seen* (in the Regret Feedback treatment) or *Good offers rejected* (in the No Feedback treatment). We found no significant effects. See Appendix 4, section 9. This is further evidence that comparisons between responses to price search and lottery tasks are not distorted by order effects.

about – and perhaps to worry about – the distinction between certainty and uncertainty.

2.6 Conclusion

Existing industrial organisation models of the effects of time-limited offers assume that, in their search behaviour, consumers are rational and risk-neutral. Our main objective was to investigate whether the search-deterring effects of time-limited offers was intensified or mitigated by behavioural factors. Our conclusion is that these effects are intensified, particularly (and surprisingly) when consumers are *not* subject to high time pressure. If one is thinking about why a 'behavioural' consumer might be particularly attracted by a time-limited offer, the most obvious source of that attraction is the perception that to reject such an offer is to accept uncertainty rather than certainty. Our findings call into question the generality of the assumption that, under time pressure, individuals use heuristics that impart a bias towards certainty. One should be wary of extrapolating too directly from laboratory behaviour to real markets, but our results raise the possibility that time-limited offers may be less likely to be accepted if they are presented in ways that subject consumers to extreme time pressure.

Also surprising was the degree of consistency between consumers' decisions about time-limited offers, made even under intense time pressure, and their responses to problems of binary choice between lotteries. The context-dependence of preferences is a recurring theme in behavioural economics, but this is a case in which the context-*independence* of preferences require a psychological explanation, and its implications for consumer decision making need thinking through.

2.7 Appendix

2.7.1 Appendix 1: Instructions for experiment

Welcome to today's experiment and thanks for coming. In this experiment, you will need to make a series of choices. You will receive your earnings from this experiment at the end.

I shall say more about what will be involved in the experiment soon. Before I do this, I need to set some ground rules, which you must all observe. There must be no talking during the experiment unless you want to ask a question. In this case, simply raise your hand and I will come to you. You must not attempt to look at what other people are doing.

Please keep to these simple rules, because anyone breaking them may be asked to leave the experiment without payment.

I will now describe the nature of the tasks within the experiment.

Tasks

This experiment contains two parts. In both parts, you can earn a certain number of experimental points depending on the decision you made. At the end of the experiment, these points may be converted to money earnings. Details will be given later. You will be paid the sum of your money earnings in these two parts.

Part 1

In Part 1, there are 30 tasks. In each task, you will be given 10 points as your initial budget, and your job is to buy a good with these points. You will have the chance to see 6 offer prices for this good. All these prices range from 0.00 to 10.00 points. Therefore, none of them will exceed your initial budget.

During the task, you need to choose one of these offer prices. The price of your chosen offer is what you will pay. Your point earnings will be equal to 10 points minus this price. Please remember that you cannot keep all these 10 points as your earnings, because you always have to choose one offer price in each task.

Now, I am going to describe the offer prices in detail.

In each task, these 6 offer prices will be presented on 6 separate cards, 5 blue and 1 red. The offer price on each blue card will be in the range from 0.00 to 10.00 points. The actual price for each blue card will be generated randomly by the computer. Each price in the range will be equally likely. The offer price on the red card will be in the range from 0.00 to 4.00 points. The actual price will be generated randomly by the computer. Each price in the range will be equally likely.

At the beginning of each task you will only be able to see the backs of these 6 cards. You will not know which card is the red card. The picture below shows you how the screen will look.

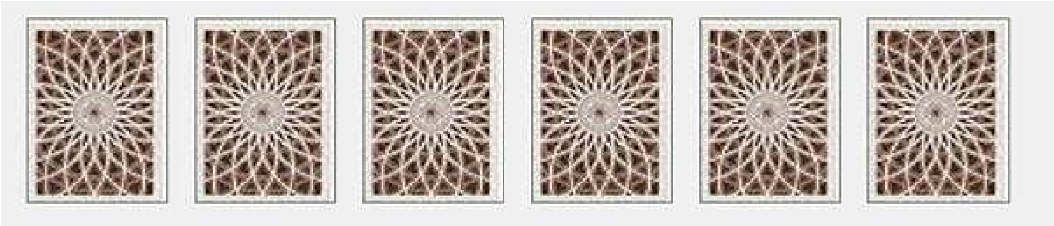

As time goes on, cards will be turned over one by one in order from left to right. The time interval between cards being turned over will be fixed within each task, but may vary between tasks.

Once a blue card is turned over, it will stay turned over. This means that the offer on the card will stay available throughout the task.

There are two types of red card. One is called a standard red card. Once a standard red card is turned over, it will stay turned over, in the same way that a blue card does, and so the offer on the card will stay available throughout the task. The other type of red card is called a time-limited red card. The offer on a time-limited red card will be available only for a certain number of seconds. When this time is out, the card will be turned back. Once it is turned back, you cannot choose that offer any more.

Until the red card is turned over, you will not know whether it is a standard red card or a time-limited red card. You will discover this only when the card is turned over. As shown in the example below, the message under the red card will tell you which type of card it is.

If the message below the red card says ‘offer available throughout the task’, this means that it is a standard red card. The picture below is an example of this type of red card and the corresponding message.

The picture below is an example of how the standard red card looks if you wait until the very end of the task before choosing an offer.

If the message under the red card says 'offer available for x seconds', this means that it is a time-limited red card and will be turned back after these x seconds. There will be a countdown clock below the message reminding you how many seconds are left before it will be turned back. The picture below is an example of this type of red card, the message, and the countdown clock.

The picture below shows how the time-limited red card looks after it is turned back, if you wait until the very end of the task before choosing an offer.

So if you meet a time-limited red card, you need to decide whether to choose the offer on the card before it is turned back.

[**For *NF (No feedback)* treatment**]

To choose an offer, you need to click the grey button with the words 'Click to choose this offer' on it below the card that you want to choose. If you accept an offer, the task ends. The offer price on the card you picked will then be shown on the screen. This is the final price you have chosen to pay for the good in the task. After you click the 'Next task' button below this price, the computer will move on to the next task.

[**For *RF (Regret feedback)* treatment**]

To choose an offer, you need to click the grey button with the words 'Click to choose this offer' on it below the card that you want to choose. If you accept an offer, the task ends. If you chose an offer before all the cards are turned over, the remaining cards will then be turned over to show you the offers on these cards. You will not be able to change your decision at this stage. The offer price on the card you picked will then be shown on the screen. This is the final price you have chosen to pay for the good in the task.

At the end of the experiment, two tasks will be picked at random from these 30 tasks by the computer. Your final point earnings in Part 1 will be the sum of the points you earned in these two selected tasks. Finally, your point earnings will be converted to money earnings at the exchange rate of £1 for every 2.5 experimental points. Please raise your hand if you have any questions. Before starting to take decisions, we ask you to answer some questions in the next several screens. The purpose of these questions is to check whether you have understood these instructions. Any mistake you may make in doing these questions will not affect your final money earnings. When you have finished Part 1, please remain seated. I will distribute the instruction for Part 2 after everyone has finished Part 1.

Part 2

In this part of the experiment, you will have the opportunity to earn an additional amount of experimental points according to the decisions you make.

There are 15 lottery tasks, each of which will require you to choose between two lotteries. At the end of the experiment, one of these lottery tasks will be picked by the computer at random and played out for real. Your point earnings in Part 2 will be determined by your decision in this lottery task.

For each lottery, the boxes beside the option label represent the possible outcomes. Points you can earn from each outcome are shown in the box. For each lottery and outcome, the number of chances out of 100 that you will get this outcome if you choose this lottery is shown at the bottom of the outcome box as a percentage. The picture below is an example of how these lottery tasks look.

In every lottery task, Lottery 1 has only one outcome, as in this example. So there is only one box beside the lottery label, and the percentage chance shown below the box is 100%.

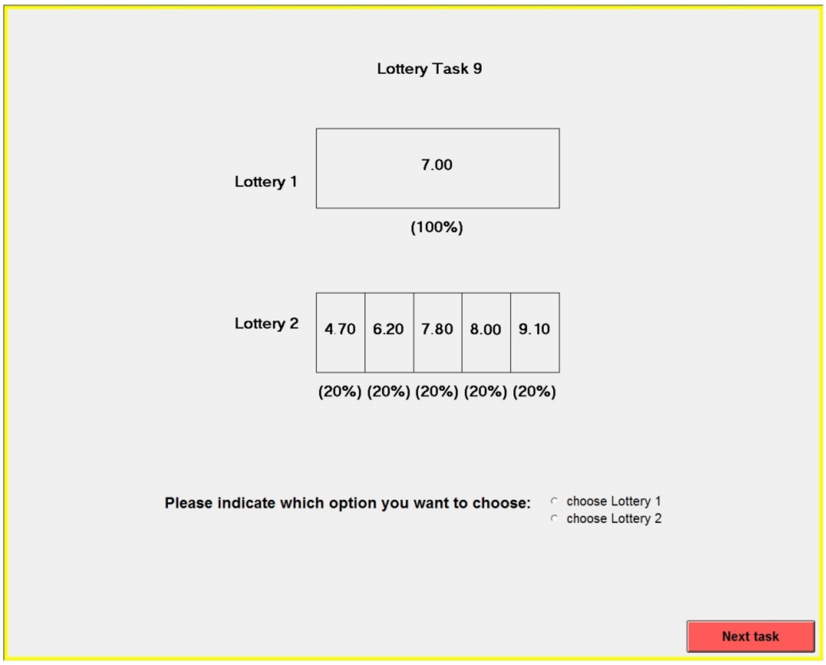

This means that, if you choose Lottery 1, you are certain to earn the number of points shown in the box.

In every lottery task, Lottery 2 has five outcomes, all of which are equally likely, as in this example. So there are five boxes beside the lottery label, and the percentage chance shown below each box is 20%. This means that if you choose Lottery 2, there are 20 chances out of 100 that you will earn the number of points shown in the first box, 20 chances out of 100 that you will earn the number of points shown in the second box, and so on.

To choose a lottery, you need to click 'Choose Lottery 1' or 'Choose Lottery 2'. After you click the 'Next task' button at the bottom of the screen, the computer will move on to the next task.

[For *NF (No feedback)* treatment]

At the end of the experiment, one lottery task will be picked at random from these 15 tasks by the computer. If you chose Lottery 1, your point earnings will be the number of points shown in the box for Lottery 1, and this will be the end of Part 2 of the experiment. If you chose Lottery 2, the computer will then determine the outcome of that lottery by picking one of the five boxes in that lottery at random. The box that the computer has picked will be highlighted on the screen. Your point earnings will be the number of points shown in the highlighted box. The picture below is an example of how the outcome of Lottery 2 is shown.

[For *RF (Regret feedback)* treatment]

At the end of the experiment, one lottery task will be picked at random from these 15 tasks by the computer. The computer will then determine the outcome of Lottery 2 in that task by picking one of the five boxes in that lottery at random. The box that the computer has picked will be highlighted on the screen. You will always see which box the computer has picked, but this will affect your earnings only if you chose Lottery 2 in this task. If you chose Lottery 1, your point earnings will be the number of points shown in the box for Lottery 1. If you chose Lottery 2, your point earnings will be the number of points shown in the highlighted box. The picture below is an example of how the outcome of Lottery 2 is shown.

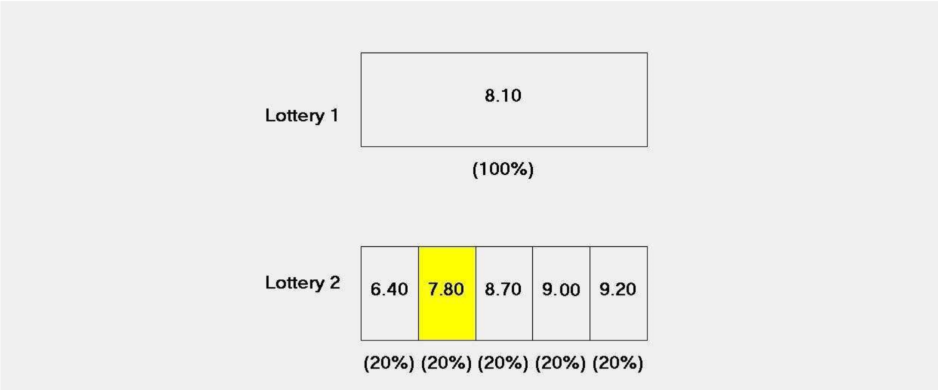

Finally, your point earnings will be converted to money earnings at the exchange rate of £1 for every 2.5 experimental points. These will be added to your earnings from Part 1.

Please raise your hand if you have any questions. Before starting to take decisions, we ask you to answer some questions in next several screens. The purpose of these questions is to check whether you have understood these instructions. Any failure you may make in doing these questions will not affect your final money earnings. When you have finished Part 2, please remain seated and wait patiently until you are paid.

2.7.2 Appendix 2: Impatience and dominance violations in card tasks

We define a response as the behaviour of a specific subject in a specific task. In a price search task, every possible response involves a decision to accept one of the six cards presented in that task. Each task has possible *stages* $s = 1, ..., 6$, some of which may not be reached. Stage 1 starts when the first card is turned over and ends when an offer is accepted or when the second card is turned over (whichever comes first). Stages $2, ..., 5$ are defined similarly. Stage 6 starts when the last card is turned over and ends when an offer is accepted. Thus, for each response there is exactly one stage in which an offer is accepted.

Table 2.3 reports responses separately for four types of task, classified according to whether the task is free-recall or time-limited, and whether the time interval is 4s or 12s. We use the following mutually exclusive and exhaustive classification of responses, which is slightly different for free-recall and time-limited tasks.

Table 2.3: Dominance violations, impatient responses and rational responses

	Number of responses				Average loss			
	time-limited tasks		free-recall tasks		time-limited tasks		free-recall tasks	
	4s	12s	4s	12s	4s	12s	4s	12s
Rational responses								
free-recall offer accepted at stage 6	771	630	1221	1116	–	–	–	–
	49.27%	40.13%	78.12%	70.99%				
time-limited offer accepted	650	749	–	–	–	–	–	–
	41.53%	47.71%						
Impatient responses								
free-recall offer accepted at stage 1	23	35	68	99	0.121	0.459	0.223	0.286
	1.47%	2.23%	4.35%	6.30%				
free-recall offer accepted at stage 2: (i) free-recall tasks	–	–	81	88	–	–	0.127	0.154
			5.18%	5.60%				
(ii) time-limited offer seen at stage 1	18	19	–	–	0.074	0.196	–	–
	1.15%	1.21%						
(iii) time-limited offer not seen at stage 1	11	22	–	–	0.113	0.269	–	–
	0.70%	1.40%						
free-recall offer accepted at stage 3	23	22	81	107	0.15	0.128	0.108	0.17
	1.47%	1.40%	5.18%	6.81%				
free-recall offer accepted at stage 4	28	33	32	67	0.13	0.112	0.054	0.116
	1.79%	2.10%	2.05%	4.26%				
free-recall offer accepted at stage 5	21	42	47	59	0.086	0.085	0.063	0.111
	1.34%	2.68%	3.01%	3.75%				
Dominance violations								
Accepted offer dominated by previous offer	20	18	33	36	2.551	2.033	2.947	3.173
	1.28%	1.15%	2.11%	2.29%				
Total	1565	1570	1563	1572				

Notes:Percentages do not always add up to 100% due to rounding.

Classification of responses to free-recall tasks

Starting with the set of all responses to the relevant class of free-recall tasks, we first pick out *dominance violations.* A response is classified as a dominance violation if, at the stage at which the subject accepted an offer, another offer with a strictly lower price was visible.

Considering only responses that are not dominance violations, we then pick out *impatient responses.* A response is classified as impatient if the subject accepted an offer (with non-zero price) at any of the stages $1, ..., 5$. Impatient responses are then disaggregated according to the stage at which the offer was accepted.

Responses that are neither dominance violations nor impatient are classified as rational. These are responses in which the best (or an equal-best) of the six offers was accepted at stage 6, or if zero-price offer was accepted at any stage. Notice that if a response is consistent with expected utility theory for *some* attitude to risk, it is classified as rational.

Classification of responses to time-limited tasks

Starting with the set of all responses to the relevant class of time-limited tasks, we first pick out *dominance violations.* A response is classified as a dominance violation if, at the stage at which the subject accepted an offer, another offer with a strictly lower price was visible.

Considering only responses that are *not* dominance violations, we then pick out *impatient* responses [18]. A response is classified as impatient if at any of stages 1, ..., 5, the subject accepted a free-recall offer. Impatient responses are then disaggregated according to the stage (1, ..., 5) at which an offer was accepted. Notice that if an impatient response was made at stage 1, the subject ended the task *before* seeing the time-limited offer. If an impatient response was made at stages 3, 4 or 5, the subject ended the task *after* seeing the time-limited offer. But impatient responses at stage 2 can be further disaggregated according to whether, at this stage, the subject had seen the time-limited offer (i.e. the time-limited offer was the first or second card) or had not seen it.

Responses that are neither dominance violations not impatient are classified as *rational.* These responses are then subdivided according to whether the accepted offer was free-recall (in which case, by our definitions, it was the best or equal-best free-recall offer and was accepted at stage 6) or time-limited.

Average loss

For each response that was a dominance violation or was impatient, we find the expected value of the price that the subject would have paid if, instead of behaving as she did, she had waited to stage 6 and then chosen the best available offer. This expectation is the *counterfactual price.* In free-recall tasks, the counterfactual price is always the price that the subject would have paid, had she behaved in a fully rational way. In these tasks it is therefore natural to define the subject's *loss* (i.e., the loss resulting from dominance violation or impatience) as

[18] As noted in Section 2.4.1, a response is also classified as impatient if the subject chose a time-limited offer with a price exactly equal to the best previous offer. This possibility was never observed.

the price she actually paid minus the counterfactual price. For each of our classifications of dominance-violating or impatient responses to free-recall tasks, Table 2.3 shows the average loss to subjects who made the relevant responses.

Now consider time-limited tasks. Consider impatient responses made at stages 3, 4 and 5, and impatient responses made at stage 2 when the time-limited offer had been seen (necessarily, at stage 1). In all these cases, the subject has chosen to reject the time-limited offer. The counterfactual price is therefore the expected value of the price that the subject would have paid, conditional on rejecting the time-limited offer, had her behaviour been fully rational. For each of these classifications of responses, the difference between actual and counterfactual prices can be interpreted as a loss resulting from impatience. The corresponding average losses are shown in Table 2.3. For completeness, the 'average loss' column for time-limited tasks also shows the average differences between actual and counterfactual prices for (a) impatient choices made before the time-limited offer was seen and (b) choice of dominated offers. In case (a), the entries in the table may understate the true loss from impatience, as they do not take account of the option of choosing the time-limited offer. In case (b), the entries can be interpreted as losses resulting from impatience and dominance violation.

In Section 2.4.1, we noted that in free-recall tasks, impatient responses were more common in 12s tasks than in 4s tasks, and that average losses due to impatience were very low. Table 2.3 shows the same regularities in time-limited tasks.

2.7.3 Appendix 3: Expected value difference

First consider the following abstract lottery. An individual will make m independent draws from a uniform distribution of values with support $[0, k]$. She will then receive the maximum of (i) the highest-valued of these draws and (ii) some constant value z. The expected value of this lottery is:

$$\frac{k(m + z^{m+1})}{m + 1} \tag{2.1}$$

Now consider a subject in the experiment who faces a consequential time-limited task. Assume that she does not accept any offer before the stage $s \in 1, 2, 3$ at which the time-limited offer is revealed. If $s > 1$, let b be the lowest free-recall price so far revealed. If $s = 1$, define $b \equiv 10$ (i.e., the highest possible price, in EP units). Let p be the time-limited price. Since the task is consequential, $p < b$. Suppose the subject rejects the time-limited offer and waits until all offers have been revealed. She will see $5 - s$ additional free-recall offers. The payoff associated with each of these offers is an independent random draw from a uniform distribution with support $[0, 10]$. (Recall that the subject has a budget of 10EP for each task; her payoff is the budget *minus* the price of the offer she accepts.) If she makes a rational choice at this stage, her payoff is the maximum of (i) the payoff associated with the best of these additional offers and (ii) $10 - b$. Thus, as viewed at stage s, the option of rejecting the time-limited offer and then acting rationally is formally equivalent to the lottery described in the first paragraph. The expected value of this option is given by (1) when $k = 10, m = 5 - s$ and $z = 10 - b$. The *expected value difference* is defined as the payoff from accepting the time-limited offer (i.e., $10 - p$) *minus* the expected value of the payoff from rejecting it. Equivalently, it is the expected value of the price paid if the time-limited offer is rejected *minus* the time-limited price.

2.7.4 Appendix 4: Additional statistical analysis

1. Treatment effect

Table 2.4: Treatment effect

	(1)	
	β	ME
Expected value difference	1.874***	0.452***
	(0.088)	(0.021)
Treatment	0.253	0.061
	(0.156)	(0.038)
Constant	0.355**	
	(0.118)	
Observations	2526	
Pseudo R^2	0.403	
LR chi2	463.454	
Prob > chi2	0.000	
Baseline predicted probability	0.416	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 % level. Cluster-robust standard errors in parentheses. Cluster-robust standard errors in parentheses. The dependent variable in the model is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. For the model, the left column contains coefficients, and the right column report marginal effects. Results for the model are based on logit estimations with subject-level clustering.

2. Effect of time pressure

Table 2.5: Effect of time pressure

	(1)	
	β	ME
Expected value difference	1.879***	0.453***
	(0.090)	(0.022)
Available time	0.540***	0.129***
	(0.111)	(0.027)
Constant	0.220**	
	(0.100)	
Observations	2526	
Pseudo R^2	0.409	
LR chi2	433.274	
Prob > chi2	0.000	
Baseline predicted probability	0.417	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 % level. Cluster-robust standard errors in parentheses. The dependent variable in the model is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. For the model, the left column contains coefficients, and the right column report marginal effects. Results for the model are based on logit estimations with subject-level clustering.

3. Lottery choices: No Feedback and Regret Feedback

Figure 2.9: Lottery choices: No Feedback and Regret Feedback

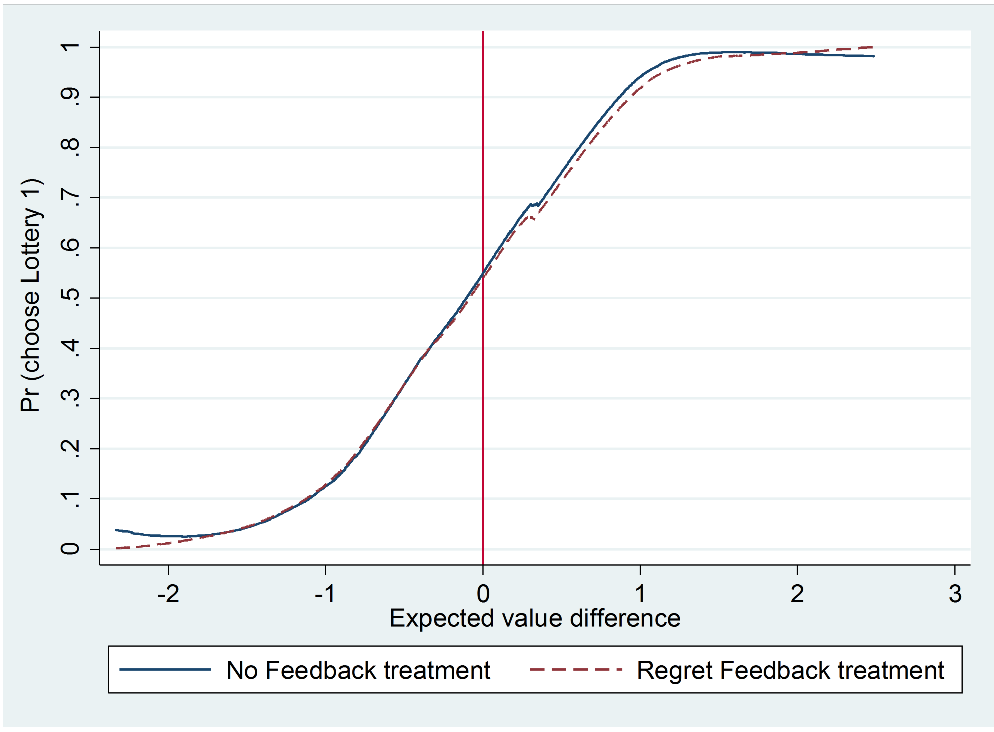

4. Logit model combining card and lottery tasks

Table 2.6: Logit model combining card and lottery tasks

	(1)	
	β	ME
Expected value difference	-2.288^{***}	-0.560^{***}
	(0.101)	(0.025)
4s	0.007	0.002
	(0.133)	(0.033)
12s	0.616^{***}	0.146^{***}
	(0.128)	(0.028)
Constant	0.258^{***}	
	(0.068)	
Observations	5052	
Pseudo R^2	0.486	
LR chi2	520.405	
Prob > chi2	0.000	
Baseline predicted probability	0.428	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 % level. Cluster-robust standard errors in parentheses. The dependent variable in the model is a dummy equal to 1 if the subject chose the time-limited offer/Lottery 1 and 0 if the subject rejected the time-limited offer/Lottery 1 in the task. For the model, the left column contains coefficients, and the right column report marginal effects. Results for the model are based on logit estimations with subject-level clustering.

5. Comparisons between first and second fifteen periods

Figure 2.10: Comparisons between first and second fifteen periods

(a) 12s card task

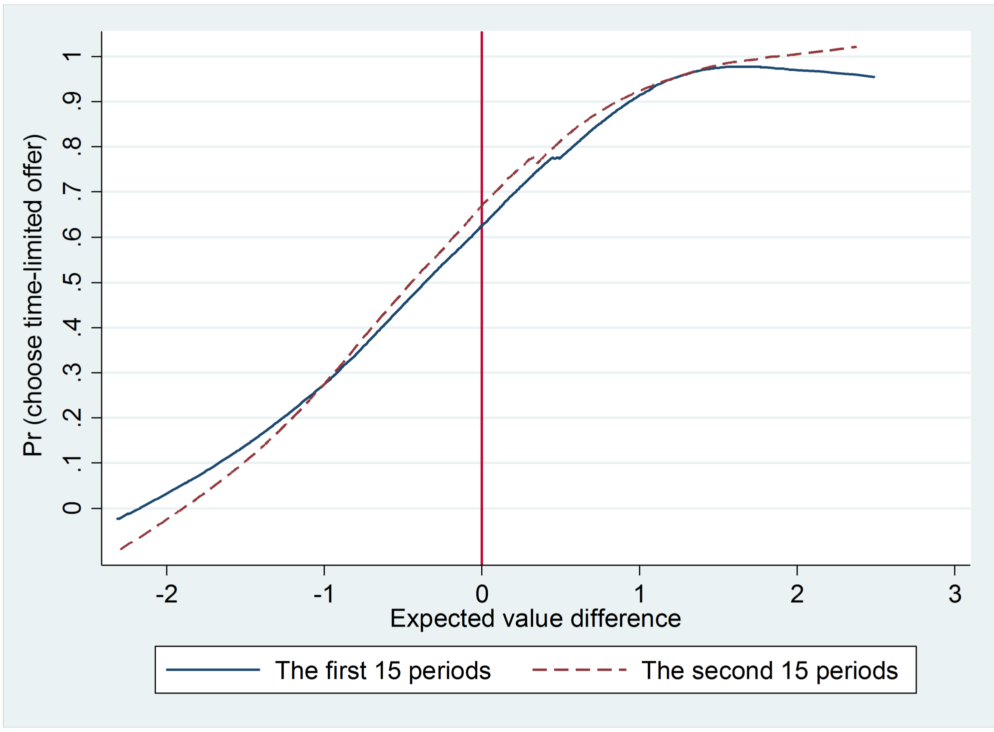

(b) 4s card tasks

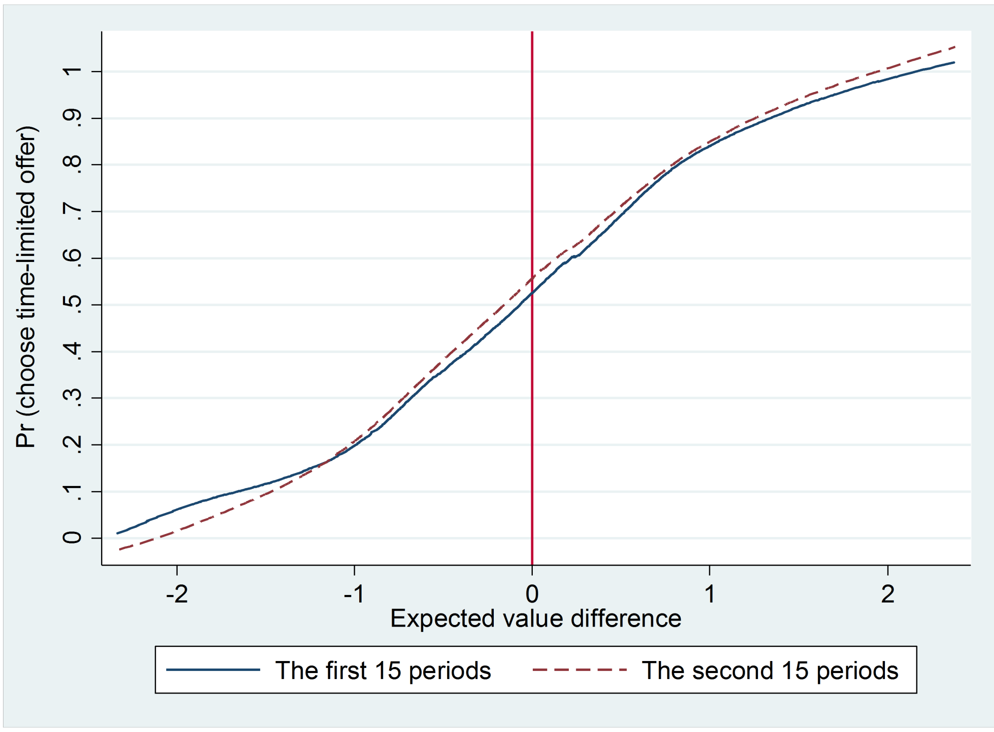

Table 2.7: Regression

	4s		12s	
	(1)		(2)	
	β	ME	β	ME
Expected value difference	1.752***	0.438***	2.038***	0.434***
	(0.107)	(0.027)	(0.130)	(0.026)
Period dummy	0.082	0.020	0.107	0.023
	(0.157)	(0.039)	(0.163)	(0.035)
Constant	0.168		0.748***	
	(0.117)		(0.129)	
Observations	1294		1232	
Pseudo R^2	0.379		0.433	
LR chi2	266.155		251.460	
Prob > chi2	0.000		0.000	
Baseline predicted probability	0.381		0.453	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 % level. Cluster-robust standard errors in parentheses. The dependent variable in the model is a dummy equal to 0 if the task is in period 1 to period 15, and 1 if the task is in period 16 to period 30. For the model, the left column contains coefficients, and the right column report marginal effects. Results for the model are based on logit estimations with subject-level clustering.

6. Models using Expected value difference

Table 2.8: Interactions between *Available time* and *Good offers seen* or *Good offers rejected*: models using *Expected value difference*

	Overall (1) β	NF (2) β	RF (3) β
Treatment	0.330^{*}		
	(0.200)		
Expected value difference	2.397^{***}	2.692^{***}	2.138^{***}
	(0.104)	(0.168)	(0.128)
Available time	0.762^{***}	0.812^{***}	0.909^{***}
	(0.171)	(0.258)	(0.327)
Good offers rejected	2.224^{***}	3.268^{***}	
	(0.331)	(0.540)	
Available time x Good offers rejected	-0.271	-0.686	
	(0.425)	(0.648)	
Good offers seen			1.625^{***}
			(0.448)
Available time x Good offers seen			-0.441
			(0.626)
Period	0.004	0.006	0.004
	(0.008)	(0.012)	(0.011)
Constant	-0.585^{***}	-0.823^{***}	-0.476^{*}
	(0.210)	(0.291)	(0.277)
Observations	2456	1246	1210
LR chi2	537.186	257.388	280.652
Prob > chi2	0.000	0.000	0.000
Baseline predicted probability	-0.796	-1.070	-0.643

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. We used panel data to estimate all these models. The data used to estimate model 2 contains 1246 observations from 105 subjects. The data used to estimate model 3 contains 1210 observations from 104 subjects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

7. Models using offer variables

Table 2.9: Interactions between *Available time* and *Good offers seen* or *Good offers rejected*: models using *offer variables*

	Overall	NF	RF
	(1)	(2)	(3)
	β	β	β
Treatment	0.307		
	(0.225)		
Red card	-2.146^{***}	-2.461^{***}	-1.895^{***}
	(0.128)	(0.215)	(0.155)
Position	0.828^{***}	0.763^{***}	0.861^{***}
	(0.178)	(0.265)	(0.237)
Best blue card	0.097^{**}	0.074	0.129^{**}
	(0.040)	(0.061)	(0.052)
Available time	0.653^{***}	0.694^{**}	0.670
	(0.223)	(0.330)	(0.436)
Good offers rejected	1.826^{***}	2.806^{***}	
	(0.442)	(0.702)	
Available time x Good offers rejected	0.078	−0.051	
	(0.548)	(0.845)	
Good offers seen			1.260^{*}
			(0.653)
Available time x Good offers seen			−0.052
			(0.845)
Period	−0.001	−0.013	0.011
	(0.010)	(0.016)	(0.014)
Constant	1.205^{**}	2.079^{**}	0.455
	(0.604)	(0.915)	(0.804)
Observations	1444	738	706
LR chi2	292.001	134.398	155.430
Prob > chi2	0.000	0.000	0.000
Baseline predicted probability	−0.196	−0.224	−0.212

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. We used panel data to estimate all these models. The data used to estimate model 2 contains 738 observations from 105 subjects. The data used to estimate model 3 contains 706 observations from 104 subjects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

8. Interactions between Available time and offer variables

Table 2.10: Interactions between *Available time* and *offer variables*

	Overall	NF	RF
	(1)	(2)	(3)
	β	β	β
Treatment	0.305		
	(0.226)		
Red card	-2.060^{***}	-2.286^{***}	-1.879^{***}
	(0.161)	(0.256)	(0.210)
Position	0.633^{**}	0.395	0.825^{**}
	(0.258)	(0.385)	(0.350)
Best blue card	0.120^{**}	0.044	0.200^{**}
	(0.059)	(0.090)	(0.080)
Available time	0.308	-0.613	1.188
	(1.065)	(1.637)	(1.406)
Good offers rejected	1.846^{***}	2.717^{***}	
	(0.338)	(0.565)	
Good offers seen			1.251^{***}
			(0.464)
Available time x Red card	-0.172	-0.360	-0.041
	(0.199)	(0.309)	(0.265)
Available time x Position	0.385	0.685	0.095
	(0.354)	(0.535)	(0.476)
Available time x Best blue card	-0.042	0.055	-0.125
	(0.079)	(0.122)	(0.105)
Period	-0.001	-0.013	0.011
	(0.010)	(0.016)	(0.014)
Constant	1.389^{*}	2.860^{**}	0.133
	(0.842)	(1.297)	(1.115)
Observations	1444	738	706
LR chi2	291.363	134.392	155.252
Prob > chi2	0.000	0.000	0.000
Baseline predicted probability	-0.179	-0.172	-0.226

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the time-limited offer and 0 if the subject rejected the time-limited offer in the task. We used panel data to estimate all these models. The data used to estimate model 2 contains 738 observations from 105 subjects. The data used to estimate model 3 contains 706 observations from 104 subjects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

9. Choice of Lottery 1

Table 2.11: Choice of Lottery 1

	Overall (1)		NF (2)		RF (3)	
	β	ME	β	ME	β	ME
Good offers rejected	−0.107	−0.027	−0.605	−0.149		
	(0.471)	(0.117)	(0.712)	(0.175)		
Good offers seen					0.262	0.065
					(0.964)	(0.240)
Expected value difference	3.381***	0.839***	3.670***	0.904***	3.146***	0.784***
	(0.155)	(0.039)	(0.248)	(0.062)	(0.198)	(0.050)
Constant	0.377**		0.624**		0.138	
	(0.176)		(0.264)		(0.480)	
Observations	2526		1281		1245	
LR chi2	477.709		219.624		252.986	
Prob > chi2	0.000		0.000		0.000	
Baseline predicted probability	−0.789		−0.788		−0.794	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the Lottery 1 and 0 if the subject chose Lottery 2 in the Lottery task. We used panel data to estimate all these models. The data used to estimate model 2 contains 1281 observations from 105 subjects. The data used to estimate model 3 contains 1245 observations from 104 subjects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

Chapter 3

Does strategy fairness make inequality more acceptable?

3.1 Introduction

Existing theories and experimental evidence on social preferences has identified at least five types of fairness considerations: inequality aversion (Güth *et al.*, 1982; Camerer and Thaler, 1995; Roth, 1995; Camerer, 2011; Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000), intention-based reciprocity (Blount, 1995; Falk *et al.*, 2003; Offerman, 2002; Charness and Rabin, 2002; Charness, 2004; Falk *et al.*, 2008), social welfare preferences (Charness and Grosskopf, 2001; Andreoni and Miller, 2002; Charness and Rabin, 2002), desert-based fairness (Hoffman *et al.*, 1994; Konow, 2000; Cappelen *et al.*, 2007), and procedural fairness as randomness (Bolton *et al.*, 2005; Cox and Deck, 2005). By taking into account people's preference for fairness, these theories can explain many seemingly puzzling behaviours for which conventional economic theories cannot give an explanation, such as the resistance to 'unfair' outcomes in the ultimatum game and cooperative behaviour in the trust game. People's attitudes to fairness have also been considered in the literature on social norms (Fehr and Fischbacher, 2004; Bicchieri, 2008; Bicchieri and Chavez, 2010; Xiao and Bicchieri, 2010; Krupka and Weber, 2013). The literature on social norms focuses on almost the same dimensions of attitudes to fairness as the literature on social preference. However, both literatures have neglected a type of procedural fairness that may be particularly important in market environments and in public choice: fairness of a framework of rules within which individuals pursue self-interest. The conjecture is that inequalities will tend to be seen as acceptable if they come about through the workings of fair rules, even though they are the result of self-interested intentions. To fill in this gap of literature on social preferences, we propose a new concept of fairness: strategy fairness. This is tested using an experiment.

The ultimatum game has been used as a standard experimental design for studying all these different types of fairness preferences (see Section 3.2 for a comprehensive review). In the ultimatum game, subjects are assigned to be either the proposer or the responder. The proposer makes a decision on how to split an endowment between himself and the responder, and then the responder chooses to either accept or reject this proposal. If the responder accepts the offer, the payoff is distributed according to the proposal. If the responder rejects the offer, both players get a zero payoff. The choices of the proposer and the responder reflect their ideas of fairness.

The ultimatum game provides a natural environment for studying fairness considerations regarding the allocation of social welfare. However, one limitation of using the ultimatum game is that it does not provide a measure of fairness considerations in a competitive context. Competitions are ubiquitous in our everyday life. We have sports competitions (e.g. football or basketball), competitions in politics (e.g. elections), competition in business, and competition in education. The distinction between the situation which is described by the ultimatum game and competitions is that participants play different roles in the ultimatum game and move sequentially, but in fair competitions, the positions of participants are completely symmetric, so that each participant has the same strategy set, therefore equal opportunity to win. There is evidence suggesting that people may hold different fairness preferences in situations where they have equal opportunity to compete and in the ultimatum game. In their bargaining games, Isoni *et al.* (2014) find that subjects are more willing to settle on efficient but unequal allocations rather than equal but inefficient allocations when they have equal opportunity to compete. In this paper, we will be concerned with people's fairness preferences in a competition environment. We will introduce a new concept of fairness: strategy fairness. Strategy fairness in competitions implies that every one in the competition has equal opportunity to win.

In the ultimatum game, responders reveal their attitudes towards fairness by rejecting unequal allocations or allocations which signal unkind intention of proposers. However, the right to reduce other people's payoff is asymmetric between the proposer and responder. The proposer is given no opportunity to react to the decision made by the responder. As a result, the cost of making a payoff reducing move is also asymmetric between proposers and responders. Many recent studies have suggested that findings of experiments with games involving asymmetric opportunities of punishment lack external validity. Fehr and Gächter (2000, 2002) show that the existence of punishment opportunities increases the contribution level in public good games. However, Nikiforakis (2008) shows that when both punishment and counter-punishment are allowed in the public good game, cooperators' willingness to punish decreases, which leads to the breakdown of cooperation. In the presence of counter-punishment opportunities, people also reveal their strong desire to reciprocate punishments in the public goods game (Cinyabuguma *et al.*, 2006; Denant-Boemont *et al.*, 2007). Given these facts, an experiment providing equal opportunities to players to punish (or take resources from) one another has obvious advantages in studying fairness preferences. The experiment reported in this paper was designed with that objective.

The experiment had two parts. In Part 1, subjects were paired to compete with one another in either a effort task or a series of card games. In each case, the winner of this competition was given nine lottery tickets, and the loser was given three tickets. One of these lottery tickets entitled the holder to a prize, but the subjects did not know which ticket this was until after the end of Part 2. In Part 2, the same pair of subjects faced a version of the vendetta game, originally introduced by Bolle *et al.* (2014). In this game, subjects were given opportunities to take lottery tickets from their co-players, in alternating turns, to increase their number of tickets by a fraction of what they took. By using or not using these opportunities to take tickets, subjects were able to reveal their attitudes to equality and fairness.

In the card game, players were dealt a card each. Each card had a number of points on it. The player who held the card with the higher number of points at the end of each card game was the winner of game. Both players were given some opportunities to replace cards. The number of replacement opportunities varied between treatments. In the Fair Rule treatment, players were given equal numbers of replacement opportunities. In the Unfair Rule treatment, one player was allowed to change more cards than the other player. The Encryption Task presented in Erkal *et al.* (2011) was used as the real-effort task. In the Encryption Task, subjects were asked to encrypt words by substituting letters with numbers using a given table. The subject who encoded more words or encoded the same number of words in a shorter time was the winner.

The Part 1 tasks were designed to generate unequal outcomes between players. In the series of card games, the rules of the game satisfied strategy fairness in the Fair Rule treatment but not in the Unfair Rule treatment. This feature of the design allowed us to use the vendetta game to make between-treatment comparisons of individuals' attitudes to inequalities generated by fair and unfair rules. Effort tasks are commonly used to test the concept of desert-based fairness (e.g. Burrows and Loomes, 1994; Bosman *et al.*, 2005; Oxoby and Spraggon, 2008). The task in the Real Effort treatment was used as another way of generating an unequal distribution of the lottery tickets. By comparing the Real Effort and Fair Rule treatments, we were able to compare attitudes to strategy fairness and desert-based fairness.

In this chapter, we are not interested in how subjects play Part 1 games – these games just generate inequality by fair or unfair rules. But these games are of independent interest, because they are 'search competitions'. Subjects' behaviour in these games will be analysed in Chapter 4.

We begin by reviewing the existing concepts of fairness and measures of these concepts in more detail (Section 3.2). We then propose a model of strategy fairness (Section 3.3). We describe the basic principles of experimental design (Section 3.4) and the application of the model and our hypotheses (Section 3.5). This is followed by a more detailed description of the experimental design and its implementation (Section 3.6), the results (Section 3.7), and their implications (Sections 3.8 and 3.9).

3.2 Literature on the concept of fairness

Researchers have generally studied a number of fairness considerations: inequality aversion, intention-based reciprocity, social-welfare preferences, desert-based fairness and procedural fairness as randomness in the allocation procedure. In this section, we will review these concepts of fairness in more detail, outlining the measures of them.

Studies show that when making decisions in the ultimatum game, responders are not only concerned about their own payoff, but also care about the distribution of payoffs. Responders frequently reject offers of less than 20 percent (Güth *et al.*, 1982; Camerer and Thaler, 1995; Roth, 1995; Camerer, 2011). The two best-known models of inequality aversion are developed by Fehr and Schmidt (1999) and Bolton and Ockenfels (2000). By adding concerns about the distributive consequence of outcomes, these models explain responders' tendency to reject unequal distribution of payoffs at the cost of efficiency.

A large number of experimental works show that responders are more likely to reject a given offer with an unequal distribution of payoffs when the intentions signalled by proposers' actions are unkind, which implies that inequality aversion is not sufficient to explain responders' punishment behaviour. Blount (1995) finds that responders indicate significantly higher minimum acceptable outcomes when the payoff is allocated by an interested party than by a random device. Falk *et al.* (2003) find that responders are more likely to reject a given offer with an unequal distribution of payoffs when proposers have other alternatives, which offer more equal distribution of payoffs. These pieces of evidence suggest that intention-based reciprocity does better than inequality aversion in explaining responders' different responses to identical offers. Studies with other sequential games also confirm that people tend to reward kind actions and punish unkind actions, even when the choice of rewarding or punishment is costly (Offerman, 2002; Charness and Rabin, 2002; Charness, 2004; Falk *et al.*, 2008).

Most studies of social welfare preference use dictator games or distribution games. These studies show that many subjects are concerned with social efficiency. Charness and Rabin (2002) show that in their dictator games, proposers are willing to sacrifice their money in order to increase efficiency, especially when these sacrifices are inexpensive. They suggest that social welfare preferences provides a better explanation than inequality aversion for helpful sacrifice behaviour. Andreoni and Miller (2002) and Charness and Grosskopf (2001) find similar results with distribution games.

Regarding the process by which the payoff distribution is generated in ultimatum games, Bolton *et al.* (2005) find that in response to a given unequal distribution of payoffs, responders' behaviour differs between when the offer is proposed by a co-player and when it is generated by a random device. Studies show that responders' decisions also vary depending on the fairness of random devices (Bolton *et al.*, 2005; Cox and Deck, 2005). Responders are less likely to accept unfair random offers than fair random offers, which cannot be explained by intention-based reciprocity. All this evidence suggest that responders have preferences for procedural fairness.

Studies of desert-based fairness focus on various entitlement conditions in both ultimatum games and dictator games. There are two dimensions of entitlement: the first dimension implies how the role of first mover is assigned; the second dimension refers to how the initial endowment is produced. Hoffman *et al.* (1994) find that in both ultimatum games and dictator games, the origin of initial entitlements matters – first movers who earn the right to the role by answering more questions correctly offer significantly less to their opponents than first movers who have it assigned randomly. It is also interesting that in the ultimatum game, they did not find any detectable difference in the rejection rate of second movers between treatments with different entitlement conditions. Regarding the second dimension of entitlement, first movers who make more contribution to the initial endowment than their opponents are found to keep more for themselves than first movers who make less contribution (Konow, 2000; Cappelen *et al.*, 2007). Both studies also find that people treat the entitlement earned through factors within individual control (e.g. effort) differently from that earned through factors beyond individual control (e.g. talent) – people find it more fair to allocate payoffs according to factors within individual control than factors beyond individual control. Ruffle (1998) shows that in the dictator game, allocators offer significantly more to recipients who exert effort to create a large pie than to lucky recipients who create a large pie as a result of a coin toss.

3.3 Model of strategy fairness

3.3.1 Concept of strategy fairness

Consider any normal form game. Each player $i \in \{1, 2, ..., n\}$ has a strategy space S_i containing m_i pure strategies, indexed by the numbers $1, ..., m_i$. The set of all strategy profiles, i.e. $S_1 \times ... \times S_n$, is denoted by S. Typical profiles are denoted by $\boldsymbol{s}$, $\boldsymbol{s'}$. For each player i, material payoffs are described by a function $\pi_i : S \to \mathbb{R}$. If there is some m such that $m_i = m$ for all i, the game has *identical strategy spaces*, i.e. $S_i = \{1, ..., m\}$ for all i. To develop a concept of strategy fairness, we need to make comparisons between the opportunities faced by different players in the same game. Consider any game in which, for two distinct players i and j, $m_i = m_j$. We will say that two strategy profiles $\boldsymbol{s}$, $\boldsymbol{s'}$ for this game *differ only by an* $\{i,j\}$ *transposition* if $s'_i = s_j$ (i.e. i's component of $\boldsymbol{s'}$ has the same index number as j's component of $\boldsymbol{s}$), $s'_j = s_i$, and for all $k \neq i, j$, $s_k = s'_k$.

Definition 1 *(Direct fairness).* A game is *directly fair with respect to* two distinct players i and j if $m_i = m_j$ and if, for all pairs of strategy profiles $\boldsymbol{s}$, $\boldsymbol{s'}$ that differ only by an $\{i, j\}$ transposition, $\pi_i(\boldsymbol{s}) = \pi_j(\boldsymbol{s'})$.

Definition 2 *(Direct bias).* A game is *directly biased towards* some player i *relative to* another player j if $m_i = m_j$ and if, for all pairs of strategy profiles $\boldsymbol{s}$, $\boldsymbol{s'}$ that differ only by an $\{i, j\}$

transposition, $\pi_i(\boldsymbol{s}) \geq \pi_j(\boldsymbol{s}')$, with a strict inequality for at least one such pair.

Two corollaries of these definitions will turn out to be significant:

Corollary 1. In any game that is directly fair with respect to two players i and j, $\sum_{\boldsymbol{s}\in S} \pi_i(\boldsymbol{s}) = \sum_{\boldsymbol{s}\in S} \pi_j(\boldsymbol{s})$.

Corollary 2. In any game that is directly biased towards one player i relative to another player j, $\sum_{\boldsymbol{s}\in S} \pi_i(\boldsymbol{s}) > \sum_{\boldsymbol{s}\in S} \pi_j(\boldsymbol{s})$.

However, $\sum_{\boldsymbol{s}\in S} \pi_i(\boldsymbol{s}) = \sum_{\boldsymbol{s}\in S} \pi_j(\boldsymbol{s})$ does not imply that the game is directly fair, and $\sum_{\boldsymbol{s}\in S} \pi_i(\boldsymbol{s}) > \sum_{\boldsymbol{s}\in S} \pi_j(\boldsymbol{s})$ does not imply the existence of direct bias.

'Direct' fairness and bias are defined with respect to arbitrary assignments of index numbers to players' strategies. To remove this limitation, we use the idea that a game can be 'relabelled' by changing this assignment. Consider any game G in which, for two distinct players i, j, $S_i = S_j = \{1, ..., m\}$. A *relabelling* of this game with respect to i and j is a pair of one-to-one mappings $f_i : \{1, ..., m\} \to \{1, ..., m\}$, $f_j : \{1, ..., m\} \to \{1, ..., m\}$. The relabelled game G' is identical to G in all respects except that, for each $k \in \{1, ..., m\}$, the strategy for i that is indexed by k in G is indexed by $f_i(k)$ in G', and the strategy for j that is indexed by k in G is indexed by $f_j(k)$ in G'.

Notice that, because index numbers are transformed separately for the two players, a pair of strategies (one for i and one for j) that have the same index in G may have different indices in G'. However, relabelling cannot affect the value of $\sum_{\mathbf{s}\in S} \pi_i(\mathbf{s})$ or $\sum_{\mathbf{s}\in S} \pi_j(\mathbf{s})$, i.e. the total of all possible payoffs to each player, summing over all strategy profiles. The following theorem can be proved by combining this fact with Corollaries 1 and 2:

Theorem 1. For any game G, for any distinct players i and j, no more than one of the following propositions is true:

[i] There is some relabelling of G with respect to i and j such that the re-labelled game G' is directly fair with respect to i and j.

[ii] There is some relabeling of G with respect to i and j such that the re-labelled game G' is directly biased towards i relative to j.

[iii] There is some relabeling of G with respect to i and j such that the re-labelled game G' is directly biased towards j relative to i.

Theorem 1 legitimates the following definitions:

Definition 3 *(Label-independent fairness)*. A game G has *label-independent fairness with respect to* two distinct players i and j if $m_i = m_j$ and if there is some relabelling of G with respect to i and j such that the re-labelled game G' is directly fair with respect to i and j.

Definition 4 *(Label-independent bias)*. A game G has *label-independent bias towards* some player i *relative to* another player j if $m_i = m_j$ and if there is some relabelling of G with respect to i and j such that the re-labelled game G' is directly biased towards i relative to j.

Table 3.1: Battle of the Sexes 1

		Wife	
		Boxing	Opera
Husband	Boxing	3, 2	0, 0
	Opera	0, 0	2, 3

These definitions can be illustrated by using the 'Battle of the Sexes 1' game 1 in Table 1. Husband gets more payoff than Wife if both of them choose Boxing, while Wife gets more than Husband if both of them choose Opera. If, for each player, we index Boxing by 1 and Opera by 2, we find that the game is neither directly fair with respect to the two players, nor biased towards either of them. However, we can relabel the game for assigning the index 1 (and the name 'Preferred') to Boxing for Husband and to Opera for Wife, and by assigning the index 2 (and the name 'Less Preferred') to Opera for Husband and to Boxing for Wife. The relabelled game, 'Battle of the Sexes 2' is given in Table 2. In Battle of the Sexes 2, both Husband and Wife get zero payoff if they both choose Preferred or both choose Less Preferred. Husband gets 2 if he chooses Less Preferred and Wife chooses Preferred, which is as same as the payoff that Wife gets if she chooses Less preferred and Husband chooses Preferred. Similarly, Husband gets 3 if he chooses Preferred and Wife chooses Less Preferred, which is as same as the payoff that Wife gets if she chooses Preferred and Husband chooses Less Preferred. So, Battle of the Sexes 2 is directly fair with respect to its two players under Definition 1. Therefore, we can say that Battle of the Sexes 1 has label-independent fairness with respect to its players under Definition 3.

Table 3.2: Battle of the Sexes 2

		Wife	
		Preferred	Less preferred
Husband	Preferred	0, 0	3, 2
	Less preferred	2, 3	0, 0

We now extend our concept of bias to cases in which the players' strategy spaces are not identical.

Consider any normal form game G for n players. Consider any players i, j with strategy spaces $Si = \{1, ..., m_i\}$, $S_j = \{1, ..., m_j\}$ where $m_i > m_j$. By removing any $m_i - m_j$ strategies from player i's strategy space and then re-labelling the game so that $Si = \{1, ..., m_j\}$, we can create a game G' that is a *reduction of* G *with respect to* i and j and in which i and j have identical strategy spaces.

Definition 5 *(Reduced-game bias).* A game G has *reduced-game bias towards* one player i *relative to* another player j if $m_i > m_j$ and if there is some game G' that is a reduction of G with respect to i and j such that G' has label-independent bias towards i relative to j.

3.3.2 Utility function

In this subsection, we propose a method of incorporating attitudes to strategy fairness into players' utility functions. We model these attitudes as modifying attitudes to inequality via *strategy fairness factors.* Consider any game G, defined by its profile $(S_1, ..., S_n)$ of strategy spaces and its profile $(\pi_1, ..., \pi_n)$ of payoff functions. These properties determine whether, with respect to each pair of players, the game is strategically fair or biased. Let $\mathbf{x} = (x_1, ..., x_n)$ be profile of actual payoffs in the game, as determined by players' strategy choices.[1] The strategy fairness factor for the comparison between x_i and x_j, as viewed by player i, is given by

$$\varphi_i(G, x_i, x_j) = \begin{cases} r_i & \text{if } x_i \geq x_j \text{ and the game is biased towards player } i \\ 1 & \text{if the game is fair, no matter } x_i \geq x_j \text{ or } x_i < x_j \\ s_i & \text{if } x_i \geq x_j \text{ and the game is biased towards player } j \\ p_i & \text{if } x_i < x_j \text{ and the game is biased towards player } i \\ q_i & \text{if } x_i < x_j \text{ and the game is biased towards player } j \end{cases}$$

The parameters r_i, s_i, p_i and q_i represent i's attitudes to strategy fairness. We impose the restrictions (i) $r_i > 0$, $s_i > 0$, $p_i > 0$, and $q_i > 0$, (ii) if $r_i > 1$ then $s_i < 1$, (iii) if $r_i < 1$ then $s_i > 1$, (iv) if $p_i > 1$ then $q_i < 1$, (v) if $p_i < 1$ then $q_i > 1$. We will explain the motivation for these restrictions after we have explained the role of the strategy fairness factor in the utility function.

The utility function of player $i \in \{1, ..., n\}$ is given by

$$U_i(\mathbf{x}) = x_i - \alpha_i(\frac{1}{n-1}) \sum_{j \neq i} max\{x_j - x_i, 0\} \varphi_i(G, x_i, x_j) - \beta_i(\frac{1}{n-1}) \sum_{j \neq i} max\{x_i - x_j, 0\} \varphi_i(G, x_i, x_j)$$

where $0 \leq \beta_i < 1$ and $\beta_i \geq \alpha_i$

In the case of a two-player game, this simplifies to

$$U_1(\mathbf{x}) = x_1 - \alpha_1 max\{x_2 - x_1, 0\} \varphi_1(G, x_1, x_2) - \beta_1 max\{x_1 - x_2, 0\} \varphi_1(G, x_1, x_2)$$

[1] We use x_i rather than π_i in this context, to distinguish ex post payoffs from payoff functions.

This utility function is based on the Fehr and Schmidt (1999) model of inequality aversion. It takes players' preferences for strategy fairness into account by modifying players' attitudes to payoff inequality according to the fairness of the game in which the inequality has come about. Consistent with the Fehr-Schmidt model, $\beta_i \leq \alpha_i$ means that a player suffers more from disadvantageous inequality than advantageous inequality. $0 \leq \beta_i$ rules out the possibility that a player likes to be better off than others. $\beta_i < 1$ means that a player's utility is always increasing in her own material payoff in a fair game.

Returning to the restrictions placed on the parameters r_i, s_i, p_i and q_i, the restrictions $r_i > 0$, $s_i > 0$, $p_i > 0$ and $q_i > 0$ capture the idea that strategy fairness can partially influence a player's attitude to payoff inequality, but cannot change it completely. For example, suppose that player i is inequality averse and has a lower monetary payoff than player j (i.e. $x_i < x_j$). If the game is biased towards player i, this can make player i feel better about the payoff inequality or worse about the payoff inequality, but it cannot make her completely indifferent to inequality or inequality seeking. Restrictions (ii) to (v) express the idea that disadvantaged strategy unfairness and advantaged strategy unfairness affect utility in opposite directions. If player i feels better about disadvantaged inequality when a game is biased towards her than when it is fair, then she feels worse about disadvantaged inequality when the game is biased against her.

3.4 Basic principles of experimental design

The experiment consists of two stages, Stage 1 with either a series of card games or an effort task, and Stage 2 with a vendetta game. Before the experiment began, subjects were randomly assigned to seats with numbers on them. The randomness of the seat allocation was salient to subjects, but because it took place before the experiment began, we expected it to be separated from the game in subjects' minds when they were thinking about the fairness of the game.

At the beginning of the experiment, subjects were told that those with odd seat numbers became participant As and those with even seat numbers became participant Bs. Each participant A was randomly and anonymously matched with a coparticipant B. This matching stayed the same throughout the experiment. During the experiment, participant As and participant Bs competed for some lottery tickets. The results of the series of card game or the effort task in Stage 1 determined the distributions of the lottery tickets between participants that was carries over to Stage 2. The winner got more lottery tickets than the loser. The vendetta game in Stage 2 gave the two participants opportunities to change the distribution of the tickets. At the end of the experiment, one tickets was picked out randomly as the winning ticket. If the winning ticket was held by one of the participants, that participant got a prize.

In order to investigate people's preference for strategy fairness in competition environments, we wanted a Stage 1 game with the following features. First, the game had to provide an environment in which subjects interact with each other. Second, the game needed to create fixed inequality as the outcome. Third, within the interaction, we also wanted subjects to reveal their intentions of trying to win in the game. Therefore, random devices would not be suitable for the purpose. Finally, the rules of the game needed to be easily understood as fair or unfair.

We designed a card game which has all these features. At the start of each game, participants were dealt a card each. Each card had a number of points. In each game, both participant As and participant Bs were offered opportunities to replace cards. The numbers of replacement opportunities available to participant As and participant Bs varied between treatments. By allowing participant As and participant Bs to have equal/unequal replacement opportunities, the rules of the game were deliberately made fair/unfair between participants. In the Fair Rule treatment, both participant As and participant Bs had the same number of replacement opportunities. In the Unfair Rule treatment, participant As had less replacement opportunities than participant Bs. At the end of the card game, the participant who held a card with a higher number of points won the game.

In the interest of comparing strategy fairness with other widely tested concept of fairness, we chose to have a Real Effort treatment in which unequal outcomes in Stage 1 were generated by a real-effort task. In the literature on attitudes to fairness, real-effort tasks are commonly used to test the concept of desert-based fairness (e.g. Hoffman *et al.*, 1994; Ruffle, 1998; Fahr and Irlenbusch, 2000; Oxoby and Spraggon, 2008). Comparisons between the Fair Rule treatment and the Real Effort treatment allows us to analyse similarities or differences between strategy fairness and desert-based fairness.

In Stage 2, we wanted a game which can pick out people's attitudes towards the inequality created in Stage 1, i.e. how willing they are to accept this inequality. The game was selected with the following considerations in mind. First, the game had to provide a natural environment in which both participants are allowed to make moves. As discussed in Section 3.1, games with a single punishment round may lack external validity. We wanted a game which provides both participants with the same opportunities to change the distribution of payoffs, at the same cost. The idea of 'having the same opportunities' in the Stage 2 game mirrors the role of equality of opportunities in the concept of strategy fairness that we use in analysing Stage 1. Second, the game should offer participants opportunities to redistribute the lottery tickets rather than to punish each other. In the card game, participants interact with each other under rules that can be fair or unfair. If these rules are biased towards one of the players, there seems no reason for the disadvantaged player to *punish* a co-participant who did not set up the rule. But there is a reason to want to *redistribute* the outcomes of an unfair game.

We used the vendetta game created by Bolle *et al.* (2014), which meets all the criteria. To make the game easier to understand and to fit better with Stage 1 of our experiment, we made some changes to Bolle et al.'s original design. Details of game will be described in the next section. In the vendetta game, participant As and participant Bs took turns to choose whether to stay with the current distribution of lottery tickets or to choose a *taking* move – that is, a move that takes lottery tickets from the co-participant at a cost in terms of 'wasted' tickets. The loser in stage 1 made a choice first. Participants always had the option of not taking any tickets. The game ended when both participants had not taken for two successive turns each, or when no taking moves remained for either participant.

3.5 Application of the model and hypotheses

In section 3.3, we showed how to incorporate the concept of strategy fairness into a more complete model. Although the experiment was not designed to test the model, the model still can provide us some insight into the behaviours that we may observe in our experiment.

Table 3.3 summarises the initial positions in the vendetta game. The initial positions are determined by the type of tasks that were carried out in Stage 1 and their outcomes. In both the Fair Rule treatment and the Unfair Rule treatment, the series of card games was used to determine the allocation of lottery tickets between participant A and participant B. In the Fair Rule treatment, participant A and participant B had the same number of replacement opportunities; i.e., the game provided them with opportunities of competing fairly. Therefore, subjects can be classified as either fair winners or fair losers. In the Unfair Rule treatment, participant B was allowed to change more cards than participant A. This rule is obviously biased towards participant B. (In terms of the analysis in Section 3.3.1, there is reduced-game bias.) However, according to the design of the card game in the Unfair Rule treatment, participant A still has a chance to win. Hence, the outcome of the series of games can either be that the advantaged participant B wins and the disadvantaged participant A loses, or (but with much lower probability) the opposite. In the Real Effort treatment, the effort task is used to determine the allocation of lottery tickets between participant A and participant B in each pair. Subjects can be classified as either high effort winners or low effort losers.

Table 3.3: Initial positions in the vendetta game

	Loser	**Winner**
Fair Rule treatment	Fair loser	Fair winner
Unfair Rule treatment		
Advantaged player wins	Disadvantaged loser	Advantaged winner
Disadvantaged player wins	Advantaged loser	Disadvantaged winner
Real Effort treatment	Low effort loser	High effort winner

3.5.1 Application of the model

Figure 3.1 shows the component of player i's utility that derives from inequality of material payoffs. The baseline case of utility from inequality in a strategically fair game is shown by the solid lines. The other lines are drawn on the assumption that that $q_i > 1 > p_i$ and $r_i > 1 > s_i$. Our prior intuition was that most people's attitudes to unfairness would be well represented by these assumptions. The interpretation of $q_i > 1 > p_i$ is that if player i has a lower monetary payoff than player j, she would feel less bad about the payoff inequality if the game was biased towards her than if the game was fair, while she would feel worse about it if the game was biased towards player j. The interpretation of $r_i > 1 > s_i$ is that if player i has a higher monetary payoff than player j, she would feel less bad about the payoff inequality if the game was biased towards player j, while she would feel worse about it if the game was biased towards her.

These conjectures are about what attitudes to unfairness are *most common*. One would expect to find a lot of heterogeneity of attitudes in any population.

Applied to our experiment, these conjectures suggest that disadvantaged losers would be more likely to take than fair losers, and that advantaged winners would be less likely to take than fair winners.

Figure 3.1: Utility from payoffs

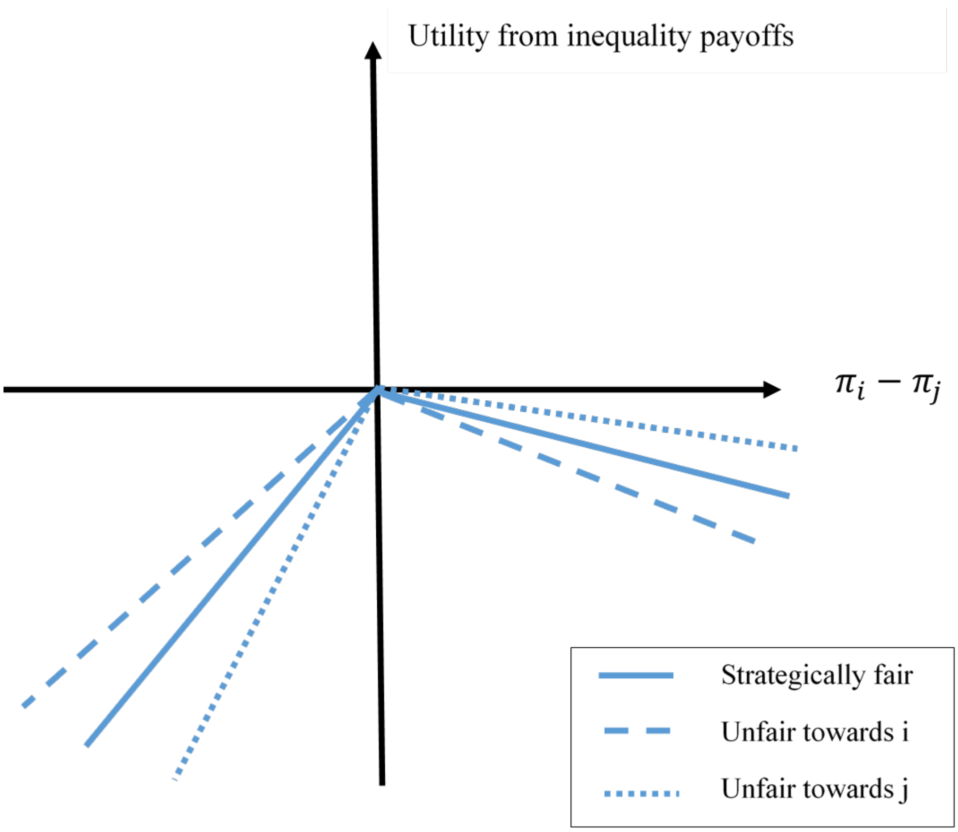

3.5.2 Hypotheses

The existing literature on social preferences suggests that people have a desire for fairness. When people are treated unfairly, they have the desire to rectify the unfair situation by means of punishment (e.g. Rabin, 1993; Charness and Rabin, 2002; Falk *et al.*, 2003; Segal and Sobel, 2007). In our experiment, the Fair Rule treatment produces inequality as the result of a game that is strategically fair in a way that is likely to be salient to players. The Unfair Rule treatment uses the same framing as the Fair Rule treatment but deliberately generates unfairness. Hence, if the fairness of the 'equal opportunities to change cards' rule is salient in the card games, one would expect more desire to rectify the inequality between disadvantaged losers and advantaged winners in the Unfair Rule treatment than to rectify the inequality between fair losers and fair winners in the Fair Rule treatment. In the Real Effort treatment, the framing of the effort task makes people think of a situation where effort ought to be rewarded. If one participant is to get more lottery tickets than the other, it ought to be the one who makes more effort. It is consistent with the idea of desert that people who put in more effort should get a greater reward as they are relatively deserving (e.g. Hoffman and Spitzer, 1985; Burrows and Loomes, 1994; Bosman *et al.*, 2005). Hence, less taking behaviour should be observed in the Real Effort treatment than in the Unfair Rule treatment. We therefore test the following hypothesis about efficiency loss:

> *Hypothesis 1:* More taking behaviour occurs in the Unfair Rule treatment with disadvantaged losers and advantaged winners than in the Fair Rule treatment or in the Real Effort treatment.

Apart from differences in efficiency loss between treatments, we were also interested in individual players' propensities to take. First, we consider the first moves of the losers. In the Fair Rule treatment, the fair loser and the fair winner got the unequal outcome by going through a fair competition. In the Real Effort treatment, the high effort winner wins the game by making more effort. For the disadvantaged losers in Unfair rule treatment, the advantaged winner won the game through making use of a rule that was biased in her favour. Therefore, one would expect that unfair losers in the Unfair Rule treatment should be less willing to settle on the initial distribution of lottery tickets than fair losers in the Unfair Rule treatment or low effort losers in the Real Effort treatment. For the first moves of losers, we therefore test the following hypotheses:

> *Hypothesis 2:* Comparing the first moves of losers, disadvantaged losers in the Unfair Rule treatment are more likely to take than fair losers in the Fair Rule treatment or low effort losers in the Real Effort treatment.

Moreover, any taking behaviour by advantaged winners in the Unfair Rule treatment can be expected to provoke more counter-taking behaviour by disadvantaged losers than the corresponding behaviour of winners in the other two treatments. Therefore, for the propensity to steal of losers in these three treatments, we also want to test the following hypotheses:

> *Hypothesis 3:* At any point in the vendetta game at which the loser moves, disadvantaged losers in the Unfair Rule treatment are more likely to take than fair losers in the Fair Rule treatment or low effort losers in the Real Effort treatment.

Furthermore, it is also interesting to compare winners' propensity to take in these three treatments. In both the Fair Rule treatment and the Real Effort treatment, fair winners and high effort winners won more lottery tickets either by playing and winning a strategically fair card game or by making more effort than their co-participants. In the Unfair Rule treatment, advantaged winners instead were favoured by the rules of the card game. One might expect that within an ongoing vendetta, advantaged winners in the Unfair Rule treatment would be more tolerant of the taking behaviour of the disadvantaged losers than either the fair winners or high effort winners in the other two treatments. Our final hypothesis is, therefore:

> *Hypothesis 4:* At any point in the vendetta game at which the winner moves, advantaged winners in the Unfair Rule treatment are less likely to take than fair winners in the Fair Rule treatment or high effort winners in the Real Effort treatment.

As Table 3.3 shows, there is a fourth type of initial position for the vendetta game – the position that can arise in the Unfair Rule treatment when the series of card game is won by participant A. In this case, the vendetta game is played between a disadvantaged winner and an advantaged loser. The experiment was not designed to investigate this case; it occurs only as a necessary by-product of setting up a genuine game with unfair rules. Precisely because of the unfairness of the game, this case occurs relatively rarely, and so our design produces relatively few observations of it. (17 (22.67 per cent) of the vendetta games in the Unfair Rule treatment were of this type.) But, while recognising the small number of observations, it is still interesting to compare the propensity to take of disadvantaged winners and advantaged losers in the Unfair Rule treatment with winners and losers in the other two treatments.

We do not state any formal hypotheses regarding differences between taking behaviour in the Fair Rule treatment and the Real Effort treatment. The Fair Rule treatment gets inequality by incorporating a certain kind of fairness: strategy fairness. The Real Effort treatment incorporates a different kind of fairness: desert-based fairness, which is more commonly tested in the literature of social preference. Existing theories of attitudes to fairness do not provide unambiguous predictions about the direction that differences between the two treatments might take. Still, it is interesting to compare taking behaviour in these two treatments.

3.6 Design details and implementation

3.6.1 Overall structure of experiment

Each session of the experiment was randomly assigned to one of the three treatments (Fair Rule, Unfair Rule or Real Effort). Stage 1 of each session consisted of either a series of card games or an effort task; Stage 2 consisted of a vendetta game. At the beginning of each stage, each subject received a copy of the instructions for that part; these instructions were read aloud by the experimenter. These instructions are reproduced in Appendix 1. Each subject then completed a computerised questionnaire which tested her understanding of the tasks. If a subject made a mistake, the computer would show her the correct answer and the relevant part of the instructions. Subjects were invited to ask the experimenter for clarification.

In Stage 1, subjects competed for twelve lottery tickets numbered 1 to 12. The results of the series of card game or the effort task determined the initial allocation of these tickets. At the end of Stage 1, the computer picked nine of these numbered tickets at random and assigned them to the winner. The remaining three tickets were assigned to the loser. The vendetta game in Stage 2 gave subjects the opportunities to change the initial distribution of the tickets.

At the end of each session, the experimenter put twelve numbered tickets into a bag. One of the participants was asked to come forward and pick one ticket from the bag. The number on this ticket was the number of the winning ticket. In each pair of subjects, if either member of that pair held a ticket with the winning number, she got the prize of £24. If the winning ticket had been wasted during the vendetta game, neither member of the pair got the prize. In all cases, both members of a pair also received a participation fee of £3.

3.6.2 The Series of Card Games

In both the Fair Rule treatment and the Unfair Rule treatment, subjects played a series of card games in Stage 1. The basic structure of the game was described in Section 3.4 above; here we fill in the details.

At the start of each card game, participants were dealt a card each. Each card had a number of points, which could be any of the whole numbers in the range from 1 to 100. Each of these numbers was equally likely at each 'deal'.

In each game, participant A and participant B were offered opportunities to replace cards. In the Fair Rule treatment, both participants were allowed to replace cards up to three times in each card game. In the Unfair Rule treatment, participant A was allowed to replace cards up to one time in each card game, while participant B was allowed to replace cards up to three times. During the game, participants could decide to stick with the card that they had been dealt or to replace it with a new card. If a participant decided to replace her card with a new card, then the computer would randomly draw a new one for her. Each number in the range from 1 to 100 was still equally likely at this 'deal'. However, the participant could not go back to the replaced card again. Participants could decide to stick with any card that was dealt to them. Once they used up all the replacement opportunities, they could not make any further replacement and had to stick with the last card that they had been dealt. During this stage of the game, neither participant could see what cards her co-participant was being dealt, or whether the co-participant was using replacement opportunities.

After both participants had made the decision of sticking with a card they had been dealt or had used up all their opportunities for replacing cards, their cards were compared. At this stage, they could both observe the points on both participants' cards. They were also shown how many replacement opportunities their co-participants had used. A participant who held a card with a higher number of points than the card held by her coparticipant won the game. If both of them had the same number of points, the game was a draw. The first participant to win 4 games was the overall winner of the series of games. Draws were not counted.[2] We used a series of card games rather than just one because we wanted the game to be fairly simple (and so not to contain too many replacement opportunities) but also wanted the advantaged player to win with high probability.

3.6.3 Real-effort task

In the Real Effort treatment, participants faced the Encryption Task presented in Erkal *et al.* (2011). In this task, participants were given an encryption table which assigned a number to each letter of the alphabet in a random order. Each participant was then presented with words in a predetermined sequence and was asked to encrypt them by substituting the letters with numbers using the encryption table. All participants were given the same words to encode in the same sequence.

After a participant encoded a word, the computer would tell her whether the word had been encoded correctly or not. If the word had been encoded wrongly, the participant would be asked to check her codes and correct them. Once the participant encoded a word correctly, the computer then prompted her with another word which she was asked to encode. This process continued for six minutes.

After both participants had finished the task, the number of words they had encoded was counted as their scores for the task. At the end of the task, the participant with the higher score was the winner. If both participants got the same score, then the person who encoded the words in a shorter time was the winner.

[2] In all sessions, draw happened in only 18 games (1.23% of all card games played).

3.6.4 Vendetta Game

In the second stage of the experiment, we used a modified version of the vendetta game created by Bolle *et al.* (2014). Participants took turns to choose whether to stay with the current distribution of lottery tickets or to change it. The loser in stage 1 made a choice first, starting from the distribution of lottery tickets determined in Stage 1. When it was a participant's turn to move, she was asked to choose whether she wanted to take lottery tickets from her co-participant, and if so, how many lottery tickets to take. Amounts taken had to be in blocks of three (so the number of tickets taken could be three, six or nine), and up to as many as the co-participant had at the time. The *transfer rate*, which denotes the marginal gain per unit taken, was 1/3, implying an efficiency loss from taking tickets. Therefore, for every block of three tickets that the participant took from her co-participant, she gained one ticket and two were wasted. Participants always had the option of not taking any tickets. All the feasible points of the vendetta game are given in Figure 3.2 below. The initial starting point is (3, 9), i.e. the Stage 1 loser (and first mover in the vendetta game) held three tickets and the Stage 1 winner held nine.

Figure 3.2: The Feasible Points of the Vendetta Game

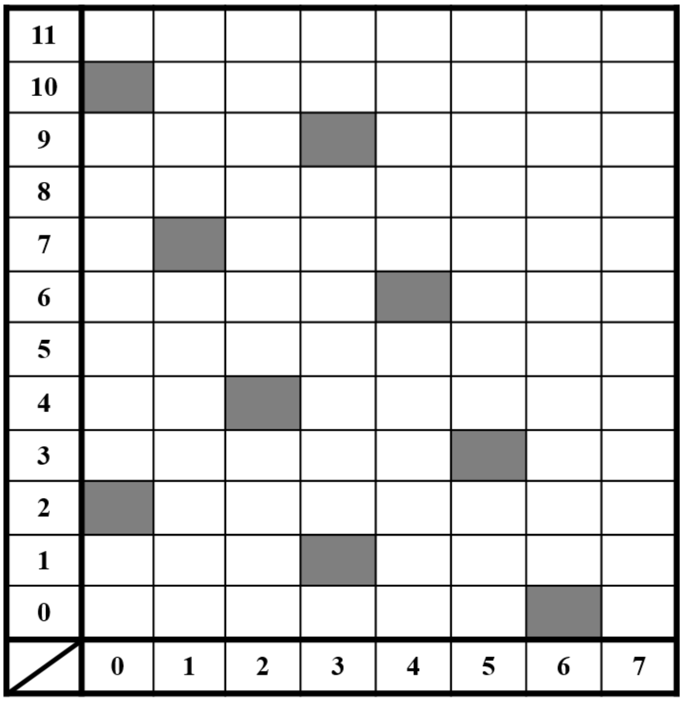

Notes: Shaded areas indicate all feasible points of the vendetta game. The horizontal axis refers to the number of lottery tickets owned by the Part 1 loser and the vertical axis refers to the number of lottery tickets owned by the Part 1 winner.

The game ended if one of two cases applied. The first case occurred if one or both participants could still take tickets from their co-participants but the participant(s) who were able to do this had chosen not to do so for two consecutive times. The second case occurred if both participants held less than three tickets, so no positive multiple of three tickets could be taken from either of them. Therefore, (0, 2) is the terminal point, i.e. the (only) distribution of tickets at which no further taking moves are possible.

Figure 3.3: Sample computer display of the vendetta game

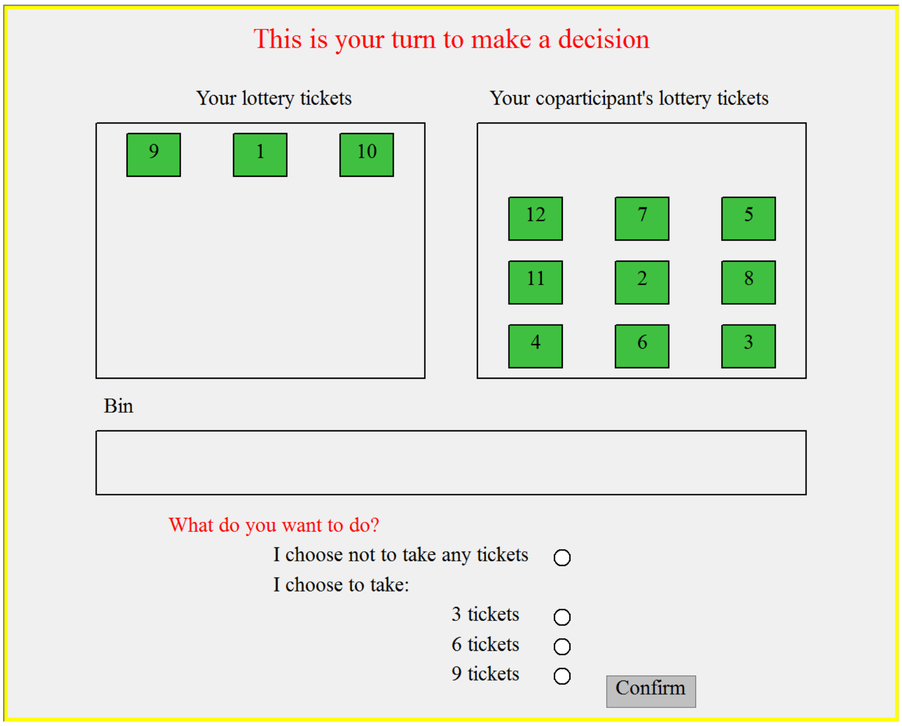

Notes: The sample computer display shows what subjects saw at the beginning of the vendetta game.

The computer display is shown in Figure 3.3. On this display, the participant could see three baskets. One contained the lottery tickets that she held at present, one contained the lottery tickets that her coparticipant held at present, and one was the bin. Before making any decision, she could try different possible numbers of blocks of lottery tickets to take away from her coparticipant. The computer would show her how many of these lottery tickets would be moved from her coparticipant's basket into her basket and how many of these lottery tickets would be moved from her coparticipant's basket into the bin. After the participant made her decision, the baskets and bin were updated to show her the outcome of her decision and the location of all the lottery tickets at the time. After her coparticipant had chosen, the baskets and bin were updated again to show her the location of all the lottery tickets at the time as a result of her coparticipant's choice.

The vendetta game replicates the 'mini-vendetta' game in Bolle *et al.* (2014), in which all possible sequences of taking moves are relatively short (the longest possible game has five taking moves). The advantage of having such a design is that it minimises the degree of reasoning required from the subjects, in terms of the number of steps required to backward induce to subgame perfection. It makes the game simpler for subjects, and therefore reduces the likelihood that a vendetta might be caused by confusion. Compared to the original mini-vendetta game, the major changes that we made in our experimental design are as follows. An important difference is that we used lottery tickets instead of describing outcomes in terms of numerical probabilities of winning the game. Lottery tickets are more concrete objects and easier to understand, while still making it easy for subjects to read off probabilities (in our design, as multiples of 1/12). On the computer display, each participant was able to see very clearly the current distribution of the tickets; i.e. how many tickets were in her basket, how many were in her co-participants' basket, and how many were in the bin. By allowing participants try out different possible actions before making a final decision, our design enabled participants to see vividly the consequences that alternative actions would produce. The existence of the bin helped participants to have a clearer sense of the waste that takes place when they take each other's tickets, as they can see the number of tickets that go into the bin.

3.6.5 Implementation

The experiment was conducted between November 2015 and January 2016 at the CBESS Experimental Laboratory at the University of East Anglia. Participants were recruited from the general student population via the CBESS online recruitment system (Bock *et al.*, 2012). The experiment was programmed and conducted with the experimental software z-Tree (Fischbacher, 2007). We ran 18 sessions in total: four for the Real Effort treatment, six for the Fair Rule treatment and eight for the Unfair Rule treatment. A total of 326 subjects participated in the experiment, of whom 72 were in the Real Effort treatment, 104 in the Fair Rule treatment, and 150 in the Unfair Rule treatment. We needed more participants in the Unfair Rule treatment than the Fair Rule treatment is because some series of card games would be won by disadvantaged players in the Unfair Rule treatment and we wanted the number of games won by advantaged winners in the Unfair Rule treatment to be close to the number of games won by fair winners in the Fair Rule treatment. Fewer participants were recruited in the Real Effort treatment, as we were more interested in investigating possible differences in behaviour between strategically fair and biased games. 123 subjects were male and 196 were female.[3] Most of the participants were students from a wide range of academic disciplines and with an age range from 18 to 63. The experiment lasted about 50 minutes. Average earnings were £10.67 per person, including a show-up fee of £3.00. The lowest earning was £3.00, the highest was £27.00.

[3]Seven subjects selected 'prefer not to say' in the gender question.

3.7 Results

3.7.1 Efficiency Loss and Vendetta behaviour

We begin by looking at the outcomes of the vendetta games across treatments.

Figure 3.4 provides an overview of vendetta game outcomes in the Real Effort treatment, the Fair Rule treatment, and the Unfair Rule treatment with advantaged winners. It shows that in the Real Effort treatment, subjects did not take in 14 out of 36 pairs (i.e. the vendetta game ended at the initial starting points), while 9 of the 36 pairs ended at the terminal point (0,2). In the Fair Rule treatment, 23 out of 52 pairs settled on the initial starting points, and 11 out of 52 pairs ended at the terminal point. In the Unfair Rule treatment with advantaged winners, subjects did not take in 13 out of 58 pairs, and 16 out of 58 pairs ended at the terminal point.

Figure 3.4: Outcomes of the Vendetta Game
(a) Real Effort treatment

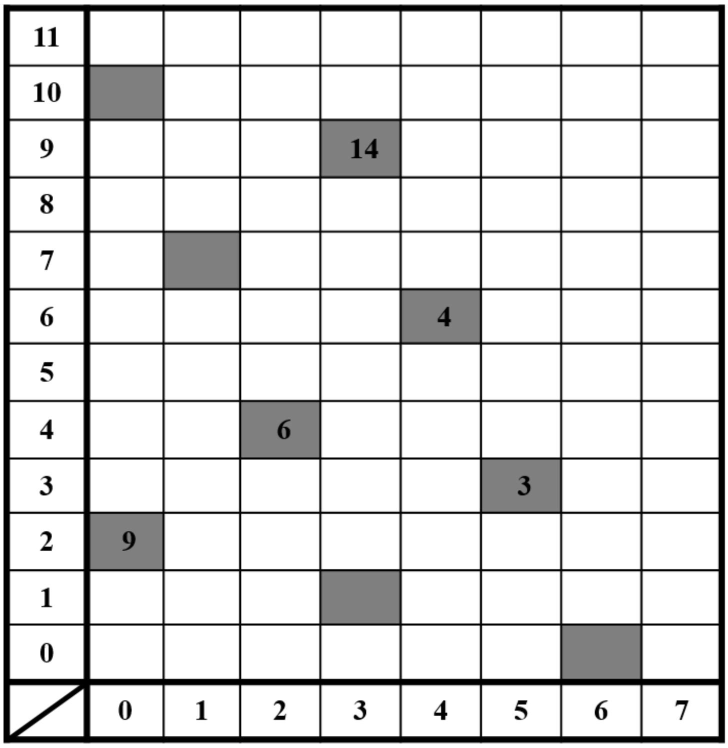

(b) Fair Rule treatment

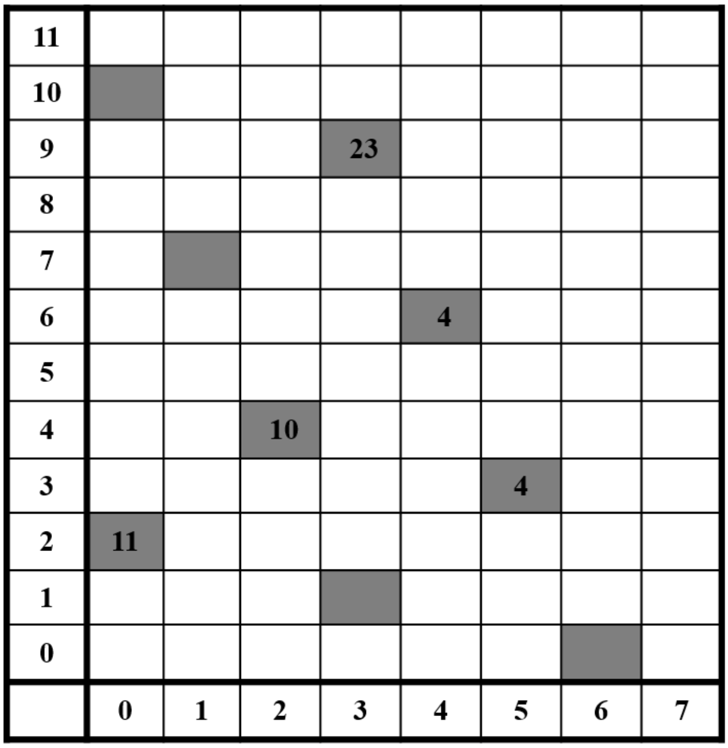

(c) Unfair Rule treatment with advantaged winners

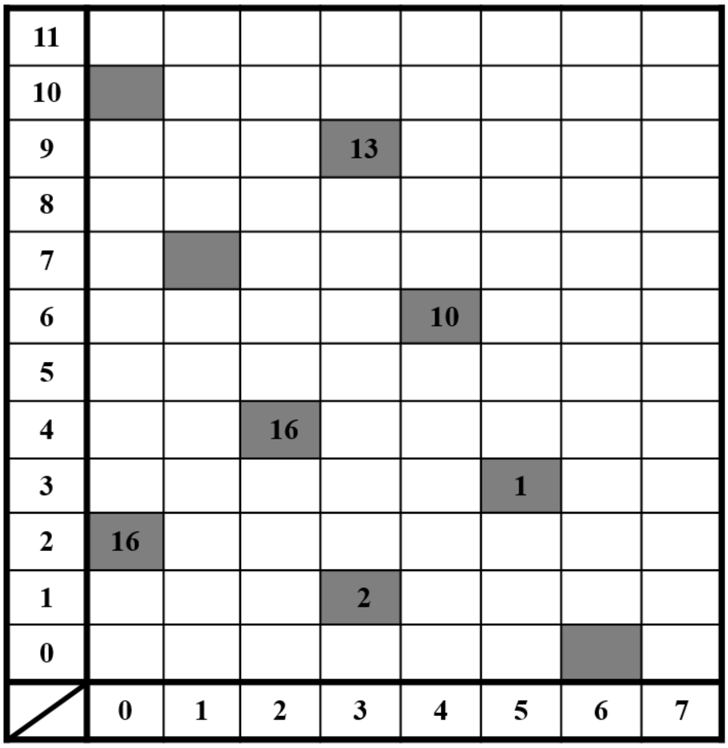

Notes: Shaded cells correspond to the points that can be reached in the game. Numbers on each grid represent the number of times a point (x, y) was obtained as the final point in the experiment.

The cumulative distributions of the outcomes of the vendetta games across these three treatments are shown in Figure 3.5. For any given pair of subjects, *final total holdings* is the total number of lottery tickets that two co-participants hold at the end of vendetta game. This is a measure of the efficiency of the outcome or, equivalently, an inverse measure of the extent of taking and counter-taking during the game. It takes its maximum value of 12 if neither participant chooses to take anything. It takes its minimum value of 2 if taking and counter-taking continues until no more taking moves are possible. On average, pairs in the Fair Rule treatment and Real Effort treatment ended up with more lottery tickets than pairs in the Unfair Rule treatment. The mean value of final total holdings is 7.94 in the Real Effort treatment, 8.27 in the Fair Rule treatment and 6.90 in the Unfair Rule treatment with advantaged winners. The distributions of final total holdings are significantly different between the Fair Rule treatment and the Unfair Rule treatment (Mann-Whitney p=0.046)[4]. No significant difference in the distributions of final outcomes of vendetta games is found either between the Fair Rule treatment and the Real effort treatment (Mann-Whitney p=0.679) or between the Real effort treatment and the Unfair Rule treatment (Mann-Whitney p=0.177).

Figure 3.5: Cumulative distributions of final total holdings

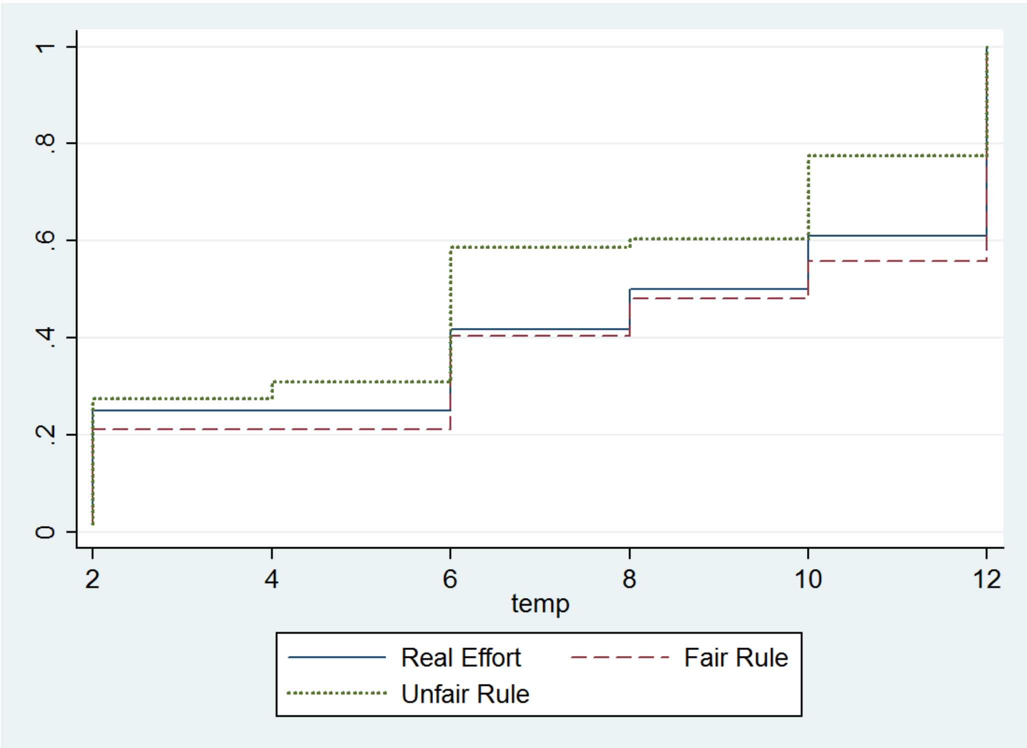

Result 1: As implied by Hypothesis 1, significantly greater efficiency losses were observed in the Unfair Rule treatment between disadvantaged losers and advantaged winners than in the Fair Rule treatment between fair losers and fair winners. No significant difference in efficiency losses was found either between the Fair Rule treatment and the Real Effort treatment or between the Real Effort treatment and the Unfair Rule treatment.

[4]In this chapter, all reported p values are two-sided.

3.7.2 The first moves of losers and winners

In the light of the results that we have reported so far, it would be interesting to know whether there were systematic differences between the taking behaviour of losers and winners across treatments. However, the problem with comparing individual players' decisions across treatments is that apart from the initial point of the game, different subjects may move to different feasible points in the vendetta game. Therefore, we have to test separately for each feasible point in the game. However, except for the initial point of the game (the loser's first opportunity to take), the number of observations is necessarily lower – often much lower – than the total number of participants in the relevant role (i.e. winner or loser). For most points in the game, we do not have enough data for informative statistical tests about behaviour at individual points.

To solve this problem, we use several methods. First, we test for differences in losers' first moves across treatments. 47.2% of low effort losers in the Real Effort treatment and 46.2% of fair losers in the Fair Rule treatment made a taking decision in their first move. In the Unfair Rule treatment, 60.3% of disadvantaged losers chose to take in their first moves. These data suggests that low effort losers in the Real Effort treatment and fair losers in the Fair Rule treatment are more willing to settle on the initial unequal distribution of lottery tickets than disadvantaged losers in the Unfair Rule treatment. However, no statistically significant difference in taking behaviour is found either between fair losers and disadvantaged losers in their first moves (Mann-Whitney p=0.138) or between low effort losers and disadvantaged losers in their first moves (Mann-Whitney p=0.216).

If the loser made a non-taking decision in her first move and the winner did not take in her first move either, the loser then would get another chance to make a taking decision at the initial point (3, 9). Therefore, if we want to check losers' willingness to settle on the initial unequal distribution of lottery tickets, it is reasonable for us to also take into account losers' second moves. Our data show that, conditioning on losers not having taken in their first moves, there are only two winners who made a taking move in their first moves: one fair winner and one advantaged winner. After we exclude these two cases, the data shows that after the first two successive turns of losers, 38.9 % of pairs in the Real Effort treatment and 47.1% of pairs in the Fair Rule treatment still held their initial amount of lottery tickets, while 30.4% of pairs with disadvantaged losers and advantaged winners in the Unfair Rule treatment were still at the initial point. We find evidence that more fair losers in the Fair Rule treatment chose to settle on the initial unequal distribution of lottery tickets than disadvantaged losers in the Unfair Rule treatment (Mann-Whitney p=0.077). But there is no significant difference in willingness to settle on the initial unequal distribution between low effort losers in the Real Effort treatment and disadvantaged losers in the Unfair Rule treatment (Mann-Whitney p=0.401).

Result 2: In line with our Hypothesis 2, disadvantaged losers in the Unfair Rule treatment were more likely to take in their first moves than fair losers in the Fair Rule treatment

In order to get a rough idea of winners' attitude towards losers' taking behaviour, we compare winners' first moves across treatments. After winners had seen losers choosing to take in their first moves, 47.1% of high effort winners in the Real Effort treatment and 50.0% of fair winners in the Fair Rule treatment made a taking decision in their first moves. In the Unfair Rule treatment, 54.3% of advantaged winners chose to take in their first moves after their co-participants had made a taking decision. There is no evidence of less taking behaviour by advantaged winners in their first moves in the Unfair Rule treatment than by fair winners in the Fair Rule treatment (Mann-Whitney p=0.746).

3.7.3 Index method

To get further insight into the propensity to steal of losers and winners in the vendetta games, we need a method which can solve two problems – the problem of having small numbers of observations for most feasible points, and the problem of dependence between losers' moves and winners' moves. Bolle *et al.* (2014) investigate the dynamics of taking behaviour in their vendetta game using a regression method. However, we believe that this regression method is not sufficient to deal with our second problem. It cannot fully disentangle subjects' decisions from their co-participants' decisions. If subjects in one role (winner or loser) in one treatment make more taking moves (or take more tickets in total) than subjects in the same role in another treatment, we would not be able to tell from the regression results whether this was because the subjects in the first treatment had a stronger desire to take, other things being equal, or because they we retaliating against co-participants who had taken more.[5] To solve these problems, we use an 'index' method.

For each subject, we want to construct an index which represents the subject's propensity to take, relative to the overall behaviour of all subjects with the same role, controlling for differences in the points in the game that are reached by different subjects. The index for an average subject should be equal to zero. Subjects with indexes above zero are the subjects who have a greater propensity to take than an average subject, while subjects with indexes below zero are the ones who have a lower propensity to take than an average subject.

[5]We run a regression analysis. The results can be found in Appendix 1. Table 3.10 and 3.11 both contain three overall regressions. Model 1 tests for a treatment effect. Model 2 tests taking behaviour in later rounds by controlling for the taking decisions that the co-player made in the most recent round in which that player had an opportunity to take. Model 3 tests the effect of the difference in the number of lottery tickets (*Lottery ticket difference*) on taking behaviour. Lottery ticket difference is the difference between the number of lottery tickets held by the subject and the number of lottery tickets held by her co-player. The treatment variables are *Real Effort* (= 1 if the subject is in the Real Effort treatment) and *Unfair Rule* (=1 if the subject is in the Unfair Rule treatment). The results show that there is no significant treatment effect. *Lag Taken* has positive and significant effect on both losers' decisions and winners' decisions. *Lottery ticket difference* has negative and significant effect on players' decision.

The first step in defining the index is to define, for each of the two roles in the vendetta game, the set of possible *taking opportunities.* Each taking opportunity is a point in the game at which a player in the relevant role can choose between taking and not taking, and which can be reached by some feasible combination of previous decisions by the two players. A 'point in the game' for a given role is defined in terms of the two players' current holdings of tickets and (in cases where both players hold three or more tickets) whether, conditional on those holdings, this is the first or second opportunity for player in the relevant role to take. Taking opportunities for the two roles are shown in Table 3.4.

Table 3.4: Taking opportunities for losers and winners

Loser	**Winner**
(3,9) opportunity 1	(3,9) opportunity 1
(3,9) opportunity 2	(3,9) opportunity 2
(4,6) opportunity 1	(4,6) opportunity 1
(4,6) opportunity 2	(4,6) opportunity 2
(4,6) opportunity 1	(5,3) opportunity 1
(4,6) opportunity 2	(5,3) opportunity 2
(0,10)	(6,0)
(1,7)	(3,1)
(2,4)	

For example, consider taking opportunities for the loser. '(3, 9): opportunity 1' is the initial point in the game: the current distribution is (3, 9), and the loser has her first opportunity to choose between taking and not taking. '(3, 9): opportunity 2' occurs if the loser chooses not to take in her first move and if the winner then chooses not to take in his first move. '(1, 7)' occurs if the loser takes three tickets at the initial distribution and the winner then takes three tickets. It also occurs if the winner takes three tickets at the initial distribution and the loser then takes three tickets. Notice that, because tickets can be taken only in blocks of three, it is not possible for the winner to take any tickets at (1, 7). Thus, there is no strategic difference between the loser's first and second opportunity to take at this point in the game. Accordingly, we treat (1, 7) as a single taking opportunity for the loser.

We now define an index for losers. (We use the same method to define an index for winners.) We assume that a player's behaviour at any given point in the game is independent of how that point was reached. Given this assumption, any given (mixed) strategy for a loser implies a *taking probability* p_l (i.e. the probability that a taking move is chosen) at each taking opportunity l in the set L of all possible taking opportunities for losers (i.e. all those listed in Table 3.4.[6] Intuitively, the propensity to take of a given player in the role of loser can

[6]As in analyses of subgame-perfect Nash equilibrium, we allow for the possibility of low-probability 'trembles'. Thus, for example, even if a strategy assigns zero probability to taking moves at (3, 9), it is meaningful to define taking probabilities at (4, 6).

be measured by some weighted sum of that player's taking probabilities. Provided that the weights are fixed, such a measure can in principle be used to compare different players' propensities to take, independently of differences in the behaviour of their co-players. However, there is no uniquely correct way of assigning these weights. For the purposes of our statistical tests, we adopt the convention of weighting each taking opportunity l by the proportion of vendetta games in our experiment which the loser faced that opportunity.[7] This proportion is denoted by q_l. Thus, for any given strategy, the *taking propensity* is:

$$\sum_{l \in L} p_l q_l.$$

Intuitively, a player's taking propensity is the expected number of taking moves that she would make in the vendetta game if she faced each taking opportunity with the same probability as an 'average' player.

However, our experimental design does not allow us to observe complete strategies. For each participant, we observe behaviour only at those taking opportunities that she in fact reached. Consider any given participant playing in the role of loser. Let L^* be the set of taking opportunities that she in fact reached. For each l in L^*, let a_l be the *actual decision* of that player, where $a_l = 0$ denotes 'not take' and $a_l = 1$ denotes 'take'. Let e_l be the *expected proportion* of taking moves at opportunity l (i.e. considering all those vendetta games in the experiment in which opportunity l was reached, the proportion in which a taking move was made at that opportunity). We define the *index of excess taking* for that player as:

$$\sum_{l \in L^*} (a_l - e_l) q_l.$$

Notice that if all players follow the same mixed strategy, the expected value of this index is zero. Intuitively, the value of this index for a given player can be thought of as an estimate of the difference between this player's taking propensity and the taking propensity of an 'average' player, based only on actual observations.

The 'index' approach is useful for the following reasons. First, it solves the problem of lack of observations of behaviour at taking opportunities that are not reached. Second, it takes individual subjects as independent units of observations and allows us to combine all the moves of each individual subject. Thus, it allows us to do statistical tests at the level of the individual subject. Third, it controls the problem of dependence between losers' moves and winners' moves, so that, for each role (loser and winner) separately, we can compare the distributions of indexes across treatments. If we find a difference in the behaviour of

[7] In defining each q_l, we aggregate across the three treatments in our experiment, giving each observation equal weight. Although different treatments had different numbers of participants, this procedure is legitimate for tests where the null hypothesis is that behaviour does not differ across treatments.

losers (winners) between treatments, we are able to say that the difference is not caused by differences in the behaviour of winners (losers) between treatments.

As a variant of the index of excess taking defined above, we also defined an *index of excess taking with value*. The only difference between these indexes is that, while the index of excess taking is a measure of taking *moves* (i.e. of all moves in which three, six or nine tickets were taken from the co-player), the index of excess taking with value is a measure of the *number* of tickets taken. Defining q_l as before, let s_l be the actual number of tickets taken at opportunity l and let t_l be the expected number of tickets taken at that opportunity (i.e. considering all those vendetta games in the experiment in which opportunity l was reached, the average number of tickets stolen at that opportunity). The index of excess taking with value is:

$$\sum_{l \in L^*} (s_l - t_l) q_l.$$

The distributions of these two indexes in the three treatments are summarised in Table 3.5. From Table 3.5 we can see that low effort losers in the Real Effort treatment and fair losers in the Fair Rule treatment are more likely to take than disadvantaged losers in the Unfair Rule treatment. This is indicated by both indexes. However, the difference between fair losers and disadvantaged losers is not statistically significant either for the index of excess taking (Mann Whitney p= 0.146) or for the index of excess taking with value (Mann-Whitney p= 0.177).

Table 3.5: Distributions of the indexes

		Loser		**Winner**	
		Mean	Std. Dev	Mean	Std.Dev
Index of excess taking	RE	-0.014	0.543	-0.032	0.119
	FR	-0.077	0.531	-0.016	0.132
	UR	0.078	0.519	0.034	0.158
Index of excess taking with value	RE	-0.035	3.090	-0.092	0.354
	FR	-0.310	3.129	-0.051	0.398
	UR	0.300	3.036	0.102	0.485

Notes: RE represents the Real Effort treatment, FR represents the Fair Rule treatment, and UR represents the Unfair Rule treatment.

Surprisingly, the distributions of both indexes show that high effort winners in the Real Effort treatment and fair winners in the Fair Rule treatment are *less* likely to take than advantaged winners in the Unfair Rule treatment. The difference between fair winners and advantaged winners is significant both for the index of excess taking (Mann-Whitney p= 0.073) and the index of excess taking with value (Mann-Whitney p= 0.057).

Result 3: As predicted by Hypothesis 3, disadvantaged losers in the Unfair Rule treatment are more likely to take than losers in other two treatments. However, this difference is not statistically significant. Contrary to Hypothesis 4, there is some evidence that advantaged winners in the Unfair Rule treatment are more likely to take than winners in the Fair Rule treatment.

3.7.4 Gender difference in the propensity to take

The original propose of this experiment was not to test for gender differences in attitudes towards strategy fairness. However, many studies suggest the existence of differences between male and female attitudes towards rules and competitions. A study of children's social behaviour has shown that boys play rule-based games more often than girls, such as sports games which are governed by a set body of rules and aim at achieving an explicit goal, and consequentially boys gain more experience in the judicial process (Lever, 1976). Piaget (1932 (1968)) observed that in the games played by children, boys were more explicit about agreements and more concerned with legal elaboration than girls. Gilligan (1982) claims that for men, fairness is more of a matter of principle, while for women, fairness does not appear to be a moral imperative. We might conjecture that male participants are predisposed to care more than female participants about the rules of the game that generates inequality in our experiment. We test for gender differences in propensities to take by losers and winners using the index method.

Table 3.6: Distributions of indexes for losers by gender

		Male		**Female**	
		Mean	Std. Dev	Mean	Std.Dev
Index of excess taking	RE	0.037	0.544	-0.023	0.552
	FR	-0.311	0.464	0.078	0.520
	UR	0.122	0.518	0.067	0.507
Index of excess taking with value	RE	0.196	3.354	-0.062	2.967
	FR	-1.048	3.281	0.211	3.002
	UR	0.183	2.881	0.512	3.155

Notes: RE represents the Real Effort treatment, FR represents the Fair Rule treatment, and UR represents the Unfair Rule treatment.

Table 3.6 summarizes the distributions of the indexes for losers by gender. First, we compare the indexes of male losers and female losers by treatments. There is no significant difference between male and female behaviour in the Real Effort and Unfair Rule treatments. However, females are significantly more likely than males to take in the Fair Rule treatment, whether this is measured by the index of excess taking (Mann-Whitney $p = 0.009$), or by the index of excess taking with value (Mann-Whitney $p = 0.026$). We also consider cross-treatment differences in losers' taking behaviour separately for males and females. We find that male losers are significantly more likely to take in the Unfair Rule treatment than in the Fair Rule treatment, which is indicated by both indexes[8], while female losers' behaviour is relatively more consistent between treatments.

Table 3.7: Distributions of indexes for winners by gender

		Male		**Female**	
		Mean	Std. Dev	Mean	Std.Dev
Index of excess taking	RE	-0.013	0.076	-0.043	0.138
	FR	-0.019	0.113	-0.007	0.145
	UR	0.007	0.134	0.047	0.167
Index of excess taking with value	RE	-0.033	0.229	-0.125	0.410
	FR	-0.069	0.339	-0.016	0.435
	UR	0.026	0.401	0.137	0.520

Notes: RE represents the Real Effort treatment, FR represents the Fair Rule treatment, and UR represents the Unfair Rule treatment.

Table 3.7 summarizes the distributions of the indexes for winners by gender. No significant difference between male winners and female winners is found in any of the treatments. Nor is there any significant difference between the behaviour of male winners in the Fair Rule and Unfair Rule treatments. Surprisingly, female winners are significantly more likely to take in the Unfair Rule treatment than in the Fair Rule treatment[9] or in the Real Effort treatment[10].

3.7.5 The taking behaviour of disadvantaged winners and advantaged losers

Although the Unfair Rule treatment was designed to produce data about the taking behaviour of advantaged winners and disadvantaged losers, it also generates a relatively small number of vendetta games between disadvantaged winners and advantaged losers. So, it is natural to ask whether there were systematic differences in the behaviour of winners and losers between the two forms of the Unfair Rule treatment.

There are 75 pairs of subjects who participated in the Unfair Rule treatment sessions. Of these 75 pairs, 58 pairs ended up with advantaged players winning the series of card games and 17 pairs ended up with disadvantaged players winning the series of card games.

[8]Mann-Whitney p=0.006 for index of excess taking; Mann-Whitney p= 0.038 for index of excess taking with value.

[9]Mann-Whitney p=0.091 for index of excess taking; Mann-Whitney p= 0.098 for index of excess taking with value.

[10]Mann-Whitney p=0.026 for index of excess taking; Mann-Whitney p= 0.024 for index of excess taking with value.

Figure 3.6 shows the outcomes of vendetta games carried out by these 17 pairs. It shows that 3 out of 17 pairs settled on the initial starting points, and 8 out of 17 pairs ended at the terminal point. No significant difference in the distribution of final outcomes of vendetta games is found between these 17 pairs with disadvantaged winners and the other 58 pairs with advantaged winners (Mann-Whitney p=0.446).

Figure 3.6: Outcomes of the Vendetta Game with disadvantaged winners in the Unfair Rule treatment

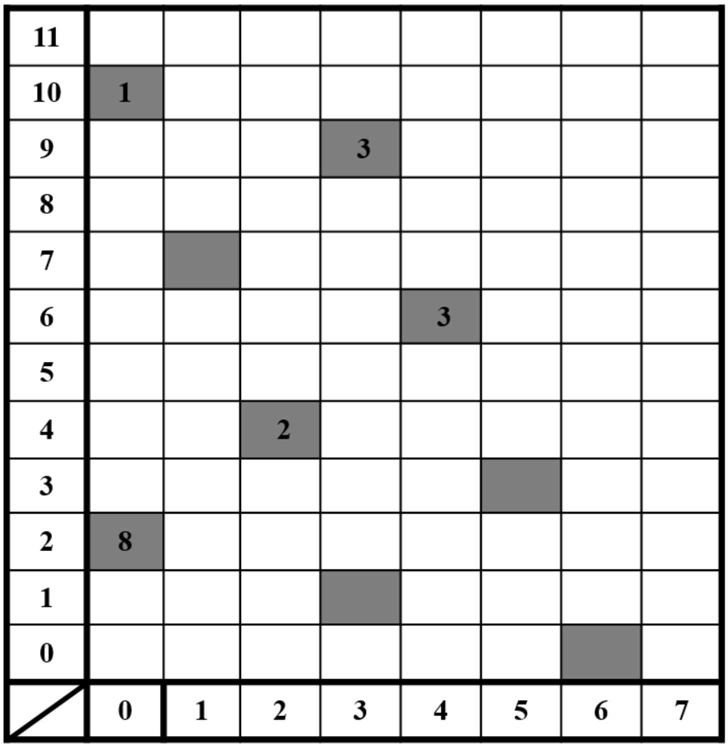

In these 17 pairs, 70.6% of advantaged losers made a stealing decision in their first move, which is slightly higher than the proportion of disadvantaged losers who stole lottery tickets from their coparticipant in their first move (60.3%). However, the difference is not statistically significant (Mann-Whitney p=0.414).

Table 3.8: Distributions of indexes of advantaged losers and disadvantaged winners

		Loser		**Winner**	
		Mean	Std. Dev	Mean	Std.Dev
Index of excess taking	Disadvantaged	-0.026	0.521	0.048	0.387
	Advantaged	0.088	0.501	-0.014	0.397
Index of excess taking with value	Disadvantaged	-0.328	3.027	0.132	1.186
	Advantaged	1.120	3.568	-0.039	1.242

A summary of indexes for disadvantaged losers, advantaged losers, advantaged winners and disadvantaged winners in the Unfair Rule treatment is shown in Table 3.8. We can see from

Table 3.8 that the mean values of both indexes are higher for disadvantaged winners than for advantaged winners and, more surprisingly, are higher for advantaged losers than for disadvantaged losers. None of these differences is statistically significant.[11] However, it is interesting that the direction of the observed difference for losers parallels the surprising part of Result 3 – that advantaged winners are more likely to take than fair winners. These findings raise the possibility that competing with unfair rules might have a *general* tendency to induce taking behaviour, even on the part of individuals who have benefited from the unfairness.

3.8 Discussion

Our findings arise a number of interesting issues.

A. *Does the fairness of competition matter?*

Previous research shows that people care about equality of outcome, intention-based reciprocity, social welfare preferences, desert-based fairness, and procedural fairness as randomness. However, none of these theories can well explain the results that we find in the experiment.

In our Fair Rule treatment and Unfair Rule treatment, the allocations of lottery tickets depended on the results of the series of card games in Part 1. In the Real Effort treatment, subjects competed in effort tasks. Winners got more lottery tickets then losers, which is the same across all three treatments. It is obvious that neither the effort task nor the series of card games can be counted as a random procedure, which suggests that the theory of procedural fairness as randomness cannot play any role in explaining our results.

As the distributions of lottery tickets were the same among all three treatments, the models of inequality aversion (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000) cannot explain the difference in degrees of willingness to accept inequalities between the Fair Rule treatment and the Unfair Rule treatment.

During the card game, a player's self-interested intention of trying to win the game was revealed if she used any of her replacement opportunities. According to our data, only 3 subjects in the Fair Rule treatment and 4 subjects in the Unfair Rule treatment did not use any replacement opportunities in the series of card games, which implies that almost all of the subjects in both the Fair Rule and Unfair Rule treatments revealed their self-interested intentions in the game. Therefore, we cannot use the theory of intention-based reciprocity (Blount, 1995; Offerman, 2002; Falk *et al.*, 2003) to explain the finding that people are more tolerant of the inequality in the Fair Rule treatment than in the Unfair Rule treatment.

[11] There is no significant difference in the index of excess taking between advantaged losers and disadvantaged losers (Mann Whitney p= 0.280) or between advantaged winners and disadvantaged winners (Mann-Whitney p= 0.693). If the index of excess taking with value is used, the corresponding tests give $p = 0.165$ and $p = 0.818$ respectively.

Although the card games were not random procedures, winning or losing in the Fair Rule treatment was mainly determined by luck. Even in the Unfair Rule treatment, where one player had three times as many replacement opportunities as the other and a player had to win four games in order to be the winner of the series, 23 per cent of the series were won by the disadvantaged player.[12] Clearly, the game involved an element of skill, but winning was not obviously a matter of effort. Working out the equilibrium strategy (or a best response to a given belief about the behaviour of an opponent) is a difficult mathematical problem; it is unlikely that any subject would have been able to solve this problem while taking part in the experiment. Therefore, it is hard to see how a theory of desert-based fairness could explain *both* the similarity in taking behaviour between the Fair Rule and Real Effort treatments *and* the dissimilarity between the Fair Rule and Unfair Rule treatments.

Our theory of strategy fairness instead provides a simple explanation of our findings. People's tolerance of inequality is sensitive to strategy fairness in the competition. To much the same extent that people are willing to accept inequalities that result from differences in effort, they are are willing to accept inequalities that are the result of fair procedures, even if individuals reveal self-interested intentions in the competition and even if the inequality does not reward effort or ability. This means in particular that the acceptability of a given allocation as a result of fair/unfair competition cannot be captured by any theory of fairness consideration in the existing literature. Strategy fairness is conceptually distinct from distribution fairness, fairness intention, desert-based fairness, social welfare preferences or procedural fairness as randomness. My strategy fairness model meshes strategy fairness with distribution fairness, and demonstrates how such an approach can explain people's tolerance to unequal outcomes that occur as a result of fair competition.

B. *Why is it important to make the competition fair?*

Our experiment provides a simple benchmark to test the role of strategy fairness in competitions. In our experiment, subjects are matched up and take part in competitions. The winners of the competitions get higher chances to win a prize. By using a treatment in which the rules of this competition are biased, we deliberately make some subjects able to win the competition more easily than others.

In our Fair Rule treatment, 44.2% of pairs chose to settle on the unequal distribution of lottery tickets. In the Unfair Rule treatment, this occurred in 22.4% percent of the games. These results suggest that when people compete in a fair competition, they are more tolerant of ex post inequality. This finding also supports the idea of Isoni *et al.* (2014) that procedural fairness reduces the salience of considerations of distribution fairness. On the other hand, when the rules of the competition are unfair, disadvantaged participants are more likely to try to re-establish fairness, even if doing so is costly.

[12] This is almost exactly the proportion that would occur in Nash equilibrium. In Nash equilibrium, the probability that the advantaged player wins a single game is 0.630 (see Chapter 4). The probability that the advantaged player wins a series of seven games is 0.766.

It is not surprising that disadvantaged losers in the Unfair Rule treatment were more likely to take than fair losers in the Fair Rule treatment. However, the results also show that advantaged winners in the Unfair Rule treatment were more likely to take than fair winners in the Fair Rule treatment. It is reasonable to expect that competing under unfair rules would make disadvantaged losers reluctant to accept the unequal outcome and make retaliatory moves. A similar line of thought would lead to the expectation that, as advantaged winners are favoured by the rule of the competition, they would be more tolerant to taking behaviour of disadvantaged losers. Surprisingly, the evidence is against that expectation. It seems that, when inequality is generated unfairly, the person who has benefited from the unfairness feels entitled to get even more. One possible explanation is that some people enjoy being in an advantaged position, and they do not feel bad about earning more than others by using their advantages. An alternative explanation is that advantaged winners might have a self-biased expectation about disadvantaged losers' beliefs. They might expect disadvantaged losers to believe that they (the advantaged winners) are not the one who set up the biased rules of the competition and that therefore, they should not be punished for the unequal outcome.

As subjects in the Unfair Rule treatment engaged in more taking behaviour, we observed greater efficiency losses in that treatment than in the Fair Rule treatment. This result indicates that as people care about strategy fairness and are ready to involve themselves in costly vendettas if they are treated unfairly, competitions with unfair rules would result in significant social inefficiencies.

C. *Procedural fairness as fair competition vs desert-based fairness*

My experimental design makes it possible to compare the influence of two different concepts of fairness: strategy fairness and desert-based fairness. We find no evidence that subjects in the Real Effort treatment behave differently from subjects in the Fair Rule treatment. Moreover, there is no significant difference in efficiency losses between these two treatments. The implication is that if the outcome has to be unequal, giving people equal opportunities to compete in competitions can have a similar tendency to mitigate resistance to inequality as offering them equal opportunities to put in effort.

However, even when subjects are given equal opportunity to put in effort or equal opportunity to compete, we find that more than 45% of the losers made a taking decision in their first moves in all three treatments. This result may suggest that people care both about strategy fairness and equality of distributions in competitions, and they would be willing to rectify the unequal outcomes even if the competition offer them equal opportunities to compete. Bolle *et al.* (2014) find that there is no significant difference between stealing ratios in the vendetta games with equal initial winning probability and stealing ratios in the vendetta games with unequal initial winning probability. They provide two explanations. One explanation assumes that individuals are motivated by pure nastiness (such as preferences with strong spite). The other explanation assumes that vendettas are triggered by the boundedly rational temptation of immediate gains from taking behaviour.

D. *Implications*

Although our experimental setup is simple and abstract, it provides a stylised representation of many real world situations. For instance, competition in markets often generates unequal outcomes between competitors. If people's willingness to tolerate inequality is influenced by strategy fairness in the market competition, then maybe policy makers should focus more on how to make the market competition more fair by ensuring that individuals have equal opportunities to compete, instead of just trying to equalize the final outcomes.

The finding that competing in an unfair environment makes both the disadvantaged party and the advantaged party behave more aggressively seems to be surprising, but if one looks more closely into various psychological and economic studies, one can find evidence in the direction of this finding. Milgram (1963) carried out one of the most famous studies of obedience in psychology. In the studies, participants were divided into two groups: learners and teachers. Teachers were asked to administer increasingly severe electric shocks to learners when they provided a wrong answer. Shock levels were labelled from 15 to 450 volts. Although most subjects were uncomfortable about doing this, all subjects continued to 300 volts. 65% of participants in the teacher group continued to give shocks up to the highest level of 450 volts. In 1971, Zimbardo and his team conducted the Stanford Prison Experiment (Zimbardo, 2009). Participants were recruited and told they would participate in a two-week prison simulation. Participants were assigned the role of either prisoners or guards. In the end, the experiment had to be terminated after only six days because the brutality of the Guards and the suffering of the Prisoners was way too intense. Zimbardo suggests that the behaviour of subjects who acted guards was significantly influenced by the situation that was created by the experiment, such as the roles, the norms, conformity pressures, and group identity. Karakostas and Zizzo (2016) conducted an experiment where participants were ordered directly or indirectly to destroy half of another participant's earnings at a cost to their own earnings. They find that around 60% of participants decide to comply with the orders. They suggests that the occurrence of the high destruction rate is due to the existence of the norm of compliance towards authority. The finding about the behaviour of advantaged players in our experiment can be also driven by the norm of compliance or obedience. One conjecture is that the advantaged winners in the Unfair Rule treatment may believe that they are picked by the 'authority' to have the right to earn more than others, and therefore they should try to earn more than their co-participants as 'ordered' by the 'authority'. Although future studies are required to check the robustness of our finding that advantaged players behave more aggressively in fair competitions than in unfair competitions, our finding draws attention to large potential cost of unfair competitions which can be caused by the decisions of both advantaged and disadvantaged parties in the society.

3.9 Conclusion

The main objective of our experiment was to explore how strategy fairness of competitions affects the willingness of individuals to accept inequality. We proposed a utility model of individual preferences for strategy fairness which complements the Fehr-Schmidt model for inequality aversion. The model assumes that strategy fairness influences fairness perceptions of outcomes.

We designed a novel card game which creates fixed inequality as an outcome. The rules of the card game could be easily understood as fair or unfair. When playing the game, subjects also need to reveal their self-interested intentions to win the game. This card game allowed us to explore whether people are more willing to accept inequalities that result from fair competitions than competitions with unfair procedures, even if individuals reveal self-interested intentions in the competition. We used a vendetta game (Bolle *et al.*, 2014) as an instrument to measure people's attitude towards the status of fairness.

Overall, the evidence shows that people are more tolerant of inequalities that result from fair competitions than competitions with unfair rules. We find a tendency for players to settle on initial inequalities when the card game gives them equal opportunities to compete in the game. Significantly more efficiency losses are observed when the rules of the game are biased. Surprisingly, we also find that in the unfair competition, not only are disadvantaged players more likely to take from their co-players, but so too are advantaged players. The results also show that males and females hold different beliefs about fairness norms or have difference preferences about strategy fairness.

3.10 Appendix

3.10.1 Appendix 1: Instructions for experiment

Welcome to today's experiment and thanks for coming. This is an experiment in decision-making. At the end of the experiment you will be paid the earnings you obtained from this experiment plus a participation fee of £3.

It is important that you remain silent and do not look at other people's work. If you have any questions, or need assistance of any kind, please raise your hand and an experimenter will come to you. If you talk, laugh, exclaim out loud, etc., you will be asked to leave and you will not be paid. We expect and appreciate your cooperation.

I will now describe the nature of the tasks within the experiment.

Tasks

This experiment contains two parts. At the beginning of this experiment, individuals with odd seat numbers will become participant As and individuals with even seat numbers will become participant Bs. Each participant A will be randomly matched with a coparticipant B. This matching will stay the same throughout the experiment. You will never be told who your coparticipant is.

During the experiment, you and your coparticipant will compete for 12 lottery tickets numbered 1 to 12. At the end of the experiment, the experimenter will put 12 tickets with the numbers 1 to 12 on them into a bag. One of you will be asked to come forward and pick one ticket from the bag. The number on this ticket will be the number of the winning ticket. If you hold the winning ticket, you will get £24. If your coparticipant holds the winning ticket, he or she will get £24.

Part 1

[For Effort treatment]

In this part, you will be given a task and your coparticipant will be given the same task. You and your coparticipant will do the task independently. After you both have finished the task, your score will be compared with your coparticipant's score. **At the end of the task, the winner will get 9 lottery tickets and the loser will get 3 lottery tickets.**

In the task, you will be presented with a number of words and your task will be to encode these words by substituting the letters of the alphabet with numbers using Table 1 below.

Table 3.9

A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
8	12	14	10	9	6	24	22	7	5	11	3	18	1	21	16	23	2	13	19	25	4	26	17	20	15

Example 1: You are given the word FLAT. The letters in Table 1 show that F=6, L=3, A=8, and T=19.

In the task, Table 1 will also be shown on each screen. The picture below shows you how the computer screen will look.

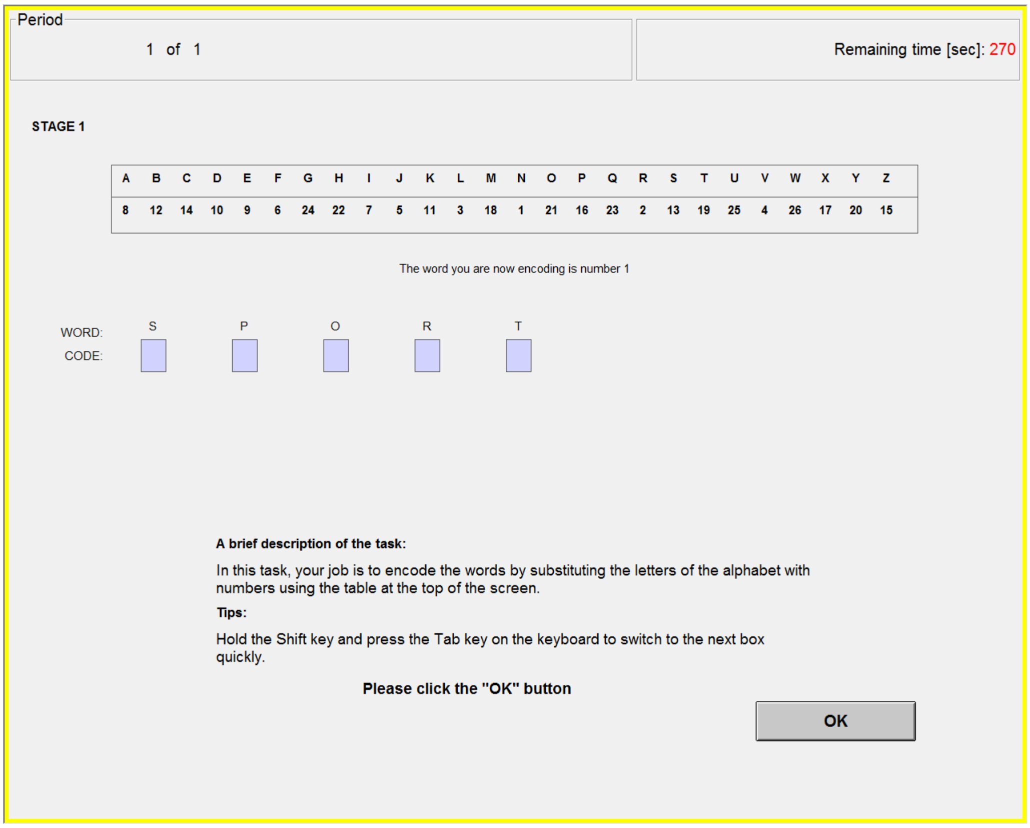

All the codes need to be entered into the boxes under the letters of the word that you are asked to encode. You can shift among boxes by clicking the boxes. After you encode a word, you need to click the 'OK' button to verify your codes. The computer will tell you whether the word has been encoded correctly or not. If the word has been encoded wrongly, you need to check your codes and correct them. Then, you need to click the 'OK' button again to verify the codes.

Once you encode a word correctly, the computer will prompt you with another word which you will be asked to encode. Once you encode that word, you will be given another word and so on. This process will continue for 6 minutes (360 seconds).

You and your coparticipant will be given the same words to encode in the same sequence.

After you both have finished the task, the number of words you have encoded will be your score for the task. Your score will be compared with your coparticipant's score. If you and your coparticipant get the same score, then the computer will compare the total amount of time that you used encoding these words (i.e. the time between the start of the task and when the OK button was clicked after you finished the last word) with the total amount of time that your coparticipant used.

At the end of the task, the person with the higher score will be the winner. If you and your coparticipant get the same score, then the person who encodes the words in the shorter time will be the winner. The winner will get 9 lottery tickets and the loser will get 3 lottery tickets.

Please raise your hand if you have any questions.

Before you start to take decisions, we ask you to answer some questions in the next several screens. The purpose of these questions is to check whether you have understood these instructions. Any mistake you may make in doing these questions will not affect your final money earnings.

When you have finished Part 1, please remain seated. When everyone has finished Part 1, I will distribute the instructions for Part 2.

[**For Fair Rule treatment**]

In this part, you will play a series of card games with your coparticipant. **The winner of the series of card games will get 9 lottery tickets and the loser of the series of card games will get 3 lottery tickets.**

At the start of each card game, you will be dealt a card and your coparticipant will also be dealt a card. Each card has a number of points, which can be any of the whole numbers in the range from 1 to 100. Each of these numbers is equally likely at each 'deal'. You and your coparticipant will have some opportunities to replace cards. To win the game, you need to hold a card with a higher number of points than the card hold by your coparticipant.

You can decide whether to **stick** with the card you have been dealt or **replace** it with a new card. If you decide to replace the card with a new card, then the computer will randomly draw a new one for you. Each number in the range from 1 to 100 is still equally likely at this 'deal'. However, if you decide to replace this card, you cannot go back to it again.

In each game, if you are participant A, you are allowed to replace cards up to 3 times. If you are participant B, you are allowed to replace cards up to 3 times. During the game, you can decide to stick with any card that is dealt to you. Once you have used up all these replacement opportunities, you will not be able to make any further replacement and have to stick with the last card that you have been dealt. On the computer display, there will be a message reminding you how many replacement opportunities you have left. The picture below shows you how the computer screen might look in the first game, before you had made any decision.

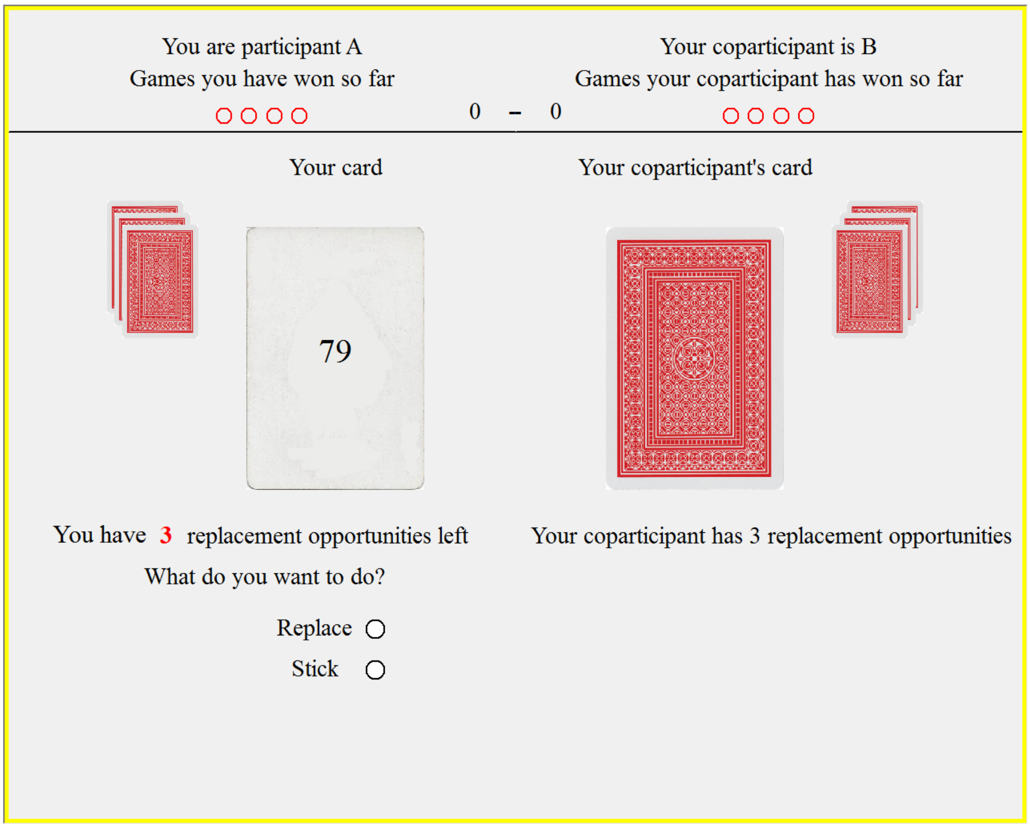

While you are making your decisions about whether to stick or to use replacement opportunities, you will not know what decisions your coparticipant is making. Nor will you know the numbers on the cards that are dealt to him or her. Similarly, your coparticipant will not know what decisions you are making, or the numbers on the cards that are dealt to you.

After you and your coparticipant have made the decision of sticking with a card you have been dealt or have used up all your opportunities for replacing cards, your coparticipant's card will be turned over. You can observe the points on your card and the points on your coparticipant's card. You will also be shown how many replacement opportunities your coparticipant has used. Whoever has the card with the higher number of points on it wins the game. If you both have the same number of points, the game is a draw. The first participant to win 4 games will be the overall winner of the series of games. Draws will not be counted. The overall winner will get 9 lottery tickets and the overall loser will get 3 lottery tickets.

Please raise your hand if you have any questions.

Before you start to take decisions, we ask you to answer some questions in the next several screens. The purpose of these questions is to check whether you have understood these instructions. Any mistake you may make in doing these questions will not affect your final money earnings. When you have finished Part 1, please remain seated. When everyone has finished Part 1, I will distribute the instructions for Part 2.

[For Unfair Rule treatment]

In this part, you will play a series of card games with your coparticipant. **The winner of the series of card games will get 9 lottery tickets and the loser of the series of card games will get 3 lottery tickets.**

At the start of each card game, you will be dealt a card and your coparticipant will also be dealt a card. Each card has a number of points, which can be any of the whole numbers in the range from 1 to 100. Each of these numbers is equally likely at each 'deal'. You and your coparticipant will have some opportunities to replace cards. To win the game, you need to hold a card with a higher number of points than the card held by your coparticipant.

You can decide whether to **stick** with the card you have been dealt or **replace** it with a new card. If you decide to replace the card with a new card, then the computer will randomly draw a new one for you. Each number in the range from 1 to 100 is still equally likely at this 'deal'. However, if you decide to replace this card, you cannot go back to it again.

In each game, if you are participant A, you are allowed to replace cards up to 1 time. If you are participant B, you are allowed to replace cards up to 3 times. During the game, you can decide to stick with any card that has been dealt to you. Once you have used up all these replacement opportunities, you will not be able to make any further replacement and will have to stick with the last card that you have been dealt. On the computer display, there will be a message reminding you how many replacement opportunities you have left. The picture below shows you what you might see on the computer screen in the first game, if you were participant A and before you had made any decision.

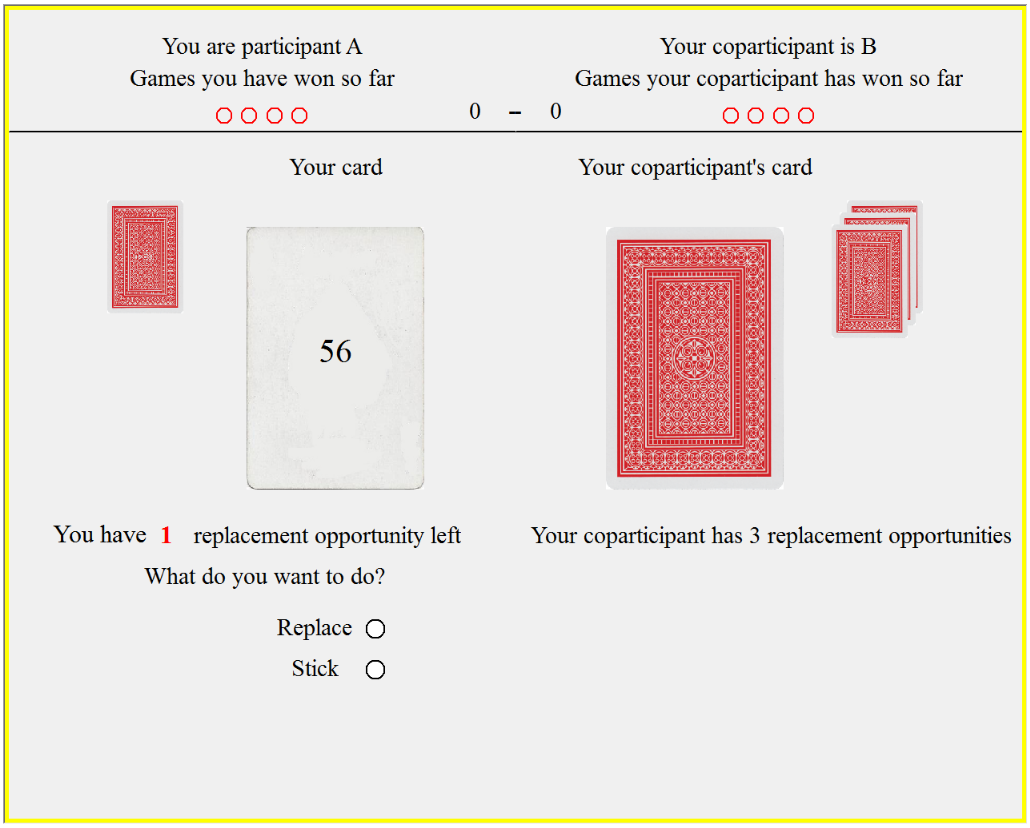

Participant B would see a similar screen, showing the card that he or she had been dealt and saying that he/she had 3 replacement opportunities left. While you are making your decisions about whether to stick or to use replacement opportunities, you will not know what decisions your coparticipant is making. Nor will you know the numbers on the cards that are dealt to him or her. Similarly, your coparticipant will not know what decisions you are making, or the numbers on the cards that are dealt to you.

After you and your coparticipant have made the decision of sticking with a card you have been dealt or have used up all your opportunities for replacing cards, your coparticipant's card will be turned over. You can observe the points on your card and the points on your coparticipant's card. You will also be shown how many replacement opportunities your coparticipant has used. Whoever has the card with the higher number of points on it wins the game. If you both have the same number of points, the game is a draw. The first participant to win 4 games will be the overall winner of the series of games. Draws will not be counted. The overall winner will get 9 lottery tickets and the overall loser will get 3 lottery tickets.

Please raise your hand if you have any questions.

Before you start to take decisions, we ask you to answer some questions in the next several screens. The purpose of these questions is to check whether you have understood these instructions. Any mistake you may make in doing these questions will not affect your final money earnings.

When you have finished Part 1, please remain seated. When everyone has finished Part 1, I will distribute the instructions for Part 2.

Part 2

At the end of Part 1, 12 lottery tickets were allocated between you and your coparticipant based on the result of the tasks you carried out. The winner in Part 1 got 9 lottery tickets and the loser in Part 1 got 3 lottery tickets. The tickets are numbered 1 to 12. The computer has picked 9 of these numbered tickets at random and assigned them to the winner. The remaining 3 tickets have been assigned to the loser. One of these lottery tickets will be the winning ticket, which gives a prize of £24. At the end of Part 2, the number of the winning ticket will be picked at random. Therefore, each lottery ticket gives a 1/12 chance of winning the prize. At the end of Part 2, if you hold the winning ticket, you will get the prize. If your coparticipant holds the winning ticket, he or she will get the prize.

In this part of the experiment, you and your coparticipant will take turns to choose whether to stay with the current distribution of lottery tickets or to change it. The loser in Part 1 will make a choice first.

On the screen there will be three baskets. One contains the lottery tickets that you currently hold, one contains the lottery tickets that your coparticipant currently holds, and one is a bin.

When it is your turn to choose, you will be asked to decide whether you want to take some lottery tickets from your coparticipant, and if so, how many lottery tickets to take. Amounts taken have to be in blocks of three (so if you choose to take tickets, the number you take can be 3, 6 or 9, up to as many as your coparticipant has at the time). You always have the option of not taking any tickets.

For every block of 3 lottery tickets that you take away from your coparticipant, one ticket from the block will be moved into your basket and the other two will be moved into the bin. If at the end of Part 2 the winning ticket is in the bin, neither you nor your coparticipant gets the prize.

The picture below shows how the computer screen would look at the start of Part 2 if you had been the loser in Part 1. Your basket is on the left, containing the three tickets that you earned in Part 1. Your coparticipant's basket is on the right, containing the nine tickets that he or she earned in Part 1. The bin is at the bottom. At the top of the screen you are told that it is your turn to make a decision.

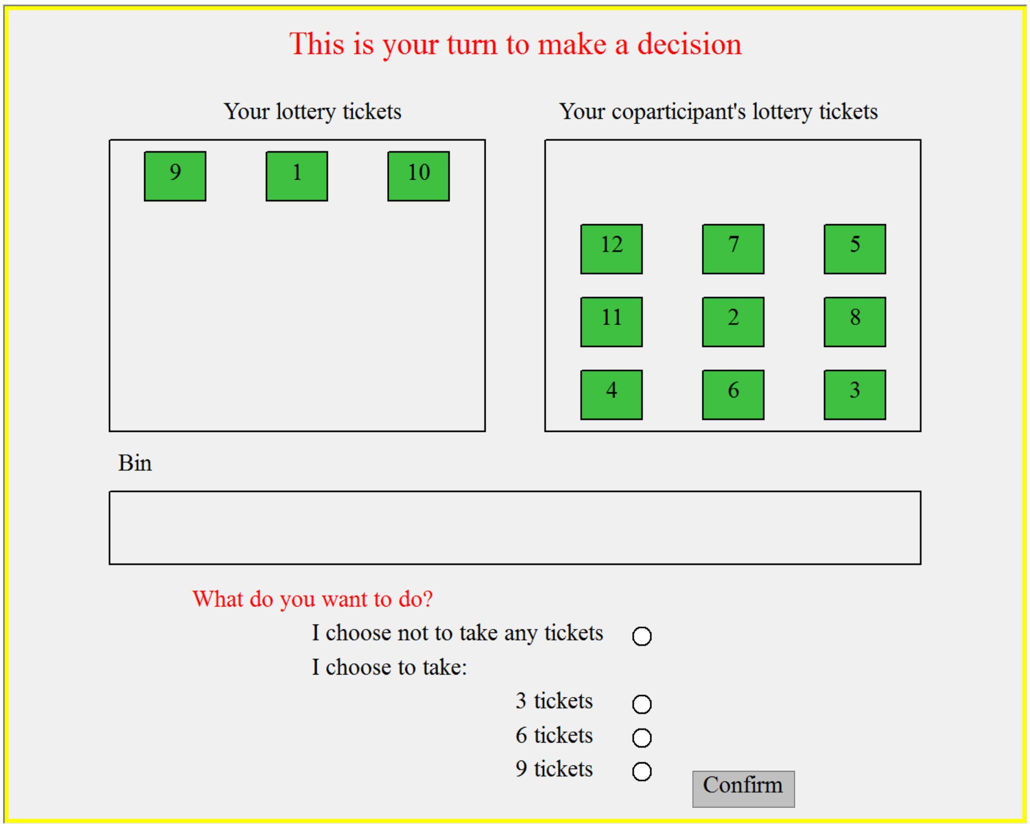

Before making any decision, you are allowed to try different possible numbers of blocks of lottery tickets to take away from your coparticipant. The computer will show you how many of these lottery tickets will be moved from your coparticipant's basket into your basket and how many of these lottery tickets will be moved from your coparticipant's basket into the bin. These lottery ticket will be shown in a different colour.

For example, if you clicked the option 'I choose to take 3 tickets', the picture below shows you what the computer would display. From the picture, you can see that three lottery tickets have been moved from your coparticipant's basket. One of these (number 12) has been moved into your basket. The other two (numbers 3 and 6) have been moved into the bin. All these tickets (numbers 3, 6 and 12) are in yellow.

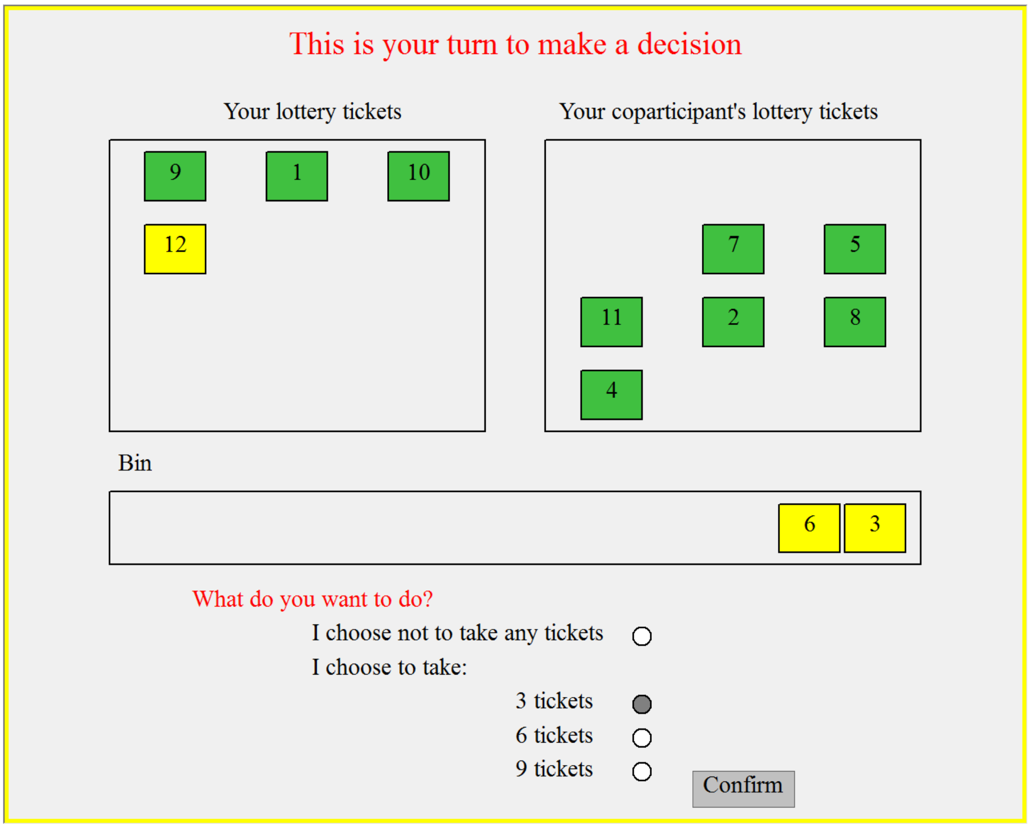

After your decision is made, you need to click the 'Confirm' button. The baskets and bin will be updated to show you the outcome of your decision and the current location of all the lottery tickets. All the lottery tickets will come back to being coloured green. Then it will be your coparticipant's turn to make decisions on whether to take lottery tickets from you, and if so, how many lottery tickets to take. After your coparticipant has chosen, the baskets and bin will be updated again to show you the current location of all the lottery tickets as a result of your coparticipant's choice. Then it will be your turn to choose again, and so on.

If there are four turns in a row (two for you and two for your coparticipant) in which neither of you takes lottery tickets, then Part 2 will end. Because tickets can be taken only in blocks of three, Part 2 will also end if your basket and your coparticipant's basket both contain less than three tickets.

The experimenter will then put 12 numbered tickets into a bag. One of you will be asked to come forward and pick one ticket from the bag. The number on this ticket will be the number of the winning ticket. You will see whether this winning ticket is in your basket, or in your coparticipant's basket, or in the bin. If the winning ticket is not in the bin, whoever holds it will get the prize of £24. If the winning ticket is in the bin, then neither you nor your coparticipant gets the prize. In all cases, both of you will also get a £3 participation fee.

Please raise your hand if you have any questions.

Before you start to take decisions, we ask you to answer some questions in the next several screens. The purpose of these questions is to check whether you have understood these instructions. Any mistake you may make in doing these questions will not affect your final money earnings.

When you have finished Part 2, please remain seated until everyone has finished Part 2.

3.10.2 Appendix 2: Regression results

Table 3.10: Estimation for losers' decisions

	Overall		NF		RF	
	(1)		(2)		(3)	
	β	ME	β	ME	β	ME
Real Effort	0.218	0.050	0.371	0.076	0.453	0.091
	(0.284)	(0.067)	(0.462)	(0.098)	(0.470)	(0.098)
Unfair Rule	0.361	0.083	0.174	0.035	0.158	0.030
	(0.248)	(0.057)	(0.412)	(0.082)	(0.426)	(0.082)
Lag Taken			2.066***	0.448***	2.016***	0.430***
			(0.362)	(0.074)	(0.376)	(0.077)
Lottery ticket difference					−0.218***	−0.041***
					(0.073)	(0.013)
Constant	−0.819***		−1.745***		−2.610***	
	(0.187)		(0.337)		(0.476)	
Observations	390		195		195	
LR chi2	2.123		33.478		35.508	
Prob > chi2	0.346		0.000		0.000	
Baseline predicted probability	−0.602		−1.303		−2.223	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the steal and 0 if the subject chose not to steal. We used panel data to estimate all these models. The data used to estimate model 1 contains 390 observations from 146 subjects. The data used to estimate model 2 and model 3 contain 195 observations from 134 subjects. For each model, the left column contains coefficients, and the right column report marginal effects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

Table 3.11: Estimation for winners' decisions

	Overall		NF		RF	
	(1)		(2)		(3)	
	β	ME	β	ME	β	ME
Real Effort	0.074	0.014	−0.097	−0.018	−0.332	−0.044
	(0.711)	(0.134)	(0.412)	(0.076)	(0.752)	(0.095)
Unfair Rule	1.079^{+}	0.209^{+}	0.420	0.081	0.871	0.132
	(0.641)	(0.127)	(0.359)	(0.071)	(0.663)	(0.104)
Lag Taken			2.754^{***}	0.525^{***}	0.447	0.064
			(0.331)	(0.050)	(0.635)	(0.098)
Lottery ticket difference					-0.632^{***}	-0.089^{***}
					(0.180)	(0.021)
Constant	-1.574^{***}		-2.411^{***}		−0.575	
Observations	289		283		283	
LR chi2	3.403		71.500		28.799	
Prob > chi2	0.182		0.000		0.000	
Baseline predicted probability	−1.059		−1.616		−0.280	

Notes: $*$ 5% level, $**$ 1% level, $***$ 0.1 %. Standard errors in parentheses. The dependent variable in these three models is a dummy equal to 1 if the subject chose the steal and 0 if the subject chose not to steal. We used panel data to estimate all these models. The data used to estimate model 1 contains 289 observations from 146 subjects. The data used to estimate model 2 and model 3 contain 283 observations from 146 subjects. For each model, the left column contains coefficients, and the right column report marginal effects. Results for all three models are based on random effects logit estimations in which subject-specific random effects are controlled.

Chapter 4

Search competition: a theoretical and experimental investigation

4.1 Introduction

This paper introduces a new type of game, a search competition game. This is related to two theoretical and experimental literatures. The contest literature studies games in which players compete for a fixed prize by expending costly effort; players who expend relatively more effort have a higher probability of winning. Thus the decision about how much effort to expend is made under strategic uncertainty. The search literature studies situations in which an individual searches sequentially among a set of alternative offers of unknown value; there is either an explicit cost of inspecting each offer, or the maximum number of inspections is fixed and offers must be accepted or rejected sequentially without recall (or both). In either case, the decision about when to terminate the search is made under (non-strategic) uncertainty. Our paper is about a competition for a fixed prize when the competition takes the form of parallel searches without recall. Each player's objective is to find an offer that is better than the offer found by the other player(s). There are many real world examples of this kind of situation, such as football managers searching for players, competing firms searching for managers, and firms in patent races searching for researchers. Little work has been done on this type of game. Existing research suggests opposing intuitions about actual behaviour, relative to Nash equilibrium play. Experiments on contests have found a general tendency for players to expend more effort than the equilibrium solution (intuitively: trying too hard to win). Experiments on search have found a general tendency for individuals to search less than is optimal given risk-neutrality (intuitively: not trying hard enough to find the best offer). Our experiment is aimed to find out which (if either) of these intuitions applies to search competition.

We implement the search competition as a card game. In the search competition, players compete in fixed pairs. At the beginning of the game, players are dealt a card each. There is a number of points on each card, which is randomly drawn from a discrete uniform distribution. In the game, each player is given some opportunities to replace cards. If a player decides to use a replacement opportunity, she will be dealt with a new card, which is drawn randomly from the same distribution. All draws are stochastically independent. Rejected cards cannot be recalled. There is no cost of replacing cards. Once a player has used up all replacement opportunities, she has to stick with the last card dealt to her. Players cannot see each other's cards or know their replacement decisions until the end of the game. After both players have made a decision of sticking to a card or have used up all their replacement opportunities, the game ends. The player who holds a card with a higher number wins the game.

The general picture of our search competition is similar to games considered in contests literature. In both situations, there is a single prize to be awarded. Multiple players compete for the prize, but only one player receives the prize. However, almost all games in the contests literature are developed based on three models: rent-seeking contests (Tullock, 2001), rank-order tournaments (Lazear and Rosen, 1981), and all-pay auctions (Hirshleifer, 1978; Nalebuff and Stiglitz, 1983; Hillman and Riley, 1989). Players in these games compete by exerting costly irreversible effort. Each player's probability of winning depends not only on her own effort level but also on other players' effort levels. In contrast, players in our game compete for the prize by playing a pure strategic game. There is no monetary cost of exerting effort. But players face opportunity costs when making decisions on whether or not to reject a card and ask for a new one.

Our search competition should also interest economists who study individual search behaviour because it resembles a dynamic search situation arising in many important economic contexts. In our search competition, players search for a card sufficiently good for winning the game from a known distribution. Each player has a limited number of replacement opportunities, which means she can only observe a finite number of cards. A new card is drawn for the player only if she rejected the previous card. The decision problem that players face in our search competition is similar to many search problems in the real world, just as buyers must decide whether or not to search for another house, and a job market candidate must decide whether or not to reject the current offer and search for another job. A large number of studies have been done to investigate variations of individual search problems, such as the price search problem and the secretary problem[1] (Stigler, 1961; Schotter and Braustein, 1981; Hey, 1981, 1982, 1987; Kogut, 1990, 1992; Sonnemans, 1998; Schunk and Winter, 2009a,b). To our knowledge, all the existing studies on search problems focus on individual search behaviour. Our paper is the first to study a search problem in a competitive setting. There are many

[1]In the secretary problem, the decision maker is presented with a set of items, one at a time, in a random order. At any period, the decision maker is able to rank order all the items that have been observed in terms of their desirability. In each period, the decision maker must either accept the presented item , in which case the observation process terminates, or reject it, in which case the decision maker is presented the next item in the randomly determined order. See for example Rapoport and Tversky (1970) and Seale and Rapoport (1997, 2000) for detailed discussions of the secretary problem.

real world situations involving search with competitions. For instance, football managers compete on signing better players than others. Firms compete on hiring good managers. So, it is important to investigate how people behave in search competitions and what factors may influence their decisions.

In the paper, we show the method of deriving the unique subgame perfect Nash equilibrium to the search competition. Using our method, we derive the unique subgame perfect Nash equilibrium for competitions with up to three replacement opportunities for each player. The equilibrium solutions show that players should use reservation value strategies. That is, for any given replacement opportunity, there is a reservation value such that that opportunity should be used if and only if the number of points on the card currently held is lower than the reservation value. A player's probability of winning depends on the number of replacement opportunities of both players. The more replacement opportunities the opponent has, the higher reservation values a player should hold.

To gain insight into players' behaviour in the card game compared to the theoretical prediction, we designed a laboratory experiment. In the experiment, subjects competed for lottery tickets by playing a series of card games in fixed pairs. The first player to win 4 card games was the overall winner. The winner got more lottery tickets than the loser. At the end of the experiment, one of these lottery was randomly drawn as the winning lottery ticket. The player who held the winning lottery ticket got the prize.

The experimental literature on contests has generally found that the average effort level is significantly higher than theoretical predictions.[2] However, existing experimental evidence on search behaviour show that subjects tend to stop searching too early, compared to the optimal strategy. As our search competition consists elements of both individual search problem and competition, it is hard to predict whether subjects in our experiment would search too little or too much in comparison to the equilibrium prediction.

As far as we are aware, there is only one paper studying search related competition using data from laboratory experiment. Tenorio and Cason (2002) analyse contestants' behaviour in *The Wheel* game in the television game show *The Price is Right* using data from both the television show and a matching laboratory experiment. In *The Wheel* game, three contestants take turns to spin a wheel with twenty equal partitions labelled from 5 to 100. Points that a contestant gets depend on where the wheel stops. Each contestant is given the option of spinning the wheel twice, and her score equals the sum of her spin(s). The contestant whose score is closest to but not more than 100 wins. They find that contestants frequently make decisions inconsistent with the equilibrium prediction. The rate of deviation increases when the decision problem becomes more difficult. Deviations from optimal play often happens because of failing to spin again when contestants should have chosen to continue. They suggest two potential explanations for contestants' under-spinning bias: computational errors and omission bias. There are several differences between our search competition and *The Wheel*

[2] See Section 4.2 for a comprehensive review

game. First, *The Wheel* is a sequential game. The fact that contestants in later positions can observe choices made by contestants in earlier positions before they make any decision makes the information asymmetric between contestants. In contrast, players make decisions simultaneously in our search competition. Information about the result of the game and other players' decisions can only be learnt at the end of the game. Moreover, the decision problem faced by contestants in *The Wheel* game is much more complicated than the problem in our search competition, in the sense that contestants in *The Wheel* game not only need to get a score high enough too win with two spins but also need to make decisions conditional on the previous decisions of other contestants and the results of those contestants' spins. Furthermore, *The Wheel* game is an artificial game which is not directly related to any real world scenario. Our search competition game instead combines competition and search which are two topics of interest to economists.

We begin by providing an overview of the related literature in Section 4.2. Section 4.3 introduces the card game and the theoretical model. Section 4.4 details the experimental design and implementation. In section 4.5 , we construct the econometric model for the data analyses. We then report the experimental results (Section 4.6), discuss their implications (Section 4.7), and conclude (Section 4.8).

4.2 Literature review

4.2.1 Overbidding and heterogeneity in contests

Over the last few decades, a rapidly growing number of controlled laboratory experiments have been conducted to test contest theory. The advantage of using controlled experiments is that it allows researchers to measure the actual effort made by the contestants, while controlling for confounding effects. Experimental literature on contests focuses on a relatively small range of games, which are based on three models: rent-seeking contests (Tullock, 2001), rank-order tournaments (Lazear and Rosen, 1981), and all-pay auctions (Hirshleifer, 1978; Nalebuff and Stiglitz, 1983; Hillman and Riley, 1989). The models of all these three games assume that players expend costly efforts while competing for a prize and each individual player's probability of winning the prize of depends on the proportion of her expenditure relative to the other individuals' expenditures.

Overbidding and heterogeneity in the behaviour of individual contestants are the two main phenomena observed in contest experiments: average effort levels are significantly higher than the risk-neutral Nash equilibrium prediction, and there is a large variation in effort levels between and within subjects; see Dechenaux *et al.* (2015) for a comprehensive review. There are many different explanations that have been proposed for overbidding in contests literature (Sheremeta, 2013). The first most commonly cited explanation is bounded rationality: subjects are prone to mistakes in contests (Sheremeta, 2011; Chowdhury *et al.*, 2014; Lim *et al.*, 2014).

The second explanation is that in addition to the monetary value of the prize, subjects also get a non-monetary utility of winning (Sheremeta, 2010; Price and Sheremeta, 2011; Mago *et al.*, 2016). The third explanation is judgemental biases, such as probability distortions and the hot hand fallacy (Amaldoss and Rapoport, 2009; Sheremeta and Zhang, 2010; Sheremeta, 2011). The final explanation is relative payoff maximization. Studies shows that subjects care about their relative payoffs as well as the utility of winning (Herrmann and Orzen, 2008; Mago *et al.*, 2016).

Between-subject heterogeneity in contests is usually attributed to heterogeneous preferences and demographic differences (Sheremeta, 2013). Studies suggest that gender and religion are the two most important demographic differences that can explain heterogeneous behaviour between contestants (Mago *et al.*, 2013; Price and Sheremeta, 2015). Individual differences in preferences towards risk, losses and winning can also explain heterogeneous behaviour of the subjects in contests. Several experimental studies show that compared with risk-neutral or risk-seeking subjects, risk-averse subjects expend less effort in lottery contests (Millner and Pratt, 1991; Sheremeta and Zhang, 2010; Sheremeta, 2011; Mago *et al.*, 2013). More loss-averse subjects tend to exert lower efforts in contests than less loss-averse subjects (Shupp *et al.*, 2013). Finally, subjects' other-regarding preferences, such as inequality aversion, are also correlated with effort levels in contests (Balafoutas *et al.*, 2012; Mago *et al.*, 2016).

4.2.2 Feedback information and learning

Many experimental studies of contests point out that within-subject heterogeneity of effort in contests can be explained by learning and by the hot hand fallacy. In this section, we focus on the learning behaviour and the effect of feedback on subjects' performance in contests.

Studies show that in experiments with repeated contests, as subjects become more experienced overtime, they learn to lower their effort levels (Davis and Reilly, 1998; Gneezy and Smorodinsky, 2006; Price and Sheremeta, 2011; Brookins and Ryvkin, 2014; Mago *et al.*, 2016). However, all of these studies find that subjects' learning in contests is not sufficient – average effort levels remain significantly higher than the Nash equilibrium prediction, even in the last periods of the experiment.

Sheremeta (2011) finds that, given information about the outcomes of previous contests, subjects are more likely to expend higher effort if they won the previous game. Sheremeta and Zhang (2010) report similar results. They point out that the hot hand fallacy might be an interpretation of this phenomenon, which was first described by Gilovich *et al.* (1985) as a belief that a person who has experienced success with a seemingly random event has a greater chance of further success in additional attempts.

Several experimental studies compare the effort level chosen by subjects in contests under different feedback conditions. Kuhnen and Tymula (2012) examine the effect of private feedback about relative performance on performance in a real effort tournament. After receiving feedback, subjects who got better ranks than expected decrease output, but expect even better ranks in the future, whereas subjects who ranked lower than expected increase their output and lower their future expectations on ranks. Ludwig and Lünser (2012) find that in two-stage tournaments with feedback information on relative effort levels, subjects who lead tend to reduce their effort levels in the second stage, while subjects who lag behind tend to increase their effort levels. Nevertheless, they find that there is no significant difference in average stage effort levels between the feedback and no feedback treatments. In a rent-seeking lottery contest, Mago *et al.* (2016) also find that providing information about others' effort levels decreases the within-subjects heterogeneity of effort, but it does not affect the average effort level in contests. Subjects who expended higher effort than the winning effort reduce their effort in the future, whereas the opposite is true of subjects who exerted lower effort than the winning effort. Similarly, Fallucchi *et al.* (2013) compare the effects of information feedback in share contests and lottery contests. Additional feedback about rivals' choices and earnings increases average expenditures in the share contest. However, unlike the other studies, they find that providing subjects with additional feedback about rivals' effort reduces aggregate effort in a lottery contest.

4.2.3 Search problem

Many experimental studies have been done to investigate variations of individual search problems, mainly the price/wage search problem and the secretary problem (Stigler, 1961; Schotter and Braustein, 1981; Hey, 1981, 1982, 1987; Kogut, 1990, 1992; Sonnemans, 1998; Schunk and Winter, 2009a,b). In the price/wage search problem, a decision maker needs to search for a price/wage which can maximize her payoff. The secretary problem is a more complicated decision problem, in which a decision maker has to pick the option with the highest rank based only on the relative ranks of the presented alternatives.[3]

One way of classifying these studies on search problems is by the opportunity of recall – whether the rejected offer can be recalled or not. Most experimental studies on individual search behaviour focus on cases where recall is allowed. In these cases, there is always a search cost, otherwise a subject can always decide to check all the offers and choose the best one, and therefore the problem would be trivial. Much less work has been done on search problems where recall is not possible. Rapoport and Tversky (1970) and Schotter and Braustein (1981) study standard search problem both with recall and without recall in various cost conditions. Seale and Rapoport (1997, 2000) examine the secretary problem with assumptions that recall of rejected applicants is not possible and there is no cost of searching.

Existing experimental evidence on search problems both with and without recall suggests

[3]See for example Rapoport and Tversky (1970); Seale and Rapoport (1997, 2000) for detailed discussions of the secretary problem.

that subjects tend to search too little compared to the risk neutral prediction of the optimal strategy and there is heterogeneity in subjects' search behaviour (Schotter and Braustein, 1981; Hey, 1987; Cox and Oaxaca, 1989; Sonnemans, 1998; Seale and Rapoport, 1997, 2000). However, experimental findings in relation to search with recall show that subjects' search behaviour is relatively efficient in the sense that their actual earnings are close to those they would have achieved had they used the optimal strategy, even though the subjects' searching strategy is different from the theoretical optimal strategy (Schotter and Braustein, 1981; Harrison and Morgan, 1990; Kogut, 1990, 1992; Sonnemans, 1998).

Several explanations have been proposed for under-search in search problems with recall. Schotter and Braustein (1981) find that risk averse subjects set lower reservation wages than risk neutral subjects, and therefore search less. Cox and Oaxaca (1989) test a finite horizon model of sequential search by individual agents. They suggest that the subjects' tendency to search too little can be explained by their risk-averse behaviour. Similarly, Kogut (1990) also finds that in an infinite-horizon search experiment, subjects stopped searching even though the expected gain was greater than the marginal cost. Evidence suggests that a large percentage of this behaviour can be explained by risk aversion. In contrast, Schunk and Winter (2009a) show that behavioural heterogeneity in search can be linked to heterogeneity in individual preferences, such as loss aversion and reference point updating. But, no evidence supports the hypothesis that risk aversion is related to subjects' search behaviour. Schunk and Winter also find that there is no correlation between individual risk attitudes and their reservation prices and numbers of searches, whereas the correlation between loss aversion and these search parameters is significant.

Camerer (2006) suggests that the phenomenon of searching too little can be explained by heuristic rules. Many experimental studies find that subjects' preferences, especially risk attitudes, do not help to predict their search behaviour, while a few simple decision heuristics can (Sonnemans, 1998; Houser and Winter, 2004; Schunk and Winter, 2009a). Sonnemans (1998) finds that only a small amount of early stopping behaviour can be explained by risk aversion, while subjects' learning processes can partly explain the early stopping tendency. Two different kinds of learners are considered: naive learners and sophisticated learners. Naive learners tend to repeat behaviour that results in good outcomes. Sonnemans finds that naive subjects tend to lower their reservation price after experiencing negative regret that they could have earned more if they had stopped earlier. There is no evidence of learning from positive regret. Moreover, naive learners who stop early have less opportunity to learn than naive learners who stop late, which may also explain the early stopping behaviour. Sophisticated learners instead tend to test different strategies first, and then follow the most successful one. Systematic biases of sophisticated learners can also cause early stopping. Schunk and Winter (2009a) show that simple heuristics, like the constant reservation price heuristic and the satisfier heuristic can describe observed search behaviour which cannot be explained by individual risk references.

The theory of optimal search strategy predicts that in search contexts with recall, a rational subject should never decide to return to an offer drawn earlier, and the decision on which offer to accept should depend only on the period of the observation but not on the history of the process. However, experimental evidence shows that there is a tendency to recall rejected offers (Schotter and Braustein, 1981; Kogut, 1990, 1992). Kogut (1990) suggests that subjects' decisions might be influenced by sunk costs, which can explain the fact that recall occurs. In studies of the secretary problem, Seale and Rapoport (1997, 2000) find that subjects' decisions to either accept or reject a candidate depends on observed patterns in the sequence. Zwick *et al.* (2003) study sequential search behaviour in a secretary problem. They find that the rate of recall is influenced by periods since the last candidate was encountered and the average rate of candidate arrival. The observed sequence of candidate arrival affects the length of subjects' search period.

The existing studies on search problems without recall do not provide many explanations of the phenomenon of under-search in that search context. Seale and Rapoport (1997) and Seale and Rapoport (2000) sugguest that in the secretary problem, subjects' behaviour of stopping too early can be explained by *endogenous* cost of search, i.e. subjects consider time spent on observing the applicants as an implicit search cost.

Experimental evidence of learning in various search contexts is quite mixed. Seale and Rapoport (1997) examine the secretary problem in experiments with repeated tasks, with a known number of applicants, without search cost and with no recall allowed. They find signs of learning over repeated trials – there is an increase in the proportion of decisions consistent with theoretical predictions. In contrast, Seale and Rapoport (2000) study a similar search problem, but they find only very week learning trends. Zwick *et al.* (2003), who examine the secretary problem with recall, find that learning is insignificant.

4.3 The basic idea of the card game and Nash equilibrium solution

4.3.1 The basic idea of the card game

In the card game, players compete with each other in groups. Each group contains two players, $i = \{1, 2\}$. At the start of each card game, players are dealt a card each. Each card has a number of points, which can be any of the whole numbers in the range from 1 to m. Each of these numbers are equally likely at each 'deal'. Players are informed of the distribution.

During the game, each player is offered a number of opportunities to replace cards, O_i. Each player decides individually on whether she wants to stick with the card that she has been dealt or to use the replacement opportunities. If the player decides to replace her card, then the computer will randomly draw a new one for her. Each number in the range from 1 to m

is still equally likely at this 'deal'. However, the player cannot return to the replaced card. Once the player has used up all the replacement opportunities, she has to stick with the last card that she has been dealt.

At the end of the game, the player who holds a card with a higher number of points than the card held by her co-player wins the game.

4.3.2 Subgame perfect Nash equilibrium solution

We now derive the subgame perfect Nash equilibrium solution for the card game using the backward induction method. In the theoretical model, we used continuous distributions (for ease of analysis), while in the experiment, we use discrete distributions (for ease of implementation). As the experiment used a high value of m ($m = 100$), this difference should not seriously affect the accuracy of the theoretical predictions.

Consider a card game with two players, $i = \{1, 2\}$. The numbers on the cards are drawn randomly and independently from a uniform distribution over the interval $[0, 1]$. Each player knows the distribution from which the number on the card is independently and randomly drawn. In the card game, each player i is given O_i replacement opportunities.

Assume both players are fully rational and prefer winning to losing. Therefore, they have the target of maximizing their probability of winning. There is no need to assume risk neutrality. According to expected utility theory, any attitude to risk can be represented by a utility function. The risk attitude of a player influences the curvature of her utility function, i.e. a concave utility curve indicates risk-averse behaviour, a convex utility curve indicates risk-seeking behaviour, and risk neutrality is reflected by a linear utility function. For each player, the card game has only two possible outcomes: win or lose. Picking any two points on the utility function where one is preferred to the other, the player's utility function can be normalized by assigning the worse outcome to zero and the better outcome to one. Since there are only two outcomes and there are only two numbers to be assigned to the outcomes in the utility function, the risk attitude of the player would not influence the result that the better one is always preferred to the worse. Therefore, the subgame perfect Nash equilibrium solution does not depend on the risk attitude of the players.

Let $\widehat{x_0}$ be the actual value of first card dealt to player 1. Let $\widehat{x_1}, ..., \widehat{x_{o_1}}$ be the actual value of player 1's replacement cards 1 to O_1. Assume player 1 follows a reservation value strategy, she chooses *cut-offs* $\kappa_1^1, \kappa_2^1, ..., \kappa_{O_1}^1$, i.e.

1. Stick with the first card iff $\widehat{x_0} \geq \kappa_1^1$; otherwise use the first replacement opportunity; and

2. Stick with the second card iff $\widehat{x_1} \geq \kappa_2^1$; otherwise use the second replacement opportunity;

and

$\vdots$

O_1. Stick with the O_1th card iff $\widehat{x_{O_1}} \geq \kappa^1_{O_1}$; otherwise use the O_1th replacement opportunity.

We take player 1's cut-offs as given. Given the cut-offs, we know that her final card has density $f_1(x)$, and cumulative distribution $F_1(x)$.

Let $\widehat{y_0}$ be the value of the first card dealt to player 2. Let $\widehat{y_1}, ..., \widehat{y_{O_2}}$ be the value of player 2' replacement card 1 to O_2. Assume player 2 chooses cut-offs $\kappa^2_1, \kappa^2_2, ..., \kappa^2_{O_2}$, i.e.

1. Stick at first card iff $\widehat{y_0} \geq \kappa^2_1$; otherwise use the first replacement opportunity; and
2. Stick at second card iff $\widehat{y_1} \geq \kappa^2_2$; otherwise use the second replacement opportunity; and

$\vdots$

O_2. Stick at O_2th card iff $\widehat{y_{O_2}} \geq \kappa^2_{O_2}$; otherwise use the O_2th replacement opportunity.

Let $\kappa^{2*}_1, \kappa^{2*}_2, ..., {\kappa^2_{O_2}}^*$ be the optimal values of $\kappa^2_1, \kappa^2_2, ..., \kappa^2_{O_2}$. Let π^*_0 be player 2's probability of winning (before knowing $\widehat{y_0}$), if player 2 plays optimally. Let $\pi^*_{O_2}$ be player 2's probability of winning, conditional on using the O_2th replacement opportunity (before knowing $\widehat{y_{O_2}}$).

Using backward induction, we begin by deriving ${\kappa^2_{O_2}}^*$. Suppose player 2 has used $(O_2 - 1)$ replacement opportunities, and holds card $\widehat{y_{O_2}}$. If player 2 decides to stick to card $\widehat{y_{O_2}}$, the probability of winning is $F_1\left(\widehat{y_{O_2}}\right)$; if player 2 decides to use the O_2th replacement opportunity, the probability of winning is $\pi^*_{O_2}$. Notice that given the density function of player 1's final card,

$$\pi^*_{O_2} = \int_0^1 F_1(x)dx \tag{4.1}$$

Given that player 2 is fully rational, optimality requires $F_1\left({\kappa^2_{O_2}}^*\right) = \pi^*_{O_2}$, i.e.

$${\kappa^2_{O_2}}^* = F^{-1}(\pi^*_{O_2}) \tag{4.2}$$

Now, we can derive ${\kappa^2_{O_2-1}}^*$. Suppose player 2 has used $(O_2 - 2)$ replacement opportunities, and holds card $\widehat{y_{O_2-1}}$. If player 2 decides to stick to card $\widehat{y_{O_2-1}}$, the probability of winning is $F_1\left(\widehat{y_{O_2-1}}\right)$; if player 2 decides to use the $(O_2 - 1)$th replacement opportunity, the probability of winning is $\pi^*_{O_2-1}$. Notice that given the density function of player 1's final card and $\pi^*_{O_2}$

$$\pi^*_{O_2-1} = {\kappa^2_{O_2}}^* * \pi^*_{O_2} + \int_{{\kappa^2_{O_2}}^*}^{1} F_1(x)dx \tag{4.3}$$

The first term in (4.3) is the effect of using the O_2th replacement opportunity, the second term is the effect of sticking to the current card.

So, optimality requires $F_1({\kappa^2_{O_2-1}}^*) = \pi^*_{O_2-1}$, i.e.

$${\kappa^2_{O_2-1}}^* = F^{-1}(\pi^*_{O_2-1}) \tag{4.4}$$

Using the same method, we can get that,

$$\begin{aligned} {\kappa^2_{O_2-2}}^* &= F^{-1}(\pi^*_{O_2-2}) \\ &\vdots \\ \kappa^{2*}_1 &= F^{-1}(\pi^*_1) \end{aligned}$$

Player 2's ex-ante probability of winning (before knowing $\widehat{y_0}$) if she plays optimally is,

$$\pi^*_0 = \kappa^{2*}_1 * \pi^*_1 + \int_{\kappa^{2*}_1}^{1} F_1(x)dx \tag{4.5}$$

Player 2's best response function can be determined by finding the first derivatives of her ex-ante probability of winning (π^*_0) w.r.t. $\kappa^{2*}_1, \kappa^{2*}_2, ..., {\kappa^2_{O_2}}^*$ and set them equal to zero.

Using the same procedure, we can derive player 1's best response function.

Given both players' best response functions, we can derive the combination $(\kappa^{1*}_1, \kappa^{1*}_2, ..., {\kappa^1_{O_1}}^*, \kappa^{2*}_1, \kappa^{2*}_2, ..., {\kappa^2_{O_2}}^*)$ of reservation values such that each player's action is a best response to the other player's action. This combination is the Nash equilibrium solution of the card game.

Next, we will derive the Nash equilibrium solution of a card game in which both players are given 1 replacement opportunity as an example.

Suppose that player 1 plays with a cut-off of κ^1_1 (i.e. she sticks iff her first card is $\geq \kappa^1_1$). So, player 1's final card has density $f_1(x)$ and cumulative distribution $F_1(x)$

$$f_1(x) = \begin{cases} \kappa^1_1 & \text{if } x < \kappa^1_1 \\ 1 + \kappa^1_1 & \text{if } \kappa^1_1 \leq x \leq 1 \end{cases}$$

$$F_1(x) = \begin{cases} \kappa^1_1 x & \text{if } x < \kappa^1_1 \\ x + \kappa^1_1 x - \kappa^1_1 & \text{if } \kappa^1_1 \leq x \leq 1 \end{cases} \tag{4.6}$$

Assume player 2 follows a reservation value strategy by choosing a cut-off κ_1^2, i.e. she sticks iff her first card is $\geq \kappa_1^2$. Let $\kappa_1^{2^*}$ be the optimal values of κ_1^2. Notice that given the density function of player 1's final card, player 2's probability of winning if she uses the replacement opportunity is

$$\pi^* = \int_0^1 F_1(x)dx = \frac{1}{2} - \frac{1}{2}\kappa_1^1 + \frac{1}{2}(\kappa_1^1)^2 \tag{4.7}$$

Given that player 2 is fully rational, optimality requires $F_1(\kappa_1^{2^*}) = \pi_1^*$. As player 1's final card has cumulative distribution $F_1(x)$, we know that

$$F_1(\kappa_1^{2^*}) = \begin{cases} \kappa_1^1\kappa_1^{2^*} & \text{if } \kappa_1^{2^*} < \kappa_1^1 \\ \kappa_1^{2^*} + \kappa_1^1\kappa_1^{2^*} - \kappa_1^1 & \text{if } \kappa_1^1 \leq \kappa_1^{2^*} \leq 1 \end{cases} \tag{4.8}$$

Player 2's ex-ante probability of winning (before knowing $\widehat{y_0}$, i.e. the value of the first card) if she plays optimally is

$$\begin{aligned} \pi_0^* &= \kappa_1^{2^*}\pi_1^* + \int_{\kappa_1^{2^*}}^1 F_1(x)dx \\ &= \begin{cases} \frac{1}{2} - \frac{1}{2}\kappa_1^1 + \frac{1}{2}(\kappa_1^1)^2 + \frac{1}{2}\kappa_1^{2^*} - \frac{1}{2}\kappa_1^1\kappa_1^{2^*} + \frac{1}{2}(\kappa_1^1)^2\kappa_1^{2^*} - \frac{1}{2}\kappa_1^1(\kappa_1^{2^*})^2 & \text{if } \kappa_1^{2^*} < \kappa_1^1 \\ \frac{1}{2} - \frac{1}{2}\kappa_1^1 + \frac{1}{2}\kappa_1^{2^*} + \frac{1}{2}\kappa_1^1\kappa_1^{2^*} + \frac{1}{2}(\kappa_1^1)^2\kappa_1^{2^*} - \frac{1}{2}(\kappa_1^{2^*})^2 - \frac{1}{2}\kappa_1^1(\kappa_1^{2^*})^2 & \text{if } \kappa_1^1 \leq \kappa_1^{2^*} \leq 1 \end{cases} \end{aligned} \tag{4.9}$$

Player 2's best response function can be determined by differentiating her ex-ante probability of winning (π_0^*) w.r.t. $\kappa_1^{2^*}$

$$\frac{\partial \pi_0^*}{\partial \kappa_1^{2^*}} = f_1(x) = \begin{cases} \frac{1}{2} - \frac{1}{2}\kappa_1^1 + \frac{1}{2}(\kappa_1^1)^2 - \kappa_1^1\kappa_1^{2^*} & \text{if } \kappa_1^{2^*} < \kappa_1^1 \\ \frac{1}{2} + \frac{1}{2}\kappa_1^1 + \frac{1}{2}(\kappa_1^1)^2 - \kappa_1^{2^*} - \kappa_1^1\kappa_1^{2^*} & \text{if } \kappa_1^1 \leq \kappa_1^{2^*} \leq 1 \end{cases} \tag{4.10}$$

By setting equation 4.10 equal to zero, we get player 2's best response function

$$\kappa_1^{2^*} = f_1(x) = \begin{cases} \frac{1}{2\kappa_1^1} - \frac{1}{2} + \frac{\kappa_1^1}{2} & \text{if } \kappa_1^{2^*} < \kappa_1^1 \\ \frac{1+\kappa_1^1+(\kappa_1^1)^2}{2(1+\kappa_1^1)} & \text{if } \kappa_1^1 \leq \kappa_1^{2^*} \leq 1 \end{cases} \tag{4.11}$$

Given that the game is symmetric between player 1 and player 2, player 1's best response function is

$$\kappa_1^{1^*} = f_2(x) = \begin{cases} \frac{1}{2\kappa_1^2} - \frac{1}{2} + \frac{\kappa_1^2}{2} & \text{if } \kappa_1^{1^*} < \kappa_1^2 \\ \frac{1+\kappa_1^2+(\kappa_1^2)^2}{2(1+\kappa_1^2)} & \text{if } \kappa_1^2 \leq \kappa_1^{1^*} \leq 1 \end{cases} \tag{4.12}$$

By solving these two best response functions, we get that there is a unique pure strategy Nash equilibrium solution $(\kappa_1^{1*}, \kappa_1^{2*}) = (\frac{\sqrt{5}}{2} - \frac{1}{2}, \frac{\sqrt{5}}{2} - \frac{1}{2})$. It means that both players should hold the cut-off $\frac{\sqrt{5}}{2} - \frac{1}{2}$, i.e. they should use the replacement opportunity if the value of their first card is below $\frac{\sqrt{5}}{2} - \frac{1}{2}$, and stick if the value of the first card is above $\frac{\sqrt{5}}{2} - \frac{1}{2}$. At the equilibrium, each player's ex ante probability of winning is equal to 0.5.

4.3.3 General features of the subgame perfect Nash equilibrium solutions

Using the same method, we derive the subgame perfect Nash equilibrium solutions for games in which either player has up to three replacement opportunities (including cases that both players have the same/different number of replacement opportunities).

The optimal cut-off is a function of three variables: *replacement opportunity (x)*, *total number of replacement opportunities (y)*, and *number of replacement opportunities for co-participant (z)*. A change of any of these three variables would change the cut-off. The solutions are shown in Table 4.1, in which $\kappa(x, y, z)$ represents the optimal cut-off given the values of x, y, and z. The probability of winning of a subject depends on the total number of replacement opportunities and the number of replacement opportunities for the co-participant, i.e. variable y and variable z. Table 4.2 shows the probability of winning for each player in each game, in which $p(y, z)$ represents the probability of winning given the values of y and z.

Table 4.1: Subgame perfect Nash equilibrium solutions

$\kappa(1,1,0) = 0.500$	$\kappa(1,1,1) = 0.618$	$\kappa(1,1,2) = 0.652$	$\kappa(1,1,3) = 0.682$
$\kappa(1,2,0) = 0.625$	$\kappa(1,2,1) = 0.698$	$\kappa(1,2,2) = 0.743$	$\kappa(1,2,3) = 0.763$
$\kappa(1,3,0) = 0.695$	$\kappa(1,3,1) = 0.749$	$\kappa(1,3,2) = 0.781$	$\kappa(1,3,3) = 0.827$
$\kappa(2,2,0) = 0.500$	$\kappa(2,2,1) = 0.593$	$\kappa(2,2,2) = 0.657$	$\kappa(2,2,3) = 0.676$
$\kappa(2,3,0) = 0.625$	$\kappa(2,3,1) = 0.705$	$\kappa(2,3,2) = 0.739$	$\kappa(2,3,3) = 0.762$
$\kappa(3,3,0) = 0.500$	$\kappa(3,3,1) = 0.574$	$\kappa(3,3,2) = 0.632$	$\kappa(3,3,3) = 0.671$

Table 4.2: Probability of winning

$p(1,0) = 0.625$	$p(2,0) = 0.695$	$p(3,0) = 0.742$
$p(1,1) = 0.500$	$p(2,1) = 0.577$	$p(3,1) = 0.630$
	$p(2,2) = 0.500$	$p(3,2) = 0.555$
		$p(3,3) = 0.500$

Table 4.1 and Table 4.2 indicate some general features of the subgame perfect Nash equilibrium solutions and the probabilities of winning for different games. First, we notice that given the total number of replacement opportunities of the subject, for a given replacement opportunity

(e.g. first, second, third for a player with three opportunities), the cut-off increases with the number of replacement opportunities for the co-participant. Second, given the number of replacement opportunities for the co-participant, for a given replacement opportunity, the cut-off increases with the total number of replacement opportunities. Third, given the total number of replacement opportunities of the subject and the number of replacement opportunities for the co-participant, the cut-off decreases as 'replacement opportunity' increases. The last feature is that given the number of replacement opportunities of the subject's co-participant, the more replacement opportunities she is given, the higher the probability of winning she has. For example, given that the subject's co-participant has no replacement opportunity, her probability of winning increases from 0.5 to 0.742 when the number of her replacement opportunity increases from 0 to 3.

4.4 The experimental design and implementation

4.4.1 Overall structure of experiment

The experiment contains two treatments. In the first treatment, subjects with 3 replacement opportunities competed with subjects with 3 replacement opportunities. In the second treatment, subjects with 3 replacement opportunities competed with subjects with 1 replacement opportunity. The same two subjects played a series of games. The first subject to win 4 games is the 'winner' of the sequence. This experimental setting was used for the following reasons (see Chapter 3 for details). First, we wanted to generate given inequality by fair or unfair rules. Having subjects with 3 replacement opportunities playing against subjects with 3 replacement opportunities makes the rule of the game fair, while having subjects with 3 replacement opportunities playing against subjects with 1 replacement opportunity makes the rule of the game unfair. Second, we needed an overall winner for the competition. Using best of seven series is a way to generate the overall winner. This feature of the experiment does not affect the subgame perfect Nash equilibrium solutions derived in Section 4.3.3. These solutions were derived for a single game on the assumption that each player tries to maximize her probability of winning. For any player in the series of card games, each game has only two outcomes, win or lose. A fully rational player should try to maximize her probability of winning in each card game, irrespective of how many previous games have been played and irrespective of the results of those games.

This experiment was run as a part of the experiment in Chapter 3. The original experiment consisted of two stages, Stage 1 with either a series of card games or an effort task; Stage 2 with a vendetta game. In this paper, we only consider the series of card games.

Each session of the original experiment was randomly assigned to one of the three treatments (Fair Rule, Unfair Rule or Real Effort). At the beginning of each session, subjects were randomly allocated to numbered seats. Subjects with odd seat numbers were assigned to be participant As and subjects with even seat numbers were assigned to be participant Bs. Each participant A was randomly and anonymously matched with a co-participant B. This matching stayed the same throughout the experiment.

At the beginning of each part, each subject received a copy of the instructions for that part; these instructions were read aloud by the experimenter. These instructions are reproduced in Appendix 1, Chapter 3. Each subject then completed a computerised questionnaire which tested her understanding of the tasks. If a subject made a mistake, the computer would show her the correct answer and the relevant part of the instructions. Subjects were invited to ask the experimenter for clarification.

In stage 1, subjects competed for twelve lottery tickets numbered 1 to 12. The results of the series of card game or the effort task determined the initial allocation of these tickets. At the end of Stage 1, the computer picked nine of these numbered tickets at random and assigned them to the winner. The remaining three tickets were assigned to the loser. The vendetta game in Stage 2 gave subjects the opportunities to change the initial distribution of the tickets, but at a cost: for every three tickets that one player took from the other, she received only one, the other two being wasted.[4]

At the end of each session, the experimenter put 12 numbered tickets into a bag. One of the participants was asked to come forward and pick one ticket from the bag. The number on this ticket was the number of the winning ticket. The subject who held the winning ticket got the prize of £24. If the winning ticket was wasted during the vendetta game, then neither subjects for the pairing got the prize. In all cases, both of them also received a participation fee of £3.

4.4.2 The rules of the card game

In both the Fair Rule treatment and the Unfair Rule treatment, participants played card games.

At the start of each card game, participants were dealt a card each. Each card had a number of points, which could be any of the whole numbers in the range from 1 to 100. Each of these numbers was equally likely at each 'deal'.

In each game, both participant As and participant Bs were offered some opportunities to replace cards. The numbers of replacement opportunities that participant As and participant Bs had varied between treatments. In the Fair Rule treatment, both participant As and participant Bs were given 3 replacement opportunities. In the Unfair Rule treatment, participant As had 1 replacement opportunity and participant Bs had 3 replacement opportunities. During the game, participants could decide to stick with the card that they had been dealt or to replace it with a new card. If the participant decided to replace her card with a new card,

[4] The detail of the vendetta game is described in Chapter 3.

then the computer would randomly draw a new one for her. Each number in the range from 1 to 100 was still equally likely at this 'deal'. However, the participant could not go back to the replaced card again. Participants could decide to stick with any card that was dealt to them. Once they used up all the replacement opportunities, they could not make any further replacement and had to stick with the last card that they had been dealt.

After both participants had made the decision of sticking with a card they had been dealt or had used up all their opportunities for replacing cards, their cards were compared. They both could observe the points on their cards and the points on their co-participants' cards. They were also shown how many replacement opportunities their co-participants had used.

At the end of the card game, the participant who held a card with a higher number of points than the card held by her co-participant won the game. If both of them had the same number of points, the game was a draw.

4.4.3 Subgame perfect equilibrium solution of the card game

In section 4.3.3, we derived the subgame perfect Nash equilibrium solutions for the game in which both players are given 3 replacement opportunities and for the game in which one player has 3 replacement opportunities and the other player has 1 replacement opportunity, when the number of points on the card is drawn form a uniform distribution over the interval $[0, 1]$. In the card games in our experiment, the number of points on the card could only be the whole numbers in the range from 1 to 100. Therefore, we adjust the subgame perfect Nash equilibrium solutions to reflect this.

According to the subgame perfect Nash equilibrium solution, subjects in the Fair Rule treatment should hold cut-offs 83, 76, and 67, i.e. use the first replacement opportunity if the number of points on first card is below 83, use the second replacement opportunity if the number of points on the second card is below 76, and use the third replacement opportunity if the number of points on the third card is below 67. If these cut-offs are used, both players have the same probability of winning, i.e. 0.5.

In the Unfair Rule treatment, the one replacement opportunity player should hold a cut-off 68 and the cut-offs for the three replacement opportunities player should be 75, 71, and 57 for the first, second and third card respectively. The probability of winning for the three replacement opportunities player and the one replacement opportunity player are 0.63 and 0.37.

4.4.4 Implementation

The experiment was conducted between November 2015 and January 2016 at the CBESS Experimental Laboratory at the University of East Anglia. Participants were recruited from the general student population via the CBESS online recruitment system (Bock *et al.*, 2014). The experiment was programmed and conducted with the experimental software z-Tree (Fischbacher, 2007). We ran 18 sessions in total: four for the Real Effort treatment, six for the Fair Rule treatment and eight for the Unfair Rule treatment. Subjects were not allowed to participate in more than one session. A total of 326 subjects participated in the experiment, of whom 72 were in the Real Effort treatment, 104 in the Fair Rule treatment, and 150 in the Unfair Rule treatment. 123 of the subjects were male and 196 were female.[5] Most of the participants were students from a wide range of academic disciplines and with an age range from 18 to 63. Each session lasted approximately 50 minutes. Average earnings were £10.67 per person, including a show-up fee of £3.00. The lowest earning was £3.00, the highest was £27.00.

4.5 The Econometric Model

In this section, we introduce the econometric model that simultaneously estimates subjects' cut-offs and variables that we believe may influence their choices of cut-offs.

4.5.1 General Setup

Each subject i plays t games, $i \in \{1, ..., n\}$ and $t \in \{1, ..., T_i\}$. In each game, t, subject i observes up to J_i card *draws* with values $v_{1,t}$ to $v_{J_i,t}$. For subjects with 3 replacement opportunities, J_i equals 4; for subjects with 1 replacement opportunity, J_i equals 2.

If subject i accepts a draw, her participation in the game ends. If subject i gets as far as draw J_i, she must accept draw J_i. Define $a_{i,t}$ to represent draw accepted by subject i in game t: $a_{i,t} = j$ if subject i accepts jth draw in game t, $j = 1, ..., J_i$.

Subject i wins game t ($w_{i,t} = 1$) if her accepted draw is higher than her opponent's accepted draw.

4.5.2 Econometric Model

Consider the decisions of subject i. But for simplicity, subscript i is suppressed in the rest of this section.

In game t, the subject has a 'cut-off' for each draw: $\kappa_{1,t}$ to $\kappa_{J-1,t}$. Conditional on having rejected previous draws (if any), the subject accepts draw j if $v_{j,t} - \kappa_{j,t} + \varepsilon_{j,t} > 0$, Where $\varepsilon_{j,t} \sim N(0, \sigma^2)$, represents computational error.

[5] 7 subjects selected 'prefer not to say' in the gender question.

If $J > 2$, the subject's cut-off declines between draws according to:

$$\kappa_{j,t} = \kappa_{1,t} - d_t(j-1), j = 2, ..., J-1$$

d_t therefore represents the rate of decline of the cut-off. d_t does not enter the model when $J = 2$.

Both cut-offs and the rate of decline of the cut-off are allowed to change with experience (t). A reciprocal specification is adopted here, since this implies convergence to an equilibrium [6] with experience (see Moffatt, 2015, Chapter 4). The specification is:

$$\kappa_{1,t} = \kappa_{1,e} + \tau_\kappa(\frac{1}{t}), t = 1, 2, 3, ...$$

$$d_t = d_e + \tau_d(\frac{1}{t}), t = 1, 2, 3...$$

Note that $\kappa_{1,e}$ and d_e are respectively the equilibrium cut-off and equilibrium rate of change of the cut-off to which subjects converge with experience. τ_κ and τ_d represent the distance away from these respective equilibria with no experience (i.e. when $t = 1$).

Furthermore, the cut-off is assumed to shift in response to the values of previous draws (for subjects with $J = 3$):

$$\kappa_{2,t} = \kappa_{1,t} - d_t + \gamma_1 v_{1,t}$$

$$\kappa_{3,t} = \kappa_{1,t} - 2d_t + \gamma_2 v_{2,t} + \gamma_3 v_{1,t}$$

Finally, we allow the cut-offs to depend on whether the subject won the previous game (w_{t-1}), and also on whether the subject is playing against an opponent who only has one opportunity to replace a card (i.e. an opponent for whom $J = 2$). The latter variable is labelled 'opp1'. Hence the cut-off in draw 1 becomes:

$$\kappa_{j,t} = \kappa_{1,e} + \tau_\kappa \frac{1}{t} + \beta_w w_{t-1} + \beta_{opp1} opp1, t = 1, 2, 3...$$

4.5.3 Construction of Likelihood function

Estimation is by Maximum Simulated Likelihood (MSL).

Recall that $a_t = j$ if subject accepts jth draw in game t, $j = 1, ..., J$

[6] This need not be a Nash equilibrium.

The probability of accepting draw 1 in game t is:

$$P(a_t = 1|\kappa_{1,e}) = \Phi(\frac{v_{1,t} - \kappa_{1,t}}{\sigma})$$

Where $\Phi(.)$ is the standard normal cdf.

If $J = 2$, the probability of accepting draw 2 is:

$$P(a_t = 2|\kappa_{1,e}) = 1 - \Phi(\frac{v_{1,t} - \kappa_{1,t}}{\sigma})$$

If $J = 4$, the probability of accepting draw 2 is:

$$P(a_t = 2|\kappa_{1,e}) = [1 - \Phi(\frac{v_{1,t} - \kappa_{1,t}}{\sigma})]\Phi(\frac{v_{2,t} - \kappa_{2,t}}{\sigma})$$

The probability of accepting draw 3 is:

$$P(a_t = 3|\kappa_{1,e}) = [1 - \Phi(\frac{v_{1,t} - \kappa_{1,t}}{\sigma})][1 - \Phi(\frac{v_{2,t} - \kappa_{2,t}}{\sigma})]\Phi(\frac{v_{3,t} - \kappa_{3,t}}{\sigma})$$

The probability of accepting draw 4 (note that this is just the probability of rejecting all of the first three draws) is:

$$P(a_t = 4|\kappa_{1,e}) = [1 - \Phi(\frac{v_{1,t} - \kappa_{1,t}}{\sigma})][1 - \Phi(\frac{v_{2,t} - \kappa_{2,t}}{\sigma})][1 - \Phi(\frac{v_{3,t} - \kappa_{3,t}}{\sigma})]$$

Between-subject heterogeneity

The cut-off for the first draw with experience ($\kappa_{1,e}$) is assumed to vary between subjects, according to:

$$\kappa_{1,e} \sim N(\mu_\kappa, \eta_\kappa^2)$$

When $J = 4$, there are eleven parameters to estimate: $\mu_\kappa, \eta_\kappa, d_e, \sigma, \gamma_1, \gamma_2, \gamma_3, \tau_\kappa, \tau_d, \beta_w, \beta_{opp1}$.

When $J = 2$, there are five parameters to estimate: $\mu_\kappa, \eta_\kappa, \sigma, \tau_\kappa, \beta_w$

Likelihood function

The per-subject (subject i) likelihood contribution is given by:

$$L_i = \int_{-\infty}^{\infty} \left[\prod_{t=1}^{T} P(a_{t,i}|\kappa_{1,e}) \right] f(\kappa_{1,e}; \mu_\kappa, \eta_\kappa) d\kappa_{1,e}$$

where $P(a_{t,i}|\kappa_{1,e})$ is the (conditional) probability of subject i's observed decision in game t (given by one of the formulae above), and $f(\kappa_{1,e};\mu_\kappa,\eta_\kappa)$ is the normal density function evaluated at $\kappa_{1,e}$.

The integral appearing in (.) is evaluated using Halton draws. The integral is converted into the sum:

$$\widehat{L_i} = \frac{1}{R}\sum_{r=1}^{R}\prod_{t=1}^{T} P(a_t|\kappa_{1,e,r})$$

Where $\kappa_{1,e,r}$, $r = 1, ..., R$, are the transformed Halton draws for subject i.[7] The draws are transformed to represent draws from a $N(\mu_\kappa, \eta_\kappa^2)$ distribution. The draws are fixed over t for a given i.

Finally, the sample log-likelihood function is obtained as:

$$LogL = \sum_{i=1}^{n}\widehat{L_i}$$

The LogL is maximised with respect to all parameters of the model in order to obtain MLEs.

For further details of the MSL technique see Chapter 10 of Moffatt (2015).

4.6 Results

In this section, we estimate the model constructed in the previous section, using the method of maximum simulated likelihood. Our sample consist of 254 subjects, of whom 104 are in the Fair Rule treatment and 150 are in the Unfair Rule treatment. On average, each subject in the Fair Rule treatment played 5.81 rounds of card games and each subject in the Unfair Rule treatment played 5.76 rounds.

Table 4.3 shows the separate results of the model for players with 3 replacement opportunities and players with 1 replacement opportunity. Remember that μ_κ and η_κ represent the mean and the standard deviation of the equilibrium cut-offs for the first card, that is, the cut-off to which they converge with experience. To find the cut-offs of subjects with no experience, we need to find $\mu_\kappa + \tau_\kappa$.

[7] $\kappa_{1,e}$ is the parameter that is assumed to vary between subjects. The method of MSL involves using simulations of this parameter in order to evaluate $\widehat{L_i}$. Halton draws are used as the simulated values. Halton draws are uniformly distributed. They need to be transformed to normality so that they represent the variation in $\kappa_{1,e}$.

Table 4.3: Estimates from the model

	3 opportunities player (J=4)	1 opportunity player (J=2)
μ_κ	68.55(1.86)	60.13(2.83)
η_κ	9.72(0.98)$***$	10.37(1.75)$***$
d_e	8.06(2.01)$***$	
σ	13.12(0.65)$***$	15.00(1.57)$***$
γ_1	0.08(0.04)$*$	
γ_2	0.30(0.07)$***$	
γ_3	$-0.09(0.07)$	
τ_κ	$-18.44(2.88)***$	$-15.09(4.87)***$
τ_d	0.62(2.86)	
β_{opp1}	3.13(1.94)	
β_w	$-1.64(0.72)**$	$-1.97(1.67)$
n	179	75
T (mean of)	5.79	5.76
$LogL$	-549.76	-136.86

Notes: The estimation technique used is maximum simulated likelihood. Standard errors are given within parentheses. The first column presents estimates for players with three replacement opportunities; the second column presents estimates for players with one replacement opportunity. * p<0.1; ** p<0.05; ***p<0.01. It is only logical to test whether a parameter equals zero if that is an interesting hypothesis. We know that μ_κ is non-zero. It makes no sense to apply stars to μ_κ.

4.6.1 Actual cut-offs versus subgame perfect Nash equilibrium predictions

We first analyse the extent to which subjects make decisions consistent with subgame perfect Nash equilibrium solutions. Table 4.4 lists the posterior estimates of the actual cut-offs for subjects without experience and the equilibrium predictions.

Table 4.4: Summary of actual play without experience and equilibrium predictions

	3 opportunities player				1 opportunity player	
	Against 3 opportunities opponent		Against 1 opportunity opponent			
	Actual	Predicted	Actual	Predicted	Actual	Predicted
Card 1	50.11	83	53.24	75	45.04	68
Card 2	42.05	76	45.18	71		
Card 3	33.99	67	37.12	57		

In our experiment, there are two types of 3 replacement opportunities players: players who play against players with 3 replacement opportunities in the Fair Rule treatment and players who play against players with only 1 replacement opportunity in the Unfair Rule treatment. The estimates show that the cut-off for the first card chosen by an average subject with 3 replacement opportunities who plays against a 3 replacement opportunities player is 50.11 in the first game. The cut-off for the first card chosen by an average subject with 3 replacement opportunities who plays against a 1 replacement opportunity player is 53.24 in the first game. Within the game, the average rate of decline of the cut-off is 8.06, which is the same for both types of subjects. Therefore, the cut-offs chosen by an average subject with 3 replacement opportunities who plays against a 3 replacement opportunities player are 50.11, 42.05, and 33.99 in her first game. The cut-offs chosen by an average subject with 3 replacement opportunities who plays against a 1 replacement opportunity player are 53.24, 45.18, and 37.12 in her first game. Figure 4.1 and Figure 4.2 display the distributions of posterior estimates of cut-offs chosen by subjects with 3 replacement opportunities in their first card games. The vertical line in each histogram indicates the subgame perfect Nash equilibrium prediction. It is striking to see that, all the cut-offs chosen by both types of subjects with 3 replacement opportunities in their first games are below subgame perfect Nash equilibrium predictions.

Figure 4.1: Posterior estimates of cut-offs - 3 vs 3 without experience

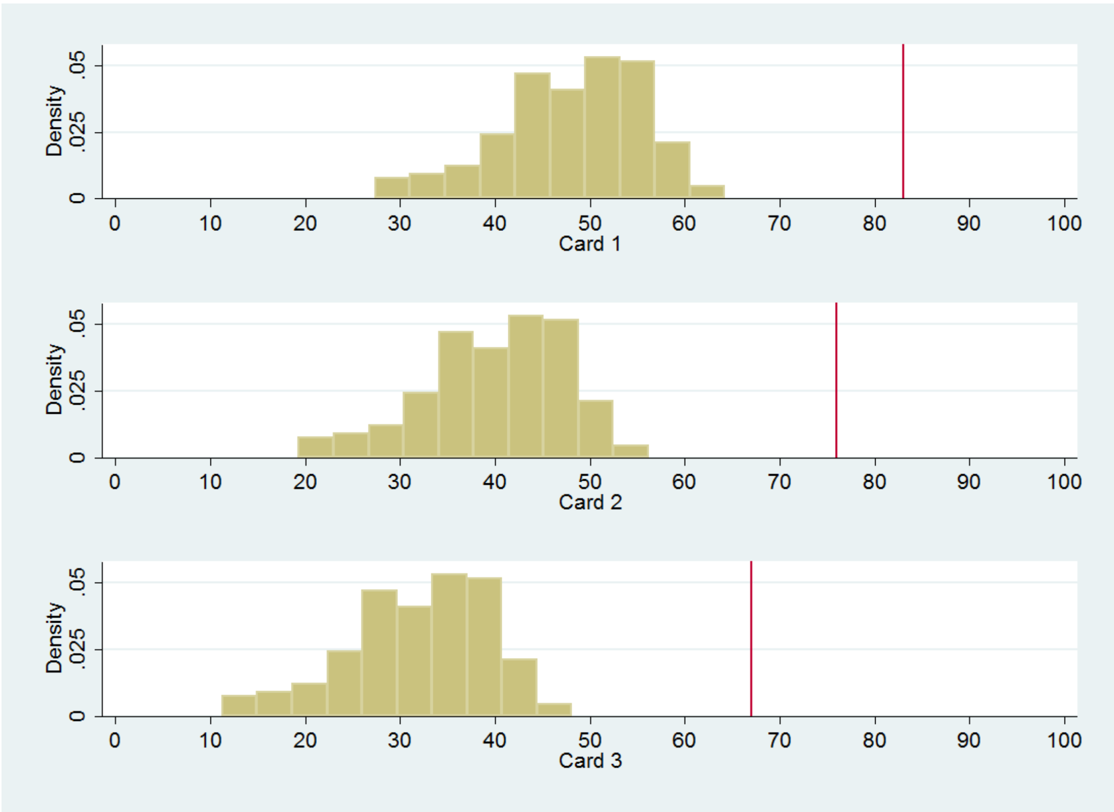

Notes: The histograms display the distributions of posterior estimates of cut-offs chosen by 3 replacement opportunities subjects who play against 3 replacement opportunities players. The vertical lines indicate subgame perfect Nash equilibrium predictions.

Figure 4.2: Posterior estimates of cut-offs - 3 vs 1 without experience

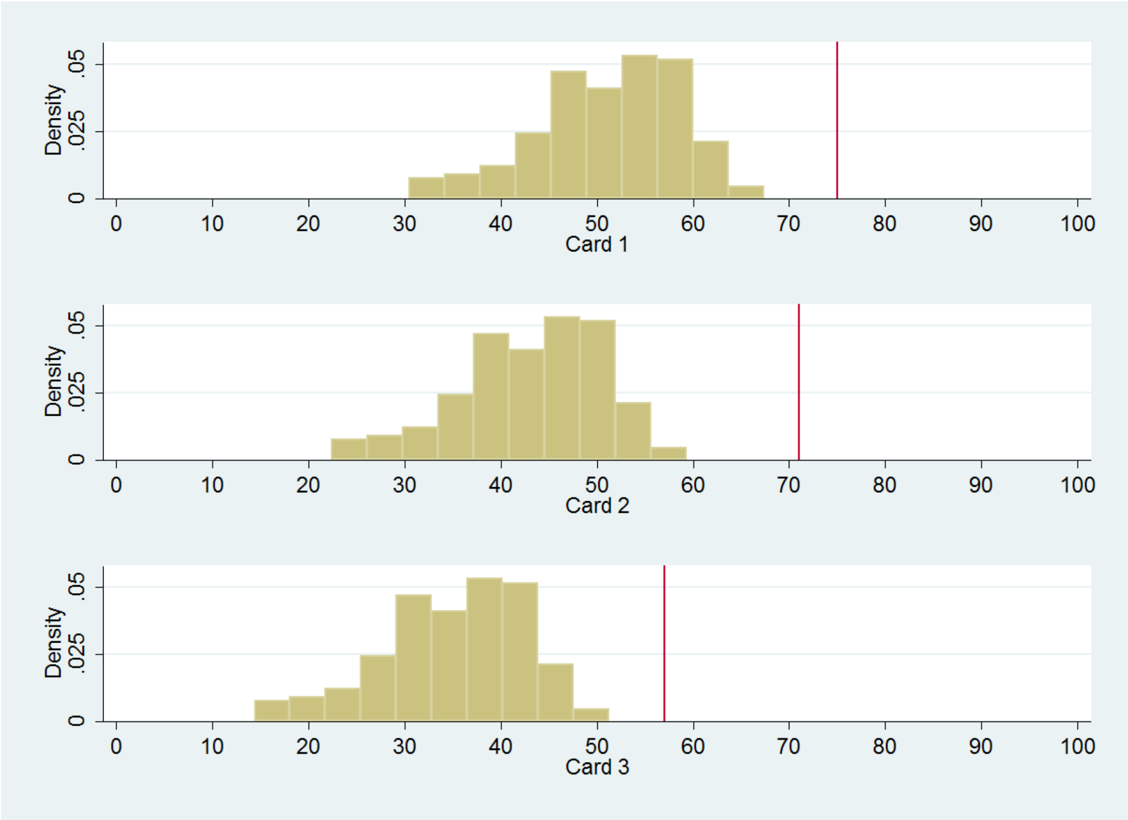

Notes: The histograms display the distributions of posterior estimates of cut-offs chosen by 3 replacement opportunities subjects who play against 1 replacement opportunity players. The vertical lines indicate subgame perfect Nash equilibrium predictions.

An average subject with 1 replacement opportunity chose a cut-off of 45.04 for the first card in the first game. Figure 4.3 shows the distribution of posterior estimates of cut-offs chosen by subjects with 1 replacement opportunity in their first games. It is clear that none of the cut-offs chosen by subjects with 1 replacement opportunity is consistent with or above the subgame perfect Nash equilibrium prediction.

Figure 4.3: Posterior estimates of cut-offs - 1 vs 3 without experience

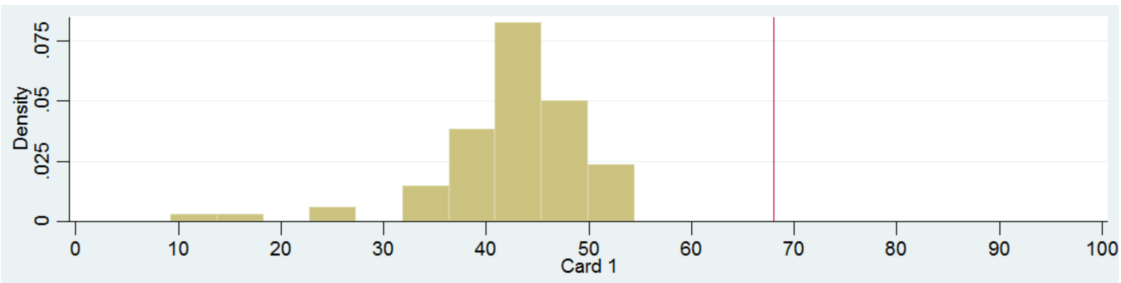

Notes: The histogram displays the distribution of posterior estimates of the cut-offs chosen by 1 replacement opportunity subjects. The vertical line indicates the subgame perfect Nash equilibrium prediction.

Result 1. The actual cut-offs chosen by subjects without experience are lower than subgame perfect Nash equilibrium predictions.

According to the subgame perfect Nash equilibrium prediction, the cut-offs for 3 replacement opportunities players who play against 3 replacement opportunities players should be higher than the cut-offs for 3 replacement opportunities players who play against 1 replacement opportunity players. Our estimation results show that, contrary this prediction, playing against an opponent with 1 replacement opportunity has a small positive effect on cut-offs chosen by subjects, but the effect is not significant.

Result 2. For subjects with 3 replacement opportunities, cut-offs are not significantly affected by the number of replacement opportunities available to their opponents.

Result 3. Although the actual cut-offs chosen by subjects do not match with the subgame perfect Nash equilibrium prediction, the decline of subjects' cut-offs between cards is significant and in the same direction as predicted.

When a player with three replacement opportunities makes decisions on the rational strategy for the card game, it involves two levels of reasoning. The question that arises at the first level of reasoning is what strategy she should choose when playing against an opponent with a given strategy, implying a given distribution of the value of the opponent's final card. The intuition is that, for any such distribution, the player should decrease cut-offs between cards. The question that arises at the second level of reasoning is what strategy the opponent will choose, and therefore what the distribution of the opponent's final card will be. The intuition is that the player should adjust her cut-offs based on the number of replacement opportunities held by her opponent, i.e. the more replacement opportunities held by her opponent, the higher cut-offs she should hold. From Result 2 and Result 3, we can see that subjects are able to conduct the first level of reasoning, but their reasoning is not deep enough for them to realize that cut-offs should vary with the number of replacement opportunities held by opponents.

Moreover, from Figure 4.1 to Figure 4.3 we can see that, instead of following a unique pure strategy subgame perfect Nash equilibrium, subjects' cut-offs are distributed over a wide range. For example, the cut-offs chosen by subjects with 3 replacement opportunities who play against 3 replacement opportunities players should be concentrated at 83, but instead they range from 27.40 to 64.26. Similar behaviour is observed of subjects with 3 replacement opportunities who play against 1 replacement opportunity players. For subjects with 1 replacement opportunity, their cut-offs for the first card should be 68, but the actual cut-offs instead cover a wide range from 9.27 to 67.39. The estimates show that η_κ is significantly greater than zero in both models, implying heterogeneity in the cut-offs .

Result 4. The actual cut-offs are distributed over a wide range, instead of concentrating at the equilibrium value.

4.6.2 Feedback information

In our experiment, subjects were provided with the same feedback information in both the Fair Rule treatment and the Unfair Rule treatment. At the end of each card game, subjects saw the card held by their opponents and learned whether they had won or lost that game. In this section we focus on the impact of feedback on subjects' performance in the card games.

In the previous section, we showed that at the beginning of the series of card game subjects chose cut-offs significantly lower than predicted. Thus, one might expect that cut-offs would increase as subjects gained experience. Table 4.5 shows estimated actual cut-offs for *experienced* players – that is, the estimated values to which cut-offs converge as t increases. The results of our model show that as subjects become more experienced, cut-offs increase. Note that τ_d, which measures the difference between the cut-off decline rates for subjects with and without experience, is not significant, which suggests that the rate of decline of cut-offs is stable with experience.[8]

Table 4.5: Summary of actual play with experience, equilibrium predictions, and empirical best responses

	3 opportunities player						1 opportunity player		
	Against 3 opportunities opponent			Against 1 opportunity opponent					
	Actual	Predicted	Empirical	Actual	Predicted	Empirical	Actual	Predicted	Empirical
Card 1	68.55	83	76	71.68	75	74	60.13	68	67
Card 2	60.49	76	72	63.62	71	69			
Card 3	52.43	67	67	55.56	57	61			

Our model implies that, for an average subject with 3 replacement opportunities, each cut-off is 18.44 higher when the subject is experienced than in the first game. For an average subject with 1 replacement opportunity, the cut-off is 15.09 higher when the subject is experienced than in the first game. Figure 4.4 and Figure 4.5 display the distributions of posterior estimates of cut-offs chosen by 3 replacement opportunities subjects with experience. In Figure 4.4 and Figure 4.5, the solid vertical lines indicate the subgame perfect Nash equilibrium predictions. (The dashed vertical lines in these figures, and the data in the 'Empirical' columns of Table 4.5, will be explained in Section 4.6.3)

[8] As the coefficient of τ_d is relatively small and insignificant, we do not take it into account when we estimate actual cut-offs.

Figure 4.4: Posterior estimates of cut-offs - 3 vs 3 with experience

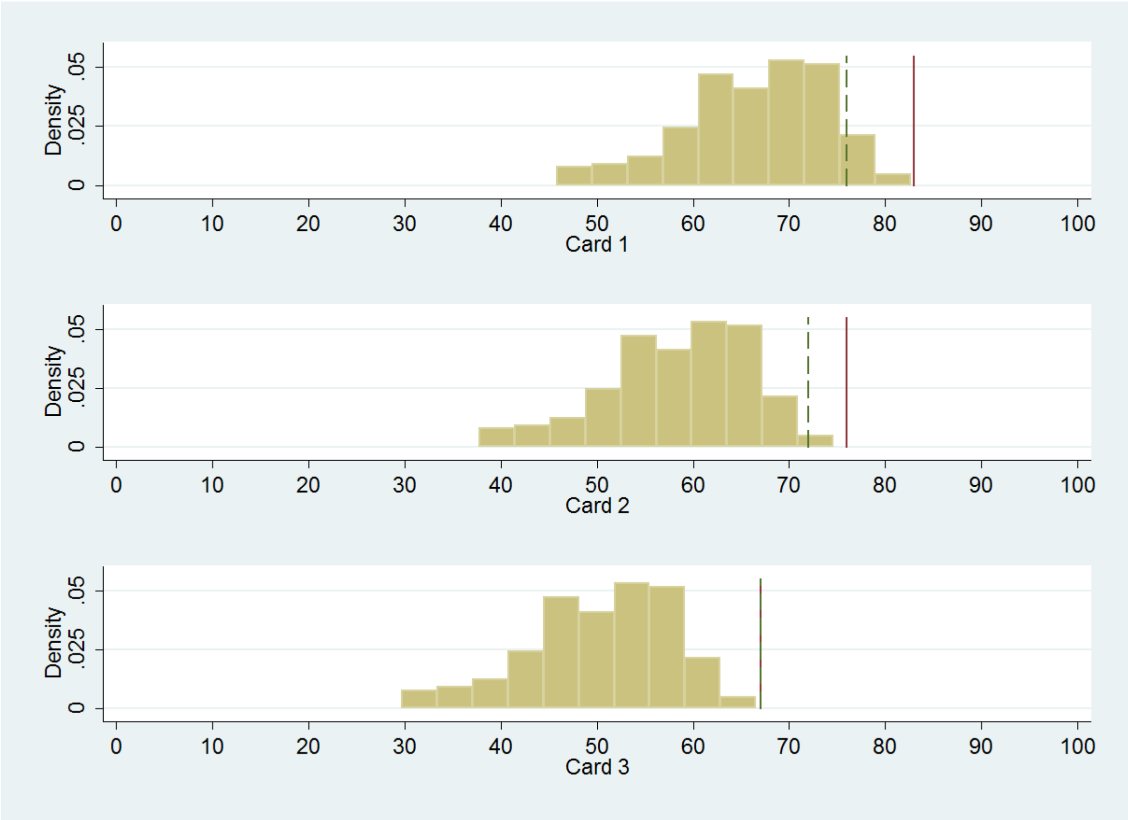

Notes: The histograms display the distributions of posterior estimates of cut-offs chosen by 3 replacement opportunities subjects who play against 3 replacement opportunities players. The solid vertical lines indicate subgame perfect Nash equilibrium predictions. The dashed vertical lines indicate empirical best responses.

Figure 4.5: Posterior estimates of cut-offs - 3 vs 1 with experience

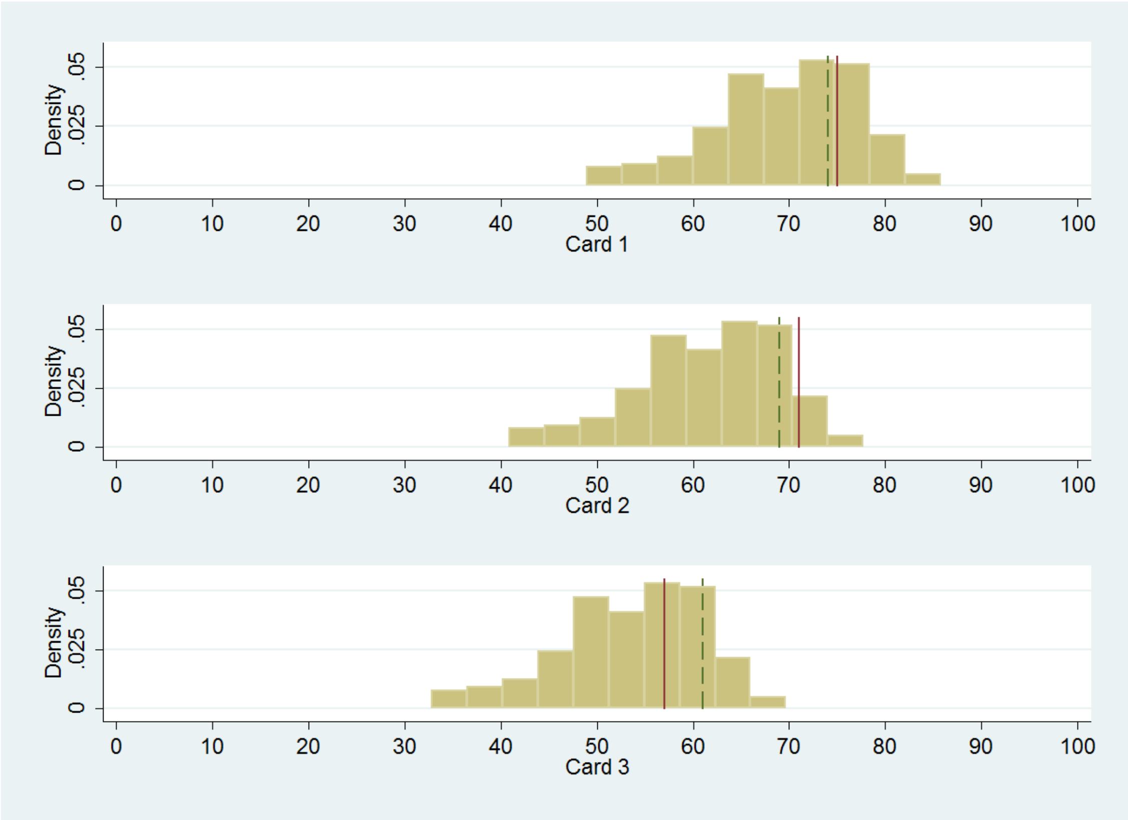

Notes: The histograms display the distributions of posterior estimates of cut-offs chosen by 3 replacement opportunities subjects who play against 1 replacement opportunity players. The solid vertical lines indicate subgame perfect Nash equilibrium predictions. The dashed vertical lines indicate empirical best responses.

It appears that subjects learn to increase their cut-offs over the course of the experiment. With experience, all of the cut-offs chosen by subjects with 3 replacement opportunities who play against 3 replacement opportunities players remain lower than the subgame perfect Nash equilibrium predictions. For subjects with 3 replacement opportunities who play against 1 replacement opportunity players, 27.93 % of their cut-offs for the first card are higher than the equilibrium prediction, 7.26% of their cut-offs for the second card are higher than the equilibrium prediction, and 40.22% of their cut-offs for the third card are higher than the equilibrium prediction

Figure 4.6 shows the distribution of posterior estimates of cut-offs chosen by subjects with 1 replacement opportunity after they became experienced with the card game. From the histogram, we can see that with experience, almost all of the cut-offs chosen by subjects with 1 replacement opportunity who play against 3 replacement opportunities players remain lower than the subgame perfect Nash equilibrium prediction. Only 1.33% of the cut-offs are higher than the equilibrium prediction.

Figure 4.6: Posterior estimates of cut-offs - 1 vs 3 with experience

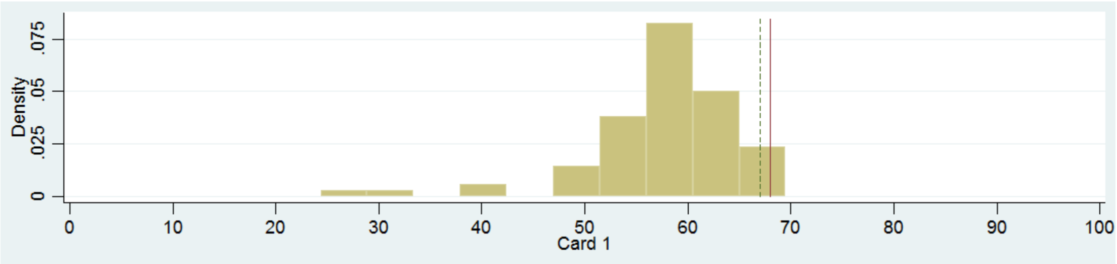

Notes: The histogram displays the distribution of posterior estimates of cut-offs chosen by 1 replacement opportunity subjects. The solid vertical line indicates the subgame perfect Nash equilibrium prediction. The dashed vertical line indicates the empirical best response.

Result 5. Subjects learn to increase their cut-offs over time, but for experienced subjects, most cut-offs still remain lower than predicted.

According to the theoretical prediction, subjects should use the strategy of maximizing the winning probability in each card game. Therefore, the results of the preceding game should have no impact on the cut-offs chosen by subjects in the later games. However, it may be possible that the cut-offs chosen by subjects in our experiment were influenced by the outcomes of previous games. In our model, we include variable β_w, which takes into account the influence of the outcome of the previous game on cut-offs chosen by subjects in the later games. It turns out that for subjects with 3 replacement opportunities, β_w is positive and statistically significant (p-value< 0.05). If the previous game was a win, on average subjects with 3 replacement opportunities lower their cut-offs in the following game by 1.64. For subjects with 1 replacement opportunity, β_w is positive but not significant.

Result 6. Subjects' decisions are influenced by the outcome of the preceding game.

4.6.3 Actual cut-offs versus empirical best responses

In previous sections, we have shown that although there is an increasing trend, the average cut-offs chosen by subjects remain lower than the subgame perfect equilibrium prediction. Next, we examine to what degree the players learn to choose the empirical best response. We define the empirical best response as the strategy that is ex post optimal given the strategy played by an average player in the game. A similar concept is used by Weizsäcker (2010), who defines an *empirically optimal action*, the action that is ex post optimal in the majority of cases under identical conditions. Weizsäcker suggests that empirically optimal action is a better measure of the success of social learning than Perfect Bayesian Equilibrium or Quantal Response Equilibrium, as it reflects the true behaviour of other players.

Given the cut-offs of an average player with experience presented in Section 4.6.2, we calculate the empirical best responses to these cut-offs. The empirical best responses are shown in

Table 4.5. For players with 3 replacement opportunities who play against players of the same type, the empirical best response is to choose cut-offs 76, 72, and 67. For players with 3 replacement opportunities who play against 1 replacement opportunity players, the empirical best response predicts cut-offs 74, 69, and 61. The empirical best response cut-off for players with 1 replacement opportunity is 67.

We measure how well subjects learn to play the empirical best responses by comparing the cut-offs estimated for experienced subjects with the predictions of the empirical best response. Figure 4.4, 4.5, and 4.6 show the distributions of posterior estimates of cut-offs chosen by experienced subjects. The dashed vertical lines indicate the empirical best response cut-offs in each case.

Figure 4.4 and Figure 4.6 clearly show that, for experienced subjects with 3 replacement opportunities playing against players of the same type and for experienced subjects with 1 replacement opportunity, posterior estimates of cut-offs are almost always lower than empirical best responses. Of all players with 3 replacement opportunities who play against 3 replacement opportunities players, 7.26% are estimated as having cut-offs for the first card that are higher than the empirical best response. For the second card, the corresponding percentage is 1.12%. For the third card, all cut-offs are lower than the empirical best response. For subjects with 1 replacement opportunity, only 1.33% have estimated cut-offs higher than the empirical best response. In contrast, the average cut-offs estimated for subjects with 3 replacement opportunities who play against 1 replacement opportunity players are closer to the empirical best responses. For these subjects, 33.52% of the cut-offs for the first card and 15.64% of the cut-offs for the second and third cards are higher than the empirical best responses.

> *Result 7.* For experienced subjects, most cut-offs are lower than the empirical best responses.

4.7 Discussion

Our search competition takes the form of parallel searches without recall. Players compete for a fixed prize by finding an offer that is better than the offer found by the other player(s). Existing research suggests opposing intuitions about actual behaviour, relative to Nash equilibrium play. The experimental literature on contests has generally found that the average effort level is significantly higher than the theoretical prediction. Experimental evidence on search suggests that subjects tend to stop searching too early, compared to the risk neutral benchmark.

In our search competition, we observe a significant tendency for subjects to set lower cut-offs than is optimal. This finding is similar to the general phenomena of under-search reported by experiments on search. In literature on search problem with recall, early termination of search is usually attributed to risk aversion or loss aversion (Schotter and Braustein, 1981; Cox and Oaxaca, 1989; Kogut, 1990; Schunk and Winter, 2009a). However, because our card game has only two possible outcomes, either win or lose, and because we used the binary lottery mechanism as the incentive system in our experiment, the behaviour of a player who acted in accordance with expected utility theory would be independent of her attitude to risk. For the same reasons, the behaviour of a player who acted in accordance with prospect theory would be independent both of her attitudes to gain and loss and of the properties of her probability weighting function.[9]

Similar to the finding in contest experiments that subjects learn to reduce their efforts over the course of the experiment, but their effort levels remain significantly higher than the Nash equilibrium prediction (Davis and Reilly, 1998; Gneezy and Smorodinsky, 2006; Price and Sheremeta, 2011; Brookins and Ryvkin, 2014; Mago *et al.*, 2016), we find that as subjects become more experienced, they tend to choose higher cut-offs, but their cut-offs remain lower than the subgame perfect Nash equilibrium predictions, especially the cut-offs chosen by subjects with 3 replacement opportunities who play against 3 replacement opportunities players and the cut-offs chosen by subjects with 1 replacement opportunity. To get more insight into how successful subjects learn to play the best response of other players' strategies, we compare subjects' cut-offs with the predictions of the empirical best response. We find that although subjects are given full feedback information, cut-offs chosen by subjects with experience are still different from cut-offs predicted by the empirical best response.

We also find that as subjects become experienced with the card game, the average cut-offs chosen by 3 replacement opportunities subjects who play against 1 replacement opportunity players are closer to both subgame perfect Nash equilibrium solutions and empirical best response predictions than the average cut-offs chosen by 3 replacement opportunities subjects who play against 3 replacement opportunities players and the average cut-offs chosen by 1 replacement opportunity subjects. One possible explanation to this finding is that compared to the decision problems faced by subjects with 3 replacement opportunities, the decision problem for subjects who play against 1 replacement opportunity players is relatively easier than the decision problem for subjects who play against 3 replacement opportunities players. Therefore less effort may be required for subjects who play against 1 replacement opportunity players to learn to play optimally.

As subjects in the competition get full feedback information about the outcome of each game, our experimental design allows us to study more about subject's learning process. We find that subjects who win the previous game tend to decrease their cut-offs in later rounds, while subjects who lose the previous game tend to increase their cut-offs in later rounds. Many experimental studies on tournaments and multi-battle contests report similar results. For instance, Mago *et al.* (2016) find that in multi-battle contests with feedback information

[9]However, there is some evidence that, when the binary lottery incentive system is used in experiments, subjects treat probabilities of winning as if they were material consequences (Selten *et al.*, 1999).

on relative effort levels, subjects who put effort lower than the winning effort increase their effort in the consecutive battle, while subjects who put effort higher than the winning effort decrease their effort in the consecutive battle. In contests, given the information about the other players' effort levels, subjects may able to increase their payoffs by reducing effort levels if their effort levels are higher than the winning effort level in the previous round. However, as in our search competition there is no cost of replacing cards, adjusting cut-offs only influences subjects' probability of winning the competition. Subjects got feedback on whether they won the game or not, but not on the probability of winning given the cut-offs that they chose. It is difficult for them to learn what is the best strategy given the limited feedback information. One possible explanation is that subjects tend to play 'safe' psychologically by lowering their cut-offs and therefore changing cards less often after having won a few games, while subjects tend to play 'risky' psychologically by raising their cut-offs and therefore changing cards more often after having lost a few games.

Subjects' reluctance to replace cards, especially after winning a game(s) might be explained by some behavioural biases. Omission bias is one of the possible biases that might influence the card replacing behaviour of subjects. Many psychological studies have shown that when a decision leads to a bad outcome, relative to what might have been, people think that the decision was worse if the outcome resulted from action than if it resulted from inaction (Ritov and Baron, 1992; Spranca *et al.*, 1991; Ritov and Baron, 1992; Baron and Ritov, 1994). For subjects in our search competition, if the game ends up with a loss, they might feel worse if they could have won by not using the replacement opportunity and keeping a card that they were dealt, especially when losing the game means losing the leading position or even the competition.

4.8 Conclusion

We introduce a new type of game, a search competition. The game describes a situation in which players compete for a fixed prize by searching for an offer that is better than the offer found by the other player(s). It combines features of both games in the contest literature and games in the search literature. The search competition differs from games in the contest literature: exerting effort in contests is costly and the cost is irreversible, while exerting effort in the search competition is costless. Compared to games studied in the contest literature, our search competition focuses more on the strategic interaction between players in the competition. We conducted a laboratory experiment to study individual behaviour in the search competition, and compared our results with the main findings in the contest literature and search literature.

We find that in the search competition, subjects set reservation prices too low relative to theoretical predictions, and therefore search is less than optimal. Our results indicate that individual behaviour in the search competition is more consistent with the search behaviour reported in the search literature than with competitive behaviour in contests.

Our experiment results show that subjects' learning behaviour in the search competition is similar to behaviour reported in contest literature. Subjects learn to increase their cut-offs over the course of the experiment, but their reservation prices remain lower than equilibrium predictions. Although feedback information on previous games and opponents' choices should not influence subjects' decisions, subjects in our experiments tend to adjust their reservation prices based on this information in the competition. Subjects' learning behaviour can be explained to some extent by behavioural biases.

References

Amaldoss, W. and Rapoport, A. (2009). Excessive expenditure in two-stage contests: Theory and experimental evidence. *Game Theory: Strategies, Equilibria, and Theorems. Hauppauge, NY: Nova Science Publishers.*

Andreoni, J. and Miller, J. (2002). Giving according to garp: An experimental test of the consistency of preferences for altruism. *Econometrica*, **70** (2), 737–753.

Armstrong, M. and Zhou, J. (2016). Search deterrence. *Review of Economic Studies*, **83** (1), 26–57.

Balafoutas, L., Kerschbamer, R. and Sutter, M. (2012). Distributional preferences and competitive behavior. *Journal of Economic Behavior & Organization*, **83-334** (1), 125–135.

Baron, J. and Ritov, I. (1994). Reference points and omission bias. *Organizational Behavior and Human Decision Processes*, **59** (3), 475–498.

Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, **30** (5), 961–981.

Ben Zur, H. and Breznitz, S. J. (1981). The effect of time pressure on risky choice behavior. *Acta Psychologica*, **47** (2), 89–104.

Bicchieri, C. (2008). The fragility of fairness: An experimental investigation on the conditional status of pro-social norms 1. *Philosophical Issues*, **18** (1), 229–248.

— and Chavez, A. (2010). Behaving as expected: Public information and fairness norms. *Journal of Behavioral Decision Making*, **23** (2), 161–178.

Blount, S. (1995). When social outcomes aren′t fair: The effect of causal attributions on preferences. *Organizational Behavior and Human Decision Processes*, **63** (2), 131–144.

Bock, O., Baetge, I. and Nicklisch, A. (2014). hroot: Hamburg registration and organization online tool. *European Economic Review*, **71**, 117–120.

—, Nicklisch, A. and Baetge, I. (2012). *hroot: Hamburg registration and organization online tool*, *WiSo-HH working paper series*, vol. 1. Hamburg: WiSo-Forschungslabor.

Bolle, F., Tan, J. H. and Zizzo, D. J. (2014). Vendettas. *American Economic Journal: Microeconomics*, **6** (2), 93–130.

Bolton, G. E., Brandts, J. and Ockenfels, A. (2005). Fair procedures: Evidence from games involving lotteries*. *The Economic Journal*, **115** (506), 1054–1076.

— and Ockenfels, A. (2000). Erc: A theory of equity, reciprocity, and competition. *The American Economic Review*, **90** (1), 166–193.

Bone, J. (2006). *The hard sell: An ethnographic study of the direct selling industry.* Aldershot Hampshire England and Burlington VT: Ashgate.

Bosman, R., Sutter, M. and van Winden, F. (2005). The impact of real effort and emotions in the power-to-take game. *Journal of Economic Psychology*, **26** (3), 407–429.

Brookins, P. and Ryvkin, D. (2014). An experimental study of bidding in contests of incomplete information. *Experimental Economics*, **17** (2), 245–261.

Brown, A., Viriyavipart, A. and Wang, X. (2014). Exploding offers with experimental consumer goods.

Burrows, P. and Loomes, G. (1994). The impact of fairness on bargaining behaviour. In J. D. Hey (ed.), *Experimental Economics*, Heidelberg: Physica-Verlag HD, pp. 21–41.

Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, **100** (3), 432–459.

Camerer, C. (2006). Individual decision making. In J. H. Kagel (ed.), *The handbook of experimental economics*, Princeton, NJ: Princeton Univ. Press.

— and Thaler, R. H. (1995). Anomalies: Ultimatums, dictators and manners. *The Journal of Economic Perspectives*, **9** (2), 209–219.

Camerer, C. F. (2011). *Behavioral Game Theory: Experiments in Strategic Interaction.* The roundtable series in behavioral economics, Princeton: Princeton University Press.

Cappelen, A. W., Hole, A. D., Sørensen, E. Ø. and Tungodden, B. (2007). The pluralism of fairness ideals: An experimental approach. *American Economic Review*, **97** (3), 818–827.

Charness, G. (2004). Attribution and reciprocity in an experimental labor market. *Journal of Labor Economics*, **22** (3), 665–688.

— and Grosskopf, B. (2001). Relative payoffs and happiness: an experimental study. *Journal of Economic Behavior & Organization*, **45** (3), 301–328.

— and Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, **117** (3), 817–869.

Chowdhury, S. M., Sheremeta, R. M. and Turocy, T. L. (2014). Overbidding and overspreading in rent-seeking experiments: Cost structure and prize allocation rules. *Games and Economic Behavior*, **87**, 224–238.

Cialdini, R. B. (2003). *Influence: Science and practice*. Boston: Allyn & Bacon, 4th edn.

Cinyabuguma, M., Page, T. and Putterman, L. (2006). Can second-order punishment deter perverse punishment? *Experimental Economics*, **9** (3), 265–279.

Cox, J. and Oaxaca, R. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, **2** (3).

Cox, J. C. and Deck, C. A. (2005). On the nature of reciprocal motives. *Economic Inquiry*, **43** (3), 623–635.

Davis, D. D. and Reilly, R. J. (1998). Do too many cooks always spoil the stew? an experimental analysis of rent-seeking and the role of a strategic buyer. *Public Choice*, **95** (1/2), 89–115.

Dechenaux, E., Kovenock, D. and Sheremeta, R. M. (2015). A survey of experimental research on contests, all-pay auctions and tournaments. *Experimental Economics*, **18** (4), 609–669.

Denant-Boemont, L., Masclet, D. and Noussair, C. N. (2007). Punishment, counterpunishment and sanction enforcement in a social dilemma experiment. *Economic Theory*, **33** (1), 145–167.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory*, **3** (2), 156–168.

Erkal, N., Gangadharan, L. and Nikiforakis, N. (2011). Relative earnings and giving in a real-effort experiment. *American Economic Review*, **101** (7), 3330–3348.

Fahr, R. and Irlenbusch, B. (2000). Fairness as a constraint on trust in reciprocity: earned property rights in a reciprocal exchange experiment. *Economics Letters*, **66** (3), 275–282.

Falk, A., Fehr, E. and Fischbacher, U. (2003). On the nature of fair behavior. *Economic Inquiry*, **41** (1), 20–26.

—, — and — (2008). Testing theories of fairness—intentions matter. *Games and Economic Behavior*, **62** (1), 287–303.

Fallucchi, F., Renner, E. and Sefton, M. (2013). Information feedback and contest structure in rent-seeking games. *European Economic Review*, **64**, 223–240.

Fehr, E. and Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior*, **25** (2), 63–87.

— and Gächter, S. (2000). Cooperation and punishment in public goods experiments. *The American Economic Review*, **90** (4), 980–994.

— and Gächter, S. (2002). Altruistic punishment in humans. *Nature*, **415** (6868), 137–140.

— and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, pp. 817–868.

Finucane, M. L., Alhakami, A., Slovic, P. and Johnson, S. M. (2000). The affect heuristic in judgments of risks and benefits. *Journal of Behavioral Decision Making*, **13** (1), 1–17.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, **10** (2), 171–178.

Gilligan, C. (1982). *In a different voice: Psychological theory and women's development*. Cambridge, Mass.: Harvard University Press.

Gilovich, T., Vallone, R. and Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive psychology*, **17** (3), 295–314.

Gneezy, U. and Smorodinsky, R. (2006). All-pay auctions?an experimental study. *Journal of Economic Behavior & Organization*, **61** (2), 255–275.

Güth, W., Schmittberger, R. and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior & Organization*, **3** (4), 367–388.

Harrison, G. W. and Morgan, P. (1990). Search intensity in experiments. *The Economic Journal*, **100** (401), 478–486.

Heidhues, P. and Kőszegi, B. (2014). Regular prices and sales. *Theoretical Economics*, **9** (1), 217–251.

Herrmann, B. and Orzen, H. (2008). The appearance of homo rivalis: Social preferences and the nature of rent seeking. *CeDEx discussion paper series*, **2008-10**.

Hertwig, R., Barron, G., Weber, E. U. and Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological science*, **15** (8), 534–539.

Hey, J. D. (1981). Are optimal search rules reasonable? and vice versa? (and does it matter anyway?). *Journal of Economic Behavior & Organization*, **2** (1), 47–70.

— (1982). Search for rules for search. *Journal of Economic Behavior & Organization*, **3** (1), 65–81.

— (1987). Still searching. *Journal of Economic Behavior & Organization*, **8** (1), 137–144.

Hillman, A. L. and Riley, J. G. (1989). Politically contestable rents and transfers. *Economics and Politics*, **1** (1), 17–39.

Hirshleifer, J. (1978). *Elements of the theory of auctions and contests*, *Working papers Department of Economics, University of California*, vol. 118. Los Angeles, rev edn.

Hoffman, E., McCabe, K., Shachat, K. and Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior*, **7** (3), 346–380.

— and Spitzer, M. L. (1985). Entitlements, rights, and fairness: An experimental examination of subjects' concepts of distributive justice. *The Journal of Legal Studies*, **14** (2), 259–297.

Houser, D. and Winter, J. (2004). How do behavioral assumptions affect structural inference? evidence from a laboratory experiment. *Journal of Business & Economic Statistics*, **22** (1), 64–79.

Huck, S., Brian Wallace and in association with London Economics (2010). *The Impact of Price Frames on Consumer Decision Making*. Report OFT 1226: London: Office of Fair Trading.

Humphrey, S. J. (2004). Feedback-conditional regret theory and testing regret-aversion in risky choice. *Journal of Economic Psychology*, **25** (6), 839–857.

Ibanez, M., Czermak, S. and Sutter, M. (2009). Searching for a better deal – on the influence of group decision making, time pressure and gender on search behavior. *Journal of Economic Psychology*, **30** (1), 1–10.

Isoni, A., Poulsen, A., Sugden, R. and Tsutsui, K. (2014). Efficiency, equality, and labeling: An experimental investigation of focal points in explicit bargaining. *American Economic Review*, **104** (10), 3256–3287.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47** (2), 263.

Karakostas, A. and Zizzo, D. J. (2016). Compliance and the power of authority. *Journal of Economic Behavior & Organization*, **124**, 67–80.

Kocher, M. G., Pahlke, J. and Trautmann, S. T. (2013). Tempus fugit: Time pressure in risky decisions. *Management Science*, **59** (10), 2380–2391.

Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, **14** (3), 381–392.

— (1992). Recall in consumer search. *Journal of Economic Behavior & Organization*, **17** (1), 141–151.

Konow, J. (2000). Fair shares: Accountability and cognitive dissonance in allocation decisions. *American Economic Review*, **90** (4), 1072–1092.

Krupka, E. L. and Weber, R. A. (2013). Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association*, **11** (3), 495–524.

Kuhnen, C. M. and Tymula, A. (2012). Feedback, self-esteem, and performance in organizations. *Management Science*, **58** (1), 94–113.

Lazear, E. P. and Rosen, S. (1981). Rank-order tournaments as optimum labor contracts. *Journal of Political Economy*, **89** (5), 841–864.

Lever, J. (1976). Sex differences in the games children play. *Social Problems*, **23** (4), 478–487.

Lim, W., Matros, A. and Turocy, T. L. (2014). Bounded rationality and group size in tullock contests: Experimental evidence. *Journal of Economic Behavior & Organization*, **99**, 155–167.

Loomes, G., Moffatt, P. G. and Sugden, R. (2002). A microeconometric test of alternative stochastic theories of risky choice. *Journal of Risk and Uncertainty*, **24** (2), 103–130.

— and Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *The Economic Journal*, pp. 805–824.

Ludwig, S. and Lünser, G. K. (2012). Observing your competitor – the role of effort information in two-stage tournaments. *Journal of Economic Psychology*, **33** (1), 166–182.

Mago, S. D., Samak, A. C. and Sheremeta, R. M. (2016). Facing your opponents: Social identification and information feedback in contests. *Journal of Conflict Resolution*, **60** (3), 459–481.

—, Sheremeta, R. M. and Yates, A. (2013). Best-of-three contest experiments: Strategic versus psychological momentum. *International Journal of Industrial Organization*, **31** (3), 287–296.

Milgram, S. (1963). Behavioral study of obedience. *The Journal of Abnormal and Social Psychology*, **67** (4), 371–378.

Millner, E. L. and Pratt, M. D. (1991). Risk aversion and rent-seeking: An extension and some experimental evidence. *Public Choice*, **69** (1), 81–92.

Moffatt, P. G. (2015). *Experimetrics: Econometrics for experimental economics*. UK: Palgrave Macmillan.

Nalebuff, B. J. and Stiglitz, J. E. (1983). Prizes and incentives: Towards a general theory of compensation and competition. *The Bell Journal of Economics*, **14** (1), 21.

Nikiforakis, N. (2008). Punishment and counter-punishment in public good games: Can we really govern ourselves? *Journal of Public Economics*, **92** (1-2), 91–112.

O'Donoghue, T. and Rabin, M. (2001). Choice and procrastination. *The Quarterly Journal of Economics*, **116** (1), 121–160.

Offerman, T. (2002). Hurting hurts more than helping helps. *European Economic Review*, **46** (8), 1423–1437.

Office of Fair Trading (2010). Advertising of prices. **Report OFT 1291. London: Office of Fair Trading**.

Oxoby, R. J. and Spraggon, J. (2008). Mine and yours: Property rights in dictator games. *Journal of Economic Behavior & Organization*, **65** (3-4), 703–713.

Payne, J. W., Bettman, J. R. and Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **14** (3), 534–552.

Piaget, J. (1932 (1968)). *The Moral judgement of the child*. International library of psychology, philosophy and scientific method, London: Routledge and K.Paul.

Price, C. R. and Sheremeta, R. M. (2011). Endowment effects in contests. *Economics Letters*, **111** (3), 217–219.

— and — (2015). Endowment origin, demographic effects, and individual preferences in contests. *Journal of Economics & Management Strategy*, **24** (3), 597–619.

Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American Economic Review*, **83**, 1281–1302.

— (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, **68** (5), 1281–1292.

Rapoport, A. and Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, **5** (2), 105–120.

Richard, R., vanÂ derÂ Pligt, Joop and Vries, N. K. d. (1996). Anticipated regret and time perspective: Changing sexual risk-taking behavior. *Journal of Behavioral Decision Making*, **9** (3), 185–199.

Ritov, I. and Baron, J. (1992). Status-quo and omission biases. *Journal of Risk and Uncertainty*, **5** (1), 49–61.

Roth, A. E. (1995). Bargaining experiments. In J. H. Kagel and A. E. Roth (eds.), *The handbook of experimental economics*, Princeton, N.J.: Princeton University Press, pp. 253–348.

Ruffle, B. J. (1998). More is better, but fair is fair: Tipping in dictator and ultimatum games. *Games and Economic Behavior*, **23** (2), 247–265.

Salop, S. and Stiglitz, J. (1977). Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies*, **44** (3), 493–510.

Schotter, A. and Braustein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry*, **19** (1), 1–25.

Schunk, D. and Winter, J. (2009a). The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. *Journal of Economic Behavior & Organization*, **71** (2), 347–360.

— and — (2009b). The relationship between risk attitudes and heuristics in search tasks: A laboratory experiment. *Journal of Economic Behavior & Organization*, **71** (2), 347–360.

Seale, D. A. and Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the 'secretary problem'. *Organizational Behavior and Human Decision Processes*, **69** (3), 221–236.

— and — (2000). Optimal stopping behavior with relative ranks: the secretary problem with unknown population size. *Journal of Behavioral Decision Making*, **13** (4), 391–411.

Segal, U. and Sobel, J. (2007). Tit for tat: Foundations of preferences for reciprocity in strategic settings. *Journal of Economic Theory*, **136** (1), 197–216.

Selten, R., Sadrieh, A. and Abbink, K. (1999). Money does not induce risk neutral behavior, but binary lotteries do even worse. *Theory and Decision*, **46** (3), 213–252.

Sheremeta, R. M. (2010). Experimental comparison of multi-stage and one-stage contests. *Games and Economic Behavior*, **68** (2), 731–747.

— (2011). Contest design: An experimental investigation. *Economic Inquiry*, **49** (2), 573–590.

— (2013). Overbidding and heterogeneous behavior in contest experiments. *Journal of Economic Surveys*, **27** (3), 491–514.

— and Zhang, J. (2010). Can groups solve the problem of over-bidding in contests? *Social Choice and Welfare*, **35** (2), 175–197.

Shupp, R., Sheremeta, R. M., Schmidt, D. and Walker, J. (2013). Resource allocation contests: Experimental evidence. *Journal of Economic Psychology*, **39**, 257–267.

Simonson, I. (1992). The influence of anticipating regret and responsibility on purchase decisions. *Journal of Consumer Research*, pp. 105–118.

Sitzia, S. and Sugden, R. (2011). Implementing theoretical models in the laboratory, and what this can and cannot achieve. *Journal of Economic Methodology*, **18** (4), 323–343.

Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior & Organization*, **35** (3), 309–332.

Spranca, M., Minsk, E. and Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, **27** (1), 76–105.

Stewart, N., Chater, N. and Brown, Gordon D A (2006). Decision by sampling. *Cognitive psychology*, **53** (1), 1–26.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, **69** (3), 213–225.

Tenorio, R. and Cason, T. N. (2002). To spin or not to spin? natural and laboratory experiments from the price is right. *The Economic Journal*, **112** (476), 170–195.

Tullock, G. (2001). Efficient rent seeking. In A. A. Lockard and G. Tullock (eds.), *Efficient Rent-Seeking*, Boston, MA: Springer US, pp. 3–16.

Varian, H. R. (1980). A model of sales. *The American Economic Review*, **70** (4), 651–659.

Weizsäcker, G. (2010). Do we follow others when we should? a simple test of rational expectations. *The American Economic Review*, **100** (5), 2340–2360.

Xiao, E. and Bicchieri, C. (2010). When equality trumps reciprocity. *Journal of Economic Psychology*, **31** (3), 456–470.

Zeelenberg, M., Beattie, J., vanÂ derÂ Pligt, Joop and deÂ Vries, Nanne K. (1996). Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes*, **65** (2), 148–158.

— and 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.

Zimbardo, P. G. (2009). *The Lucifer effect: How good people turn evil*. London: Rider, pbk. ed edn.

Zizzo, D. J. (2005). *Transfer of knowledge in economic decision making*. Basingstoke: Palgrave Macmillan.

Zwick, R., Rapoport, A., Lo, A. K. C. and Muthukrishnan, A. V. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, **22** (4), 503–519.