AN EXPERIMENTAL INVESTIGATION OF THE IMPACT OF EXPERIENCE ON LOSS AVERSION

Luke Lindsay, MSc.

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

September 2008

Abstract

The first chapter introduces the thesis and reviews the literature on loss aversion, the endowment effect and the willingness to pay/willingness to accept gap, and the effects of experience.

The second chapter reports an extended version of Knetsch's exchange of goods experiment to explore how different types of experience influence the endowment effect. The experiment has four treatments, which compare the behaviour of subjects with experience of consuming, owning, and choosing goods to a control group. The results are consistent with earlier studies in that an endowment effect is observed; however, the strength of the effect is less than in earlier studies and differs between treatments. In particular, there is a significantly stronger endowment effect in the treatments in which the endowment is acquired in two steps rather than one step.

The third chapter reports a repeated market experiment in which subjects buy and sell lotteries under symmetric and asymmetric information. Buying and selling bids and prices are compared. A gap between buying and selling prices decays under symmetric information but persists under asymmetric information. Furthermore, there are spillover effects. When the regime switches between symmetric and asymmetric information, subjects do not immediately adjust their behaviour. The results are interpreted as evidence that behaviour is driven by heuristics.

The fourth chapter reports another repeated market experiment in which subjects buy and sell lotteries. How the lotteries' odds are presented and whether the lottery gets resolved after each trial vary between treatments. Among the findings is that the gap between buying and selling bids decays when lotteries are not resolved each trial but persists when they are.

The final chapter summarises the findings of the three experiments and identifies common patterns. Directions for future experimental and theoretical research are suggested. Finally, implications for policy are discussed.

Acknowledgements

I owe a debt of gratitude to my supervisors Dr. Steven Humphrey and Professor Chris Starmer for their criticisms, comments, and suggestions.

I would also like to thank members of the Centre for Decision Research and Experimental Economics (CeDEx) at the University of Nottingham for many useful comments and similarly participants at the joint workshops with experimental economists from the University of Amsterdam and the University of East Anglia.

I am grateful to Jacinto Braga, Harshana Kasseeah, Xiaojing Kong, Pachara Lochindaratn and Maria Ruiz-Martos for help with running experiments, to CeDEx for providing resources, and to everyone who took part in the experiments.

I am further deeply grateful to the Economic and Social Research Council and the University of Nottingham School of Economics for financial support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction and Literature Review

1.1 Introduction

Which theory of choice best describes the decisions people actually make? How do patterns in the choices people make change as people learn and gain experience? The answer to these questions have wide ranging implications. For example, different theories of choice give different answers to the following questions. When does a New York taxi driver decide to finish working for the day? What factors affect the price a house seller is willing to accept or the price of an impressionist painting at auction? How do people adjust their saving and spending when they get good or bad news about future income? How should a town council decide how many trees to plant in a public park? How might U.S. policymakers have weighed the risks when deciding how long to keep troops deployed in Vietnam and more recently Iraq.

In this chapter, I outline[1] (a) the standard theory of choice used in economics and (b) alternative theories that incorporate reference point effects including loss aversion. Theories from these two classes typically give different answers to the questions above. I assess the experimental and field evidence for the two theories. There

[1]More detailed discussions of relevant theory will be presented in later chapters as is necessary to motivate experiment designs.

is considerable evidence of replicable behavioural phenomena that are inconsistent with standard preference theory, which I refer to as *anomalies*. Many anomalies are, however, consistent with alternative theories of choice. I also consider the argument that anomalies for the standard theory when they occur may be only transient effects that are eliminated by experience.

Finally, I summarise what we know from the existing literature and what questions are left unanswered. It is these unanswered questions that motivate the new experimental studies presented in this thesis.

1.2 Standard Theory

The standard assumption used in economics is that final states alone determine preferences. That is, preferences are independent of current entitlements. This allows the textbook case where a person's preferences over two goods are represented by drawing indifference curves. A person chooses quantities of the two goods that puts them on the highest indifference curve given their budget constraint. For a given budget constraint, whatever combination of the two goods they are endowed with, they will choose the same combination of the two goods.

1.3 Reference Point Effects and Loss Aversion

There is a growing body of evidence that contradicts the assumption that current entitlements do not affect preferences and choices. The first evidence appears in the psychology literature. For example, Brehm (1956) writes

> Previous studies have found that when a person is given an object he tends subsequently to see it as more desirable. This may be called the effect of ownership.

The idea that preferences might depend on current entitlements was introduced into the economics literature by Kahneman and Tversky (1979) and Thaler (1980). When a person makes a choice, their decision is often influenced by the reference point from which they evaluate the options. This influence of the reference point follows a regular pattern. Kahneman and Tversky write:

> A salient characteristic of attitudes to changes in welfare is that losses loom larger than gains. The aggravation that one experiences in losing a sum of money appears to be greater than the pleasure associated with gaining the same amount.

That is, losses relative to the reference point carry more weight than commensurate gains. Kahneman and Tversky call this effect *loss aversion.*

1.3.1 Loss Aversion in Choice under Risk

Kahneman and Tversky's (1979) Prospect Theory is an alternative theory of choice under risk to Expected Utility Theory. A theory of choice under risk gives preferences over *lotteries*. Consider the simple case of a lottery that pays one of n possible monetary outcomes. The set of consequences, levels of monetary wealth, can be written as $X = (x_1, ..., x_n)$. Each consequence x_i occurs with probability p_i. Note that $p_i \in [0,1]$ and $\sum_{i=1}^{n} p_i = 1$. Under Expected Utility Theory, the expected utility of the lottery is $\sum_{i=1}^{n} p_i u(x_i)$ where $u(x_i)$ is the utility from outcome x_i. Under Prospect Theory, the lottery's value is $\sum_{i=1}^{n} \pi(p_i) v(x_i - r)$. There are two key differences. First, probabilities are weighted by the function $\pi(p_i)$ instead of entering the value calculation directly (Kahneman and Tversky call the values that $\pi(\cdot)$ maps probability values to *decision weights*). The weighting function has an inverted S shape. Typically, when $p < 0.3$, $\pi(p) > p$; when $p > 0.3$, $\pi(p) < p$. That is, small probabilities are over-weighted; large ones, under weighted. Second, outcomes are assessed as losses and gains rather than levels of wealth. Losses and gains are calculated relative to a reference level of wealth, r. The reference point is usually assumed to be the current level of wealth. The value function, $v(\cdot)$ is S shaped, as shown in figure 1.1. It is steeper in the domain of losses than in the domain of gains and kinked at the origin. This feature captures *loss aversion.*

As well as this, the value function is concave in the domain of gains and convex in the domain of losses. As a consequence, the theory predicts risk aversion in the domain of gains and risk seeking in the domain of losses.[2]

The original (1979) version of Prospect Theory has been refined by later studies. Tversky and Kahneman's (1992) Cumulative Prospect Theory changed the way decision weights are handled so that the weights always add up to one. Sugden (2003) presented a reference-dependent model of preferences with income effects and with the possibility of risky reference points (but left out probability weights for simplicity). Schmidt et al. (2008) combined probability weighting, income effects, and the possibility of risky reference points in a single model.

1.3.2 Loss Aversion in Riskless Choice

Prospect Theory is a theory of choice under risk. Tversky and Kahneman (1991) take two of the key features of Prospect Theory, that the carriers of value are losses and gains and loss aversion, and develop a reference dependent theory of riskless choice. Consider figure 1.2. Under the standard theory of choice (the left hand graph), preferences are independent of current entitlements. On the graph, the per-

[2] An implication of risk seeking behaviour in the domain of losses is when a person suffers a loss, if they do not adjust their reference point to the new wealth level, they will exhibit more risk seeking behaviour than before the loss occurred. Kahneman and Renshon (2007) argue that when things are going badly in a conflict, cutting the losses and withdrawing appears less attractive to policy makers than taking further risks. "U.S. policymakers faced this dilemma at many points in Vietnam and today in Iraq. To withdraw now is to accept a sure loss, and that option is deeply unattractive. The option of hanging on will therefore be relatively attractive, even if the chances of success are small and the cost of delaying failure is high." Another case of increased risk seeking over losses is observed by Rasiel et al. (2005). They found that patients with life threatening conditions make more risky decisions as their condition declines.

Figure 1.1: Prospect Theory's S-Shaped Value Function

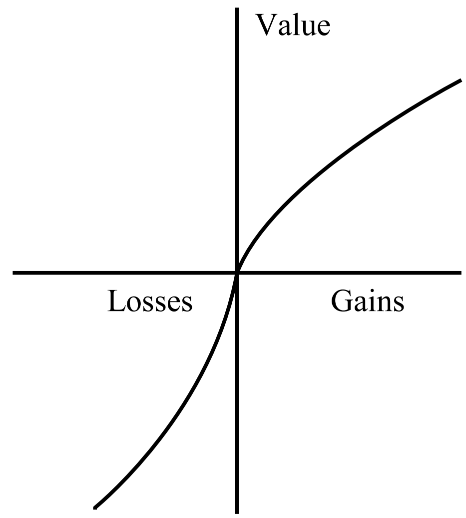

son is on a higher indifference curve at point B than at point A. Accordingly, if they are faced with a choice between A and B, they will choose B no matter what combination of goods X and Y they hold at the time they make the choice. In contrast, if preferences are reference dependent (the right hand graph), preferences over goods X and Y vary with the allocation the decision maker holds. If the decision maker holds allocation A, the reference point is A. At this reference point, the solid line represents allocations that are just as attractive as A. A is preferred to all points below this line and points above the line are preferred to A. Likewise, at point B their preferences are represented by the dashed line. B is preferred to points below the line and points above it are preferred to B. Notice that from reference point A, A is preferred to B, but from reference point B, B is preferred to A. As a consequence, someone given allocation A will not be willing to swap it for allocation B; someone given B will not be willing to swap it for A. This is known as the endowment effect. An important consequence of the endowment effect (valuing an item more once you own it) is that willingness to accept compensation (WTA), the minimum amount of money you would require to give up a good if you owned it, will exceed willingness

Figure 1.2: Alternative Models of Preferences over Two Goods

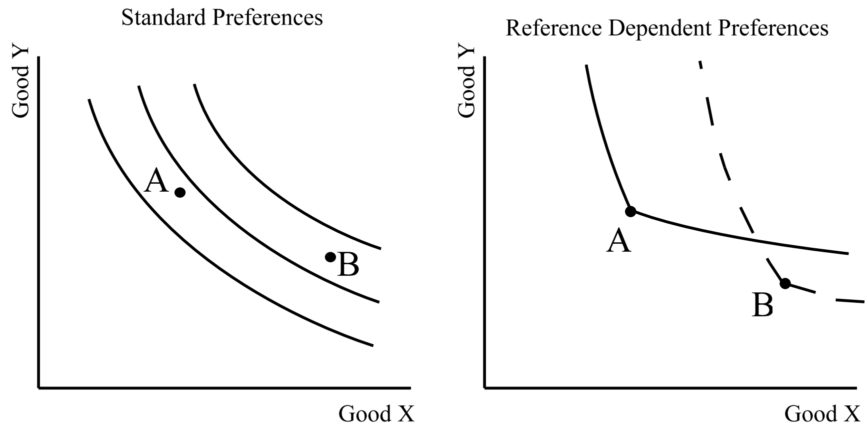

to pay (WTP), the maximum amount of money you would pay to obtain a good. This disparity occurs because once you own a good, it becomes more valuable to you, so the amount of money that has the same value to you has the good will be higher[3]. An intuitive way to think about WTA and WTP in terms of buying and selling is as follows. WTA is the minimum price we would be willing to sell an item we own for; WTP is the maximum price we be willing to pay to buy an item.

1.3.3 WTA and WTP for Lotteries

How can we model WTA and WTP for lotteries? One option is to use the riskless choice framework and let good Y be lotteries and good X be money. Alternatively, we can use the one variable value function described in section 1.3.1. Consider the lottery whose payout is determined by the colour of a ball drawn from an urn containing ten balls, two red and eight black. If a red ball is drawn, it pays out £20; otherwise, it pays out zero. If we omit probability weights for simplicity, we can calculate WTP by solving the following.

[3]For a good, WTP is the same as the Hicksian compensating variation; WTA, equivalent variation. Under the standard theory, as Willig (1976) argues, we should expect compensating and equivalent variations to be approximately equal.

$$v(0) = 0.2v(20 - wtp) + 0.8v(-wtp)$$

The solution to this equation, wtp, is the price for the lottery that would make us indifferent between (a) buying and playing the lottery and (b) not trading.

When calculating WTA, we have two options. If we take the reference point to be wealth before the lottery is played out, WTA is the solution to

$$v(wta) = 0.2v(20) + 0.8v(0)$$

In this case wta, is the amount of money that is just as attractive as playing the lottery.

Alternatively, we can allow uncertain reference points and take the reference point to be playing the lottery. In this case, if we sell the lottery, we gain the price of the lottery but lose what the lottery would have paid out to us. Accordingly, if a red ball is drawn, our net loss is £20 less the amount we sold the lottery for; if a black ball is drawn, we lose nothing and gain the amount we sold the lottery for. WTA is the solution to

$$v(0) = 0.2v(wta - 20) + 0.8v(wta)$$

The solution, wta, is the amount of money that compensates for losing what we would have received if we played the lottery.

1.3.4 Reference Dependent Theory and Risk Aversion

Reference dependent theories such as Prospect Theory can account for aversion to gambles that have both positive and negative outcomes. For instance, Kahneman and Tversky (1982) report that when faced with a 50-50 lose \$100, win X gamble, it is common for people to turn down the gamble unless X is \$200 or more. This behaviour can be explained by Prospect Theory in terms of the value function being steeper in the domain of losses than in the domain of gains.

Rabin (2000) argues that Expected Utility Theory with a concave utility of wealth function cannot explain risk aversion over modest stakes without predicting implausible degrees of risk aversion over large stakes. As a consequence, risk aversion over small stakes is an anomaly for the standard theory. Cox and Sadiraj (2006) object that Expected Utility Theory as developed by von Neumann and Morgenstern (1947) does not specify the domain of the utility function (it could be wealth, income, or vectors of commodities). Rabin's critique applies to expected utility of wealth models. They show that an *Expected Utility of Income* model can explain risk aversion over small and large stakes. One might assume that since Expected Utility of Income models can give an account of risk aversion over small stakes, then it is not an anomaly for the standard theory. I argue that this assumption would be a mistake. Expected utility of income models are a departure from the standard theory since, like Prospect Theory, they make preferences reference point dependent. As a consequence, they can predict a range of phenomena that are anomalies for the standard theory. In fact, the utility of income function Cox and Sadiraj use to demonstrate risk aversion over small stakes without implausible risk aversion over large stakes also predicts a WTA/WTP gap for lotteries, framing induced preference reversals, and choice-value preference reversals (see appendix A).

1.4 Evidence from Surveys and Experiments

There is now a large body of evidence on the effects of reference points on preferences and behaviour. Brown and Gregory (1999), Horowitz and McConnell (2002), Plott and Zeiler (2005), and Sayman and Onculer (2005) each provide useful reviews and list many of the studies.

Table 1.1: Studies Reporting Non Incentivised WTA/WTP Valuations

Study	Good	WTA[a]	WTP[a]	$\frac{WTA}{WTP}$
Hammack and Brown 1974	Waterfowl hunting	1044	247	4.2
Banford et al. 1979	Fishing pier	120	43	2.8
Bishop and Heberlein 1979	Goose hunting permit	101	21	4.8
Brookshire et al. 1980	Elk hunting quality	69	13	5.4
Rowe et al. 1980	Air visibility	24	5	5.2
Adamowicz et al. 1993	Movie ticket	9.3	4.8	2.0
MacDonald and Bowker 1994	Industrial plant odour	735	105	7.0

This table is based on Brown and Gregory (1999).
[a] Valuations are in US dollars.

In the economics literature, the early empirical evidence of reference dependent preferences comes from hypothetical valuation studies. In these studies participants were typically asked to report WTA and WTP. Table 1.1 summarises these studies that investigate WTA/WTP using surveys (as opposed to incentive compatible procedures where subjects have a incentive to give truthful answers). In all seven of the studies, WTA valuations exceed WTP and the ratio $\frac{WTA}{WTP}$ ranges from 2 to 7. Under standard theory, WTA and WTP should be approximately equal, so the results of these studies are an anomaly for the standard theory.

More recent studies have mainly used incentive compatible elicitation mechanisms instead of simply asking subjects to state valuations. In such studies subjects have an incentive to reveal their true preferences and think carefully about their decision since their payoff for participating depends on the decisions they make. The simplest example of an incentivised study is the classic chocolate and mugs experiment run by Knetsch (1989). One group of 76 students were given a coffee mug, asked to complete a short questionnaire, then asked if they wanted to swap the mug for a bar of Swiss chocolate. Another group of 87 students faced the same conditions except that they were given chocolate in the first step and asked if they wanted to swap it for the mug at the end. In the first group, 89 percent favoured the mug; in the second group just 10 percent favoured the mug. Under the standard theory, preferences do not depend on current assets, so there should be no difference between these figures.

The gap between WTA and WTP has also been extensively studied in experiments, which have used a range of incentive compatible mechanisms to elicit valuations. Tables 1.2, 1.3, and 1.4 respectively summarise one shot incentivised WTA/WTP studies, studies using lotteries instead of goods, and studies with repeated markets. Eliciting sincere valuations is more difficult than eliciting sincere choices. Obviously, if experiment participants are simply asked to state WTA and WTP valuations and then allowed to trade at these prices then they can increase their payout from the experiment by exaggerating WTA and understating WTP. To avoid this problem researchers have used pricing mechanisms where the price at which a given subject trades is not the value that the subject stated.

Most of the studies use one sided markets in the sense that subjects are buying from the experimenter or selling to the experimenter rather than trading amongst themselves. The most common mechanisms used in these studies are the Becker-DeGroot-Marschak (1964) mechanism (BDM) and versions of the Vickrey (1961) auction. One notable exception is Brookshire and Coursey (1987) who use a Smith (1980) auction to measure the value of a public good, the density of trees in a public park. A few studies use two sided markets. Such markets have been implemented in two ways. The first method is to divide the subjects into buyers and sellers, elicit WTA and WTP valuations, set the price so that supply equals demand, and execute trades that are consistent with the price. This method is used by Kahneman et al. (1990) and Morrison (1997). The second method is to have subjects simultaneously act as buyers and sellers. The method is used by Knez and Smith (1987) and Singh (1991). Subjects are endowed with different quantities of a good, then asked to state WTA to give up one unit and WTP to gain one unit, the price is set so that supply equals demand and trades consistent with the price are executed.

Table 1.2: Studies Reporting *One Shot* Incentivised WTA/WTP Valuations

Study	Good(s)	Mechanism	Ratio
Knetsch 1989, Test 2	Swiss chocolate	BDM	2.0
Bateman et al. 1997	chocolate, coke	BDM	> 1.0
Franciosi et al. 1996	mugs	TSM	2.4
Carmon and Ariely 2000	basketball game tickets	TSM	14.5
Bateman et al. 2005	luxury chocolates	MDC	1.9

In this and subsequent tables, the elicitation devices used in experiments are abbreviated as follows: MDC = multiple dichotomous choice; SPA = second price Vickrey auction; BDM = Becker-DeGroot-Marschak mechanism; RNA = random nth-price Vickrey auction; MPA= median price Vickrey auction; NPA = ninth price Vickrey auction; TSM = two sided market; DBA = double bid auction; LSA = Laboratory Smith Auction.

1.5 Evidence from Field Studies

Loss aversion and Prospect Theory in general have also been studied in non experimental settings. See Camerer (1998) for a review. Loss aversion can explain the extra return on stocks compared to bonds, the equity premium (Benartzi and Thaler, 1995). In the labour market, Camerer et al. (1997) found a tendency for New York taxi drivers to work longer hours on low-wage days; Chou (2002) found evidence of similar behaviour among taxi drivers in Singapore[4]. Goette et al. (2004) review several studies on labour supply choices and find that a reference-dependent preferences model can explain all the data. Hardie et al. (1993) find that consumers react asymmetrically to price increases and decreases: they react more to price increases than cuts. Bowman et al. (1999) find that consumers react asymmetrically to news about future income: they raise consumption following good news but fail to cut it following bad news. Not every study testing for loss aversion in the field has found positive results, however. For instance Beggs and Graddy (2005) analyse data from Impressionist and Contemporary Art auctions and find evidence of reference dependence effects but not loss aversion.

[4]Farber (2008) applies a more sophisticated econometric model to data on the labour supply of New York City taxi drivers and concludes that there may be a reference level of income on a given day that affects labour supply, but if there is, it varies from day to day giving the theory little predictive power.

Table 1.3: Studies Reporting on the WTA/WTP Gap for *Lotteries*

Study	Lotteries[a]	Mechanism[b]	Main Result[c]
Knetsch and Sinden 1984, Test 5	{$70;?}	binary choice	4.0
Harless 1989	6 binary lotteries	SPA	1.6-1.5
Singh 1991	{$1, 0.5; $2, 0,5}	TSM	1.1 to 1.57
Birnbaum et al. 1992	over 100 lotteries	survey	WTA>WTP
Kachelmeier and Shehata 1992, group 7 and 8	{$20,0.5;0,0.5}	BDM	1.8-2.1
Mellers et al. 1992	36 binary lotteries	incentivised survey[d]	WTA>WTP
Eisenberger and Weber 1995	{DM 10, ?;0, ?}	BDM	1.5-2.3
Casey 1995	12 binary lotteries	survey	WTA>WTP
Shefrin and Caldwell 2001	hypothetical scenarios	survey	2.6-18.8
Peters et al. 2003	{$100, 0.05}	RNA	7.0
Loomes et al. 2003	{£12, 0.2; 0, 0.8},{£12,0.8;0,0.2}	MPA	1.1-1.2
Schmidt and Traub 2006	lotteries: 56 risky, 2 uncertain	BDM	1.9
Loomes et al. 2007	{£18, 0.19}, {£4, 0.81}	MPA	1.1-1.4
Neilson et al. 2008	uncertain lotteries	BDM	WTA>WTP

[a] For studies that only used one or two lotteries, the actual lotteries are listed in the format {payoff, probability;...}. Where the probability is "?" it means subjects were not told the probability.
[b] The abbreviations used here are defined on the note below table 1.2.
[c] The Main Result column reports the mean WTA/WTP ratio. Where there were several lotteries, several treatments, or repeated markets, the range of the mean WTA/WTP ratio is reported. Where authors did not report the mean ratio, the relation between WTA and WTP is reported.
[d] Two of the 36 lotteries were randomly selected and the one assigned the higher price was played out.

1.6 The Scope and Magnitude of Loss Aversion

Tables 1.2 and 1.3 show that although a WTA/WTP gap is consistently observed, its size varies considerably between studies. Several hypotheses have been advanced in the literature concerning when loss aversion influences choices. Tversky and Kahneman (1991) posit that a reluctance to sell does not occur in routine commercial transactions. They ran an experiment where subjects traded tokens that could be exchanged for a fixed amount of money at the end of the experiment and found no WTA/WTP gap. In contrast, when subjects traded pens, mugs, and binoculars using the same market mechanism, there was a gap. Novemsky and Kahneman (2005) report further evidence that supports this hypothesis. They find that while there is loss aversion for goods, there is little or no loss aversion for money. Their interpretation is that money is held for the purpose of exchange, so there is no loss aversion when money is given up in exchange for goods. Bateman et al. (2005) also investigate whether there is loss aversion when people exchange money for goods, but they find evidence of loss aversion for money. The hypothesis has also been approached theoretically by Koszegi and Rabin (2006). They develop a model of reference dependent preferences where the reference point is not a person's current endowment but their rational expectation held in the recent past about what outcomes are going to occur. Accordingly, when a person expects to make a trade, their reference point will be the outcome where they make the trade and the money or goods given up in the trade are not considered a loss. When they do not expect to trade, their reference point will be the outcome where they do not trade, so anything given up in the potential trade will be considered a loss.

Another class of decision problems where loss aversion may not occur is when one person is making decisions on behalf of a third party. Knetsch and Sinden (1984) found an endowment effect using lotteries and money when people were making decisions for themselves but not when they were advising others what to do in the same decision problem. Marshall et al. (1986) find more evidence that WTA/WTP gaps occur when people are making decisions for themselves but the gap is much smaller when people act as advisers to a third party. Furthermore, Arlen et al. (2002) ran an experiment modelling an employee-employer relationship and found the employee does not exhibit an endowment effect in this setup. The evidence is not, however, universal. Birnbaum et al. (1992) conducted a survey based study investigating how people value lotteries (not the effect of agency). Subjects were asked to advise 'buyers' and 'sellers' how to price the lotteries. The study found that the advised selling prices exceeded advised buying prices even though subjects were acting as advisers.

There is also evidence that the type of good affects the degree of loss aversion. Horowitz and McConnell (2002), reviewing the data from 45 WTA/WTP studies conclude that "the farther a good is from being an *ordinary private good,* the higher the ratio". They find the WTA/WTP ratio is considerably higher for non-market goods, public goods, and health and safety, than it is for private goods and lotteries.

1.7 The Discovered Preference Hypothesis

Does the behaviour of participants in the studies discussed actually reflect their preferences? Plott (1996) argues that rational behaviour emerges from a process of reflecting, experience and practice. On this view, at least some of the behaviour observed in experimental studies appears random and inconsistent with models of preference because the participants have yet to fully appreciate what it is their best interest to do. It is only when the choice problem is repeated with appropriate incentives and feedback that people develop stable strategies. Furthermore, it is only when this has occurred and participants recognise and anticipate that other participants in the decision problem are also rational that behaviour will reflect what economic theory predicts.

If Plott's hypothesis is right, then we should expect anomalies for the standard theory to go away when decision makers have had the chance to reflect, gain experience and practice. If this occurs, it should observable in the lab. The following sections discuss evidence bearing on this. Further, note that even if Plott's hypothesis is right, it doe not mean anomalies for the standard theory are not economically significant. Even if people do eventually discover their true preference, it does not follow that they will have done so when they face economically important decisions.

1.8 Market Experience Gained in the Lab

Some of the evidence for a WTA/WTP gap and loss aversion comes from experiments where subjects made one-off decisions in non market settings. If Plott's discovered preference hypothesis is right, we have reason to be skeptical of the relevance of the results of some studies. Binmore (1999) argues we should not expect economic theory to predict behaviour in the laboratory unless (a) the problem subjects face is not only reasonably simple in itself but also framed so that it seems simple to the subjects, (b) incentives are adequate, and (c) the time allowed for trial and adjustment is sufficient. On a similar note, Knez et al. (1985) responding to early research on the WTA/WTP gap and other anomalies, write "we would urge suspension of scientific judgement until this evidence has been further examined in repetitive market-like environments."

To address concerns such as the above, researchers have investigated WTA/WTP in repeated market like environments. Table 1.4 summarises repeated experimental market studies. In the majority of studies, the initial gap closes or reduces in repeated markets. There are a number of exceptions, however.

Table 1.4: Studies Reporting on the WTA/WTP Gap in *Repeated* Markets

	Good(s)	Mechanism[a]	Trials	Main Finding
Coursey et al. 1987	tasting a bitter liquid	SPA	5-10	Gap reduces
Brookshire and Coursey 1987	trees in a public park	LSA	up to 5	Gap reduces
Harless 1989	6 resolved risky lotteries	SPA	6 buying; 6 selling	Small gap persists
Kahneman et al. 1990	induced value tokens	TSM	1-3	No gap
	pens, mugs, binoculars	TSM	4-5 per good	Gap persists
Singh 1991	unresolved risky lotteries	DBA	not reported	Gap closes
Boyce et al. 1992	pine trees	BDM	10 practise; 1 binding	Gap persists
Shogren et al. 1994	chocolate and coffee mugs	SPA	5	Gap closes
	food safety	SPA	20	Gap persists
Morrison 1997	coffee mugs	TSM	5	Gap persists
	chocolate bars	TSM	5	No gap
List and Shogren 1999	chocolate bars	SPA	4	Gap closes
	food safety	SPA	9-10	Gap reduces
Shogren et al. 2001	chocolate and coffee mugs	BDM	10	Gap persists
	chocolate and coffee mugs	SPA	10	Gap closes
	chocolate and coffee mugs	RNA	10	Gap closes
Knetsch et al. 2001	coffee mugs	SPA	6	Gap closes
	coffee mugs	NPA	6	Gap widens
Plott and Zeiler 2005	lotteries, mugs	BDM	2 practise, 15 paid	Gap closes
Loomes et al. 2003	2 unresolved risky lotteries	MPA	6	Gap closes
Loomes et al. 2007	2 unresolved risky lotteries	MPA	6	Gap closes

A risky lottery is one where the odds are know to experiment participants. A resolved lottery is resolved after each trial; an unresolved lottery is not resolved until the end of the experiment.

[a] The abbreviations used here are defined on the note below table 1.2.

1.9 Market Experience Gained in Naturally Occurring Markets

An alternative way to study the effect of market experience on behavioural anomalies is recruiting experiment subjects with different levels of market experience gained in naturally occurring markets and then comparing their behaviour. A potential advantage of this approach is that it allows the effect of experience acquired over a long time period to be investigated. Camerer and Hogarth (1999) suggest this might be important for the following reason.

> useful cognitive capital probably builds up slowly, over days of mental fermentation or years of education rather than in the short-run of an experiment (1-3 hours).

The effect of experience gained in naturally occurring markets on behaviour has been investigated in three studies by John List. Two of the studies (List, 2003, 2004b) were run at a market places for (a) sports-cards and memorabilia and (b) collector pins. The experiments involved trading unique pieces of memorabilia that subjects would not have seen before. The main findings were that the magnitude of the endowment effect decreases with the level of market experience and that the most experienced traders did not exhibit an endowment effect while less experienced ones did. Further, he found evidence consistent with the hypothesis that market experience causes a reduction in the endowment effect rather than the alternative that people with a disposition not to exhibit the endowment effect also tend to trade more intensively and so have more market experience.

The third study (List, 2004a) was similar except that it involved trading everyday consumer goods. This let him test whether experience of trading memorabilia affected behaviour when trading everyday goods. One of the tests List ran was a repetition of Knetsch's chocolate and mugs exchange experiment with subjects recruited at the sports-card show instead of university students. Some of the subjects were dealers (people who had rented a table to display their merchandise) and some non dealers (regular attendees). The dealers typically had significantly more experience of trading than non dealers. This allowed him to test whether market experience eliminates the endowment effect in general. He found that among non dealers, 77 percent of those endowed with the mug favoured the mug compared to just 19 percent of those endowed with the chocolate. Among dealers, however, the figures were 53 percent and 56 percent: the endowment effect was not observed. The results followed the same pattern as List's other studies: inexperienced traders exhibited an endowment effect; experienced ones did not.

1.10 Heuristics

There is a large body of evidence that people use a number of simplifying heuristics when making judgements and choices in the face of uncertainty.[5] An illustrative example (Tversky and Kahneman, 1974) of a heuristic is judging distance based on visual clarity. The more sharply an object is seen, the closer it appears to be. In any given scene, the nearer objects are more sharp than the more distant ones, so using the heuristic will produce reasonable estimates of distance. However, using the heuristic will also produce predictable biases in judgements: distances will be overestimated when visibility is poor because objects are seen blurred; distances will be underestimated when visibility is good because objects are seen sharply. Tversky and Kahneman argue that people use heuristics in a similar way when making judgements in the face of uncertainty. They describe the following three heuristics that explain some of the biases that are observed in the judgements people make involving the likelihood of uncertain events. First, *representativeness*: people as-

[5]The collections edited by Kahneman et al. (1982) and Gilovich et al. (2002) provide overviews and contain many of the key studies.

sess the probability that object A belongs to class B based on the degree to which A resembles B. For example, they will base their estimates of the likelihood that an individual works as a librarian on the degree to which the individual corresponds to the stereotypical librarian. Second, *availability*: people assess the frequency of a class (or probability of an event) by the ease with which instances come to mind. For example, one might assess the risk of heart attack by recalling occurrences among ones acquaintances. Third, *adjustment and anchoring*: people make estimates by adjusting an initial value and the adjustment is typically insufficient. For example, in an experiment subjects were asked to guess whether the percentage of African countries in the United Nations (N) was greater or less than X where X was a number between 1 and 100 determined by spinning a wheel of fortune, then they were asked to estimate N. The estimates were not independent of X: instead, estimates of N were higher when the wheel of fortune had produced higher values of X.

Studies investigating heuristics, such as representativeness, availability, adjustment and anchoring, provide evidence that people take short cuts when making judgements and choices. There is also some evidence that people use similar processes when making WTA and WTP valuations. Braga and Starmer (2005) and Bateman et al. (2005) argue that the observed WTA/WTP gaps may be a consequence of people using a *caution heuristic*. Bateman et al. find evidence that a *caution heuristic* in addition to loss aversion contributes to the WTA/WTP gap. The heuristic is a tendency for people to overstate *incoming* valuations and understate *outgoing* ones. An incoming valuation is where the subject states the minimum amount of a good they are willing to receive (e.g. WTA); an outgoing one is where they state the maximum amount they are willing to give up (e.g. WTP). In some environments, using this heuristic may be more advantageous than stating true valuations. For instance, in bargaining over the price of a used car or negotiating a pay rise. Furthermore, Heifetz and Segev (2004) argue such a heuristic may also be evolutionary viable. However, in experiments designed to make revealing true preferences optimal, using the heuristic results in costly errors.

Another heuristic may also explain the evolution of WTA and WTP in repeated markets. The convergence of WTA and WTP in many of the repeated market experiments could be caused by subjects adjusting their bids towards observed market prices. Loomes et al. (2003) call this effect *shaping*. This could explain the surprising result obtained by Knetsch et al. (2001). They found that the WTA/WTP gap closed in repeated second price buying and selling Vickrey auctions but widened in repeated ninth price Vickrey auctions. This results can be interpreted in terms of shaping. When second price auctions are used, the buying price is drawn from the top end of the distribution of bids (second highest) and the selling price from the bottom end (second lowest). Accordingly, if subjects adjust their bids towards the observed market price, they will tend to increase their bids in buying auctions but decrease them in selling auctions. In contrast, when the ninth price versions of the auctions are used with ten bidders, buying prices are drawn from the bottom end of the distribution of bids and selling prices from the top end, so the reverse effect occurs. Furthermore, in experiment designed to test for shaping, Loomes et al. find further evidence that shaping effects do occur.

1.11 Summary

The literature to date establishes the following. (a) There is an endowment effect: evidence comes from survey, experimental, and field studies. (b) The effect is not uniform, it varies with the good used and class of decision problem. (c) The endowment effect often but not always decays in repeated experimental markets. (d) People with relatively intense market experience gained in naturally occurring markets tend not to exhibit the endowment effect; people with less intense market experience do exhibit an endowment effect.

The following questions remain to be answered. (a) What aspects of the experience of repeated markets causes the decay in the WTA/WTP gap. (b) Why some experience outside the lab causes people not to exhibit the endowment effect, and why not everyone has this experience. (c) Why if experience of naturally occurring markets can eliminate the endowment effect do we still observe the effect in some field data.

1.12 Outline of the Thesis

This thesis presents three experimental studies that investigate how experience affects loss aversion.

The second chapter reports an extended version of Knetsch's exchange of goods experiment to explore how different types of experience influence the endowment effect. The experiment has four treatments, which compare the behaviour of subjects with experience of consuming, owning, and choosing goods to a control group. The results are consistent with earlier studies in that an endowment effect is observed; however, the strength of the effect is less than in earlier studies and differs between treatments. In particular, there is a significantly stronger endowment effect in the treatments in which the endowment is acquired in two steps rather than one step.

The third chapter reports a repeated market experiment in which subjects buy and sell lotteries under symmetric and asymmetric information. Buying and selling bids and prices are compared. A gap between buying and selling prices decays under symmetric information but persists under asymmetric information. Furthermore, there are spillover effects. When the regime switches between symmetric and asymmetric information, subjects do not immediately adjust their behaviour.

The fourth chapter reports another repeated market experiment in which subjects buy and sell lotteries. The information subjects receive about the lotteries' odds and whether the lottery gets resolved after each trial varies between treatments. Among the findings is that the gap between buying and selling bids decays when lotteries are not resolved each round but persists when they are.

The fifth and final chapter summarises the findings of the three experiments, identifies common patterns, and concludes.

Appendix A

Preference Reversals in Expected Utility of Income Models

Cox and Sadiraj (2006) use the following utility function to show that an expected utility of income model can explain risk aversion over modest stakes without implying absurd risk aversion over large stakes.

$$\mu(y) = \begin{cases} 0.9y + 0.1 & y < 1 \\ y^{0.9} & y \geq 1 \end{cases} \tag{A.1}$$

where y is income (i.e. losses or gains in wealth). Figure A.1 shows a plot of the utility function.

A.1 Reference Point Induced Preference Reversals

Imagine a person whose preferences are captured by the equation A.1 is given £200. Then they are ask to choose between options A and B from table A.1. The sure gain of £49 gives them a higher expected utility, so they choose A. What if instead they had been given £400 and asked to choose between C and D? The risky option,

Figure A.1: Cox-Sadiraj Expected Utility of Income Function

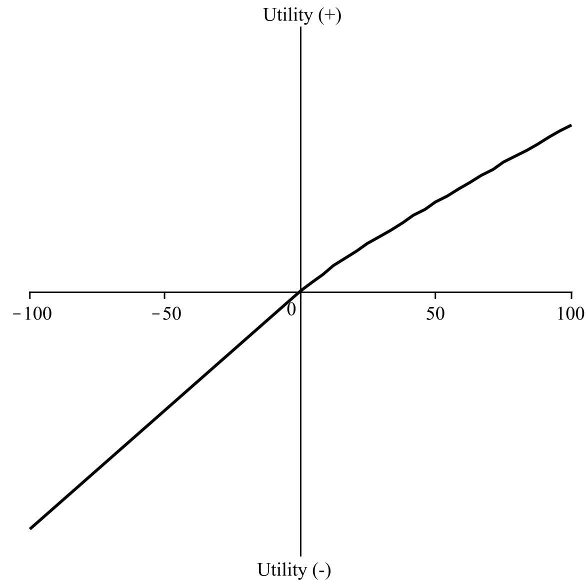

D, gives a higher expected utility, so they choose it. Notice, however, that option A gives the same final wealth as C; B gives the same as D. Changing the reference point from £200 to £400 has reversed their preference over the two outcomes defined in terms of wealth.[1]

A.2 Reference-Dependent Valuations

Bateman et al. (1997) define four measures of value: willingness to accept (WTA), willingness to pay (WTP), equivalent gain (EG), and equivalent loss (EL). They argue that the standard theory while not predicting equality between WTA and WTP does predict EG = WTA and EL = WTP. Bateman et al. write

> If losses loom larger than gains, then, in the absence of income effects, we should expect WTA_{ji} to be greater than WTP_{ji}. EG_{ji}, which expresses an equivalence between gains on the two dimensions, and EL_{ji}

[1]This example is a slight modification of one Kahneman and Tversky (1982) use. It shows how the same decision problem can be framed in different ways and how different frames can lead people to choose different options.

Table A.1: The Effect of Framing on Decisions

Endowment	Option	Income	Expected Utility	Final Wealth
£200	A	sure gain of £49	33.2	£249
£200	B	25% gain £200; 75% gain nothing	29.5	25% chance £400; 75% chance £200
£400	C	sure loss of £151	−135.8	£249
£400	D	75% lose £200; 25% lose nothing	−134.9	25% chance £400; 75% chance £200

The *Expected Utility* figures are calculated as $\sum_{i=1}^{n} p_i \mu(y_i)$ where μ is the utility function in equation A.1 and the income y_i and probability values p_i are whose specified in the *Income* column. *Final Wealth* is the sum of *Endowment* and *Income*.

Table A.2: How Lottery Valuations Depend on the Reference Point

	WTA	EL	EG	WTP
Equation	$\sum_{i=1}^{n} p_i \mu(WTA - y_i) = \mu(0)$	$\sum_{i=1}^{n} p_i \mu(-y_i) = \mu(-EL)$	$\sum_{i=1}^{n} p_i \mu(y_i) = \mu(EG)$	$\sum_{i=1}^{n} p_i \mu(y_i - WTP) = \mu(0)$
$-bet	478.2	323.0	268.7	189.3
P-bet	349.3	324.0	316.5	299.6

The $-bet is $\{1700, 0.19; 0, 0.81\}$; the P-bet, $\{0, 0.19; 400, 0.81\}$. The utility function μ is the one in equation A.1, n is the number of possible outcomes ($n = 2$ for the $ and P bets), p_i is the probability of the i^{th} outcome occurring and y_i is the income from the outcome. The reported figures are solutions to the equations for the relevant bet rounded to one decimal place. The values were calculated using the software package Maple version 11.

which expresses an equivalence between losses, should be expected to take values intermediate between WTP_{ji} and WTA_{ji}.

Table A.2 shows the result of using Cox and Sadiraj's example utility of income function to calculate the four measures for two lotteries. Notice that for both lotteries (a) WTA exceeds WTP, (b) both $EG \neq WTA$ and $EL \neq WTP$, and (c) the relative size of the four measures fits the pattern predicted by loss aversion.

What are the consequences of the differences in valuations across the four measures? Imagine a variation of Knetsch's (1989) chocolate and mugs experiment. One group of subjects is given a $-bet and then asked if the want to swap it for an amount of money, 400. Since for the $-bet, $WTA > 400$ they will refuse the trade. But what if they had been endowed with 400 instead of the lottery, would they swap it for the lottery? The WTP figure for the $-bet is $WTP < 400$, so they would not be willing to trade. We would observe an endowment effect.

A.3 Choice-Valuation Preference Reversals

Suppose the value of the P and $ bets in table A.2 are elicited as WTA valuations. The figure for the $-bet is higher, suggesting the $-bet is preferred. Now suppose the P and $ bets are offered in a straight choice, i.e. we have a choice between gaining one or the other. From the table we see that EG figure is higher for the P-bet suggesting that in a straight choice, the P-bet is preferred. The preference we inferred from the WTA valuations is reversed.

A.4 Summary

Cox and Sadiraj (2006) argue that income and wealth can enter the utility function used in Expected Utility Theory separately. I demonstrate that if they do enter separately, the resulting theory can predict preference reversals that are inconsistent with standard rational model, therefore, Expected Utility of Income is not a model of standard rational behaviour. But if under the standard model, income and wealth can not enter the utility function separately, then Rabin's critique applies and risk aversion over modest stakes is an anomaly for the standard theory.

Chapter 2

How Experience of Tasting, Choosing and Owning Goods Influence the Endowment Effect: The Crisps and Lemonade Experiment

2.1 Introduction

The standard theory of preferences used in economics suggests there should be no endowment effect. Thaler (1980) introduced the term ‘endowment effect’ to describe the tendency to under weigh opportunity costs relative to out of pocket costs. It is so called because people seem to value a good more when it is incorporated into their endowment. He illustrates the concept with several examples including the following.

> Mr. H mows his own lawn. His neighbour’s son would mow it for \$8. He wouldn’t mow his neighbour’s same-sized lawn for \$20.

In this example, the time Mr. H would gain (the opportunity cost) by not mowing his own lawn is worth less to him than $8 (a potential out of pocket cost), but the time he would lose (a potential out of pocket cost) if he mowed his neighbour's lawn is worth more to him than $20 (the opportunity cost). This phenomenon has been repeatedly observed in experimental studies of choice and valuation (see chapter 1 for a summary of the studies). For example, Knetsch (1989) conducted an experiment involving the exchange of goods. Subjects were endowed with either (a) a coffee mug or (b) a bar of Swiss chocolate. Then they were asked if they wanted to swap (a) for (b) or vice-versa. Regardless of whether they had been given (a) or (b), they tended not to swap. It appears the loss of giving up an entitlement weighs more heavily than the forgone gains of not obtaining an alternative entitlement. The endowment effect has also been observed in valuation studies. People tend to assign higher value to items when they own the items than when they do not. For example, Bateman et al. (1997) elicited preferences between money and several consumption goods using four valuation measures. The valuation figures they obtained suggested losses carry more weight than gains. In both exchange of goods experiments and valuation studies, the standard theory, as we will see in section 2.2.1, predicts there will be no endowment effect.

At this stage, it might seem one has good reason to reject the standard theory. Some studies, however, have suggested that the endowment effect may be a transient phenomenon rather than a universal feature of choice and valuation. Plott and Zeiler (2005) investigated the reported gap between 'willingness to pay' and 'willingness to accept' valuations. They found that if certain controls (an incentive compatible elicitation device, training, paid practice, and anonymity) were used together, then a gap was not observed. Likewise, many of the other repeated market experiments described in chapter 1 found that the endowment effect decayed in repeated markets.

What is the significance of the apparent discrepancy between the findings described above? First, the endowment effect and the related concept, loss-aversion have been applied to problems in numerous branches of economics. For example, Genesove and Mayer (2001) studied the Boston housing market in the 1990s and found that loss aversion determines seller behaviour. Bowman et al. (1999) propose a model of consumption and saving incorporating loss aversion and present empirical evidence from five countries that supports the novel predictions of their model. If loss aversion only affects naive or inexperienced decision makers, there are clearly implications for studies that explicitly invoke loss aversion such as the ones noted. Second, the converse is also true. The status of loss aversion also has implications for studies that invoke the assumptions of standard theory that are contradicted by loss aversion. Third, there are also implications that are less obvious. Rabin (2000) has forcefully argued that expected utility theory with any concave utility function cannot provide a satisfactory empirical theory of choice under risk because calibrating expected utility to what we know about risk aversion over modest stakes then implies absurd degrees of risk aversion over larger stakes. If we accept Rabin's argument, we need an explanation of risk aversion over modest stakes and loss aversion is a strong contender.[1] The problem is that if loss aversion is eliminated by experience, then we might expect risk aversion over modest stakes to disappear too. If so, then there are implications for studies involving the concept of risk aversion. Fourth, the status of the endowment effect also has policy implications such as assessing the welfare implications of projects that have an environmental aspect, e.g.

[1] Regret Theory (Loomes and Sugden, 1982) could also explain risk aversion over small stakes. Consider the option to play a 50-50 win £110, lose £100 gamble. Suppose that if you do not play the gamble, you are not told what it would have paid out. If you play the gamble and win, you derive utility from the extra £110 plus you gain pleasure from having made a decision that paid off (rejoice). If you play the gamble and lose, you lose utility from giving up the £100 plus you lose utility due to the feeling of regret following making a decision that made you worse off than you would otherwise have been. If the potential regret outweighs the potential rejoice, you have potentially a reason to turn down the gamble even if the utility loss from losing £100 is outweighed by the utility gain of winning £110.

building a hydroelectric plant that will involve destruction of some natural habitat. Contingent valuation studies, discussed in chapter 1, find that the value people assign to environmental goods depends on whether they are paying for the good or being compensated to give it up. If we understood how experience influences the extent to which current entitlements affect preferences, then it may be possible to develop elicitation procedures that minimise any biases caused by the endowment effect.

2.2 Theory

2.2.1 Neoclassical Preferences

Neoclassical Theory (NT) predicts that when faced with the choice of whether or not to swap one entitlement for another, a person picks the option that gives the best outcome (the one that gives the highest utility). It is irrelevant which of the two entitlements the person possessed when making the choice.

A swap decision can be modelled formally as follows[2]. There is a choice set $X = \{x, y, ...\}$. Each item, $x = (x_1, x_2)$ is a bundle of two goods. Bundle x comprises x_1 units of good 1 and x_2 units of good 2, and $x_1, x_2 \geq 0$. Preferences over X are complete, reflexive, and transitive, and are captured by the utility function $u : X \longmapsto \mathbb{R}$ ($\mathbb{R}$ is the set of real numbers). A theory of choice applied to a swap decision can either (a) predict a person will swap, (b) predict they will stick with their endowment, or (c) make no prediction. We can write the swap function for NT, f_{NT}, that maps an endowment $e \in X$ and alternative entitlement $a \in X$ to one of these three outcomes, $f_{NT} : X \times X \mapsto \{\text{swaps, indeterminate, sticks}\}$, as

[2]The same setup is used by Tversky and Kahneman (1991) to explain their model of loss aversion in riskless choice

$$f_{NT}(e,a) = \begin{cases} swaps & for\ u(a) > u(e) \\ indeterminate & for\ u(a) = u(e) \\ sticks & for\ u(a) < u(e) \end{cases} \quad (2.1)$$

The equation says that a person swaps when the alternative allocation gives a higher utility than the endowment.

2.2.2 Reference-Dependent Preferences

Kahneman and Tversky's (1979) Original Prospect Theory (OPT) was presented as a theory of choice under risk not a theory of riskless choice, but it is significant because it was the first theory that formally incorporated loss aversion. What does OPT predict will occur in the experiments described in the introduction? If it is changes in welfare that determine decisions, then whether a person swaps one entitlement for another surely depends on whether their welfare after swapping will be greater than their welfare before. Hence, OPT appears to predict the same behaviour as NT.

In order to explain the endowment effect, Tversky and Kahneman (1991) presented an extended version of Prospect Theory with two (or more) dimensions of value. This allows a potential swap to be modelled as a loss on the dimension of the item being given up and a gain on the dimension of the item being received.

A simple way to model reference dependent preferences is as follows. Preferences over gains are captured by a utility function as they are under NT. Consider the two good case described in section 2.2.1. There are three bundles of goods $x, e, a \in X$ which are illustrated in figure 2.1. A person endowed with allocation x treats x as their reference point. Suppose they are then faced with a choice between e and a. Both e and a offer more of good 1 and more of good 2 than bundle x, neither option involves giving up goods relative to the reference point. Accordingly, they choose the option that gives the highest utility. For instance, they will choose e if

Figure 2.1: Modelling an Exchange as a Loss and a Gain

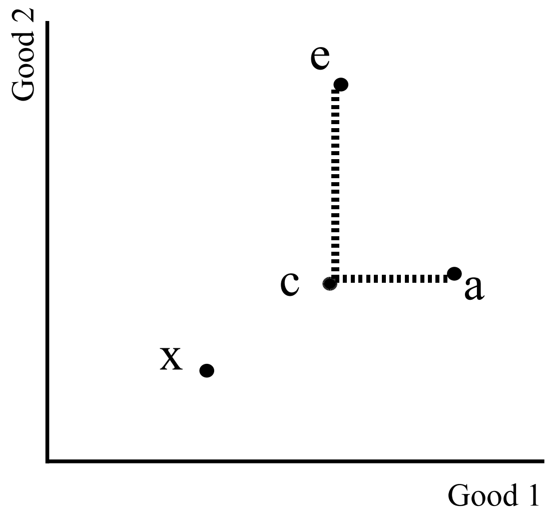

$u(e) > u(a)$.

Now suppose they are endowed with e and have the option to swap their position to a. Moving from e to a involves a loss on the dimension of good 2 represented by $e_2 - c_2$ and a gain in good 1 represented by $a_1 - c_1$. Note, $c \in X$ is the bundle of goods that is common to endowment e and the alternative entitlement a. That is, $c = (min(a_1, e_1), min(a_2, e_2))$. To model the choice, the swap is broken down into two stages: giving up part of the endowment which causes a loss in utility and gaining the alternative entitlement which causes a gain. The loss is $u(e) - u(c)$ and the gain is $u(a) - u(c)$. The person swaps if $\lambda\,(u(e) - u(c)) < u(a) - u(c)$. The constant $\lambda \geq 1$ is the coefficient of loss aversion[3]. It measures the weight assigned to losses relative to gains when weighing options.

As done for NT, let f_{RDT} be the swap function for Reference-Dependent Theory (RDT) that maps an endowment $e \in X$ and alternative entitlement $a \in X$ to one of three outcomes, $f_{RDT} : X \times X \mapsto \{\text{swaps, indeterminate, sticks}\}$. First, let κ be

[3]The coefficient of loss aversion is related to the steepness of the value function for losses relative to gains. Appendix D shows how under original prospect theory loss aversion follows from aversion to 50-50 bets and its implications.

gain/loss ratio of utilities, i.e.

$$\kappa = \frac{u(a) - u(c)}{u(e) - u(c)} \tag{2.2}$$

then the swap function is

$$f_{RDT}(e,a) = \begin{cases} swaps & \kappa > \lambda \\ indeterminate & \kappa = \lambda \\ sticks & \kappa < \lambda \end{cases} \tag{2.3}$$

The equation says that a person swaps when the gain/loss ratio exceeds the coefficient of loss aversion. NT is the special case where $\lambda = 1$.

2.3 The Experiment

2.3.1 Motivation

The experiments carried out by List (2004a), discussed in chapter 1, and Plott and Zeiler (2005) suggest experience plays some role in determining whether an endowment effect is observed. These studies do not, however, reveal what aspect or aspects of experiences matter. The following design aims to address this question. The design is similar to Knetsch's (1989) mugs and chocolate experiment in that subjects are randomly given one of two allocations, then have the option to swap. The novel features of the design is that the subjects are given different types of experience (consuming, owning, and choosing goods) before the random allocation, and that bundles of goods rather than single units of goods are used (the motivation for using bundles is explained below).

Why might these types of experience matter? First, consider experience of consuming. Suppose that the endowment effect arises when people are uncertain about the nature of goods. As a rule of thumb, *don't trade items when you don't* *know their value* could be a good defence against being exploited by better informed traders. Experience of consuming the goods gives you a better idea of their value and, all else equal, could lead to higher trading rates. In addition, consuming the goods might reveal that one of them was better or worse than expected, increasing the likelihood of that good being favoured in spite of an endowment effect. Conversely, consuming might reveal that the difference in attractiveness between the goods was in fact less than expected, thus decreasing trading rates when combined with an endowment effect.[4] Finally, the act of consuming could cause a person to decide which of two items they prefer. This could reduce the endowment effect if when faced with a swap decision they recall what they preferred before the endowment occurred.

Second, if you have already chosen between two items before being endowed with one of them, it is plausible that you could avoid the influence of the endowment effect by recalling your earlier choice, so experience of choosing would increase trading. In addition, if you have made a choice between two particular goods before, you might feel a psychological commitment to this prior choice, lessening the influence of the endowment effect. In the psychology literature, there are many studies investigating cognitive dissonance. The theory was developed by Festinger (1957). The first experimental study that found evidence[5] of the effect was run by Brehm (1956), who summarises the theory as follows. After making a choice, a person has *consonant* cognitive items (e.g. items of information) that favour the chosen alternative and *dissonant* ones that favour the unchosen alternative. All else equal, the greater the number of *dissonant* items (i.e. the greater the relative attractive-

[4] Coursey et al. (1987) ran an experiment investigating the WTA/WTP gap. For a commodity, they used a bitter tasting liquid, sucrose octa-acetate (SOA). Subjects either paid to avoid or were compensated for holding a one ounce cup of SOA in the mouth for 20 seconds. Subjects stated valuations before and after tasting a small sample. The WTA/WTP gap widened after subjects tasted the sample.

[5] Chen (2008) has recently questioned the methodology used by Brehm (1956) and many later studies. He argues the methodology involves a mistake equivalent to the one people commonly make in the three-door (or Monty- Hall) problem.

ness of the unchosen item), the greater the level of *dissonance* a person experiences. When dissonance exists, a person will attempt to reduce or eliminate it. They do this by making the chosen alternative more desirable and the unchosen one less desirable than before the choice. Furthermore, Brehm suggests the dissonance, and so attempts to reduce it, will be greater the closer the alternative choices are to being equally desirable.

If dissonance effects do occur, then a natural question is whether it is the act of choosing which triggers the effects or whether the possibility of another outcome occurring triggers them. That is, if a person is assigned one of two goods at random and they know it was a random process that determined the assignment, will it create the same dissonance effects as would have occurred if they had chosen one of the two goods?

2.3.2 Design

The experiment[6] is a modified version of Knetsch's (1989) classic exchange of goods experiment. Instead of using chocolate bars and coffee mugs, two edible goods a and b of approximately equal value were used. Subjects were randomly divided into four groups: a control group, tasters, choosers, and owners. All of the treatments had the following basic structure consisting of three stages. In the first stage, subjects (except those in the control group) gained different types of experience. In some treatments subjects gained a unit of one of the goods in this stage, in other treatments they did not. In the second stage, the subjects received a random endowment. If they had not received anything in stage one, they received aab[7] or

[6] A version of the design described here was first used in a smaller experiment (Lindsay, 2004). The experiment had two treatments, which correspond to the *choosers* and *control group*. Chocolate bars and cans of coke were used as the two goods. There were 63 subjects: 33 in the control group and 30 in the choosers group. In the control group the swap rate was 0.30; in the choosers group it was 0.47. The endowment effect was statistically significant in the control group ($p = 0.02$) but not in the choosers group.

[7] aab denotes two units of good a and one unit of good b.

abb with equal probability. If they had received a unit of good *a* in stage one, they received *ab* or *bb* with equal probability in addition to what they already had. If they had received a unit of good *b* in stage one, they received *aa* or *ab* with equal probability in addition to what they already had. As a consequence, in all treatments, subjects ended stage two either holding *aab* or *abb*. Furthermore, which of the two positions they held was independent of which of the two goods, if any, they had obtained in stage one. Finally, in stage three, those holding *aab* had the option to swap to holding *abb*; those holding *abb* had the option to swap to *aab*. The differences between treatments are described below with reference to figure 2.2.

Control Group The control treatment followed the left hand tree on figure 2.2. They did not do anything in the experience stage (stage one) so began the experiment at the node labelled *R*. They were randomly endowed with either *aab* or *abb* (stage two). This took them to one of the nodes labelled *P*. Then they had the option to swap the goods they were holding for the alternative (stage three). That is swap *aab* for *abb* or swap *abb* for *aab.* At the end of the experiment, their final allocation of goods was either *aab* or *abb* depending on the choice they made at stage three.

Tasters The tasters treatment also followed the left hand tree on figure 2.2. In stage one, at the point marked *x* the subjects tasted a small sample of the two goods. Then they followed the same sequence as the control group: they received a random endowment (stage two) and then had the option to swap it for the alternative (stage three).

Choosers The choosers treatment followed the right hand tree on figure 2.2. In stage one, at the node labelled *y* they made a choice between one unit of good *a* and one unit of *b*. They received the good they chose, taking them to one of the nodes labelled *R*. Note, at the moment subjects made the choice they knew the item they were choosing would be theirs to keep but they did not know what was going to happen in the following stages. In stage two they received a random endowment, which *topped up* the allocation of goods they held at the end of stage one to either *aab* or *abb.* If they had chosen *a* in stage one, their allocation was topped up with either *ab* or *bb.* If they had chosen *b* in stage one, their allocation was topped up with either *ab* or *aa.* As a result, at the end of stage two, they either held *aab* or *abb.* Finally, as with the other treatments, they had the option to swap the goods they were holding for the alternative (stage three).

Owners The owners treatment also followed the right hand tree on figure 2.2. The treatment was the same as the choosers treatment except for the following feature. In stage one, at the node labelled *y* instead of making a choice between one unit of good *a* and one unit of *b*, they were randomly assigned either one unit of good *a* and one unit of *b*. They received the good they were assigned, taking them to one of the nodes labelled *R*, and proceeded following the same scheme as the choosers.

For all treatments, the goods a subject holds when they make the decision to stick or swap is independent of what happened in the first stage. The right hand inverted tree shows the scheme for the choosers and owners. They acquire their endowment in two steps in contrast to the tasters and control group who acquire it in one step. Although the choosers and owners acquire the endowment in two steps, when they are faced with the swap or stick decision, they are choosing between the same two final allocations as the subjects in the one-step treatments. What is more, even in the two step treatments, whether the subject has *aab* or *abb* when they are faced with the stick or swap decision is solely determined by the output of the random device. That is, their endowment after the second step is independent of what happened in the first step. This feature of the design is made possible by using bundles of

Figure 2.2: One-Step and Two-Step Treatments

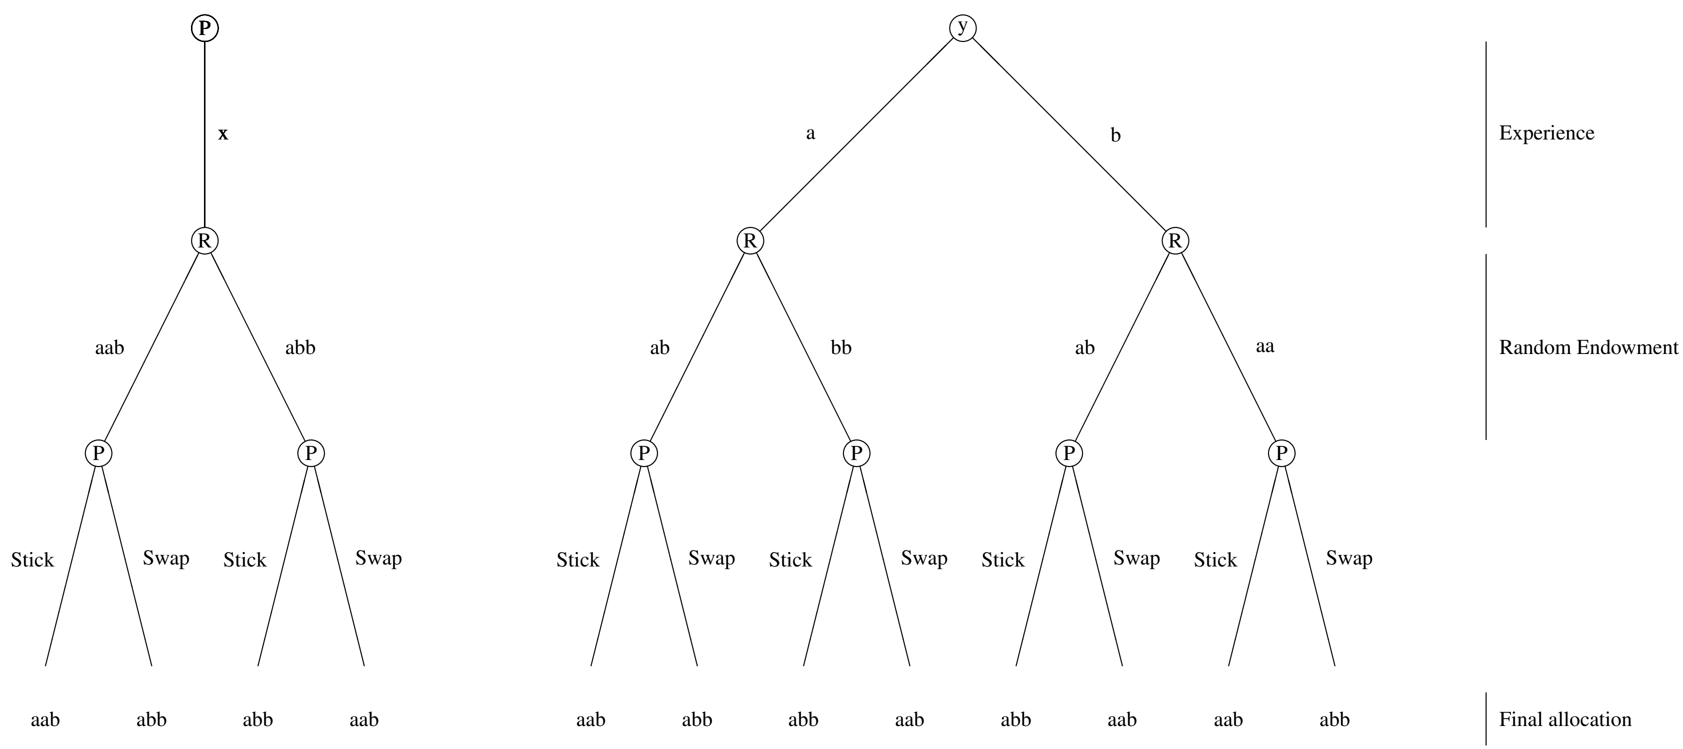

goods and is desirable for two reasons. First, it means that NT and RDT predict different swap rates. This is because when the endowment is assigned at random, NT predicts 50 percent will swap and RDT predicts less than 50 percent will do so (this is explained in more detail below). Second, it allows meaningful comparisons between treatments. The difference between the owners and choosers occurs at the first node, marked y. For the owners, the random device determines whether they receive a unit of *a* or *b*; for the choosers, the subject chooses a unit of *a* or *b*. The difference between the control group and the tasters occurs at the point marked *x* on the left hand tree. The control group do nothing; the tasters taste a sample of both goods.

These four treatments allow several comparisons. Comparing the swap rates for the tasters and control group will give a sharp test of whether experience of consuming (tasting) influences the endowment effect. Comparing the swap rates for the owners and choosers will give a sharp test of whether experience of choosing rather than merely owning influences the endowment effect.

What do the two theories predict? Under NT, preferences are independent of the random endowment. Let the proportion of subjects preferring *abb* to *aab* be $q \in [0,1]$. The probability of being endowed with *abb* is $\frac{1}{2}$. The proportion of subjects whose endowment is their preferred entitlement consists of those who prefer *abb* and received *abb* plus those who prefer *aab* and received *aab*. Assuming none of the subjects are indifferent between *abb* and *aab,* then the proportion who prefer *aab* to *abb* will be $1-q$.

$$\begin{aligned}\text{Proportion swapping} &= 1-\frac{1}{2}q-\frac{1}{2}(1-q)\\ &= \frac{1}{2}\end{aligned}$$

This means that 50 percent will not receive their preferred entitlement, so we would expect a swap rate of 50 percent. None of the differences between treatments should affect the swap rate. Note, that experience of consuming might change the value of q, since subjects expectation of the taste of the goods might be inaccurate; however, the swap rate is independent of q, so experience of consuming should have no effect on the swap rate.

Under RDT things are slightly more complex. Recall from section 2.2.2 that a subject endowed with *abb* will swap if $\kappa > \lambda$. Conversely a subject endowed with *aab* will swap if $\frac{1}{\kappa} < \frac{1}{\lambda}$. Hence, the proportion swapping will be

$$\begin{aligned}\text{Proportion swapping} &= \Pr(\kappa > \lambda).\Pr(abb) + \Pr(\frac{1}{\kappa} < \frac{1}{\lambda}).\Pr(aab) \\ &= \frac{1}{2}\left[\Pr(\kappa > \lambda) + \Pr(\frac{1}{\kappa} < \frac{1}{\lambda})\right]\end{aligned}$$

If we assume κ is the same for all subjects, then the swap rate will be zero if $\frac{1}{\lambda} < \kappa < \lambda$ and $\frac{1}{2}$ if $\frac{1}{\lambda} > \kappa$ or $\kappa > \lambda$. Now let us assume κ varies between subjects and is distributed with mean μ and variance σ^2 according to a distribution whose density function is $f(\cdot)$ and whose distribution function is $F(\cdot)$.

$$\begin{aligned}\Pr(\frac{1}{\lambda} < \kappa < \lambda) &= F(\lambda) - F(\frac{1}{\lambda}) \\ \text{Proportion swapping} &= \frac{1}{2}\left[1 - F(\lambda) + F(\frac{1}{\lambda})\right]\end{aligned}$$

If $\lambda > 1$, then $1 - F(\lambda) + F(\frac{1}{\lambda}) < 1$. It follows that the swap rate under RDT will be less than $\frac{1}{2}$. Further, it follows that factors that affect μ and σ^2 can influence the swap rate even when λ is constant. So RDT does not rule out experience of consuming influencing the swap rate if the experience alters μ and σ^2. It is plausible that the relative attractiveness of the two goods would change following tasting them. Hence tasting could alter the swap rate.

2.3.3 Experimental Procedure

Subjects were recruited by email from a database of pre-registered volunteers from the University of Nottingham student population. A total of 210 subjects took part in the experiment. The two goods used were premium organic vegetable crisps (100g packets) and handmade organic lemonade (75cl bottles). A photograph of the crisps and lemonade is in appendix F. The unit retail price for both items was approximately £2. The advantages of this pair of goods are as follows. (a) They are edible, so can be tasted. (b) They were bought from specialist wholesalers and were not available in local supermarkets, so most subjects are unlikely to have tasted them before, so letting one group taste them will produce a real difference in the experience of the goods between the subjects. (c) Since they were not available in local supermarkets, this meant subjects were also uncertain of the retail price. Furthermore, (d) subjects are unlikely to be intending to purchase either of the goods in the near future, making them unlikely to view gains in either of the goods as merely money saved.

As well as recording each subject's treatment, endowment and swap decision, some survey data was collected. Basic demographic information was collected between the experience stage and the stage where the random endowment was assigned or (in the case of choosers and owners) when the endowment was topped up. This serves two purposes, first, it is of interest whether different demographic groups behave differently. There is already some evidence that demographics matter for risk aversion[8], maybe the same is true for loss aversion. Second, it allows the owners and choosers to hold the first part of their endowment for a period of time before

[8] For example Dohmen et al. (2005) find heterogeneity in individual risk attitudes and that willingness to take risks is negatively correlated with age and being female.

they receive additional items. After subjects received their random endowment they completed an attitude to risk survey. Again, this served two purposes, first to let them keep the items they had been given for a period before they decided whether or not to swap. Second, it was used to collect additional data.

The attitude to risk survey aimed to assess subject's attitude to risk by asking them to indicate whether they would accept or reject a series of hypothetical[9] bets and to specify the size of the prize that would be sufficient to tempt them to play another series of bets. The motivation for this task is that, as noted in section 2.1 above, Rabin (2000) has argued that risk aversion over modest stakes may be explained by loss aversion. If this is true and loss aversion also explains the endowment effect, then we might expect risk aversion and a propensity to exhibit the endowment effect to be positively correlated.

At the end of the experiment the subjects completed a debrief survey. Here, among other things, they were ask to specify a hypothetical willingness to pay figure for one of the goods. This is used to shed light on whether tasting affects how subjects value the two goods.

Subjects were invited to come at 5 minute intervals so there was a stream of subjects who completed the experiment one at a time. The experiment was run by two people: experimenter A and experimenter B.

The experiment was run in a room divided in to 3 areas in such a way that those in one area could not see what people in other areas were doing. Upon arrival subjects were instructed to take a seat in Area 1 and read through a set of paper instructions for the experiment (Appendix C.1). The stages completed by subjects in each of the treatments are shown in table 2.1. Experimenter A called subjects one-at-a-time into Area 2 and assigned the subject at random to one of the four groups, then ran

[9]Given that the survey involves hypothetical payments, the results should be treated with some caution. Subjects might behave differently when facing real and hypothetical payments. For instance, Holt and Laury (2002, 2005) have found that there is no significant effect on the level of risk aversion from increasing the stakes when payments are hypothetical but there is when payments are real.

through stages 1.a to 2.b of the experiment (the script is in Appendix B.1). The demographic survey and the two attitude to risk surveys are shown in Appendices C.2, C.4, and C.3. After they had completed the attitude to risk surveys, the subject carried their goods to area 3. Here, experimenter B ran through stages 3.a to 3.c of the experiment with the subject. The option to swap was framed as follows. The experimenter read one of the statements depending on which combination of goods the subject possessed.

i You have two packets of crisps and one bottle of lemonade. If you want, you can swap one of your packets of crisps for a bottle of lemonade.

ii You have two bottles of lemonade and one packet of crisps. If you want, you can swap one of your bottles of lemonade for a packet of crisps.

The experimenter swapped the item if the subject wished to swap. The experimenter recorded the swap decision, then asked the subject to complete the debrief survey and receipt form. The debrief survey is shown in appendix C.5. Finally the experimenter gave the subject a plastic bag to carry their items. The plastic bags reduced the risk of other subjects seeing what the subject leaving the experiment had been given.

2.4 Primary Results

This section analyses the choices the 210 subjects made. A summary of the participant's demographics and further analysis of the survey responses is in section 2.6.

2.4.1 Pooled across Treatments

Result 2.1 *Across the four treatments subjects tended to stick with their endowment.*

Table 2.1: The Procedure for Each of the Treatments

Stage	Treatments			
	1-Step		2-Step	
	Control	Taster	Owners	Choosers
1.a		Taste items	Allocated item	Choose item
.b	Complete demographics survey			
2.a	Allocated items		Allocation topped-up	
.b	Complete attitude to risk surveys			
3.a	Given option to swap			
.b	Complete debrief survey			
.c	Complete receipt			

The table shows the different steps completed by subjects in each of the four treatments.

Table 2.2: Endowment Against Final Allocation for Pooled Data

	Final Allocation		
	Lemonade Rich	Crisps Rich	Total
Endowment			
Lemonade Rich	73	37	110
Crisps Rich	40	60	100
Total	113	97	210

Lemonade Rich means two bottles of lemonade and one packet of crisps. *Crisps Rich* means two packets of crisps and one bottle of lemonade. *Endowment* is what the subject held before deciding whether or not to swap. *Final Allocation* is what they held after the decision.

Support. Table 2.2 shows what subjects were endowed with and their final allocation at the end of the experiment. We can see that the majority of subjects endowed with lemonade $\frac{73}{110}$ stuck with lemonade whereas the majority endowed with crisps $\frac{60}{100}$ stuck with crisps. Running a 1-sided Fisher's Exact test gives a p-value of less than 0.001 of obtaining these results were final allocation independent of endowment. NT predicts that endowment and final allocation will be independent whereas RDT predicts that subjects will tend to stick with their endowment. This is strong evidence in favour of some variant of RDT over NT.

2.4.2 Comparisons between Treatments

RDT does not predict that there will be differences between treatments, but allows experience of consuming to alter the swap rate. Table 2.3 shows the swap rates broken down by treatment.

Result 2.2 *There was little or no endowment effect in the control group.*

Support. Table 2.3 shows the swap rate for the control group was 0.42 and that this is not statistically different from 0.5, the swap rate predicted by the standard theory.

Result 2.3 *Subjects who acquired their endowment in two steps are considerably less likely to swap than those who acquired it in one step.*

Support. Table 2.3 shows that the swap rate for treatments in which subjects acquired their endowment in one step is 0.44 compared to 0.29 for treatments where the endowment was acquired in two steps. This means that acquiring the endowment in one step made subjects $(\frac{0.44}{0.28} - 1) \times 100 = 57\%$ more likely to trade. Table 2.4 shows the difference in trading is statistically significant.

Result 2.4 *There is no evidence of tasting influencing swap rates.*

Support. The swap rate for the control group is 0.42 compared to 0.46 for the tasters. The size of the difference is small and is not statistically significant.

Result 2.5 *There is evidence that experience of choosing between goods decreases the endowment effect relative to experience of owning.*

Support. The swap rate for the owners is 0.23 compared to 0.35 for the choosers. This means that experience of choosing between the goods made subjects $(\frac{0.35}{0.23} - 1) \times 100 = 52\%$ more likely to trade. Table 2.4 reveals, however, that this difference

Table 2.3: Swap Rate by Treatment

Treatment		n	Swaps	Swap Rate	P-Value for Fisher's Exact Test
1-step	Control	50	21	0.42	0.198
	Taster	56	26	0.46	0.418
	Total	106	47	0.44	0.170
2-step	Owner	52	12	0.23	0.000
	Chooser	52	18	0.35	0.029
	Total	104	30	0.29	0.000
All		210	77	0.37	0.000

P-values are from testing the null hypothesis that endowment and final allocations are independent against the one sided alternative that subjects tend not to swap.

Table 2.4: Testing Hypotheses about Differences in Swap Rates between Treatments

H_0	H_1	P-Value
$S_{control} = S_{Taster}$	$S_{control} < S_{Taster}$	0.397
$S_{Owner} = S_{Chooser}$	$S_{Owner} < S_{Chooser}$	0.140
$S_{1-Step} = S_{2-Step}$	$S_{1-Step} \neq S_{2-Step}$	0.022
$S_{1-Step} = S_{2-Step}$	$S_{1-Step} > S_{2-Step}$	0.014

is not statistically significant.[10] The higher swap rate among the choosers relative to the owners could be explained in terms of cognitive dissonance as discussed in section 2.3.1. However, an explanation of why the choosers' swap rate is less than the ones in both the tasters treatment and the control group is needed.

2.4.3 Comparisons within Treatments

Why are the swap rates for the two step treatments lower? One explanation is that subjects have held part of their endowment for a longer period of time. Perhaps the longer someone has possessed an item the more attached to it they become. If this is true, we would expect that subjects from the owners group who were initially given

[10]The comparison between the owners and choosers is based on 104 observations while the comparison (above) between the one and two step treatments is based on all 210 observations. If the trading rates of the choosers and owners were observed in a sample of 210 observations they would be statistically significant.

Table 2.5: Order in which Items Were Acquired and Swap Rates for the Owners Treatment.

Endowment	Sequence	n	Swaps	Swap Rate
ccl	*c + cl*	12	5	41.7
	l +cc	14	1	7.1
	either	26	6	23.1
cll	*c + ll*	17	4	23.5
	l + cl	9	2	22.2
	either	26	6	23.1
all		52	12	23.1

An endowment of ccl means an endowment of two packets of crisps(c) and one bottle of lemonade(l).
A sequence of c + cl means the subject got a packet of crisps in the first step and then was given an additional packet of crisps and bottle of lemonade in the second step.

c[11] and then had their endowment topped-up to ccl would be less likely to swap to cll than those who were given l and then topped up to ccl.

Result 2.6 *When the endowment was acquired in two steps, the order in which the goods were acquired does not appear to influence the swap rate.*

Support. Table 2.5 shows swap rates against the order in which items were acquired for the owners treatment. Although it appears that sequence may influence swap rate, the variations in swap rate are not statistically significant (a χ^2 test gives a p-value of 0.188 of obtaining these results under the null hypothesis that swap rate is independent of sequence). Hence on the balance of evidence, it appears that it is the acquiring the endowment in several steps, not what is involved in these steps that causes the lower swap rates in the two-step treatments. There are, however, only on average $\frac{52}{4} = 13$ observations per sequence, so the result should be treated as tentative.

[11] c = crisps, l = lemonade

2.4.4 Comparisons with Other Studies

Result 2.7 *Swap rates were much higher than in earlier studies.*

Support. In this study the swap rate in the control group was $\frac{21}{50} = 0.42$ compared to just $\frac{10}{163} = 0.10$ in Knetsch's study.

What explains the high swap rate? There are several possible explanations. First, the most obvious difference between the studies is that Knetsch endowed subjects with one unit of a good whereas I endowed subjects with a bundles of goods. In Knetsch's study, subjects were swapping their only mug or only chocolate bar; in this study, subjects were swapping one of their two packets of crisps or one of their two bottles of lemonade. This suggests loss aversion may be more acute when one is faced with giving up the last unit of a good than when faced with giving up one of several units. Note that if this interpretation is right, it is somewhat surprising since Prospect Theory, where the concept of loss aversion has its origin, suggests that 'the carriers of value are changes in wealth or welfare, rather than final states' (Kahneman and Tversky, 1979).

Second, in this study when subjects swapped, they were gaining an additional unit of a good they already possessed. Perhaps subjects valued the item they were gaining more because they already owned one just like it.

Third, how the subjects were recruited differed between studies. In Knetsch's study, subjects were students attending a class. In this study, subjects were students who had registered themselves in a database of people interested in taking part in paid economics experiments, then they had accepted an email invitation to take part in the experiment. The differences in recruitment method could affect behaviour in two ways. First, the type of people who volunteer to take part in experiments and those who attended Knetsch's class might have different behavioural characteristics. Second, the different recruitment methods could create different expectations among subjects, which in turn might influence behaviour.

Fourth, in this study, there were two experimenters. The first experimenter endowed the subjects with the goods, then the subject carried the goods to the second experimenter, who was separated from the first experimenter by a screen. Perhaps reluctance to exchange is in part driven by a reluctance of subjects to return what they have been given, but this reluctance is reduced by moving to a different location and dealing with a different person. Note that Plott and Zeiler (2005) suggest that *anonymity* is one of the set of controls that together cause the endowment effect not to occur.

Result 2.8 *The swap rate in the control group exceeded that in the choosers group, the reverse of the pattern found by Lindsay 2004.*

Support. In this study, $\frac{21}{50} \approx 0.42$ of the control group swapped; in the earlier study, $\frac{10}{33} \approx 0.30$ of the control group swapped. In this study, $\frac{18}{52} \approx 0.35$ of the choosers group swapped; in the earlier study, $\frac{14}{30} \approx 0.47$ of the choosers group swapped. The difference in the control group swap rates between studies is not statistically significant. The difference in chooser group swap rates, however, is statistically significant.

What could account for these differences between the studies? Differing experimental procedures is one explanation. In this study, subjects in the chooser group made their choice of goods in the presence of one experimenter and later decided whether or not to swap in the presence of another experimenter. In the earlier study, both decisions were made in the presence of the same experimenter. Perhaps in the earlier study, subjects did not want to be seen to make inconsistent choices (e.g. choose a can of coke but then decline to swap if their endowment is topped up to one can of coke and two chocolate bars). In this study however, similar inconsistency in choices would not be apparent to the experimenter recording the swap decision since a different experimenter observed the earlier choice.

2.5 Interpretation

2.5.1 The Two-System Hypothesis (TSH)

Imagine a subject faced with deciding whether to swap a bottle of lemonade for a packet of crisps. There are at least two ways they could approach the decision. First, they could ask themselves whether the pain from giving up the lemonade will be made up for by the pleasure of gaining the crisps. Second, they could assess various aspects of the two goods and ask themselves which is better? E.g. which will be easier to carry home? Which is better for my health? Etc. There is some grounding in the recent literature that decisions are not always made in the same way. For instance, Kahneman (2003) distinguishes between two modes of thinking and deciding: intuition (System 1) and reasoning (System 2).

> The operations of System 1 are fast, automatic, effortless, associative, and difficult to control or modify. The operations of System 2 are slower, serial, effortful, and deliberately controlled; they are also relatively flexible and potentially rule-governed. ... the operating characteristics of System 1 are similar to the features of perceptual processes. On the other hand ...the operations of System 1, like those of System 2, are not restricted to the processing of current stimulation. Intuitive judgements deal with concepts as well as with percepts, and can be evoked by language.

Suppose the endowment effect occurs when subjects use system 1 but not when they use system 2. Further, suppose we can influence which of the two systems a subject uses when deciding whether or not to swap in an experiment. It follows that if we influence subjects differently in each treatment, then we could observe different swap rates in each treatment. Once again, we can define a swap function, this time $f_{ts} : X^2, S \mapsto \{\text{swaps, indeterminate, sticks}\}$ such that

$$f_{ts}(e,a,s) = \begin{cases} swaps & for\ \kappa > v \\ indeterminate & for\ \kappa = v \\ sticks & for\ \kappa < v \end{cases}$$

where $S = \{System\ 1,\ System\ 2\}$ and $s \in S$, and

$$v = \begin{cases} \lambda & for\ s = System\ 1 \\ 1 & for\ s = System\ 2 \end{cases}$$

How might the different experimental treatments influence which system subjects use? Perhaps it is by changing the relative accessibility of the two systems. Kahneman (2003) writes:

> The core idea of prospect theory is that changes and differences are much more accessible than absolute levels of stimulation.

The quote suggests that if we could reduce the accessibility of changes relative to outcomes, we could change how subjects make decisions. Perhaps using bundles of goods rather than single items made subjects more likely to think in terms of outcomes than changes. Likewise, perhaps acquiring an endowment in two steps focused subjects on their endowment and its composition more than if it were acquired in one step, hence making them more likely to view the swap decision in terms of changes than outcomes.

2.5.2 Varying Loss Aversion

Alternatively, we may view the imagined subject's decision as restricted to weighing the pain of the loss against the pleasure of the gain, but allow the intensity of the pain to vary. This could be modelled as λ varying across decision problems. Recall from section 2.2.2 that the swap rate was given by $\frac{1}{2}\left[1-F(\lambda)+F(\frac{1}{\lambda})\right]$. The differences in swap rates between treatments and studies could be explained by differences in experimental procedures affecting the value of λ.

2.5.3 Varying Reference Point

Another explanation for the differences between treatments is that subjects might not always take their current endowment as the reference point. Suppose subjects have reference dependent preferences, but on being endowed with items in the experiment update their reference point to include the items with probability p. The swap rate will be $\frac{1}{2}\left[1-F(\lambda)+F(\frac{1}{\lambda})\right]$ when they update their reference point and $\frac{1}{2}$ when they do not. Hence, the swap rate will be $p.\frac{1}{2}\left[1-F(\lambda)+F(\frac{1}{\lambda})\right]+(1-p).\frac{1}{2}$. Notice that the swap rate is a decreasing function of p. A possible explanation of the lower swap rate among those who acquired their endowment in two steps is that they were more likely to update their reference point to their new endowment and hence less likely to swap.

2.6 Secondary Results

The incentivised choices presented in the previous section are the main focus of this chapter. In addition to making choices subjects also completed a number of surveys, the results of which are reported in this section. The analysis of survey responses is more exploratory in nature. Nevertheless, subjects' responses to the attitude to risk and debrief surveys provide some illuminating insights. Furthermore, survey responses can predict some of the incentivised behaviour.

Table 2.6: Participants' Fields of Study

Faculty	Frequency
Arts	45
Engineering	13
Law and Social Science	64
Medicine & Health Sciences	36
Natural Sciences	39
Business School	13

The data are from the demographics survey, which is in appendix C.2.

Table 2.7: Subject's Attitudes at the End of the Experiment

	Agree 1	2	3	4	5	6	Disagree 7
Was enjoyable	34	78	65	22	5	3	2
Compensation Sufficient	50	64	47	21	11	15	1

The data are from the debrief survey, which is in appendix C.5.

2.6.1 Subjects

The demographics survey (appendix C.2) recorded subject's characteristics. There were 210 subjects: 52.9 percent were female, 49.1 percent had taken part in experiments before, and their ages ranged from 19 to 26 with a mean of 19.7. They studied courses in a range of fields as shown in table 2.6. The distance from each subject's home address to the building where the experiment took place was calculated using the address subjects supplied. The values ranged from 0.06 to 11 miles with a mean of 1.1 miles.[12]

At the end of the experiment subjects were asked to indicate whether they enjoyed taking part and whether they felt the compensation was sufficient. The answers they gave are shown in table 2.7.

[12]The bottles of lemonade were significantly heavier than the packets of crisps, so the distance the subject had to carry to goods could be a factor influencing their choice.

Table 2.8: Percentile Figures for the Coefficient of Loss Aversion

Loss	n	λ at Percentile				
		10th	25th	50th	75th	90th
£100	209	1.20	1.50	2.00	3.00	10.00
£240	209	1.25	1.33	2.00	2.92	8.33
£580	208	1.12	1.34	1.72	3.10	9.48
£1,400	202	1.29	1.43	2.14	3.57	10.71
£3,300	200	1.21	1.51	2.02	3.64	12.12
£7,900	196	1.14	1.52	2.03	3.67	12.66
£19,000	190	1.32	1.58	2.63	5.26	26.32

The value of λ were calculated using responses to the attitude to risk survey in appendix C.3. Subjects were faced by a series of 50-50 lose X/win Y bets. Values of X were as shown in the *Loss* column of the table. Subjects reported the lowest value of Y where they would still be willing to play the bet. Some subjects indicated that they would not play some of the bets whatever the value of Y. This explains why the number of observations, *n*, varies and is less than the number of participants in the experiment.

2.6.2 Attitude to Losses and Gains under Risk

The first attitude to risk survey (the form that subjects completed is in appendix C.3) aimed to estimate a coefficient of loss aversion, λ. The coefficient measures the extent to which losses loom larger than gains. Subjects were shown a series of 50/50 lose X/win Y bets with the win figure, Y, left blank. They were asked to specify the minimum value for Y where they would still be willing to play the bet.

Result 2.9 *The median coefficients of loss aversion estimated from subjects' responses to the attitude to risk survey ranged from 1.7 to 2.6.*

Support. Subjects' responses are summarised in table 2.8. The distributions of responses featured considerable skewness, ranging from 9.2 to 12.9, and kurtosis, ranging from 94 to 176, making values of the means and variance hard to interpret. For this reason, percentile figures are reported.

Result 2.10 *Subjects who swapped goods were also less loss averse when faced with 50-50 win/lose bets.*

Support. We are confronted by several problems. First, how to combine the answers to the various questions into a single measure of loss aversion. Second, how to deal with subjects who indicated that they would not play some of the bets whatever the win figure without discarding data. Third, how to prevent outliers distorting the results given that the distributions of responses feature heavy skewness and kurtosis. I ranked the answers for each of the seven 50/50 lose X/win Y bets according to the figure specified for Y. Answers indicating the bet would be rejected however large Y were assigned the highest rank. Then I constructed an overall rank by summing the rank values for the seven answers. The higher the rank value of a subject, the more loss averse they are. I ran a Wilcoxon rank-sum (Mann-Whitney) test comparing the rank of subjects who swapped to those who did not. It indicated that subjects who swapped had lower ranks than those who did not. That is, on average, those who swapped were less loss averse than those who did not. The p-value for obtaining the result under the null hypothesis of no difference in ranks is 0.092. The magnitude of the difference in loss aversion between those who swapped and those who did not is relatively small. For example, the median value of λ calculated from responses to the 50/50 lose £580/win Y question was 1.72 for those who swapped and 1.90 for those who did not.

Result 2.11 *Subjects' attitudes to risk cannot be explained by an expected utility of wealth model with diminishing marginal utility.*

Support. See appendix E.2

Table 2.9: Swap Decisions and the Ratio of Valuations (κ)

	$\frac{1}{\lambda} \leq \kappa \leq \lambda$	$\frac{1}{\lambda} > \kappa$ or $\kappa > \lambda$	All
Sticks	88	45	133
Swaps	36	41	77
All	124	86	210

2.6.3 Predicting Swap Rates

Result 2.12 *Subjects' answers to the surveys can predict whether the subject swapped.*

Support. We can use the WTP figures from the debrief survey and λ figure from the risk survey to predict swaps. Recall from section 2.2.2 that RDT predicts the swap rate will be zero if $\frac{1}{\lambda} < \kappa < \lambda$ and $\frac{1}{2}$ if $\frac{1}{\lambda} > \kappa$ or $\kappa > \lambda$. NT predicts that the swap rate will be $\frac{1}{2}$ for all values of κ. In order to test these predictions, a value of κ for each subject was calculated. This was done using the WTP figure they specified in the debrief survey and estimating the WTP figure for the other good using the model described in appendix E.2. It is assumed that the ratio of WTP figures approximates κ, which in fact is a ratio of utility values. Table 2.9 shows the outcome of the swap decision against the two relevant intervals for κ. Running a two-sided Fisher's exact test gives a p-value of 0.009 for the null hypothesis that swap decision is independent of κ. The data in the table fits neither the prediction of NT or RDT. It appears some but not all subjects are influenced by loss aversion. One interpretation is that the TSH is right and some of the subjects for whom $\frac{1}{\lambda} \leq \kappa \leq \lambda$ are using System 2 to make their swap decision so are not influenced by loss aversion. Another interpretation is that the appropriate coefficient of loss aversion λ for the swap decision is lower than that for the risk task.

2.6.4 How Subjects Explained Their Decisions

If the two system hypothesis is right, is there any way we can tell which system an individual subject used? Sloman (2002) suggests the following.

> One rule of thumb to help identify the source of an inference has to do with the contents of awareness. When a response is produced solely by the associative system [System 1], we are conscious only of the result of the computation, not the process. In contrast, we are aware of both the result and the process in a rule based computation [System 2].

After the subjects decided whether or not to swap, they were asked what factors influenced their decision. The frequencies of the more common responses are shown in table 2.10. I have calculated a measure that I have called 'decisiveness' for each of the reasons such that $d = |l - c|/(l + c)$, where d is decisiveness, l is the number of subjects who cited the reason and chose lemonade, and c is the corresponding number who chose crisps. If $d = 0$ for a reason, it means that among those who cited the reason, equal numbers chose lemonade and crisps. Conversely, if $d = 1$, it means everyone who cited the reason chose the same item.

Result 2.13 *Subjects who had a decisive reason for their choice of goods did not exhibit the endowment effect*

Support. I divided the reasons into two groups. *Decisive* reasons are those where *decisiveness* > 0.5; *indecisive* reasons are those where *decisiveness* ≤ 0.5. One interpretation is as follows. The first group, the decisive reasons, includes the responses that most clearly suggest that the subject used System 2 to make their decision. That is suggest that the subject was fully aware of how they made their choice. The second group, the indecisive reasons, is made up of the remaining responses. These responses typically involve a more subjective weighing of the goods. Of the

Table 2.10: How Subjects Explained Their Choice of Goods

Reason	Chose		Total	Decisiveness
	Lemonade	Crisps		(d)
Decisive				
Can reuse bottles	3	0	3	1.00
Easy to carry	2	23	25	0.84
Health	22	6	28	0.57
Goods in cupboards	10	3	13	0.53
Indecisive				
Quantity per unit	6	2	8	0.50
Taste of item	30	11	41	0.46
Appearance	3	7	10	0.40
Estimated monetary value	30	14	44	0.36
Positive preference	57	40	97	0.18
Current hunger or thirst	8	10	18	0.11
Negative preference	18	15	33	0.09
Good to share	9	9	18	0.00
Indifferent	1	1	2	0.00

Table 2.11: Swap Rate and Type of Explanation Subjects Gave.

	Sticks	Swaps	Swap Rate	Total
Indecisive	101	45	0.308	146
Decisive	32	32	0.500	64
All	133	77	0.667	210

210 subjects, 64 gave at least one *decisive* reason. Table 2.11 shows the swap rates for those who cited a *decisive* reason is 0.5, the neoclassical prediction, compared to 0.308 for the rest. Running Fisher's Exact Test produces a p-value of 0.007 of obtaining these results under the null hypothesis that swap rate is independent of whether a subject cites a *decisive* reason.

2.7 Discussion

Several studies have suggested that experience causes phenomena that are anomalies for NT (such as the endowment effect) to disappear, however, it is not clear from these studies what aspects of experience matter. The experiment reported in this chapter explored how different types of experience might influence the endowment effect. It had four treatments. The behaviour of subjects in a control group was compared to that of subjects with experience of consuming, owning, and choosing goods. An endowment effect was observed but it was not as strong as those reported by earlier studies and its strength differed between treatments. Notably, the effect was significantly stronger in the two treatments in which the endowment was acquired in two steps (the choosers and owners) rather than one step (the tasters and control group). The strength of the effect differed little between the tasters and control group. It was weaker among the choosers than the owners.

There are a number of possible explanations for why the strength of the endowment effect varied between treatments and differed compared to earlier studies. First, when people are faced with an option to swap, they can make their decision in different ways. They can (a) either think about the problem in terms of (i) outcomes or (ii) losses and gains. They can (b) either use (i) intuition or (ii) reasoning to make the decision. How they make decisions determines whether an endowment effect occurs, and how they make their decision is influenced by experimental procedures. Second, loss aversion is a factor in all decisions but the degree to which losses weigh heavier than gains varies between decision problems and can be influenced by experimental procedure. The high overall swap rate in the experiment could be explained by loss aversion being weaker when a person is not giving up their last unit of a good or when they already own a good like the one being gained. Third, people do not always immediately update their reference point when their endowment changes. The differences in swap rates between treatments could be caused by differences between treatments changing the likelihood that subjects updated their reference point. Fourth, other effects in addition to loss aversion are involved. The higher swap rate of the *choosers* relative to the *owners* could be explained in terms of the choosers being committed to the choice they made before the endowment occurred.

Survey data was collected in addition to the data on choices. The responses to the attitude to risk surveys suggest that subjects had reference dependent preferences and were loss averse. Swap decisions could be predicted to a degree using the survey responses. Subjects who swapped also gave responses to the risk survey that showed less loss aversion. Furthermore, the swap rate varied with the type of reason subjects gave to explain their decision.

This study found significant variation in the strength of the endowment effect. This has several broader implications. One should be cautious not to prematurely generalise the predictions of NT or RDT since this study suggests human behaviour is richer than simple forms of both these theories imply. Furthermore, the results cast doubt on some other explanations of why the endowment effect occurs. Plott and Zeiler (2005) argue that the endowment effect in the form of a gap between willingness to pay and willingness to accept is a symptom of subject's misconceptions about experimental tasks. They support their claim by presenting experiments with and without controls for misconceptions. They find that the WTA/WTP gap is not observed when a certain set of controls for misconceptions is used. In the experiment presented in this chapter, the set of controls for misconceptions is constant across treatments but the degree to which the endowment effect occurs varies between treatments. Therefore, subject's misconceptions cannot account for whether or not the endowment effect is observed.

The following are possible extensions. Exploring whether using bundles of goods rather than single items weakens the endowment effect. Testing whether acquiring an endowment in several steps rather than one step strengthens the endowment effect. Using an incentivised design to explore the link between risk aversion over modest stakes and the endowment effect (loss aversion has been cited as an explanation for both). Further tests of whether experience of choosing and tasting influence the endowment effect. A sharper test of whether the TSH can explain differing swap rates under different experimental settings.

Appendix B

Experimenter Scripts

B.1 First Experimenter

The text in italics was read out by the experimenter. Some parts of the script differed between treatments. The labels in bold indicate which treatments each part of the script applied to.

All: The subject was invited to take a seat opposite the experimenter. *This is an experiment looking at how people make decisions. During the experiment you will be asked to make decisions and answer some questions on paper forms.* Depending on which of the four treatments the subject was assigned to, the experiment proceeded as follows.

Control Group: *To start with, please could you complete this form.* The experimenter handed the subject the demographic survey and a pen.

Taster:*To start with, please could you taste these two items.* The experimenter poured the subject a small cup of lemonade and offered them a small bowl containing some vegetable crisps. The packaging of the two items were in front of the subject so they could examine them if they wished. When the subject had tasted both items, the experimenter handed the subject the demographic survey and a pen, and said

Then could you complete this form

Owner: *To start with, please could you pick an envelope and look inside.* The experimenter offered the subject a basket containing a number of envelopes. The subject picked an envelope and looked inside. It contained a paper token for either one bottle of lemonade or one packet of crisps. Depending on what the token was for, the experimenter handed the subject either a bottle of lemonade or packet of crisps and said *This is yours to keep. Then* the experimenter handed the subject the demographic survey and a pen, and said *Then could you complete this form*

Chooser: The experimenter placed a bottle of lemonade and packet of crisps in front of the subject and said *Please could you look at these two items and choose one of them: you get to keep the one you choose.* The subject chooses one of the items. Then the experimenter handed the subject the demographic survey and a pen, and said *Then could you complete this form*

All: The experimenter then presented the subject with a basket of envelopes and said *please could you pick an envelope and look inside.* Inside the envelope was a token for several units of the two goods. The experimenter handed the subject the goods specified by the token.

The experimenter then asked the subject to complete the two attitude to risk surveys. The order in which the two surveys were given to the subject was randomised. Once the subject had completed the two surveys the experimenter said *The first part of the experiment is complete.* Drawing the subject's attention to the items they have been given: *These are yours to keep. Please could you take them around the corner and [Name of second experimenter] will complete the second part of the experiment with you.*

Appendix C

Instructions and Forms Completed by Subjects

C.1 Instructions

Experiment L1

This is an experiment in individual decision making. The experiments purpose is to study various technical issues involved in decision making. The experiment is funded by The Centre for Decision Research and Experimental Economics (CeDEx) at the University of Nottingham. During the experiment you will be asked to answer a series of questions and choose between various items. At the end of the experiment you will get to keep the items you chose, so it is in your interest to carefully consider the decisions you face.

We ask that you do not communicate with other participants during the experiment. Please refrain from verbally reacting to events that occur during the experiment.

Please wait to be called to do the first stage of the experiment.

C.2 Demographics Survey

Participant number: ___

Age: ____

Subject studied: _________________________________

Year of study _____________________________________

Gender: ____

1. Have you ever bought or sold items on online auctions such as EBay?
   - No ☐
   - Yes ☐
2. Have you taken part in CeDEx experiments before?
   - No ☐
   - Yes ☐

C.3 Attitude to Risk Survey 1: Losses and Gains

Participant number: ____

The next 7 tasks are hypothetical because they involve large amounts of money. However, even though you won't win or lose money in these tasks, please think about them and try to respond to them as if they were for real. Imagine that you are given the choice between playing and rejecting each of the bets. If you play a bet, a coin is flipped. If the coin lands tails up, you lose the amount of money shown in the 'Tails' column of the table. If the coin lands heads up, you win **an unspecified amount.** If you reject a bet, no coin is flipped and you don't win or lose anything.

For each of the bets, **please enter the amount in the win column that is just enough to tempt you to play the bet – that is, the lowest figure for which you would choose to play rather than reject the bet. Please write the amount in pounds in the space provided.** For the purposes of this task, imagine that any losses could be paid back over a period of time in the same way that you might payoff a credit card debt and that the interest rate you would pay is 6%.

	Tails LOSE	Heads WIN
Bet 1	£100	?£..........
Bet 2	£240	?£..........
Bet 3	£580	?£..........
Bet 4	£1,400	?£..........
Bet 5	£3,300	?£..........
Bet 6	£7,900	?£..........
Bet 7	£19,000	?£..........

C.4 Attitude to Risk Survey 2: Small and Large Stakes

Participant number: ____

The next 9 tasks are hypothetical because they involve large amounts of money. However, even though you won't win or lose money in these tasks, please think about them and try to respond to them as if they were for real. Imagine that you are given the choice between playing and rejecting each of the bets. If you play a bet, a coin is flipped. If the coin lands tails up, you lose the amount of money shown in the 'Tails' column of the table. If the coin lands heads up, you win **the amount of money shown in the 'Heads' column of the table.** If you reject a bet, no coin is flipped and you don't win or lose anything.

For each of the bets, **please indicate whether you would play of reject it by ticking the appropriate box**. For the purposes of this task, imagine that any losses could be paid back over a period of time in the same way that you might payoff a credit card debt and that the interest rate you would pay is 6%.

	Tails LOSE	Heads WIN		
Bet 1	£100	£101	Play ☐	Reject ☐
Bet 2	£100	£105	Play ☐	Reject ☐
Bet 3	£100	£110	Play ☐	Reject ☐
Bet 4	£100	£150	Play ☐	Reject ☐
Bet 5	£100	£200	Play ☐	Reject ☐
Bet 6	£2,000	£2,320	Play ☐	Reject ☐
Bet 7	£2,000	£69,930	Play ☐	Reject ☐
Bet 8	£2,000	£12,210,880	Play ☐	Reject ☐
Bet 9	£2,000	£850,000,000,000	Play ☐	Reject ☐

C.5 Debrief Survey

Participant number ___

1. What is the maximum price you would be willing to pay for a 100g packet of vegetable crisps like the ones used in the experiment? £_._ _p

2. Towards the end of the experiment, you were asked whether you wanted to swap the items you had for another combination of items. What factors influenced your decision?

____
____
____
____
____
____

3. Taking part in the experiment was enjoyable (please circle one of the numbers below)

Strongly agree						Strongly Disagree
1	2	3	4	5	6	7

4. The items received were sufficient compensation for taking part in the experiment (please circle one of the numbers below)

Strongly agree						Strongly Disagree
1	2	3	4	5	6	7

5. Do you have any other comments?

____
____
____
____
____

Appendix D

Original Prospect Theory

Kahneman and Tversky's (1979) Original Prospect Theory (OPT) was presented as a theory of choice under risk not a theory of riskless choice, but it is significant because it incorporates loss aversion. In their exposition of the value function, they write:

> ...the carriers of value are changes in wealth or welfare, rather than final states.
>
> A salient characteristic of attitudes to change in welfare is that losses loom larger than gains. The aggravation that one experiences in losing a sum of money appears to be greater than the pleasure associated with gaining the same amount..

They observe that most people are averse to symmetric 50-50 bets. That is, given the option to either play a bet of the form $(a, 0.50; -a, 0.50)$ or abstain, most people will abstain. Further they observe that this aversion increases with the size of the stakes. That is, if $a > b \geqslant 0$, then $(b, 0.50; -b, 0.50)$ will be preferred to $(a, 0.50; -a, 0.50)$. In OPT, the overall value of the regular prospect $(x, p; y, q)$ is calculated by combining the probability weighting function π and the value function v as follows. $V(x, p; y, q) = \pi(p)v(x) + \pi(q)v(y)$. When we compare the value of the two symmetric bets above, the probability weights cancel out, leaving: $v(b)+v(-b) > v(a)+v(-a)$. If we set $b = a-h$, then we can write $\frac{v(h-a)-v(-a)}{h} > \frac{v(a)-v(a-h)}{h}$ and letting $h \to 0$ gives $v'(-a) > v'(a)$ provided the derivative exists. What this means is that the value function for losses is steeper than the value function for gains. A measure of the strength of this effect can be defined as follows: $\lambda = \frac{v'(-a)}{v'(a)}$. When λ is constant for all losses and gains, it follows that

$$\lambda v'(x) = v_i'(-x)$$
$$\lambda \int v'(x)dx = \int v'(-x)dx$$
$$\lambda \int_0^x v'(x)dx = \int_{-x}^0 v'(x)dx$$
$$\lambda v(x) = -v(-x)$$

Appendix E

Additional Results

E.1 Subject's Willingness to Pay for Crisps and Lemonade

At the end of the experiment subjects completed a debrief survey (see Appendix C.5). In this survey they were asked to specify a hypothetical Willingness To Pay (WTP) figure for one of the two goods. They were not asked to specify figures for both goods because it was thought they might attempt to rationalise their earlier decision to swap or not. Since it would be useful to have both WTP figures, the WTP figure for the good the subject did not specify was estimated using the coefficient estimates calculated from those subjects who did specify a WTP figure for that good. Then the value of the ratio $\frac{\text{WTP Lemonade}}{\text{WTP Crisps}}$ was calculated using the WTP figure specified by the subject and the estimated WTP figure. Table E.1 shows some summary statistics for the WTP figures. It shows there is significant variation between subjects and that subjects tend to assign a higher value to lemonade. Table E.2 shows estimates of how demographic properties and being assigned to the tasters treatment influence the WTP figures. It suggests that tasting the items increased the amount subjects would be willing to pay for both goods although the results are not statisti-

Table E.1: Willingness to Pay for Crisps and Lemonade

	WTP Crisps	WTP Lemonade	$\frac{\text{WTP Lemonade}}{\text{WTP Crisps}}$
Obs	100	110	210
Mean	120	186	3.13
Variance	4013	15242	203.05
Min	1	1	0.14
Max	350	800	165.29
Percentiles			
25%	70	100	0.57
50%	100	155	0.95
75%	150	200	1.61

Table E.2: What Determines WTP Values?

	WTP Crisps	WTP Lemonade	$\frac{\text{WTP Lemonade}}{\text{WTP Crisps}}$
Obs	100	110	210
Female	26.5**	5.5	−1.4
	(12.9)	(23.5)	(2.0)
Taster	23.2	18.3	−2.1
	(15.7)	(25.7)	(2.3)
Age	0.5	−11.0	0.5
	(4.4)	(7.0)	(0.6)
Constant	89.3	397.3**	−4.9
	(86.5)	(141.2)	(12.5)
R^2	0.081	0.005	0.011

The table shows coefficient estimates calculated by running a linear regression
Standard errors are shown in parentheses
** indicates estimates are statistically significant at the 5% level.

cally significant at standard levels. Further, it suggests that tasting reduces the ratio $\frac{\text{WTP Lemonade}}{\text{WTP Crisps}}$ which implies, all else equal, that tasting would make one more likely to prefer the crisps to the lemonade.

Table E.3: Attitudes to Hypothetical High and Low Stakes Bets

Bet A: 50/50 lose £100 or win...	Bet B: 50/50 lose £2000 or win...	Accepts A		Rejects A	
		Accepts B	Rejects B	Accepts B	Rejects B
101	2,320	12	13	31	154
105	69,930	21	9	162	18
110	12,210,880	34	9	159	8
150	850,000,000,000	92	5	105	8

The table reports the responses to the attitude to risk survey, which is shown in appendix C.4.

E.2 Subject's Attitude to Risk over Small and Large Stakes

At stage 1.4 (see table 2.1) subjects completed two surveys intended to assess their attitude to risk. The first survey tested whether their attitudes were consistent with expected utility theory (EUT). Subjects were presented with a list of hypothetical 50/50 bets and asked to indicate whether they would accept or reject each one. Rabin's (2000) Calibration Theorem allows low stake and high stake bets to be paired so that the pairs have the following property. If an expected utility maximiser rejects the low stakes bet (Bet A) over a range of wealth levels, then they must also reject the high stakes one (Bet B). Table E.3 shows the responses that subjects gave to four such pairs of bets. The figures in the 'Reject A, Accept B' column show a significant number of subjects gave answers inconsistent with expected utility theory. In total, 163 of the 210 subjects (77.6 percent) gave at least one pair of answers that was inconsistent with EUT. Prospect Theory, in contrast, can explain subjects rejecting Bet A but accepting Bet B in terms of a kink in the value function about the reference point.

Appendix F

The Goods Used in the Experiment

Chapter 3

How Asymmetric Information Affects the WTA/WTP Gap in Repeated Markets: The Spinning Arrow Experiment

3.1 Introduction

There is a large body of research that investigates willingness-to-accept (WTA) and willingness-to-pay (WTP). Chapter 1 summarised this research. Among the findings were (a) that a gap between WTA and WTP regularly appears in experimental and field data and (b) that this gap often but not always decays in repeated markets. One explanation for the decay is that the gap is caused by people making errors and that in repeated markets people tend to correct these errors. Another explanation is that the gap is caused in part by heuristic driven behaviour. On this view, the decay could be interpreted as a change in the heuristic used[1]. If people typically face asymme-

[1] See chapter 1 section 1.10 for a summary of the literature on heuristics.

tric information in naturally occurring markets, then using a caution heuristic that involves setting WTP<WTA may produce higher payoffs than setting WTA=WTP. This chapter explores the heuristic based explanation. I argue that observed WTA and WTP gaps may be the result of people using a *caution heuristic* that causes them to make costly errors in the lab but protects them from costly errors in naturally occurring markets. I explore the hypothesis that if using the heuristic causes a person to repeatedly make costly errors, they will stop using it. Specifically, I show that setting WTP below WTA can be optimal when faced with asymmetric information. I present an experiment in which subjects participate in repeated markets with symmetric and asymmetric information designed to test these hypotheses and whose results support this caution heuristic explanation.

Kahneman and Tversky (1979) introduced the concept of loss aversion and modelled it using a value function that is steeper for losses than gains. Thaler (1980) applied loss aversion to riskless choice. He introduced the term *endowment effect* to identify a tendency to over weigh out of pocket costs relative to opportunity costs. The endowment effect has been frequently observed as a gap between WTP and WTA: even when valuations are elicited using incentive compatible mechanisms, WTA is often significantly higher than WTP. As discussed in chapter 1, under the standard theory of choice preferences are independent of current entitlements and so, under perfect information and in the absence of wealth effects, the maximum amount a person is willing to pay to gain an item is equal to the minimum they are willing to accept as compensation for giving it up: observing a gap is an anomaly. There are, however, many studies that have found a gap.

Some may argue that while a WTA/WTP gap is commonly observed in the lab, the incentives and experience provided by natural markets would eliminate it. Furthermore, as discussed in chapter 1, some experimental studies find the gap decays in repeated markets. But if the WTA/WTP gap quickly decays in repeated markets in the lab, why has not the experience of markets that subjects receive outside the lab already eliminated the tendency to set WTA above WTP? One plausible explanation is that experience only effects relevantly similar tasks; another is that experience outside the lab does not punish setting WTA > WTP. List (2003) ran an experiment with participants recruited at a sportscard and memorabilia show. He found that subjects with relatively more intense trading experience did not set WTA > WTP while those with relatively less intense experience did. This evidence suggests that experience gained in naturally occurring markets can influence behaviour in the tasks which produce the WTA/WTP gap in the lab. The puzzle is why does a few minutes experience in the lab appear to have a greater effect than the sum of all experience outside the lab? A possible explanation is that outside the lab, setting WTA above WTP is more often than not actually optimal.

Akerlof (1970) shows how asymmetric information about the quality of a good can reduce the amount a buyer is willing to pay for it and even lead to an equilibrium at which no trade occurs. Dupont and Lee (2002) argue a variation of this process can provide a rational explanation for the WTA/WTP gap:

> With asymmetric information, the uninformed trader thinks that the informed trader with whom he might trade has private information about the true risks. When communicating the prices at which he is willing to trade, the agent takes into account that the informed trader will trade assets only at a profit. This leads to a wedge between the uninformed trader's buying and selling prices.

This could explain a gap between WTA and WTP when people are faced with asymmetric information. It does not explain, however, why we observe a gap in the lab when people are faced with symmetric information. In order to explain the WTA/WTP gap in the lab in terms of asymmetric information we need to make further assumptions, such as the following three.

1. Behaviour is determined by some type of heuristic rather than reacting optimally to each decision problem.

2. Given a set of potential heuristics, the ones used will be those that tend to produce better outcomes.

3. Most people face asymmetric information in naturally occurring markets, so the heuristic that on average produces the best outcomes is a caution heuristic that involves setting WTA above WTP.

If people have no tendency to change the heuristic used in response to incentives, then the theory cannot account for changes in behaviour in repeated market experiments. More importantly, the theory cannot explain or predict which heuristics will be used. For these reasons, assumption (2) is required. Turning to assumption (3), one might question what the natural market analogues of WTA and WTP tasks actually are. One can imagine that when a consumer is deciding whether to buy a good they ask themselves whether they are willing to pay the good's price. The natural analogues of WTA tasks are less obvious. Most people's income comes from selling labour or renting out capital and expenditure goes on buying goods and services. Accordingly, it is likely that most people execute a greater number of transactions as buyers than sellers. However, the aggregate value of a person's buying and selling transactions will balance (assuming zero net transfers and saving), and across the economy, for every buyer in a transaction there is a corresponding seller. Whether most potential transactions involve asymmetric information is hard to judge. It seems plausible that most transactions could involve a degree of asymmetric information.

What determines which heuristic is used? A possible explanation is evolution. Heifetz and Segev (2004) argue the WTA/WTP gap is an example of *toughness* and that a *toughness* bias may be evolutionary viable. They show how in an evolutionary model toughness can emerge in bargaining with asymmetric information. Further more, Chen et al. (2006) ran experiments using capuchin monkeys and find evidence of loss aversion, which suggests it may have an evolutionary origin.

Another possible explanation of which heuristics get used is learning. Friedman (1953) compared the economic agent to an expert billiard player who played shots as if they had calculated the outcome using the laws of Physics. Now, imagine a billiard player who has been practising on a table with a slight slope, they learn to hit shots harder when playing up the slope and softer when playing down it. If they switch to a perfectly flat table, we might expect them to initially hit shots as if they were still playing on the sloped table and then after a certain amount of practice correct the error. We might expect analogous behaviour from someone who is used to trading in the face of asymmetric information where setting WTA > WTP is optimal. If they switch to facing symmetric information decision problems in the lab, at first they might continue setting WTA > WTP. After the decision problem is repeated a few times, they might correct the error and start setting WTA = WTP.

We can use the theory above to reinterpret List's result that traders with relatively intense market experience do not set WTA > WTP. Perhaps these traders not only had more market experience but also had a different type[2]. That is, their experience is not just of taking part in a market: it is of taking part in a market and being better informed than their trading partners. Hence, they are less likely to have faced asymmetric information as the less informed party, so for them setting WTA > WTP is less likely to have been optimal, so they are less likely to have adopted it as a heuristic.

This chapter presents an experiment that investigates the interaction of asym-

[2]In fact, List (2004b) speculates that asymmetric information could explain why those with more market experience do not exhibit the endowment effect. His argument, however, differs from the one here. He suggests that the more experienced traders are less uncertain about the value of the auctioned item in his experimental tasks. In contrast, I suggest that experienced traders are used to being better informed, although in the experimental tasks have no informational advantage.

metric information and the WTA/WTP gap in repeated markets for lotteries. The research question addressed is (a) does experience of markets with symmetric and asymmetric information have different effects on behaviour and (b) can the WTA/WTP be explained as a caution heuristic that protects against being exploited by better informed traders. The existing literature investigates the WTA/WTP gap under symmetric information. The novel feature of the experiment reported in this chapter is that it compares the gap under symmetric and asymmetric information. I use a $2 \times 2 \times 2$ design: buying/selling, first part symmetric/asymmetric information, second part symmetric/asymmetric information. This permits the following comparisons. First within treatment comparisons, for example does the gap close under symmetric information; does it persist or widen under asymmetric information? Second, investigation of spillovers: does experience under one regime improve performance under another. The market institution used is a computerised one sided auction, run for ten repetitions per part.

The rest of this chapter is organised as follows. Section 3.2 describes how setting WTA above WTP can be optimal when bidding in a Vickrey Auction and faced with asymmetric information. Section 3.3 presents the experimental design, the hypotheses tested, and the details of how the experiment was implemented. Section 3.4 analyses the experimental results. Section 3.5 discusses the results and their implications.

3.2 Optimal Behaviour in a Vickrey Auction with Asymmetric Information

I have suggested above that asymmetric information could explain a gap between WTA and WTP under the standard theory. I now present a concrete example. It is intended to capture the key features of trading under asymmetric information in a way that can be implemented in an experiment. Suppose that there is an item that is worth 30 in one state of the world (the low state) and 70 in another (the high state). Three people are bidding to buy the item in a second price sealed-bid auction (the highest bidder receives the item and pays the second highest bid). All three are risk neutral expected utility maximisers. The two states of the world are known to obtain with equal probability. For all three bidders, it is a weakly dominant strategy[3] to bid 50, the expected value of the item [4].

Now suppose that one of the bidders (call this bidder *the informed*) observes which state of the world obtains before placing his bid. It is a weakly dominant strategy for the informed to bid 30 in the low state and 70 in the high state. The other two bidders (call them the uninformed) know the informed will observe the state of the world before bidding but cannot observe it themselves. There are no weakly dominant strategies for the uninformed. For example, if the informed player bids 40 in both states and the other uninformed player bids 41, the best response for the remaining uninformed player is to bid >41 whereas if the informed bids 30 in the low state and 70 in the high state and the first uninformed player bids 41, the second uninformed player's best response is to bid < 41. However, we can predict how the uninformed will bid using iterative deletion of weakly dominated strategies. For the informed, bidding 30 in the low state and 70 in the high state weakly dominates all other strategies, so all other strategies can be removed. In the resulting game, it is

[3]Strategy A *weakly dominates* strategy B if and only if there is at least one set of opponents strategies for which A gives a higher payoff than B and for all other sets of opponent's strategies, A gives a payoff at least as high as B. Strategy A is *weakly dominant* if and only if it weakly dominates all other strategies.

[4]Consider an auction for a lottery with expected value 50 from the point of view of bidder 1, who submits bid b. Suppose the highest competing bid is p. Since bidder 1 is risk neutral, the value of the lottery to them is 50. Bidder 1's payoff is $50-p$ if $b > p$ and zero otherwise. By bidding 50, bidder 1 will win if $p < 50$ and not if $p > 50$. Suppose bidder 1 bids $b < 50$. When $b > p$, then they still win and their profit is $50-p$. If $p > 50 > b$, they still lose. However, if $50 > p > b$, they lose but would have made a positive profit had they bid 50. Thus, bidding $b < 50$ never increases profits and in some cases decreases it. A corresponding argument applies to bidding $b > 50$.

A more detailed discussion of bidding behaviour in second price auctions can be found in Krishna (2002, p15) a weakly dominant strategy for the uninformed to bid 30. See Marx (1999) for a discussion of iterative deletion of weakly dominated strategies.

Conversely, in a second price sealed-bid selling auction, after the informed's weakly dominated strategies are deleted, it is a weakly dominant strategy for the uninformed to bid 70. So the uninformed have a reason to bid lower in buying auctions than they do in selling auctions.

3.3 The Spinning Arrow Experiment

3.3.1 Outline of Design

The following is an outline of the experiment's design. Subjects faced a variant of the decision problem described above in section 3.2. The full details of how the experiment was implemented are described in section 3.3.3. Each subject was assigned to a *trading group*. Members of a trading group bid against each other in a series of 20 auction *rounds*. Each round consisted of an auction to buy or sell lotteries after which the lotteries were played out and subjects told how much they had made or lost in the round. Each trading group and hence each subject was assigned to one of eight treatments. The organisation of the treatments is shown in table 3.1. In one half of the treatments, subjects were endowed with *credits*[5] and took part in auctions to buy lotteries from the experimenter; in the other half they were endowed with lotteries and took part in auctions to sell lotteries to the experimenter. The auctions occurred under two market regimes: symmetric information and asymmetric information. Under symmetric information, everyone had the same information when they were placing their bids. Under asymmetric information, a minority of the members of each trading group were given extra information about which lottery outcome would occur before they placed their bids. The experiment was divided

[5]The credits were exchanged for cash at the end of the experiment.

Table 3.1: The Treatments

Treatment	Regime		Type	Subjects	Trading Groups
	Rounds 1-10	Rounds 11-20			
SS	Symmetric	Symmetric	Buying	22	4
			Selling	29	5
SA	Symmetric	Asymmetric	Buying	22	4
			Selling	31	5
AS	Asymmetric	Symmetric	Buying	27	5
			Selling	31	5
AA	Asymmetric	Asymmetric	Buying	24	4
			Selling	22	4

Each row on the table represents a treatment. Each trading group consisted of 5 or 7 subjects.

into two parts each consisting of ten rounds. Table 3.1 shows that some treatments switched between symmetric and asymmetric information after ten rounds while others did not. In the rest of this chapter, the abbreviations SS, SA, AS, and AA shown in the first column of the table are used to refer to the treatments.

3.3.2 Hypotheses and Predictions

The experiment involves repeated markets under two regimes. We can make hypotheses about behaviour in at least three settings: the first round, the last round, and spillovers (when subjects move from one market regime to another). There is also the question of what the null hypothesis should be. If we are interested in anomalies for the standard theory of choice, it seems natural to take the standard theory as the null hypothesis, which in this case I take to be Expected Utility Theory. What exactly this theory is, is open to some debate. First, does the theory allow any stochastic component in behaviour? For the purposes of this chapter I allow for a stochastic component as long as the expected error in valuations is zero[6]. Second, what is the domain of the utility function used in expected utility theory? Rabin

[6] There are several ways the stochastic component in valuations can be modelled. See Loomes and Sugden (1995), Loomes (2005) and Hey (2005) for discussions of the issues involved.

(2000) argues that under expected utility theory, we should expect approximate risk neutrality over modest stakes, since even a small degree of risk aversion over small stakes would imply an absurd degree of risk aversion over large stakes. Cox and Sadiraj (2006) object that the result only holds if the domain of the utility function is final wealth; it does not hold if income and wealth enter the utility function separately. For the purposes of this chapter, I take the domain of the utility function to be final wealth. As argued in appendix A, if income and wealth are allowed to enter the utility function separately, then preference ordering of final allocations can vary with the demarcation between income and wealth. I take this to be an anomaly in the same way that preferences over final allocations varying with current entitlements is an anomaly.

First Round Data from the first round will be used to test the null hypothesis of no WTA/WTP gap against the alternative hypothesis that there is a gap.

Last Round Data from the last round will be used to test the null hypothesis of no WTA/WTP gap against the alternatives hypotheses (a) that there is a gap of the same size as in the first round and (b) that there is a gap, but it is smaller than in the first round.

Spillovers Three hypotheses about spillovers are considered. First, no spillovers, that is experience of one market regime only affects behaviour under that regime. Second, positive spillovers, that more market experience (of any market regime) reduces anomalies (including the WTA/WTP gap). Third, negative spillovers, that experience of one market regime can increase the size of anomalies that occur when the regime is changed. The caution heuristic theory described above predicts the third type of spillover effects. Participating in a market with symmetric information where setting WTP = WTA is optimal will cause subjects to tend to set WTP = WTA. When the regime switches to asymmetric information, where setting WTP < WTA is optimal, subjects who tend to set WTP = WTA will not be behaving optimally. Conversely, participating in a market with asymmetric information where setting WTP < WTA is optimal will cause subjects to tend to set WTP < WTA. Again, when the regime switches, this time to symmetric information, a tendency to set WTP < WTA will cause non optimal behaviour.

3.3.3 Experimental Procedure

The experiment is designed to discriminate between three sets of hypotheses using a variant of the decision problem described above.

Subjects were divided into trading groups of 5 or 7 and traded in one sided buying or selling markets for lotteries. The markets were one sided in the sense that all the subjects were buying (selling) lotteries from (to) the experimenter. This is an important feature of the design since it prevents the market disappearing under asymmetric information in the way Akerlof (1970) described. Market prices were generated using median price auctions as used by Loomes et al. (2003). In the buying auctions, subjects completed the sentence ‘I am willing to buy the lottery from the experimenter if the price is less than __ credits’ by typing a value. When all subjects had entered bids, the computer selected the median bid as the market price, p. Everyone who bid above p paid p and received the lottery; everyone who bid p or less did not trade. In the selling auctions, subjects completed, ‘I am willing to sell the lottery to the experimenter if the price is more than __ credits’. The median ask was selected as the price p. Everyone who asked less than p received p credits and gave up the lottery; everyone who asked p or more did not trade. The median price auction is a form of Vickrey auction and it is a weakly dominant strategy to bid ones true valuation. There are two reasons for using a median price auction. First, for a given set of bids, the price produced by a median price buying auction will be the same as the one produced by a median price selling auction. This allows meaningful comparisons between buying and selling prices. Second, as argued by Loomes et al. median price auctions control for shaping effects. This is because if bidders in a repeated market adjust their bids towards the previously observed market price, they will be adjusting their bids towards the median. This means that if buying and selling prices converge, it is not because of shaping effects.

The novel feature of the experiment is that markets with symmetric and asymmetric information are investigated. Table 3.1 shows the organisation of treatments.

The lotteries used are shown in table 3.2. The lotteries labelled *high state* and *low state* have two possible outcomes. For instance, the low state lottery pays out zero with probability 0.63 and 31 with probability 0.37. The *composite* lottery (shown on the last row of the table) is constructed by combining the low state and high state lotteries. We can think of the composite lottery as a lottery with two outcomes which are themselves lotteries. With probability 0.43 the outcome of the composite lottery is the low state lottery; with probability $1 - 0.43 = 0.57$ it is the high state lottery. The low state lottery pays out 31 with probability 0.37, so the composite lottery will pay out 31 with probability $0.37 \times 0.43 \approx 0.16$. The probability values for the other payouts are calculated in the same way.

In rounds with symmetric information, the composite lottery was traded. In rounds with asymmetric information, the minority were informed i.e. 2 in trading groups of 5, 3 in trading groups of 7. The informed traders were told whether it was a high or low state before bidding; the uninformed traders were not. So effectively, the informed were trading either the high or low state lotteries while the uninformed were trading the composite lottery. The uninformed were told that there were informed subjects in the trading group and told what the informed would have been told.

Figure 3.1 shows how the lotteries were presented to the subjects when they

Figure 3.1: The Spinning Arrow

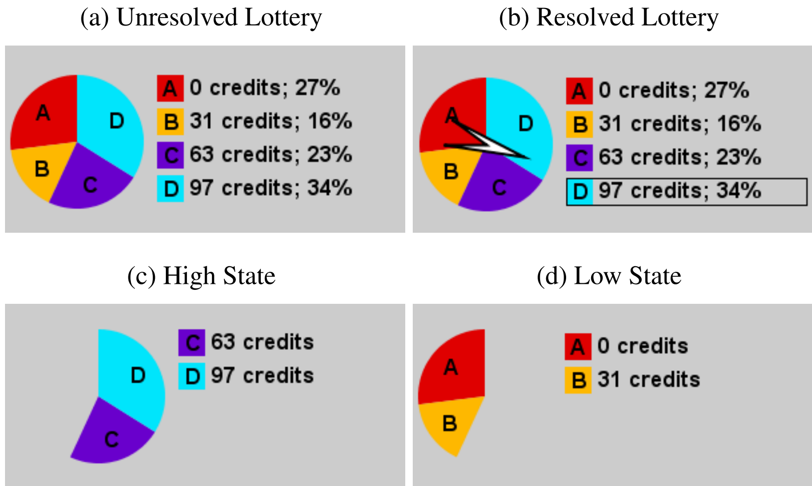

were prompted to place bids. Figure 3.1a, was shown to (a) everyone in symmetric information auctions and (b) the uninformed in asymmetric information auctions. Figure 3.1c, was shown to the informed in asymmetric information auctions when the high state occurred; figure 3.1d, was shown to them, when the low state occurred. The lottery outcomes were determined by computer generated random numbers. There was one lottery outcome per trading group per round. The outcomes were revealed to subjects after the outcome of each auction. An animated spinning arrow, figure 3.1b, was used to present the lottery outcomes. The animation lasted three seconds, initially the arrow span quickly but its speed gradually decreased until it stopped revealing the lottery payout. Examples of the complete screens subjects saw are shown in appendices H.1 to H.5.

Why is it important that only a minority of subjects in rounds with asymmetric information were informed? Suppose all the informed bid b_I and all uninformed bid b_U and $b_I \neq b_U$. Since the median bid is the price and the majority of bids are placed by the uninformed, the price will be b_U. That is the price will be determined by the uninformed. If this were not the case, then the price would be determined by the informed and the price under asymmetric information would be no different

Table 3.2: The Lotteries Used

Name		Lottery[a]	Expected Value[b]
Low state	l	$(0, 0.63; 31; 0.37)$	11.5
High state	h	$(63, 0.40; 97; 0.69)$	83.3
Composite	$(l, 0.43; h, 0.57)$	$(0; 0.27; 31, 0.16; 63, 0.23; 97, 0.34)$	47.8

[a] Each lottery is a list of consequences $(x_1, \ldots, x_n)$ and the associated probability of the consequence occurring $(p_1, \ldots, p_n)$ written in the form $(x_1, p_1; \ldots; x_n, p_n)$ where n is the number of possible outcomes of the lottery. For instance, the *lottery low* state pays out zero with probability 0.63 and pays out 31 with probability 0.37. Probability figures are rounded to two decimal places.
[b] The expected value figures are calculated as $\sum_{i=1}^{n} p_i x_i$ and rounded to one decimal place.

to what the price would be if everyone was informed. By setting the relative number of informed and uninformed so that the uninformed set the price, the effect of asymmetric information on prices can be studied.

A paper copy of the instructions was given to the subjects (see appendix G). Before the experiment started the experimenter read the instructions aloud, then gave subjects the opportunity to ask questions. All subjects were told about symmetric and asymmetric information even if they did not participate in auctions under both regimes. The motivation for this was to isolate the effect of knowing about asymmetric information from actually experiencing it.[7]

3.4 Results

A total of 208 people participated in the experiment. They were divided into 36 trading groups. Table 3.1 shows how the subjects and trading groups were divided among the 8 treatments. Several types of data are reported: prices, valuations, and costly errors.

[7] This approach is similar to the one used by Andreoni and Miller (1993) in their experiment where in some treatments there was a 50% chance of meeting a computer opponent and in others a 0.1% chance. In both treatments subjects knew about the chance of meeting a computer, but only in one was there a realistic chance of this happening.

3.4.1 Prices

In each auction the median bid or ask determined the price at which lotteries were bought or sold. There was one auction per trading group per round giving a total of $36 \times 20 = 720$ observations across all treatments. The advantage of studying prices is that they have an obvious economic meaning. If the price changes, then so do the payoffs of the subjects who traded. In contrast, if one of the bids changes without influencing the price, then the payoffs of the subjects who traded are unchanged. The subject who submits the bid that determines the price is at the margin between trading and not trading, hence there is a greater incentive for them than for non marginal bidders to bid carefully. Furthermore, since a median price auction is used, prices in buying and selling treatments are comparable. A disadvantage of studying prices is it does not allow questions about individual behaviour, e.g. do the informed and uninformed behave differently.

The evolution of prices in each of the treatments is shown in figure 3.2. The graphs show the weighted mean price. For the rounds with symmetric information, the graphs simply show the mean price across all the trading groups in the treatment. For the rounds with asymmetric information, two means were calculated separately. First, the mean price for groups where the informed traders knew the lottery was going to give a high payout, $\bar{p}_H$. Second, the mean, where they knew the lottery was going to give a low payout, $\bar{p}_L$. The weighted mean was then calculated as $\bar{p} = w_H \bar{p}_H + w_L \bar{p}_L$ where the weights are the probabilities of the high and low payout states occurring. Why is the weighted rather than simply the arithmetic mean shown? The number of trading groups with high and low payouts varies randomly between rounds. Using the weighted rather than arithmetic mean removes this random noise, making any patterns more likely to be visible.

Figure 3.2: The Evolution of Prices by Treatment

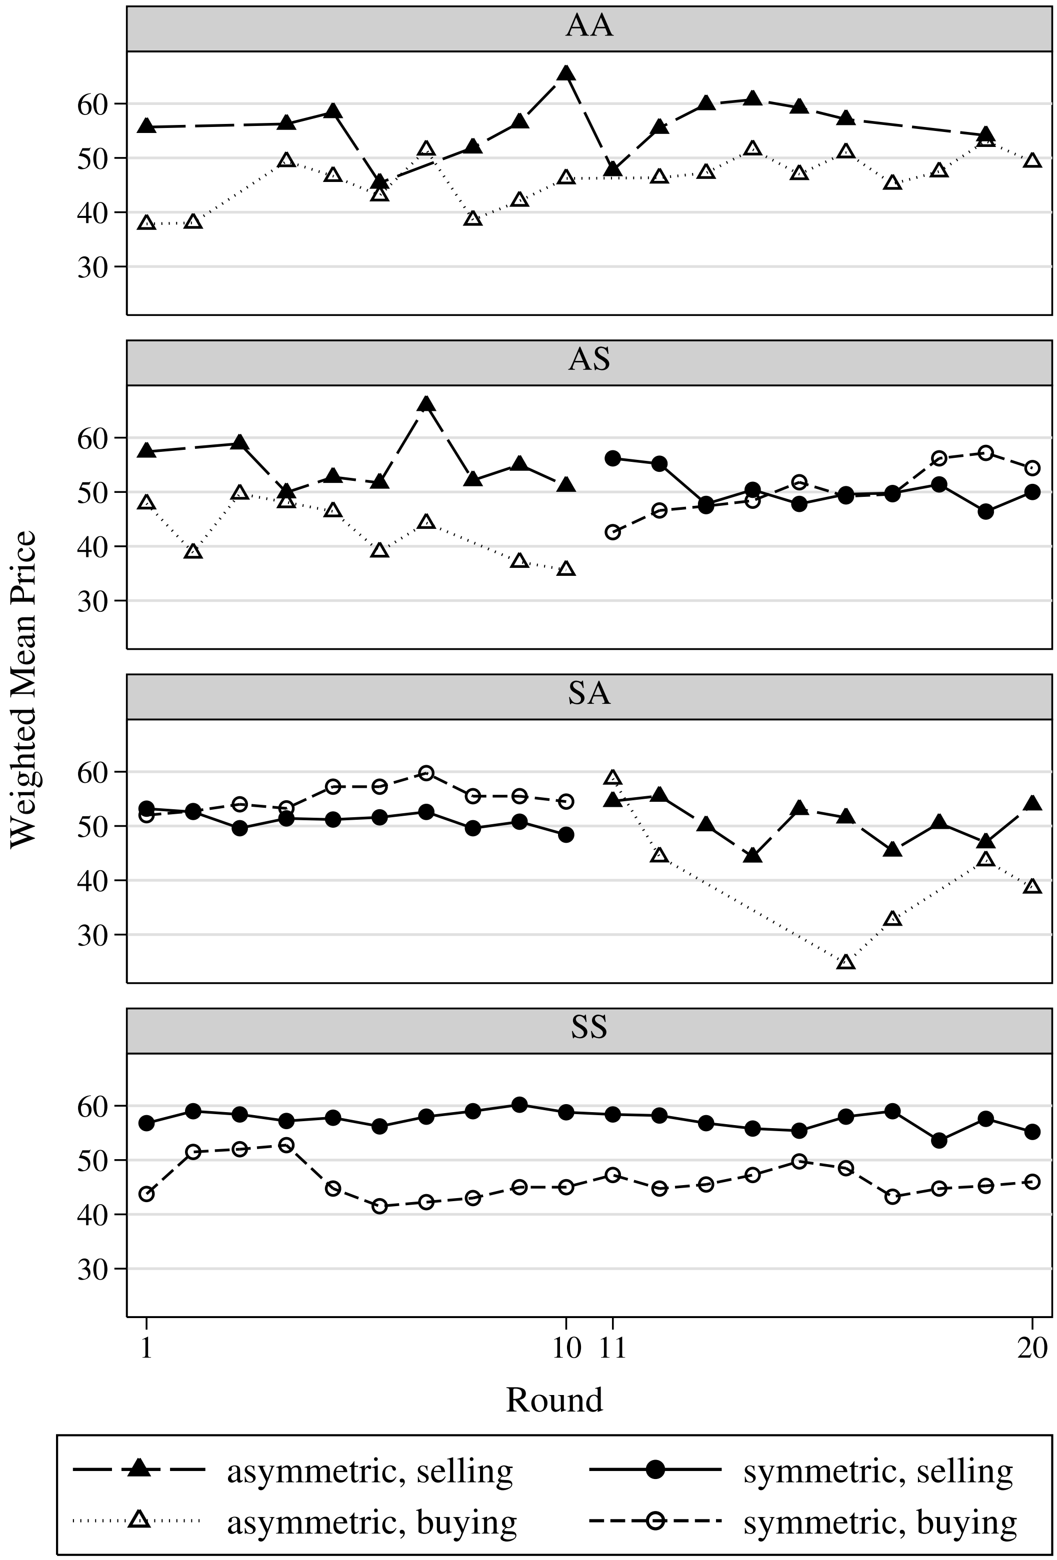

3.4.2 Valuations: Bids and Asks

In each auction, every subject submitted a bid or ask. Every subject completed 20 auctions giving a total of $208 \times 20 = 5600$ observations. The advantage of studying bids and asks is that all the data is used and questions about individual behaviour can be addressed. The disadvantage is that in most cases a subject can adjust their bid up or downwards and it will have no consequences for their payoff. Consider a trading group of 5 completing a buying auction. The third highest bid X_3 determines the price and the two bidders who bidder above X_3 buy the lottery (assuming no ties). Payoff are the same whatever value the two lowest bids take provided they are below X_3 and whatever value the two highest take provided they are above X_3. Hence, there can be patterns in the values of bids and asks that have no economic consequence.

The evolution of bids and asks in each of the treatments made by uninformed subjects is shown in figure 3.3 and the evolution of those made by informed subjects is shown in figure 3.4.[8]

3.4.3 Costly Errors

If a subject changes their bid or ask, some changes affect their payoff while other changes have no economic consequences. The market discipline hypothesis posited by Loomes et al. (2003) suggests that people have stable underlying preferences but make errors when acting on them in markets. Furthermore, people alter their behaviour if and only if they make an error which proves costly. The following definition of a costly error is intended to identify bids and asks which are costly errors for the subject who submitted them[9]. In a buying auction (a) a bidder is

[8]A subject is classed as 'informed' if for at least one part of the experiment they were an informed trader. A subject is classed as 'uninformed' if they were never an informed trader.

[9]For the purpose of defining the errors, it is assumed the subjects are expected utility of wealth maximisers. Since the stakes a low, subjects are assumed to be approximately risk neutral.

Figure 3.3: The Evolution of Bids by Treatment: Uninformed Subjects

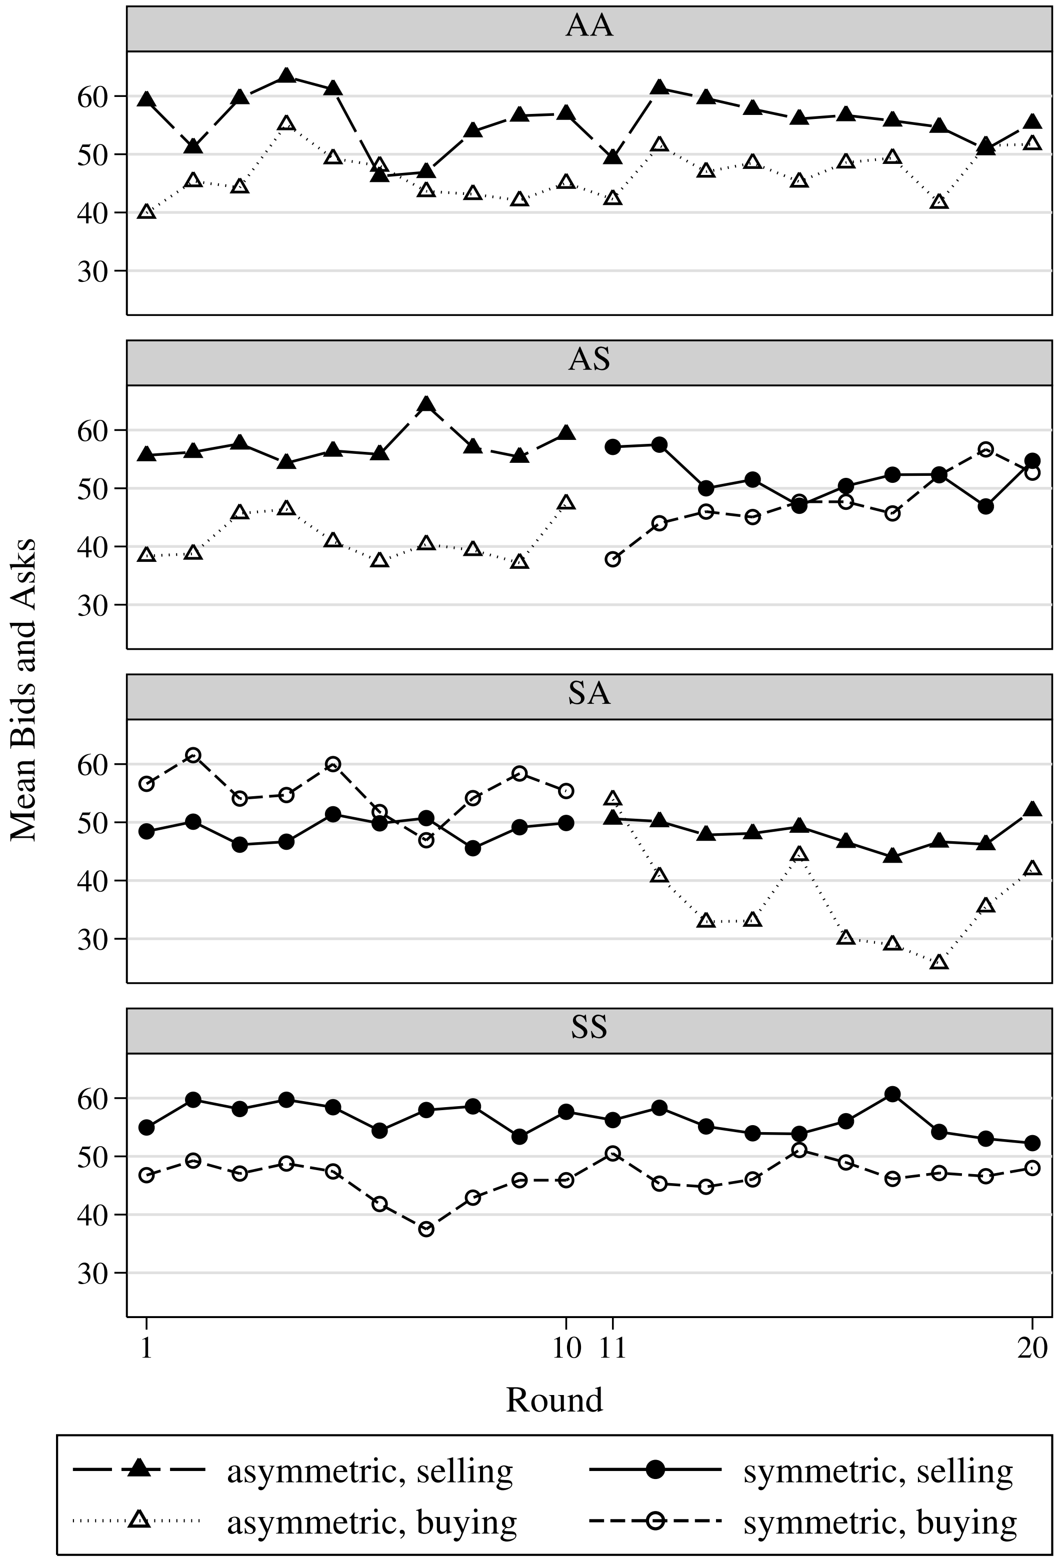

Figure 3.4: The Evolution of Bids by Treatment: Informed Subjects

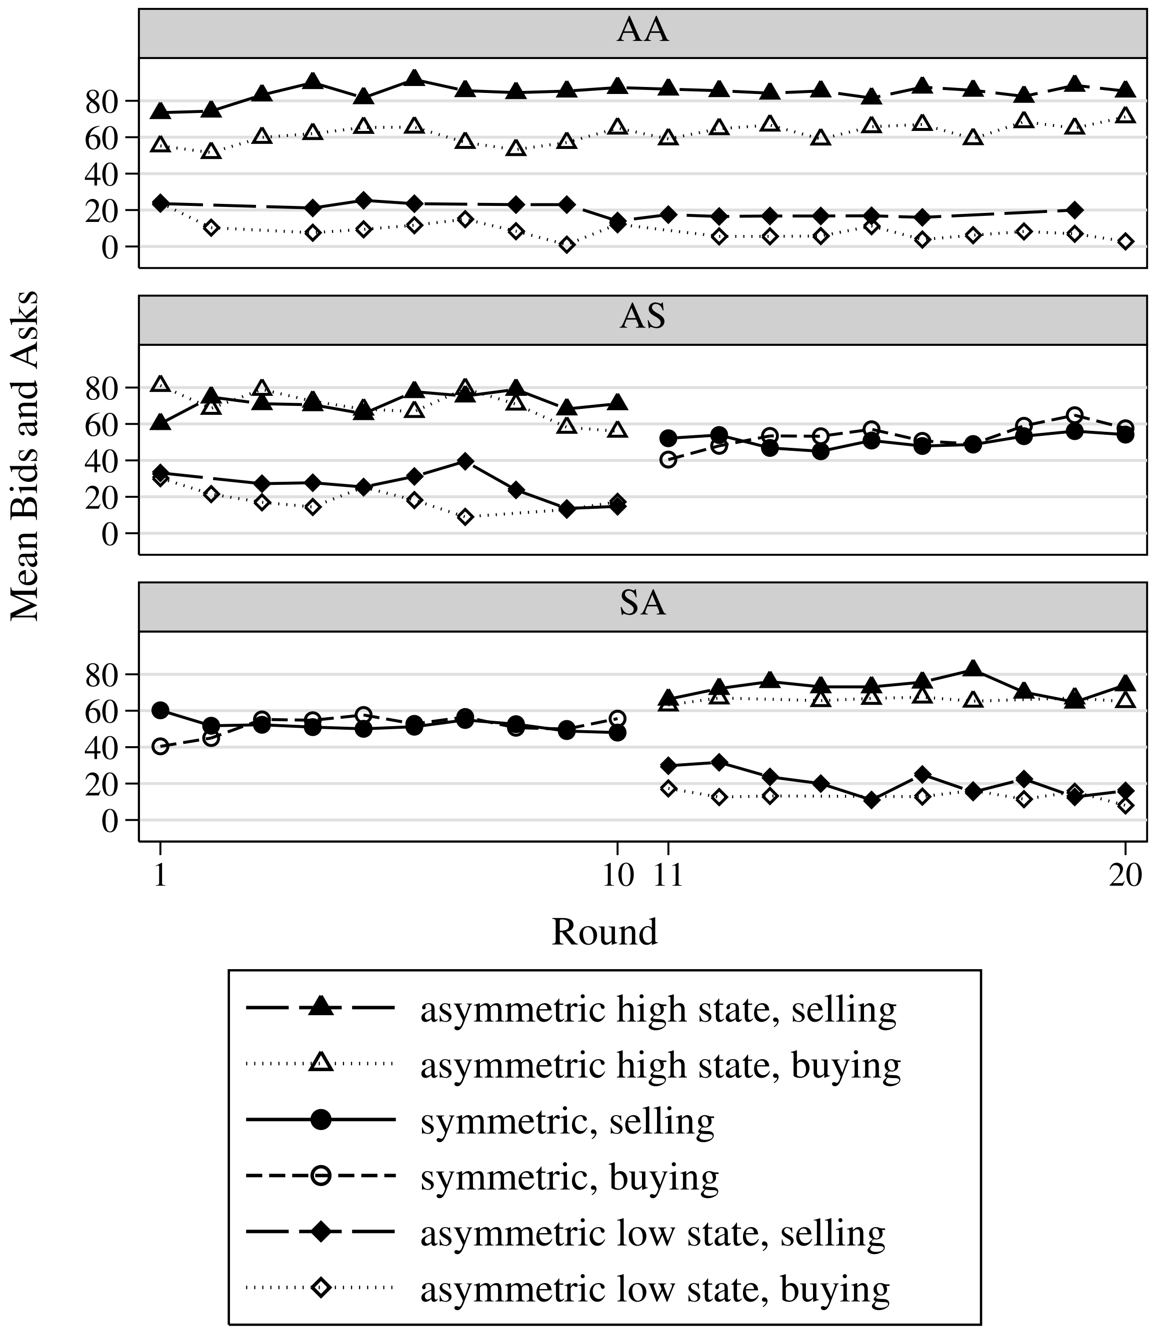

judged to have made a costly error if either (i) they bought the lottery for more than its expected value (EV) or (ii) they missed a chance to buy the lottery for less than EV; in a selling auction (b) a bidder made a costly error if either (i) they sold the lottery for less than EV or (ii) they missed a chance to sell it for more than EV. The magnitude of the error is the amount they need to adjust their bid by to avoid making the error when the bids of other subjects are held constant. A negative error indicates they bid too low; a positive error indicates they bid too high.

Costly errors are only reported for rounds with symmetric information. This is because under asymmetric information, the uninformed submit their bids before they are told whether it is a high or low payout state. In order to judge whether they made an error we need to know what their payoff would have been in both states. To determine their payoff in both states, we need to know how the informed would bid in both states, but we only observe how the informed bid in the state that occurs.

A further limitation on the use of costly errors in analysis is that they have complex statistical properties. In a given auction round, either all the errors are positive or all the errors are negative. For instance, in a buying auction where the price is less than EV, everyone who bid at or below the price made a negative costly error; everyone who bid above the price buys at less than EV so does not make an error. As a result, standard statistical tests cannot be reliably used.

3.4.4 First and Final Round Behaviour

Result 3.1 *The size of the WTA/WTP gap decreases in a repeated market with symmetric information.*

Support. Table 3.3 shows the mean buying and selling price in the first and last rounds under symmetric information. The columns labelled *gap* report the difference between mean selling and buying prices. The first row pools observations from the SS and SA treatments. Rounds 1-10 of these treatments had symmetric information, so the reported figures are from round 1 and round 10. The second row only reports data from the SS treatment where subjects traded under symmetric information for 20 rounds, so the reported figures are from rounds 1 and 20. In a given treatment, subjects were either buying for the whole experiment or selling for the whole experiment. Each subject only took part in one treatment. As a consequence, the buying and selling figures are produced by different sets of subjects but corresponding first and last round figures are produced by the same set of subjects. In the first round there is a statistically significant gap. In the final rounds the gap is smaller and not statistically significant, but still present. The Round 1 ratio of selling to buying price is $\frac{55.0}{47.9} = 1.15$ when data from the SS and SA treatments are pooled. The expected value of the lottery is 47.8. Buying prices are close to expected value whereas selling prices are a few points above it. There is little if any evidence of risk aversion.

Table 3.4 shows similar data for individual bids rather than market prices. The same pattern emerges: a statistically significant gap partially closes.

Table 3.5 shows costly errors (as defined in section 3.4.3). Again a similar pattern emerges: in the first round there is a gap that closes. In buying treatments subjects who made errors tended to bid too low; in selling treatments, they bid too high. Note the magnitude of the first round mean errors is greater in buying than selling. This can be interpreted as the error being the result of two causes: (i) the endowment effect causing a reluctance to trade and (ii) risk aversion causing a reluctance to finish the round holding a lottery. In buying rounds these two factors are pulling in the same direction while in selling rounds they are acting in opposite directions so tend to cancel each other out.

Result 3.2 *The WTA/WTP gap persists and increases slightly in size in a repeated market with asymmetric information.*

Table 3.3: Mean Market Price under Symmetric Information

	First Round			Last Round		
Rounds	buying	selling	gap	buying	selling	gap
1-10	47.9	55.0	7.1***	49.8	53.6	3.9
1-20	43.8	56.8	13.1****	46.0	55.2	9.2

The rounds 1-10 row pools the data from the SS and SA treatments; the rounds 1-20 row just includes data from the SS treatment. The null hypothesis gap = 0 is tested against the alternative gap > 0. Significance levels: *** denotes 1 percent; **** denotes 0.1 percent.

Table 3.4: Mean Individual Bids under Symmetric Information

	First Round			Last Round		
Rounds	buying	selling	gap	buying	selling	gap
1-10	48.4	54.2	5.8**	50.7	53.2	2.5
1-20	46.8	55.0	8.2**	48.0	52.3	4.3

The rounds 1-10 row pools the data from the SS and SA treatments; the rounds 1-20 row just includes data from the SS treatment. The null hypothesis gap=0 is tested against the alternative gap > 0. Significance levels: ** denotes 5 percent.

Table 3.5: Mean Individual Costly Errors under Symmetric Information

	First Round			Last Round		
Rounds	buying	selling	gap	buying	selling	gap
1-10	−9.4	2.5	11.9	−0.6	1.9	2.4
1-20	−13.1	4.6	17.7	−1.7	−1.7	0.0

The figures reported are the mean errors made by those subjects who made a costly error (see definition in 3.4.3).

Table 3.6: Weighted Mean Market Price under Asymmetric Information

	First Round			Last Round		
Rounds	buying	selling	gap	buying	selling	gap
1-10	43.9	56.3	12.4***	41.4	56.1	14.7**

Data from the AS and AA treatments are pooled. The null hypothesis $gap \leq 0$ is tested against the alternative $gap > 0$. Significance levels: ** denotes 5 percent; *** denotes 1 percent.

Support. Table 3.6 shows the weighted mean first and last round prices under asymmetric information. The table pools the round 1 and round 10 data from the AA and AS treatments.[10] The reported weighted mean[11] values are calculated as

$$weighted\ mean = mean\ high\ stated \times p(high\ state) + mean\ low\ state \times p(low\ state)$$

where high state is when the lottery pays out 63 or 97 and low state is when it pays out 0 or 31. The figures in the *gap* columns are the difference between the weighted mean selling and buying prices. First round behaviour is similar to that under symmetric information: there is a statistically significant gap in the predicted direction. However the gap persists and marginally increases in size. Note, that iterated deletion of weakly dominated strategies suggests that the buying price should fall to 11.5 (the expected value of the lottery in the low state) and the selling price should rise to 83.3 (the expected value of the lottery in the high state). Instead, although the gap persists, buying and selling prices remained relatively close to the expected value of the lottery across both states which is 47.8.

Table 3.7 shows mean bids under asymmetric information. Under asymmetric information, the uninformed did not know whether it was a high or low payout state when placing their bids but the informed did know. The bids are disaggregated

[10]The only figures reported are for rounds 1-10 unlike symmetric information where figures for rounds 1-10 and 1-20 were reported. This is because calculating the weighted figures requires at least one observation for the high and low state in both buying and selling and in both the first and last round. The data from rounds 1 and 20 of the AA treatment do not include all the required observations.

[11]Weighted means rather than arithmetic means are reported to control for variation in the frequency of high and low states between buying and selling treatments and between rounds. This was discussed in more detail in section 3.4.1.

Table 3.7: Mean Individual Bids under Asymmetric Information

	Round 1			Round 10		
	buying	selling	gap	buying	selling	gap
Uninformed	39.1	57.2	18.1	46.3	58.3	12.0
Informed						
high state	71.7	64.0	−7.7	60.5	78.2	17.7
low state	26.5	28.7	2.2	14.9	14.5	−0.4

The low state is when the lottery paid out 0 or 31; the high state when it paid out 63 or 97. The informed traders were told whether it was a high or low state before they bid; the uninformed were not. No hypothesis testing is reported on this table.

accordingly into those made by the uninformed, the informed when they knew it was a high state, and the informed when they knew it was a low state. For the uninformed, there is a gap between buying and selling bids in round 1 and round 10. The informed take advantage of the extra information they possess and bid considerably higher when they know it is a high state.

3.4.5 Between Round Spillovers

The last two sections found that the WTA/WTP gap closes under symmetric information but persists under asymmetric information. What causes the changes over successive rounds, and why do they differ between symmetric and asymmetric information? In this section I assess whether previous successful trades influence behaviour in later rounds?

Result 3.3 *Propensity to trade increases with previous trading success.*

Support. I estimate the following equation[12]:

$$bid_{it} = \alpha_i + \beta S_{i,t-1} + \gamma D_{it} + \psi_t + \varepsilon_{it} \tag{3.1}$$

[12]The equation is a two-way fixed effects model. It has a similar form to the one used by List and Shogren (1999) to determine whether previous observed market prices influence bidding in repeated second price auctions.

The variables are defined as follows: bid_{it} is the bid or ask submitted by subject i in round t of the experiment. Fixed effects are captured by α_i and ψ_t: α_i represents characteristics of subject i that influence bids but whose effects are constant across rounds; ψ_t represents factors that vary across successive auction rounds but are constant across subjects. $S_{i,t-1}$ is a measure of the relative number of profitable and loss-making trades subject i made before round t; β is the corresponding coefficient. D_{it} is a vector of dummy variables specifying the decision problem subject i faced in round t; γ is a vector of corresponding coefficients. The final term, ε_{it} captures errors, the variation in bid_{it} not accounted for by the preceding variables.

How is trading success measured? *Buying* the lottery is profitable if and only if the lottery payout *exceeds* the price. *Selling* the lottery is profitable if and only if the lottery payout *falls short of* the price. Let π_{it} indicate the outcome of subject i's trading in round t as follows:

$$\pi_{it} = \begin{cases} +1 & if\, traded\, and\, profited \\ 0 & if\, did\, not\, trade\, or\, traded\, and\, exactly\, broke\, even \\ -1 & if\, traded\, and\, made\, a\, loss \end{cases} \tag{3.2}$$

I measure subject i's relative success trading in all rounds up to and including t as follows.

$$S_{i,t} = \frac{\sum_{k=1}^{t} \pi_{ik}}{t} \tag{3.3}$$

The measure has the following properties. If the majority of trades have been profitable, then $S > 0$. If the majority of trades have been loss making, then $S < 0$. If subject i does not trade in round t, then $|S_{i,t}| < |S_{i,t-1}|$, that is the magnitude of S decreases. The motivation for this measure is that people might use a variant of the availability heuristic to judge how likely it is they will make a profit from trading. The more often they have traded and made a profit in previous auctions, the easier it will be for them to imagine that trading in the next auction will be profitable, hence they will judge that trading is more likely to be profitable[13] and accordingly they will be more willing to trade.

Table 3.8 shows the results of estimating equation 3.1 for (a) a pool of the 4 buying treatments, (b) a pool of the 4 selling treatments, and (c) each of the 8 treatments. Whichever set of bids the estimation is estimated on, increased trading success (i) increases bids in buying treatments but (ii) decreases bids in selling treatments. Bidding higher in a buying auction increases the likelihood of trading; bidding lower in a selling auction does the same. Hence the results suggest propensity to trade increases with previous trading success.

3.4.6 Spillovers between Markets with Symmetric and Asymmetric Information

Table 3.9 reports buying and selling prices for all treatments. The reported values are averages over two rounds. The reason for this is to allow meaningful comparisons of prices between treatments with asymmetric information. To calculate a weighted average price across the high and low states, in every asymmetric information round reported there must be at least one buying and at least one selling auction in both of the high and low states. In some rounds this condition is not met so averages over two rounds are used. Note, there is a lot of noise in the data with the size and sign of the gaps sometimes not being as expected. This is not surprising since the buying and selling price figures reported are the mean of just 8 or 10 observations.

[13]Kahneman and Tversky (1973) give the following example of the availability heuristic: "one may assess the divorce rate in a given community by recalling divorces among ones acquaintances. If subjects in the experiment use a similar heuristic, they may assess the probability that the next trade will be profitable by recalling what happened in previous trades.

Table 3.8: Two-Way Fixed Effects Estimates of the Relation between Amount Bid and Previous Trading Success

	Pooled		SS		SA		AS		AA	
	buying	selling	buying	selling	buying	selling	buying	selling	buying	selling
β	4.7	−7.1	4.3	−3.9	5.1	−5.1	0.8	−6.8	7.4	−14.7
	(2.1)	(1.8)	(4.4)	(2.3)	(4.8)	(3.1)	(3.5)	(4.2)	(3.9)	(5.1)
$p(\beta\|H_0)$	0.012	0.000	0.165	0.047	0.147	0.051	0.410	0.052	0.030	0.002
$F(\alpha_i = 0)$	11.2	7.8	14.7	10.1	5.9	6.4	16.9	5.6	8.7	8.2
$p(F\|\alpha_i = 0)$	0	0	0	0	0	0	0	0	0	0
$F(\psi_t = 0)$	1.9	2.6	1.6	1.5	1.2	0.7	2.8	1.1	1.1	1.2
$p(F\|\psi_t = 0)$	0.012	0.000	0.047	0.091	0.302	0.853	0.000	0.401	0.356	0.245
N	1805	2147	418	551	418	589	513	589	456	418

The table shows the results of estimating equation 3.1. β is the coefficient on the measure of previous trading success. Standard errors are shown in parentheses. $p(\beta|H_0)$ is the p value for the obtained results under the null hypothesis that (i) $\beta \leq 0$ for buying and (ii) $\beta \geq 0$ for selling. $F(\alpha_i = 0)$ is the F statistic for the null hypothesis that $\alpha_i = 0$ all for i, that is the hypothesis that there are no subject fixed effects. $p(F|\alpha_i = 0)$ is the p value for the F statistic. $F(\psi_t = 0)$ and $p(F|\psi_t = 0)$ are the corresponding figures for the fixed effects due to the number of rounds completed. N is the number of observations where each bid counts as one observation.

Recall that table 3.1 shows the number of trading groups and subjects in each of the treatments.

How do subjects with experience of trading under symmetric information behave under asymmetric information compared to (a) those with no experience and (b) those with experience of asymmetric information?

Result 3.4 *The effect of experience of markets with symmetric information spills over into markets with asymmetric information and reduces the size of the WTA/WTP gap.*

Support. From table 3.9 we see that when the regime is switched from symmetric to asymmetric information the gap is 4.7. This is (i) smaller than the gap for those with no experience (the rounds 1 & 2 gaps for the AS and AA treatments are 15.7 and 16.0) and (ii) smaller than the gap for those with experience of asymmetric information (the rounds 11 & 12 gap for the AA treatment is 8.4). Under asymmetric information, optimal behaviour results in a large gap. Hence it appears that when the regime switches to asymmetric information, experience of symmetric information is worse than experience of asymmetric information or no experience at all.

Conversely, how do subjects with experience of trading under asymmetric information behave under symmetric information compared to (a) those with no experience and (b) those with experience of symmetric information?

Result 3.5 *The effect of experience of markets with asymmetric information spills over into markets with symmetric information but its effect is similar to the effect of experience of symmetric information.*

Support. From table 3.9 we see that when the regime switches from asymmetric to symmetric information the initial gap is 11.1. This is (i) bigger than the gap for those with no experience (the rounds 1 & 2 gaps for the SS and SA treatments are 10.3 and

Table 3.9: Weighted Mean Market Price under All Treatments

	Rounds 1 & 2			Rounds 9 & 10			Rounds 11 & 12			Rounds 19 & 20		
Treatment	buying	selling	gap	buying	selling	gap	buying	selling	gap	buying	selling	gap
SS	47.6	57.9	10.3	45.0	59.5	14.5	46.0	58.3	12.3	45.6	56.4	10.8
SA	52.4	52.9	0.5	55.0	49.6	−5.4	49.3	54.0	4.7	40.9	50.0	9.1
AS	43.3	59.1	15.7	36.9	53.1	16.1	44.6	55.7	11.1	55.8	48.2	−7.6
AA	39.7	55.7	16.0	42.6	61.0	18.3	43.2	51.6	8.4	51.2	54.5	3.3

The weighed mean prices are calculated as in previous tables. No hypothesis testing is reported on this table.

0.5) but (ii) smaller than the gap for those with experience of symmetric information (the rounds 11 & 12 gap for the SS treatment is 12.3). Under symmetric information, optimal behaviour results in no gap. Hence, when the regime switches to symmetric information, it appears experience of asymmetric information is marginally better than experience of symmetric information but worse than no experience at all.

Does the presence of informed subjects and the resulting asymmetric information cause the uninformed to behave differently? If so, when the market regime changes between symmetric and asymmetric information, do those without the informational advantage adapt their behaviour immediately or is the change gradual?

Result 3.6 *The uninformed do not immediately change their behaviour when they switch between facing symmetric and asymmetric information.*

Support. I estimate the two following equations:

$$bid_{it} = \alpha_i + \delta_0 AI_{i,t} + \psi_t + \varepsilon_{it} \tag{3.4}$$

$$bid_{it} = \alpha_i + \delta_0 AI_{i,t} + \delta_{-1} AI_{i,t-1} + \ldots + \delta_{-4} SAI + \psi_t + \varepsilon_{it} \tag{3.5}$$

where bid_{it} is the bid or ask submitted by subject i in round t. As in equation 3.1, α_i and ψ_t capture fixed effects of individuals and experiment round number. $AI_{i,t}$ is a dummy variable: $AI_{i,t} = 1$ if subject i faced asymmetric information in round t of the experiment; $AI_{i,t} = 0$ if they faced symmetric information. Table 3.10 shows the estimated parameters. The equations were estimated for buying and selling auctions separately. Equation 3.4 is estimated using all bids from the SS treatments plus all the bids made by the uninformed from the AS, SA and AA treatments (i.e. bids made by the informed are excluded). Equation 3.5 is estimated using a subset of the bids used to estimate equations 3.4: the bids that are not used are those where $t < 5$ since the dummy variable $AI_{i,t-4}$ does not exist for these bids.

Table 3.10 reports the results of estimating equations 3.4 and 3.5. The estimates of the coefficients for the dummy variables show how bids are adjusted relative to bids placed under symmetric information. For instance, the estimates of equation 3.4 suggest when the uninformed face asymmetric information, they bid 12 points lower in buying auctions and 2.3 points higher in selling auctions. If the uninformed adjusted their behaviour immediately when they faced asymmetric information, we would expect (a) the AI_t coefficients to be equal in the simple model and the model with lags and (b) the coefficients on the lagged dummy variables to be zero. Instead, the coefficients on the lagged dummy variables are not zero. This indicates that the value of a bid is influenced by whether there was asymmetric information in previous rounds. Hence, it appears the uninformed do not adjust the bidding strategy immediately when the regime switches between symmetric and asymmetric information.

We can also analyse how the informed and uninformed behave after the regime switches from asymmetric to symmetric information.

Result 3.7 *After the market regime switches from asymmetric to symmetric information, the previously informed traders have a higher propensity to trade than the previously uninformed.*[14]

Support. Table 3.11 reports behaviour at the level of the individual for rounds 11 to 20 of the AS treatment (i.e. the behaviour under symmetric information of those

[14] Notice there is a contrast between result 3.5 and result 3.7. There is little difference between the round 11-20 prices under symmetric information when rounds 1-10 occurred under (a) asymmetric information and (b) symmetric information (see figure 3.2). Result 3.7, however, suggests that experiencing asymmetric information as an uninformed trader makes subjects less willing to trade. All else equal, this should lead to a larger WTA/WTP gap in round 11-20 prices in the AS treatment than the SS treatment.

Table 3.10: Two-Way Fixed Effects Estimates of the Relationship between Value Bid and Asymmetric Information

	Buying		Selling	
	Simple	Spillovers	Simple	Spillovers
Equation Estimated	3.4	3.5	3.4	3.5
AI_t	−12.0∗∗∗∗	1.3	2.3∗∗	1.1
	(1.3)	(3.1)	(1.0)	(2.4)
AI_{t-1}		−9.5∗∗		−0.4
		(4.1)		(3.2)
AI_{t-2}		−4.8		2.6
		(4.1)		(3.2)
AI_{t-3}		0.6		−0.6
		(4.1)		(3.2)
AI_{t-4}		−2.2		0.1
		(3.1)		(2.4)
$\sum_{k=0}^{4} AI_{t-k}$		−14.5		2.7
subjects N	65	65	78	78
rounds T_i	10 or 20	6, 10 or 16	10 or 20	6, 10 or 16
Observations $\sum_{i=1}^{N} T_i$	1300	1040	1560	1248

The table shows the results of estimating equations 3.4 and 3.5. The dependent variable is bid_{it}, the bid made by subject i in round t. The equations were estimated using all bids from the SS treatments and the bids of the uninformed in the other treatments. The reported figures are coefficients for dummy variables indicating whether round $t-x$ had asymmetric information. Standard errors are shown in parentheses.
Significance levels: ** denotes 5 percent; **** denotes 0.1 percent.

Table 3.11: The Effects of Experience of Asymmetric Information on Behaviour under Symmetric Information

Experience	Bid	Error	Trading Rate	Plays Lottery Rate
Uninformed				
buying	47.6	2.2	0.33	0.33
selling	52.0	−7.3	0.39	0.61
all	49.9	−2.4	0.36	0.47
Informed				
buying	53.3	−1.1	0.46	0.46
selling	50.9	0.2	0.38	0.62
all	52.0	−0.5	0.42	0.55

The table reports data from rounds 11 to 20 of the AS treatment (the rounds with symmetric information).

with experience of asymmetric information). The behaviour of the subjects who were previously uninformed traders is compared to the behaviour of those who were previously informed (rounds 1-10 of treatment AS had asymmetric information). The *Bid* column shows the mean bids. The *Error* column shows the mean costly error (as defined in section 3.4.3). The Trading *Rate* is the proportion who either buy or sell the lottery. The *Plays Lottery Rate* is the proportion of subjects who hold a lottery at the end of the round (i.e. those who buy the lottery or do not sell it). We can see that the uninformed made larger errors both when buying and selling. We see that the uninformed trade at lower rates than the informed and are less likely to finish the round holding a lottery. This can be interpreted as experience of being the uninformed party under asymmetric information having two effects: first, it increases aversion to trading (the endowment effect); second, it increases risk aversion. In buying auctions, these two factors act in the same direction; in selling auctions, they act in opposite directions. So we see a large difference between the informed and uninformed in buying auctions.

To test hypotheses about the relative trading rates of the informed and uninformed it does not make sense to look at individual level data since whether a given subject trades in an auction is not independent of whether the others trade. For example, in a trading group of 5 (and assuming no ties) exactly two members trade; in a group of 7, three trade. To get around this difficulty, we can take each auction as one observation and ask the following. Is the number of trades by informed traders more or less than expected? Is the number of informed traders holding a lottery at the end of the round more or less than expected? The expected number of trades by previously informed traders is calculated as $expected\ trades = total\ trades \times \frac{informed\ traders}{all\ traders}$. Table 3.12 reports the number of auctions in which the informed traders traded more and less than expected. If there were no difference in the behaviour of the previously informed and uninformed, then we would expect the informed to trade more than

Table 3.12: The Trading Rate of Previously Informed Traders under Symmetric Information

	Selling	Buying	Total
Number of auctions previously informed traders:			
trade less than expected	23	14	37
trade more than expected	26	36	62
Total	49	50	99

The table reports the trading rates during rounds 11-20 of the AS treatment of those who were informed traders during rounds 1-10. There were 10 trading groups in the AS treatment and each trading group completed 10 rounds under symmetric information giving 100 auctions. In one auction there were zero trades, this auction is omitted from the analysis. In each of the remaining auctions, the expected number of trades by previously informed traders is calculated as $expected\ trades = total\ trades \times \frac{informed\ traders}{all\ traders}$.

expected and less than expected with equal probability. In fact the informed traders traded more than expected in $\frac{62}{99}$ of the auctions.

Consider (i) the null hypothesis that in each auction informed traders are equally likely to trade more than expected as they are less than expected and (ii) the alternative hypothesis that informed traders trade more than expected. To test the hypotheses, I estimate the following fixed effects Probit model:

$$Pr(informed\,trades_{it} > total\ trades_{it} \times \frac{informed\ traders_i}{all\ traders_i}) = \Phi(\beta + \gamma_i + u_{it}) \tag{3.6}$$

where $informed\,trades_{it}$ is the number of trades by previously informed traders in trading group i and round t, $informed\ traders_i$ is the number of informed traders in trading group i etc. The coefficient β is positive if informed traders on average trade more than expected. The coefficient γ_i captures trading group fixed effects. It is positive if informed traders in trading group i trade more than expected more often than informed traders in the other trading groups. Finally, u_{it} is the residual. Since the model controls for trading group fixed effects, it allows meaningful hypothesis testing even though there is more than one observation per trading group. The null hypothesis $\beta = 0$ was tested against the alternative $\beta > 0$. The test produced a p-value of 0.018 for obtaining the observed results under the null hypothesis.

The same method can be used to test the hypothesis that the informed were more likely to hold a lottery after the auction. Using values from table 3.12 we can deduce that in $\frac{23+36}{99} = \frac{59}{99}$ of the auctions, more than the expected number of informed played the lottery. Estimating an equation equivalent to 3.6 and testing corresponding hypotheses gave a p-value of 0.071 for obtaining the observed results under the null hypothesis that the previously informed are no more likely to end a round holding a lottery.

3.5 Discussion

The experimental results are consistent with previous studies in that I find that under symmetric information (a) an initial gap between buying and selling prices for lotteries and (b) that this gap decays in a repeated market. The novel findings are as follows. First, under asymmetric information a gap between buying and selling prices persists in a repeated market. Second, there are spillover effects. At the level of market prices, the gap between buying and selling prices under symmetric information is greater if the subjects in the market have just been trading in a market with asymmetric information than if they have not been trading at all. Conversely, the gap under asymmetric information is smaller if the subjects have just been trading in a market with symmetric information than if they have not been trading at all. These results could be interpreted in terms of a heuristic based model of decision making in which people do not immediately adjust their behaviour when the market regime changes, although at the level of prices, the evidence is mixed. Notably, when the regimes switches from asymmetric to symmetric information, the prices observed under symmetric information are similar to those in treatments that have only traded under symmetric information. Third, there is additional evidence of spillover effects when the behaviour of previously informed and uninformed traders under symmetric information is compared. The previously uninformed are less likely to trade and less likely to end an auction holding a lottery than the previously informed traders. These results provide stronger evidence in favour of the caution heuristic based interpretation: the uninformed adopt a caution heuristic that protects them from costly errors under asymmetric information and when the regime is switched to symmetric information, they do not immediately adjust their behaviour.

One might wonder why, if using a heuristic causes suboptimal behaviour as described above, people would solve decision problems using a heuristic rather than deliberately solving each one. A possible explanation is that determining the optimal behaviour under asymmetric information is hard. For instance consider the asymmetric information problem subjects faced in the experiment. Iterative deletion of weakly dominated strategies suggests an equilibrium where the selling price is 83.3 and the buying price 11.5. Even when the same problem is repeated 10 times, the market prices were no where near these values. If people can not solve a relatively simple asymmetric information problem in the lab, even when it is repeated, can we expect them to do better in naturally occurring markets where the problems are likely less well specified and successive problems are unlikely to be identical? The most likely answer is no. Hence there is a role for heuristics.

What are the wider consequences of people using a caution heuristic? The direct consequence is that the decision of whether or not to trade is not fully determined by preferences. There will be some potential trades that will make both parties better off but will not be executed. This means that welfare gains from trade will not be fully realised. Furthermore, there are consequences to spillover effects, i.e. making a loss on a trade causes a person to be more cautious and generally more reluctant to trade. For instance, institutions that protect buyers from making losses on purchases will reduce the number of buyers suffering losses and so cautiousness among buyers.

Examples of such institutions include legal rights for buyers of goods, additional guarantees offered by some sellers, financial redress for people who were miss-sold financial products. Notice, however, if both buyers and sellers are equally prone to use a caution heuristic, institutions that transfer risk between them will not increase the gains from trade if transferring the risk makes one party less cautious but the other more cautious.

Appendix G

Instructions

Introduction

You are about to participate in an experiment investigating how people make decisions in markets. During the session, please do not talk or communicate with any of the other participants. If you have a question, please raise your hand and I will come to your desk to answer it.

Payment

You will be paid in cash at the end of the experiment. The amount you receive will depend on the decisions you and other participants make and the outcome of random events. During the course of the experiment you will gain or lose credits. At the end of the experiment, the credits you have accumulated will be exchanged for real money. You will receive £1 for every 250 credits.

Outline

The experiment involves buying lotteries. The experiment is divided into two parts: A and B. There are 10 rounds in each part, so there are 20 rounds in total. In each round you will be allocated 100 credits and have the chance to buy a lottery from the experimenter. The price of the lottery will be determined by a special type of auction (how the auction works is described in detail later). At the end of each round you will be told the result of the auction and how much the lottery paid out. How much you earn from each round is affected by some or all of the following factors: whether you buy the lottery, the price of the lottery, and how much the lottery pays out. If you do not buy the lottery, the amount you earn from the round is simply the 100 credits you were allocated at the start of the round. If you buy the lottery, your earnings are 100 credits plus whatever the lottery pays out minus the amount you paid for the lottery.

After each round your earnings from the round are banked. At the end of the experiment you will be paid based on the number of credits you have banked over the 20 rounds.

Lotteries

The lotteries will be shown to you in the following format.

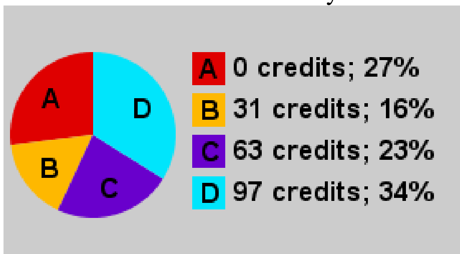

The lottery shown above pays out 0 credits with probability 27%, 31 credits with probability 16%, 63 credits with probability 23%, and 97 credits with probability 34%.

The lottery result is determined by a computer simulated spinning device as illustrated below. The amount that the lottery pays out depends on what colour the arrow is pointing to after it has been spun. In the screenshot below, the arrow is pointing to region D, so the lottery would payout 97 credits. There are no tricks. The probability figures for each of the regions are accurate. We have determined all of the lottery outcomes in advance, so we do not have to do this during the experiment. However, we can assure you that the outcomes of the lotteries were resolved in a genuinely random way. If you wish, after the experiment is over you can request a printout showing all the lottery outcomes for the session you took part in to verify this. The outcomes are then revealed to the participants in the experiment at the appropriate stage in the experiment

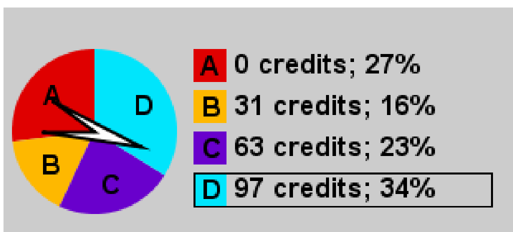

Auctions

As stated above, during the experiment you will participate in a series of auctions to buy lotteries. You will be told how many other participants are bidding in the auctions. It will be the same people bidding against you in every auction.

During the auction, you and the other participants are bidding to buy a lottery from the experimenter. Each participant can only buy one lottery per auction. However, in most auctions more than one participant will buy a lottery. The price and who buys will be determined as follows.

(1) How are bids entered? Each participant will be prompted to type a bid into the box shown on the screenshot below.

[Screen shot appeared here]

(2) How is the price determined? The computer will record the bids made by each of the participants and arrange them in order from lowest to highest. Suppose, for example, the bids were:
35, 36, **56**, 68, 72

The middle value (median) determines the price. So in this case the price would be 56.

(3) Who buys the lotteries and how much do they pay?
Each of the participants who bid above the price buys a lottery. They pay the price, not the amount they bid. So, if (as in the above example) the bids were 35, 36, 56, 68 and 72, the price would be 56 and the participants who bid 68 and 72 would each pay 56 and play the lottery.

(4) Who gets told what? After the auction, you won't be told the value of other participants' bids and they won't be told the value of your bid. However, you and the other participants will be told the price and who bought lotteries.

Informed Traders

As noted above, the experiment will be divided into two parts: A and B. Each part will consist of 10 rounds. Before the experiment starts you will be assigned to a group of 5 or 7 participants who you will play against in the auctions.

In some groups 2 or 3 participants will be selected to be *Informed Traders* in one or both parts of the experiment. Whether you are assigned to a group with *Informed Traders* and if so, whether you are selected to be an informed trader is determined at random. Everyone in the group will be told whether the group contains *Informed Traders.* If the group does contain *Informed Traders*, everyone in the group will be told how many *Informed Traders* there are in the group and whether or not they are an informed trader.

The *Informed Traders* will be given extra information about where the spinner that determines the lottery outcome stopped before they enter their bid. The screenshot below shows an example of what the informed traders might see.

[Screen shot appeared here]

The other members of the group (*The Uninformed*) will be told that there are *Informed Traders* in the group, but *The Uninformed* will not be given any extra information about where the spinner stopped before they make their bids. The screenshot below shows an example of what they might see.

[Screen shot appeared here]

How you will be told the results of the auction and the outcome of the lottery

Once you and the other participants in your group have submitted their bids, you will be told the result of the auction and the outcome of the lottery. The screenshot below shows and example of what you might be shown. (The numbers are just examples and contain no significance beyond this.)

[Screen shot appeared here]

When you click continue on the 'Round Results' screen, you will be shown a summary of the results of the experiment so far. The screenshot below shows an example of what you will see.

[Screen shot appeared here]

It shows that in round 1, you and Player #2 bought a lottery for 48 credits and it paid out 63 credits, so you both made a profit of 15 from buying the lottery. Likewise in round 2 you made a profit of 35 credits from buying the lottery. In round 3 you did not buy the lottery. Finally, in round 4 you bought the lottery for 40 credits but it paid out zero, so you made a loss of 40 credits. (These numbers are just examples and contain no significance beyond this.)

When you click continue on the 'Results so far' screen you will begin the next round.

Appendix H

Screenshots

H.1 Enter Bid: Symmetric Information

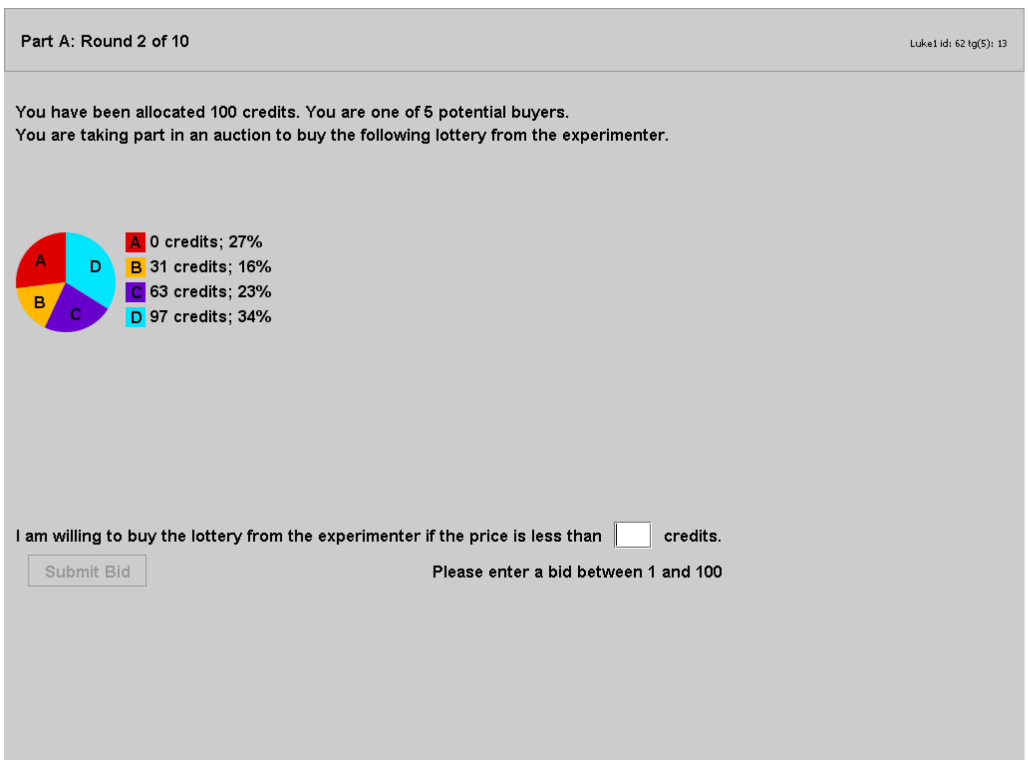

H.2 Enter Bid: Asymmetric Information Informed

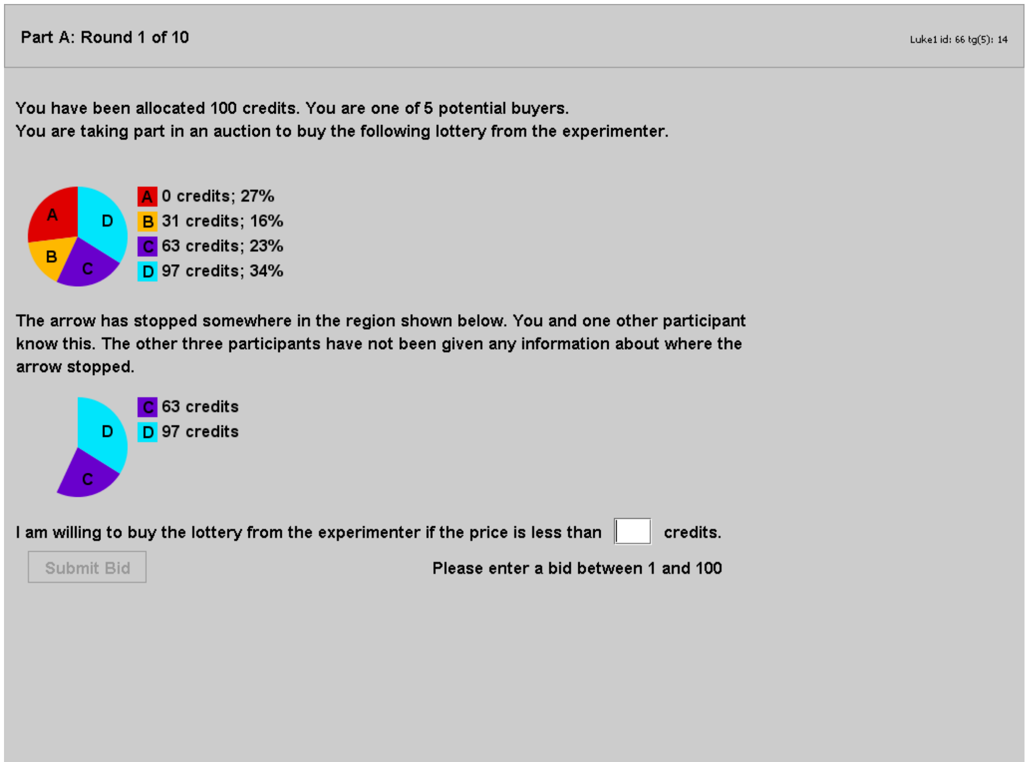

H.3 Enter Bid: Asymmetric Information Uninformed

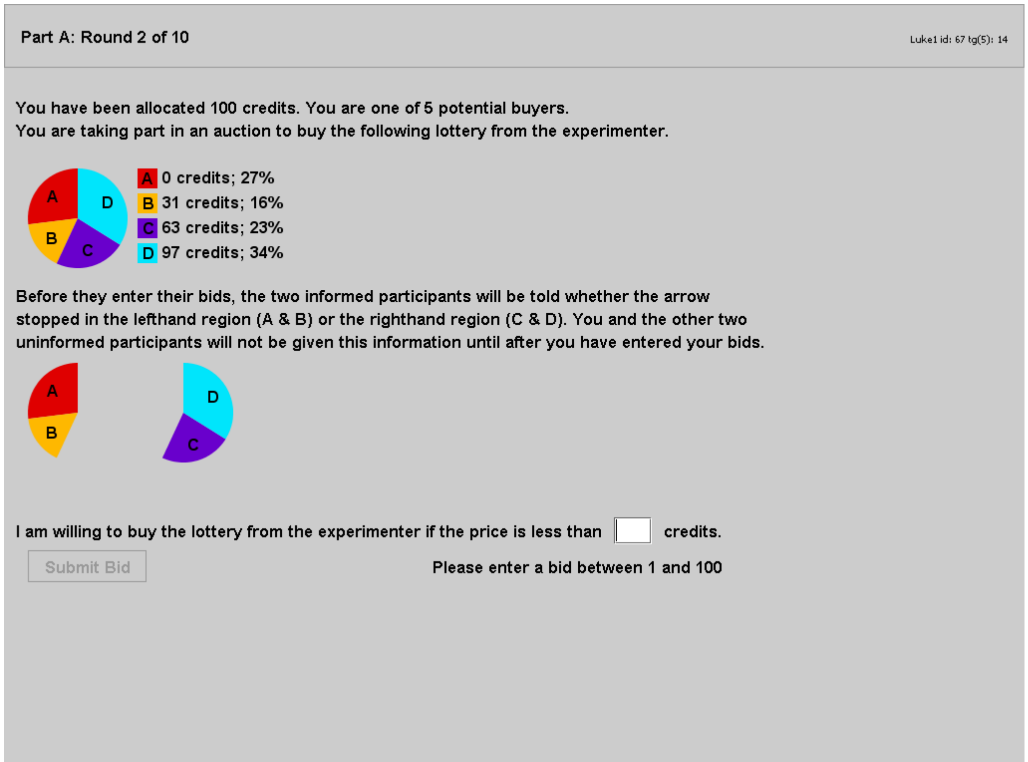

H.4 Results: Current Round

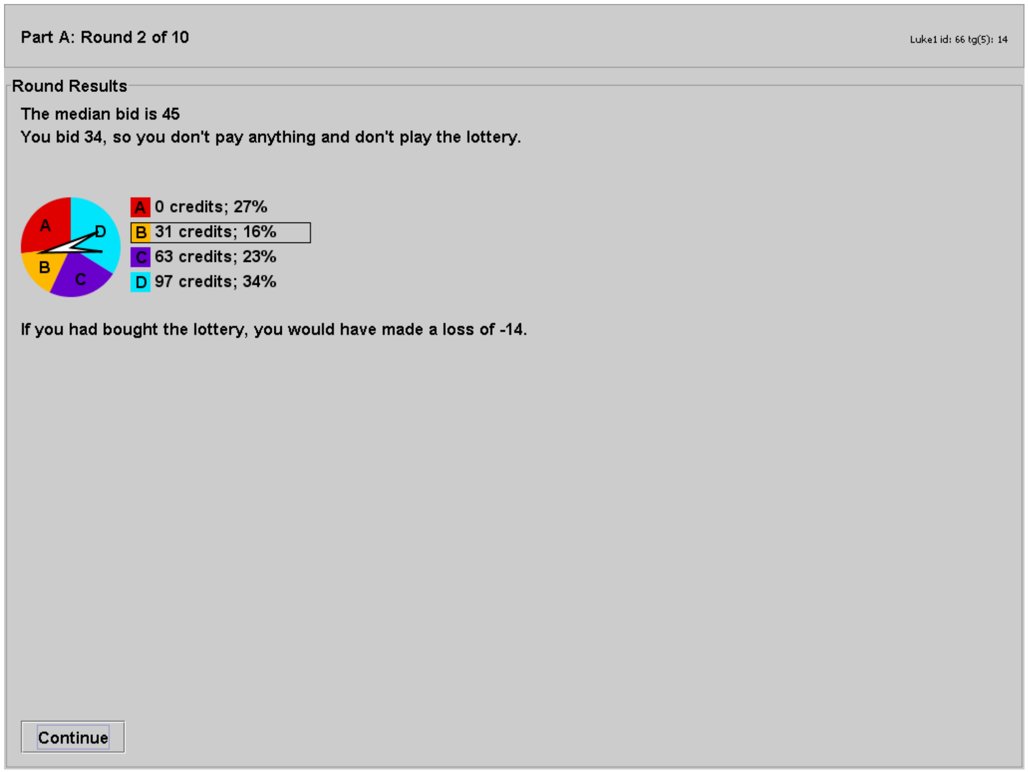

H.5 Results: All Rounds Completed

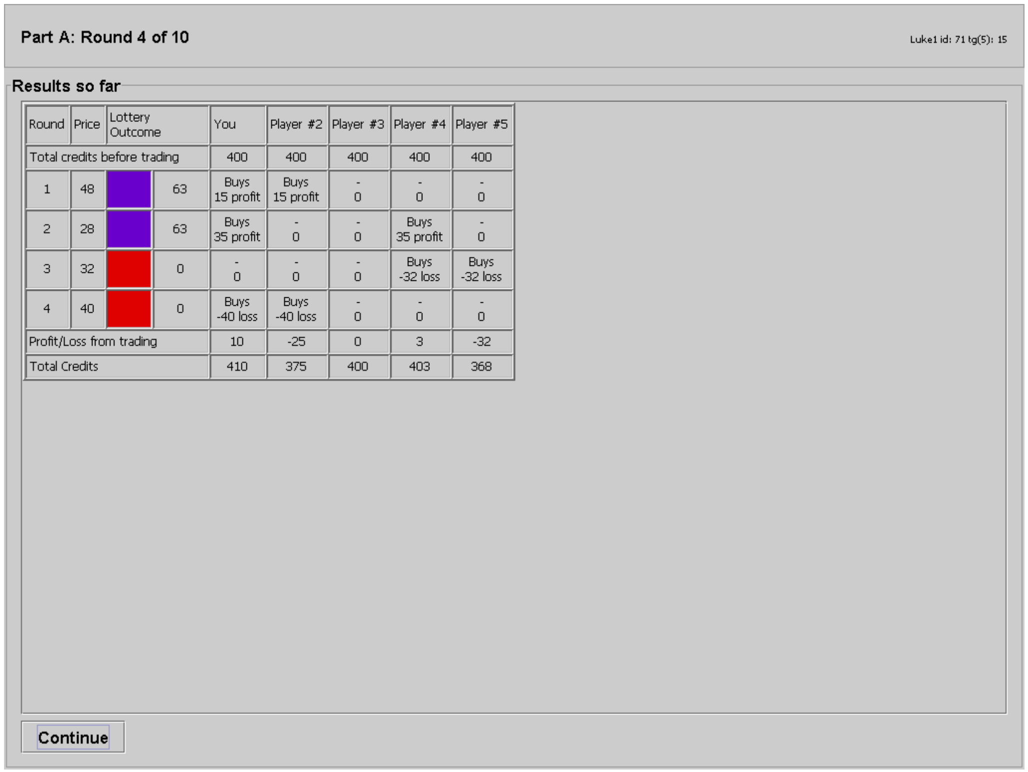

Part A: Round 4 of 10

Luke1 id: 71 tg(5): 15

Results so far

Round	Price	Lottery Outcome	You	Player #2	Player #3	Player #4	Player #5
Total credits before trading			400	400	400	400	400
1	48	63	Buys 15 profit	Buys 15 profit	- 0	- 0	- 0
2	28	63	Buys 35 profit	- 0	- 0	Buys 35 profit	- 0
3	32	0	- 0	- 0	- 0	Buys -32 loss	Buys -32 loss
4	40	0	Buys -40 loss	Buys -40 loss	- 0	- 0	- 0
Profit/Loss from trading			10	-25	0	3	-32
Total Credits			410	375	400	403	368

Continue

Chapter 4

How the Presentation of Odds and Feedback Affect Anomalies in Repeated Markets for Lotteries: The Drawing Balls from Urn Experiment

4.1 Introduction

In the experiment described in the previous chapter, I found that whether the WTA/WTP gap persisted in a repeated market depended on the market regime. Under symmetric information, the gap decayed; under asymmetric information, it persisted. When people behave optimally, we would expect no gap under symmetric information but a gap under asymmetric information, and this pattern of behaviour was observed in the later rounds of the experiment although the WTA/WTP gap under asymmetric information was smaller than standard theory predicts. There was some evidence that subjects adjusted their bids in response to previous trading success. Whether a trade was profitable depended on the price and the resolution of the lottery. There is evidence from other studies that lottery feedback plays a role in reducing anomalies. Braga et al. (2008) find that preference reversals persist in a repeated market without lottery feedback. In a repeated market with lottery feedback, the relative number of standard preference reversals decreases. With just enough feedback, standard and non standard preference reversals are equally common. With extended feedback, non standard reversals become more common.

Why might lottery feedback matter? Braga et al. propose the *loss experience hypothesis* that people will reduce their valuation of lotteries following experience of losing (playing a lottery and receiving a payout of zero or less). They find that experiment subjects reduce their valuation of lotteries in a repeated market following experiencing a loss. Another explanation is that the way probability information is presented influences valuations of lotteries. Likelihood information presented in the form of probabilities and payoffs might not give subjects a sufficient understanding of the lottery. Perhaps observing a lottery being resolved in successive rounds of a repeated market gives subjects a better understanding of the lottery. There is some experimental evidence that supports this view. Bleaney and Humphrey (2006) find that subjects assign higher values to lotteries when they are shown a representative sequence of outcomes rather than just being told the probabilities of each outcome. Furthermore, Humphrey (2006) finds that showing subjects a representative sequence of lottery outcomes makes violations of expected utility theory go away in some cases but emerge in others.

This chapter presents an experiment that investigates how lottery feedback and the presentation of likelihood information influence anomalies in a repeated market. A 2×2 design is used. Lottery feedback is varied by in some treatments resolving the lottery after each round while in others resolving the lottery at the end of the experiment. The presentation of likelihood information is varied by in some treatments informing the subjects of the exact probability figures while in others letting the subjects estimate the odds. In the first case, the lottery has known odds (a risky lottery) while in the second it has unknown odds (an uncertain lottery).

The experiment uses a novel auction mechanism designed to control for common value effects. Why might common value effects occur? First, if, as suggested above, stating probabilities and payoffs as figures does not give people sufficient information to understand lotteries, there could be common value effects even when people are trading risky lotteries. Second, in the unknown odds treatments, unless subjects correctly estimate the odds, there will be common value effects. If common value effects occur, then the standard theory no longer rules out a WTA/WTP gap.

The experiment primarily investigates preference reversals and the WTA/WTP gap. First it tests whether Braga et al.'s result is robust to changing the auction mechanism and varying the presentation of odds. Second, it tests whether the conditions that determine whether or not preference reversals get eliminated in a repeated market have the same effect on the WTA/WTP gap. Although not the main focus of the study, the experimental design also allows investigation of two other anomalies: ambiguity aversion and probability weighting effects.

There have been a number of other studies that investigate anomalies in repeated markets. Among the novel features of this study are (1) four anomalies are studied simultaneously, (2) the evolution of buying and selling prices with and without lottery feedback are compared, (3) a new auction mechanism is used to elicit buying and selling prices which controls for winners curse type effects.

The main findings are, first, that different anomalies persist and decay under different market regimes and, second, that valuations in the final rounds of the markets and choices at the end of the experiment vary considerably between treatments.

The remainder of this chapter is organised as follows. Section 4.2 argues that the second price and median price auctions used in previous repeated market studies might not produce meaningful results when the item traded has an uncertain common value. Section 4.3 presents a new auction to control for common value effects. The experiment is presented in section 4.4, the results in section 4.5, and the discussion in, 4.6.

4.2 Common Value Effects in Second and Median Price Vickrey Auctions

As outlined above, in two of the treatments lotteries with unknown odds are traded. In these treatments market participants are given a private signal that is correlated with the odds and hence the expected value of the lottery. Under these conditions, there is the potential for a *winner's curse* problem to occur. Consider the following stylised model. Five people are bidding to buy a box containing an unknown amount of money V. Each of them receives a private signal X_i such that $X_i \in \{V-2, V-1, V+1, V+2\}$. For simplicity, assume each of the signals is received by at least one bidder (and one of them is received by two bidders) and all mappings of bidders to signals that satisfy this are equally likely. Assume V is uniformly distributed between 0 and 100. This means that X_i is in the range -2 to 102. When X_i is in the range 2 to 98, $E[V|X_i = x_i] = x_i$ (appendix I shows how $E[V|X_i = x_i]$ and X_i are related). That is, the expectation of the amount of money in the box V conditional on observing one of the signals is equal to the observed signal. For the remainder of this section, assume X_i falls within this range.

What happens in a sealed bid second price buying auction (where the highest bidder buys the item for the second highest bid; ties are resolved by tossing a coin)? If each bidder simply bids their expectation of V, the highest bid will be $V+2$. A bidder who observed the signal $V+2$ will win the auction. If two bidders observed the signal $V+2$, the second highest bid and hence the price will be $V+2$. Otherwise, it will be $V+1$. In both cases, the amount of money in the box will be *less* than the amount the winning bidder expected and *less* than the amount they paid for the box. They have suffered from the winners curse! They have failed to account for the fact that winning the auction reveals that the signal they observed was one of the highest of the five signals.

In a second price selling auction (where the lowest bidder sells the item for the second lowest bid; ties are resolved by tossing a coin), the converse occurs. If bidders simply bid their signal, they will sell for $V-2$ or $V-1$. In both cases, the amount of money in the box will be *more* than the selling bidder expected and *more* than the amount they sold the box for. They have suffered from the selling version of the winners curse.

A similar effect occurs in median price auctions, although the intuition is more subtle. In a median price auction with five participants, the median bid is the price and the bidders with the two highest valuations each buy an item (again, ties are resolved by tossing a coin). What happens when bidders bid signal? The median of the signals will be $V-1$ or $V+1$ with equal probability, so the mean price will be V. On average, the amount of money in the box will be *less* than the winning bidders expected but equal to the amount they paid. The winners curse has not occurred in the same way as under the second price auction. However, everyone bidding signal is not a Nash equilibrium. If bidder 1 bids $x_1 - 1.5$ and the others bid x_i, bidder 1 now only trades at a profit. The cases where bidder 1 had previously traded at a loss were as follows: when the median bid was $V+1$ and bidder 1 observed signal $V+1$ or $V+2$ and bid signal. When they bid $x_1 - 1.5$, they no longer trade in either of these cases. The cases where bidder 1 had previously traded at a profit were as follows: when the median bid was $V-1$ and bidder 1 observed signal $V+1$ or $V+2$. Shading their bid by 1.5 in these cases does not reduce it enough to stop them trading.

The consequence of the winners curse is as follows: in buying auctions, bidders have a reason to bid below their expectation of the item's value conditional on their private signal; in selling auctions, they have a reason to bid above it. Notice that there now is a reason for buying and selling bids to diverge in the same direction as there is when bidders exhibit the endowment effect.

How should one bid in a common value auction? In a symmetric equilibrium for a common value auction, on observing a signal x you should bid an amount b such that if you were to just win the auction (i.e. the highest competing bid and hence the price were also b) then you would just break even. If the highest competing bid is equal to your bid, then by symmetry you can deduce that the bidder who placed the bid observed the same signal as you, and you should update your estimation of the object's value conditional on this new piece of information.

What information does *just* winning the auction reveal? Consider the problem from bidder 1's perspective. Let Y_1, Y_2, Y_3, Y_4 be the largest, second largest, ..., smallest from among the signals received by the other four bidders, X_2, X_3, X_4, X_5. If bidder 1 just wins a second price buying auction, then $y_1 = x_1$. If they just win a second price selling auction, then $y_4 = x_1$. Just winning the buying auction does not reveal the same information as just winning the selling auction. The same is true of median price auctions. If bidder 1 just wins a median price buying auction, then $y_2 = x_1$; if they just win the corresponding selling auction, then $y_3 = x_1$. For both second and median price auctions, just winning the buying and selling auctions reveals different information. Consequently, the expectation of the object's value conditional on observing a particular signal and just winning the auction will differ between buying and selling, as will the optimal bids.

4.3 A New Pair of Auctions to Control for Common Value Effects

I have argued that the claim that buying and selling auctions should produce the same bids does not hold for second and median price auctions when the object traded has an uncertain common value and traders receive private signals of its value. For buying and selling auctions to produce the same pattern of bids, the information revealed by winning the auction must be the same for both auctions. If the pair of auctions used comprises (a) second price buying and selling auctions or (b) and median price equivalents, this condition does not hold. It does hold, however, if the pair of auctions used comprises a k^{th} price selling auction and a $(k+1)^{th}$ price buying auction (see proof in appendix J). This is because for bidder 1 just winning the selling auction reveals that $y_k = x_1$ (where y_k is the k^{th} highest signal among those observed by the other bidders). In a $(k+1)^{th}$ price buying auction, just winning also reveals $y_k = x_1$. Hence, the expected value of the item conditional on just winning is the same in both auctions, so both auctions should produce the same pattern of bids.

What value of k should be chosen? Since the winner's curse is not the primary focus of this chapter, we should choose a value that minimises its possible effect. Ideally, we would choose k so that $E[V|X_1 = x_1] = E[V|X_1 = x_1 = y_k]$. When this condition holds, bidder 1's expectation of the auctioned object's value is independent of whether they win the auction, so it is optimal to simply bid expected value conditional on the observed signal. For an experiment, this is desirable for two reasons. First, we can meaningfully interpret bids as private valuations. Second, even participants who do not understand the winners curse and bid expected value given their signal will behave the same as those who do understand it. What value of k is most likely to have this property? Suppose V is uniformly distributed over some interval and signals are symmetrically distributed about V. Further, suppose that for some range of values, X_1 and the median of $\{X_1, ..., X_N\}$, $\tilde{X}$, are both unbiased estimators of V. That is $E[V|X_1 = x_1] = x_1$ and $E[V|\tilde{X} = \tilde{x}] = \tilde{x}$. For signals within this range, the auction has the desirable property if just winning the auction implies $\tilde{x} = x_1$. This is satisfied by auctions with the property $N = 2k$.

An auction with four participants (i.e. $N = 4$) is used in the experiment. It has the desirable property when $k = 2$. This means that in the selling auction, the second, k^{th}, highest bid is the price and the two participants with the two lowest bids sell at the price. In the buying auction, the price is the third highest bid, $(k+1)^{th}$, and the participants with the two highest bids buy at the price.

4.4 The Experiment

A total of 268 people took part in the experiment. They were recruited[1] from the student population at the University of Nottingham. For each of the four treatments four sessions were run giving a total of 16 sessions. The sessions for the treatments with unknown odds were run before those for treatments with known odds to prevent knowledge of the odds leaking. Each session had either 16 or 20 subjects who were randomly assigned to trading groups of four. The subjects were given a paper copy of the instructions (see Appendix K.1) and then the experimenter read the instructions aloud. Then the subjects completed a comprehension test (see Appendix K.2) which was checked by the experimenter before proceeding.

The two following lotteries were used: a *\$-bet* offered a 19% chance to win £17 (and a 81% chance of receiving nothing); a *P-bet* offered a 81% chance to win £4 (and a 19% chance of receiving nothing).[2]The expected value of the \$-bet is £3.23 and the P-bet, £3.24. The random device used to resolve the lotteries was an urn containing 100 balls: in some sessions 19 were orange and 81, white; in others, 19,

[1]The online recruitment management system ORSEE (Greiner, 2004) was used.

[2]Note this pair of lotteries is similar to those used by Loomes et al. (2007) in a related experiment.

white and 81, orange. The subjects were shown the balls separated by colour in two transparent containers. This let the subjects in the treatments with unknown odds arrive at a private estimate of the number of balls of each colour. The balls were then placed in a non transparent urn placed in the centre of the room and mixed thoroughly.

Each subject participated in four markets: buying and selling each of the two lotteries. After all the markets had been completed, subjects were faced with a choice between the P and $ bets. Table 4.1 shows the four possible sequences of tasks followed by subjects in the experiment. In all sessions at least one trading group followed each sequence (in sessions of 16 exactly one followed each sequence; in sessions of 20 one sequence was followed by two groups).

Each of the four markets had 10 trials.[3] Bids were entered via a computer terminal (appendix L.1 shows an example of the screen subjects saw). In selling markets, subjects were endowed with the lottery and asked to state the minimum they would be willing to accept to sell the lottery to the experimenter. In buying markets, subjects were endowed with the maximum the lottery could pay out and asked to state the maximum they would be willing to pay to buy the lottery.[4]When all subjects in a trading group had entered their bids, the market price for the round was determined. Subjects were told onscreen the outcome of the auction and whether they had traded (appendix L.2 shows an example of the screen subjects saw). In the no feedback treatments, subjects proceeded directly to the next round of the market. In the feedback treatments, once all trading groups had completed the auction, the experimenter drew a ball from the urn in the centre of the room. The result of the

[3] In some earlier studies investigating whether anomalies decay in repeated markets there were fewer than 10 trials and typically 5 or 6 e.g (Loomes et al., 2003, 2007; Grether and Cox, 1996) while others had 10 or more e.g. (Shogren et al., 2001; Braga et al., 2008). Having 10 trials provides a more rigorous test of the hypotheses.

[4]Schmidt and Traub (2006) argue that with this pattern of endowments under standard utility theory and when the degree of absolute risk aversion is non-increasing, if income effects occur at all, they will result in WTP exceeding WTA.

Table 4.1: How Tasks Were Ordered

Part	Sequence A	Sequence B	Sequence C	Sequence D
1	Buy P	Sell P	Buy $	Sell $
2	Buy $	Sell $	Buy P	Sell P
3	Sell P	Buy P	Sell $	Buy $
4	Sell $	Buy $	Sell P	Buy P
5	Choice Task			

The columns - labelled A, B, C, and D - represent sequences in which tasks could be completed. Each trading group followed one of these four sequences. In each treatment, trading groups were distributed evenly between the four sequences.

draw was recorded and subjects were shown the consequences of the draw for them on screen (see appendix L.3).

After subjects had completed each of the four markets, they faced a choice between the two bets (see appendix L.4). In the feedback treatments, once everyone had made a choice, the lottery was resolved. At this point, subjects had finished all the decision tasks in the experiment. One task (either one of the 40 auction rounds or the choice task) was selected at random to be the task for which subjects were paid. In the no feedback treatments the lottery was resolved to determine payouts in the selected task (in the feedback treatments, the lottery had been resolved immediately after each task, so the payout in the selected task was already determined). Each experimental session lasted about 70 minutes. The payments receive by subjects ranged from £0 to £33.30 with a mean of £7.12 .

4.5 Results

Section 4.5.1 describes what data was collected and then sections 4.5.2 to 4.5.5 analyse the data with respect to specific behavioural anomalies.

Figure 4.1: The Evolution of Bids: All Treatments

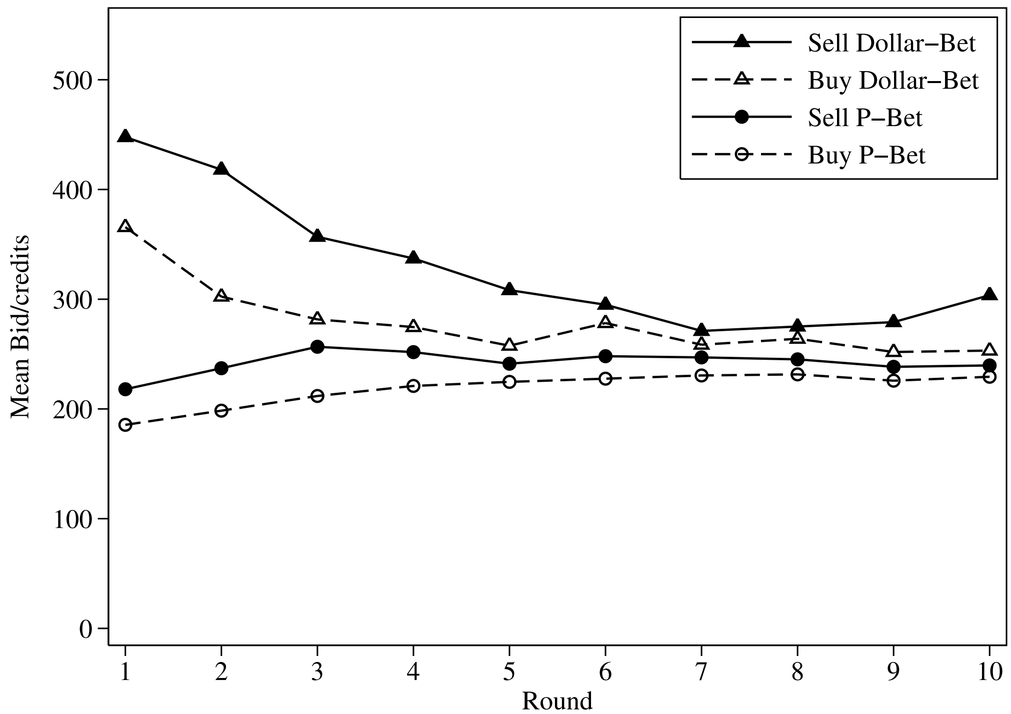

4.5.1 The Data Collected

Valuations Each of the 268 subjects took part in four markets (buying and selling the P-bet and $-bet). Each market consisted of an auction that was repeated 10 times. This gives a total of 10720 bids. The mean of bids across rounds are shown in figure 4.1. Tables 4.3 and 4.4 show the mean first and last round bids by treatment. Figure 4.8 shows mean last round bids by treatment graphically.

Choices The final task in the experiment was a choice between the P-bet and the $-bet. The results by treatment are shown in figure 4.2.

Time taken to make decisions As well as the actual decisions made by subjects, the time taken to make decisions was also recorded. This was done by recording the time (in milliseconds) at which information was displayed on subjects

Figure 4.2: Choices and Final Round Valuations by Treatment

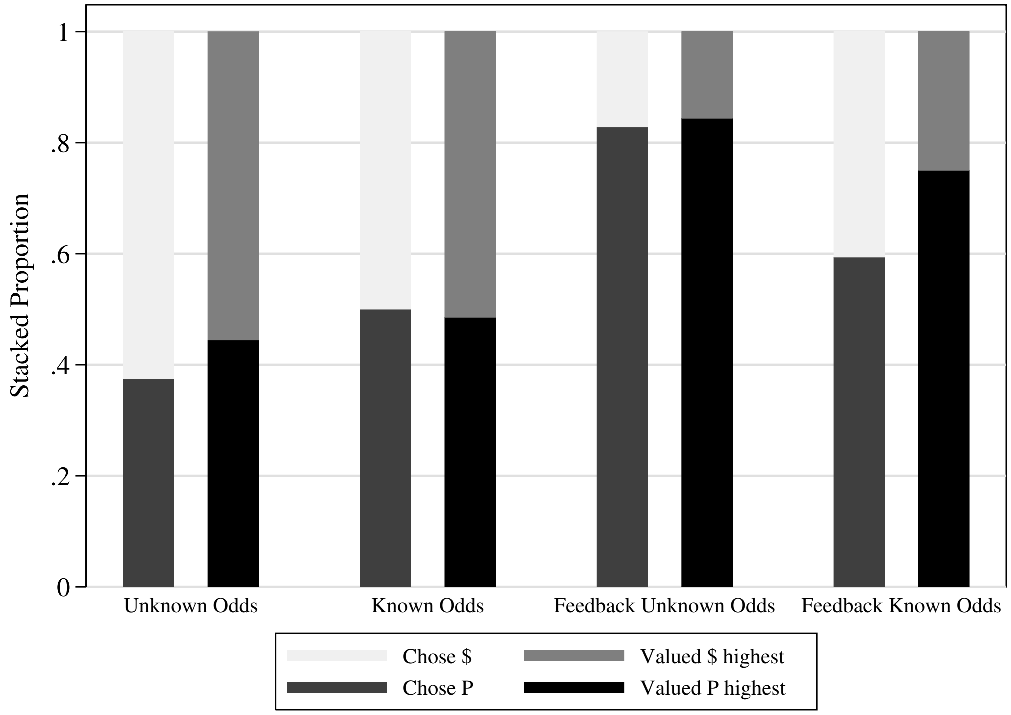

computer screens and the time at which subjects' clicked on onscreen buttons. For example, the time taken to enter bids was recorded as time between when the enter bid screen (L.1) for a given round was displayed and when the subject had entered a valid bid and clicked on *submit bid*. Figure 4.3 shows how the mean time to enter bids varied across the 40 auction tasks completed by each subject. Task 1 was the first auction subjects completed. Tasks 11, 21, and 31 are where subjects changed between trading P-bets and $-bets. Task 21 is where they changed between buying and selling as well as changing lotteries.

Comprehension Test and Debrief Survey Subjects answers to the comprehension test K.2 and debrief survey were collected. The only data reported are the estimates of the odds stated by subjects in the unknown odds treatments. Table 4.2 shows that mean estimates were close to the true value (the probability of

Figure 4.3: The Time Taken to Enter Bids in Successive Tasks

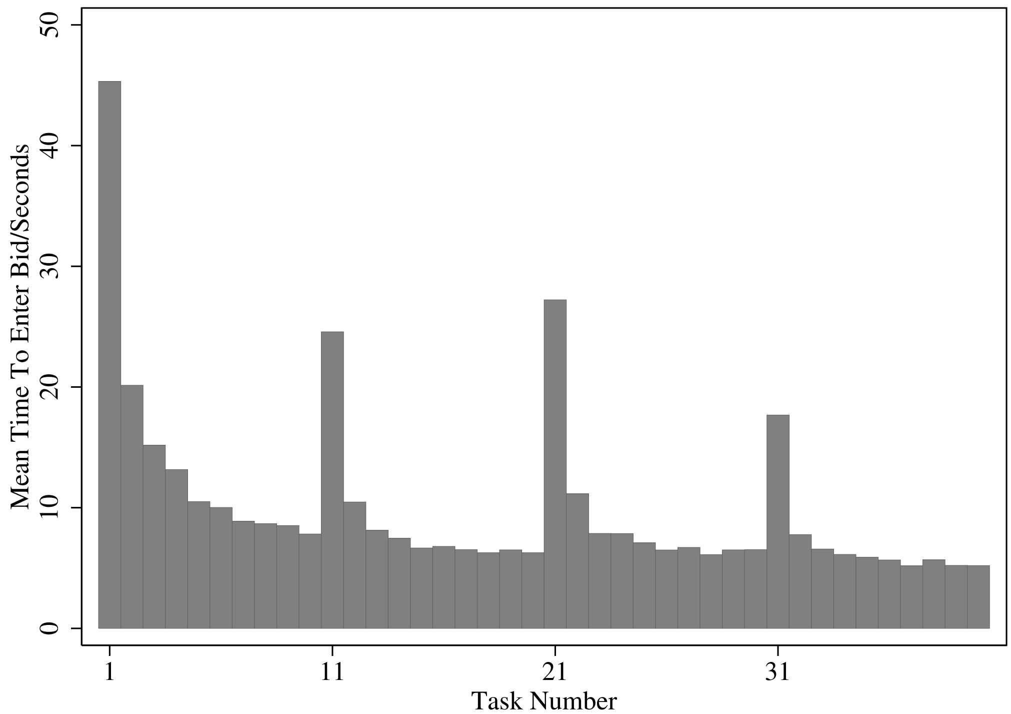

the $-bet paying out was $\frac{19}{100}$).

4.5.2 WTA/WTP Gaps

Result 4.1 *In the first round of the markets, there was a WTA/WTP gap.*

Support. In the first round of the markets for the P-bet and the $-bet bids in selling auctions exceeded those in buying auctions. Table 4.3 shows that for the P-bet, on

Table 4.2: Subject's Estimates of the Odds of Winning the $-Bet

Treatment	N	mean estimate	Std. Dev.	Min	Max
Unknown odds	64	22.5%	8.5	10.0	50.0
Unknown odds & Feedback	60	19.6%	4.6	10.0	33.3
All	124	21.1%	7.0	10.0	50.0

The reported figures are subjects' responses to the debrief survey, which is shown in appendix K.3.

average across the four treatments bids in selling auctions were 32 credits higher than in buying auctions; table 4.4 shows that for the \$-bet, on average, they were 82 credits higher. In both cases the difference is statistically significant.

Result 4.2 *In the treatments without lottery feedback, the gap tended to close over successive rounds; in the treatments with lottery feedback, it did not.*

Support. When the results of the four treatments are pooled, mean bids in buying and selling converge over successive rounds. Figure 4.1 shows this convergence. Although the gap decays, it does not close completely. The final rows of tables 4.3 and 4.4 show the size of the gap and that it was statistically significant in markets for both bets. Table 4.5 shows the correlation between the gap size and the number of rounds completed. For each subject in each of the ten rounds a measure of the gap was calculated as follows. $gap = bid\,selling\,\$ + bid\,selling\,P - bid\,buying\,\$ - bid\,buying\,P$ The resulting figures were then ranked from highest to lowest and the coefficient of correlation ρ between rank and round number calculated (the second column of the table). The null hypothesis of no correlation was tested against the alternative of either positive or negative correlation. The p-values reported in the last column of the table are the probability of obtaining correlation at least as strong as the observed results under the null hypothesis. The last row of the table reveals that the obvious convergence we see in figure 4.1 is statistically significant.

When the results of the four treatments are disaggregated, two patterns emerge. Figure 4.4 shows that on average in treatments without feedback, the buying-selling gap completely closes after 5-6 rounds. In contrast, figure 4.5 shows, in treatments with feedback, the gap persists over all 10 rounds. Returning to table 4.5, we see a similar pattern. In all treatments there is some degree of negative correlation between the gap size and round number. The correlation is stronger and statistically significant in treatments without feedback.

Table 4.3: Mean Bids in P-Bet Auctions

Treatment	First Round buy	sell	gap	Last Round buy	sell	gap
No feedback, unknown odds	187	219	32**	180	176	−4
No feedback, known odds	174	219	44**	223	206	−17
Feedback, unknown odds	189	221	33	248	290	42**
Feedback, known odds	193	213	20	274	298	24
All	186	218	32***	229	240	10**

The null hypothesis of gap = 0 is tested against the alternative hypothesis gap > 0 using a one sided t-test. Significance levels: *:10 percent; **: 5 percent; *** 1 percent.

Table 4.4: Mean Bids in $-Bet Auctions

Treatment	First Round buy	sell	gap	Last Round buy	sell	gap
No feedback, unknown odds	360	503	143**	316	388	72
No feedback, known odds	409	386	−24	348	227	−121
Feedback, unknown odds	361	445	84	116	269	153**
Feedback, known odds	331	454	123**	220	326	106**
All	366	448	82***	253	304	50**

The null hypothesis of gap = 0 is tested against the alternative hypothesis gap > 0 using a one sided t-test. Significance levels: *:10 percent; **: 5 percent; *** 1 percent.

Table 4.5: WTA/WTP Gap Trends across Rounds

Treatment	n	Spearman's ρ	p value
Unknown odds	720	-0.10	0.009
Known odds	680	-0.10	0.008
Unknown odds & feedback	640	-0.02	0.557
Known odds & feedback	640	-0.05	0.254
All	2680	-0.07	<0.001

Spearman's ρ is non-parametric measure of correlation. The p-values are the probability of obtaining ρ under the null hypothesis that there is no correlation between the size of the WTA/WTP gap and number of rounds completed.

Figure 4.4: The Evolution of Bids: Treatments without Lottery Feedback

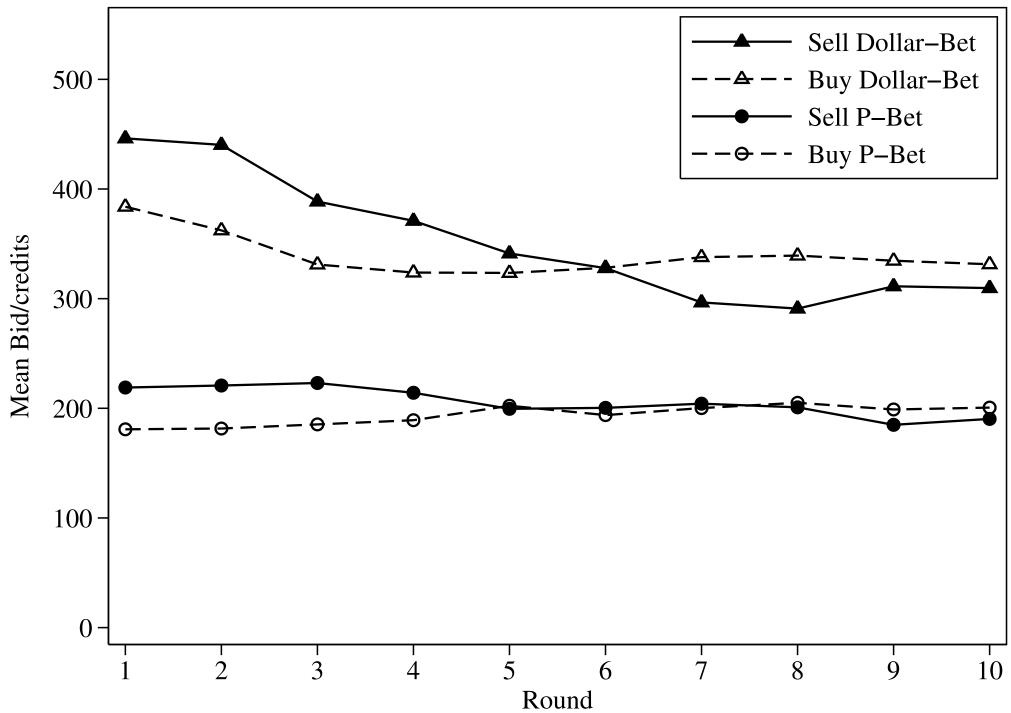

Figure 4.5: The Evolution of Bids: Treatments with Lottery Feedback

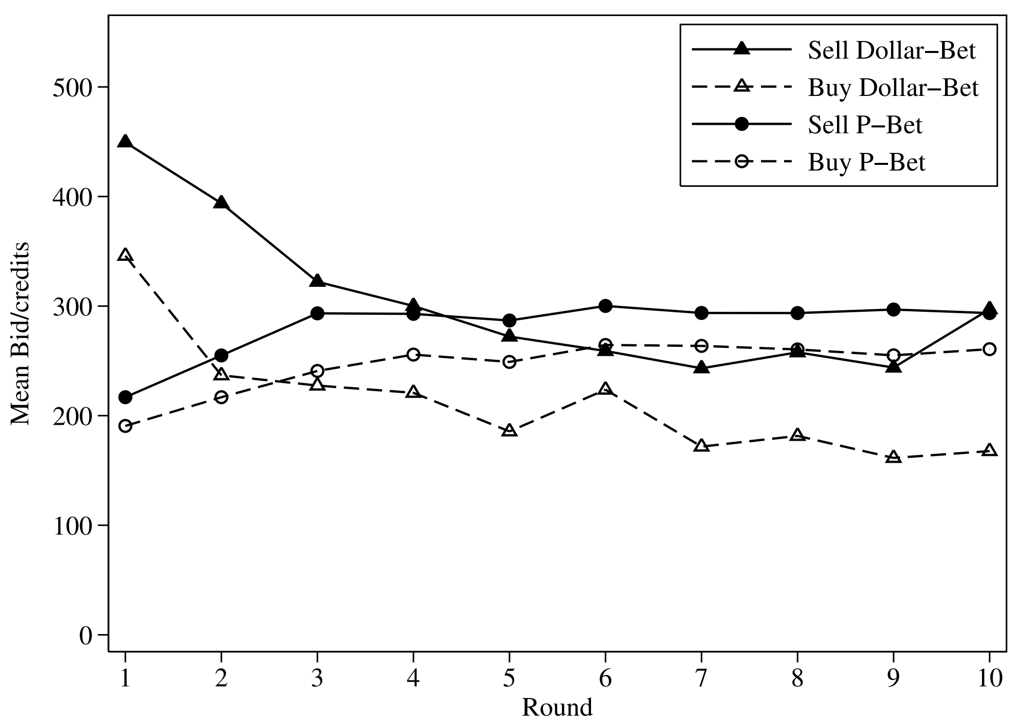

Figure 4.6: The Evolution of Bids: Treatments with Known Odds

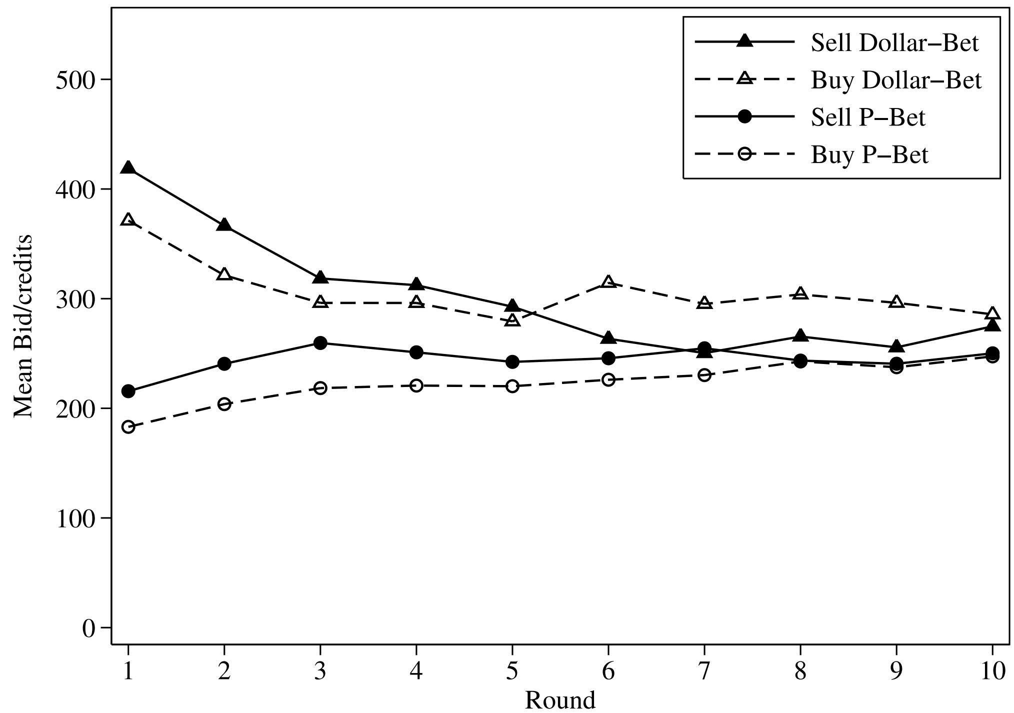

Figure 4.7: The Evolution of Bids: Treatments with Unknown Odds

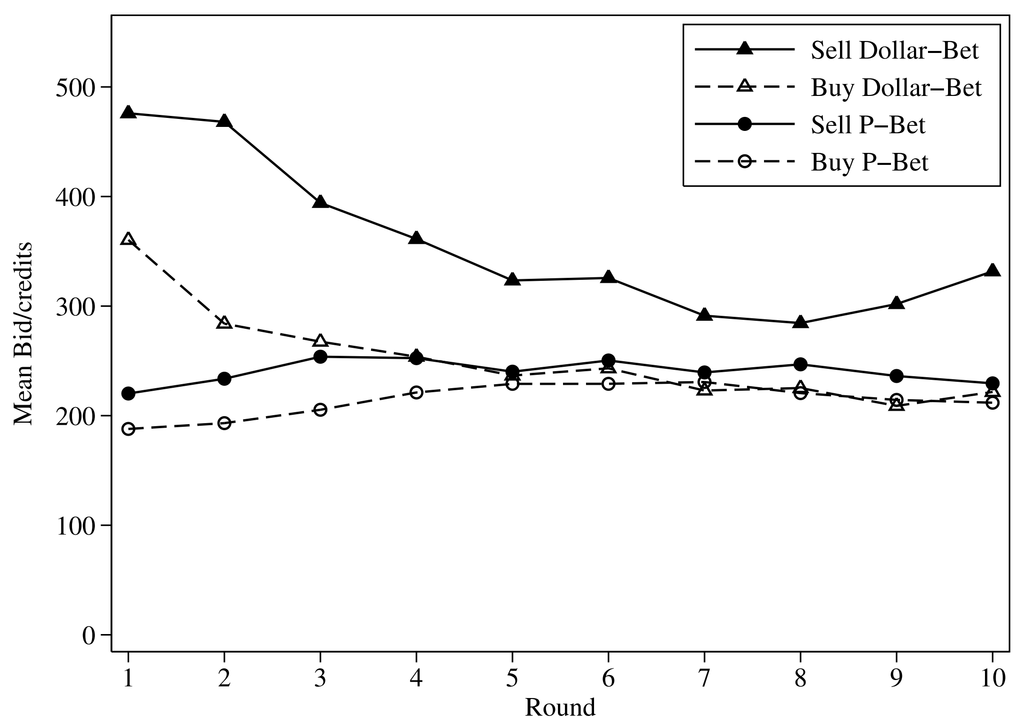

Figure 4.8: Round 10 Bids by Treatment

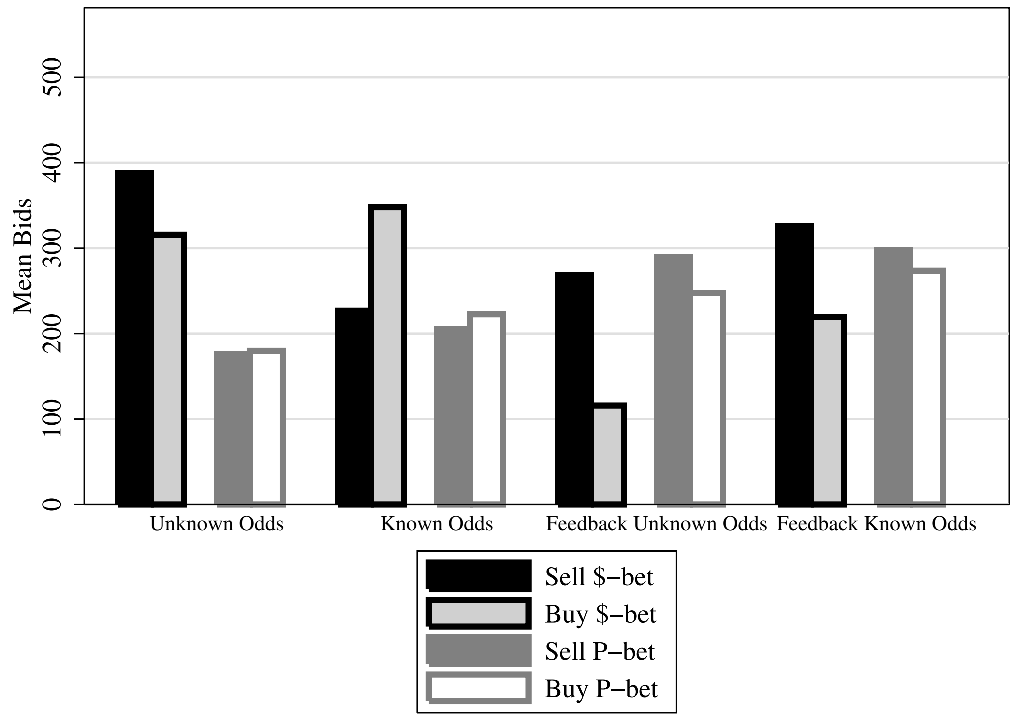

4.5.3 The Effect of Ambiguity

Figures 4.6 and 4.7 show the evolution of bids in treatments with known and unknown odds respectively. Do subjects value lotteries with known odds higher than those with unknown odds?

Result 4.3 *There is evidence of ambiguity aversion in buying but not selling.*

Support. If we just consider either the P or $ bet alone and find that valuations are higher when the odds are known, we can not distinguish between two competing explanations. Maybe the difference is due to ambiguity aversion; maybe it is due to underestimating the odds of winning the lottery with unknown odds. To isolate the effect of ambiguity the following procedure is used. Suppose a subject in the unknown odds treatment estimates that the $-bet will pay out with probability $\hat{p}$, and hence the P-bet pays out with probability $(1-\hat{p})$. If they are an expected utility maximiser with utility function $u(\bullet)$ and $u(0) = 0$, then the expected utility of the P and \$ bets are (a) $0.81u(4)$ and $0.19u(17)$ when the odds are known and (b) $(1-\hat{p})u(4)$ and $\hat{p}u(17)$ when the odds are estimated. We can write

$$v = 0.19 \times \frac{\hat{p}u(17)}{0.19u(17)} + 0.81 \times \frac{(1-\hat{p})u(4)}{0.81u(4)}$$

and for an expected utility maximiser whatever value $\hat{p}$ takes, $v = 1$. Hence, if $v < 1$, it is due to ambiguity aversion not miss-estimating the odds.

The following was used as an estimate of v (the assumptions are that bids are certainty equivalents and the ratio of certainty equivalents approximates the ratio of utilities).

$$\hat{v} = 0.19 \times \frac{Bid\,Uncertain\,\$Bet}{Mean\,Bid\,Risky\,\$Bet} + 0.81 \times \frac{Bid\,Uncertain\,PBet}{Mean\,Bid\,Risky\,PBet}$$

For each subject in the unknown odds treatments, $\hat{v}$ was calculated for rounds 1-10 of buying and selling giving a total of 20 observations per subject. Since these subjects only completed markets for lotteries with unknown odds, the mean of bids from the same round of the corresponding markets in treatments with known odds was used in calculating $\hat{v}$.

Figure 4.9 shows the mean of $\hat{v}$ for buying and selling across rounds 1-10 with and without feedback. The graph shows that subjects are ambiguity averse in buying auctions but not in selling auctions. This could be interpreted in terms of a caution heuristic having two effects. First people dislike ambiguity so value ambiguous lotteries less than risky ones. Second, when faced with uncertainty, people become more averse to trading. In buying auctions, these two effects act in the same direction, resulting in lower prices; in selling auctions, they act in opposite directions, so tend to cancel each other out.

Figure 4.9: Value of Bids for Uncertain Lotteries relative to Bids for Risky Lotteries

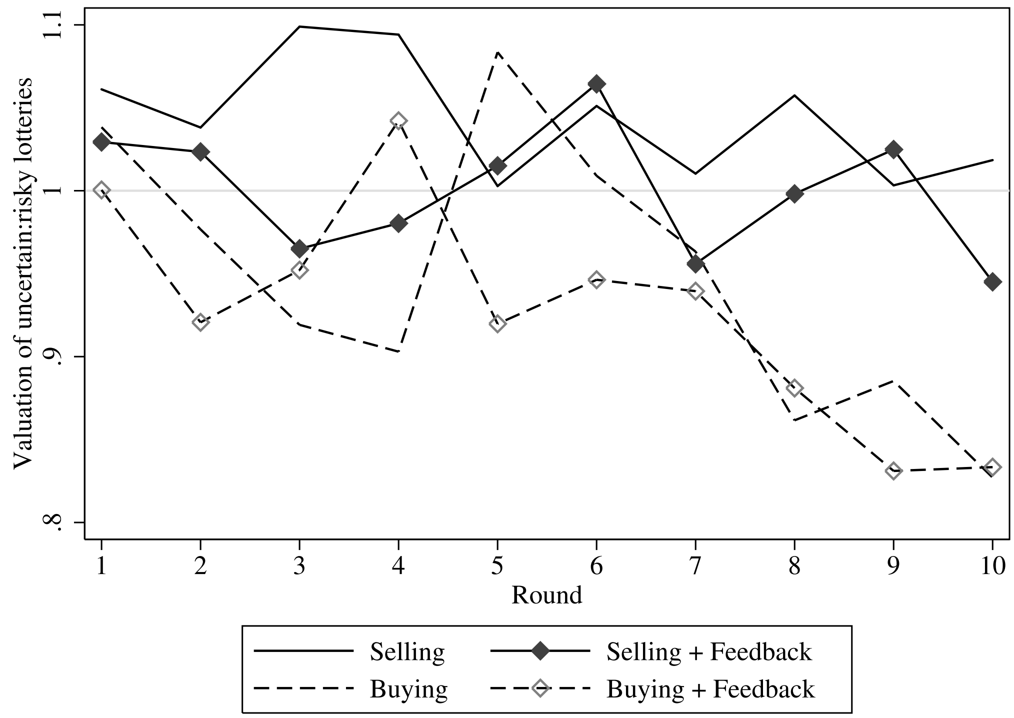

Table 4.9 shows the result of testing for trends in $\hat{v}$ across rounds. Ambiguity aversion increases in buying auctions but remains approximately constant in selling auctions. This is evidence against the hypothesis that anomalies decay in repeated markets.

Table 4.6: Testing for Trends across Rounds in the Valuations of Uncertain Lotteries relative to Risky Lotteries

Treatment	Np-trend z statistic	$Prob[Z > \|z\|]$
Selling	-0.89	0.373
Buying	-2.15	0.032
Selling + Feedback	-0.88	0.376
Buying + Feedback	-2.31	0.021

A negative z statistic indicates that valuations of uncertain lotteries relative to risky ones decrease as the number of rounds completed increases.

4.5.4 Preference Reversals

Under standard theory, if someone prefers the P-bet to the \$-bet, they should do so however the preference is elicited. There are two ways this proposition is tested. First, within subject tests: testing whether the lottery a subject valued highest is also the one they chose. This type of test allows preference reversals to be identified even when there is no systematic pattern to the reversals (i.e. the number valuing the \$-bet highest and choosing the P-bet equals the number valuing the P-bet highest and choosing the \$-bet). Second, between subject tests: testing whether the proportion choosing the \$-bet is the same when preferences are elicited differently. Differences in proportion (beyond that which can be reasonably accounted for by random variation) imply the revealed preference depends on how it is elicited. The absence of differences in proportion, however, does not imply revealed preference are independent of the elicitation method.

Result 4.4 *Standard preference reversals are more common in the early rounds; non standard ones in the later rounds.*

Support. Figure 4.7 shows the number of preference reversals in rounds 1 and 10. Figures for buying and selling are reported. The figures for selling are arguably more meaningful since in buying auctions one could in theory explain the reversals in terms of income effects. In round 1, standard reversals (valuing the \$-bet higher but choosing the P-bet) are at least twice as frequent as non standard ones. By round 10, the number of standard reversals has decreased and non standard reversals are more common.

Figure 4.10 shows how the proportions of the two types of preference reversal changed over rounds 1-10. Information is disaggregated into treatments with and without feedback. The lightest shaded regions show the proportion of subjects who valued the P-bet highest and chose the P-bet at the end of the experiment. The region above this, shaded the second lightest, shows those who valued the P-bet highest but chose the $-bet, so the sum of these two regions is the proportion who valued the P-bet highest. Similarly, the sum of the two darker shaded regions above are those who valued the $-bet highest. In round 1 (as we saw in table 4.7) the proportion who valued the P-bet highest is less than the proportion who eventually chose it over the $-bet. By round 3, the two proportions are equal. At round 10, more subjects valued the P-bet highest than chose it. Over the 10 rounds, the total number of reversals, the sum of the middle two regions remains approximately constant. The trend is that the proportion who value the P-bet highest increases over successive rounds. The trend described above is apparent in both graphs but more pronounced in the treatments with feedback.

Result 4.5 *Final choices and valuations vary across treatments suggesting exposure to a repeated laboratory market is not sufficient for people to discover their true preferences.*

Support. Figure 4.2 shows preferences implied by final round valuations and the choice task by treatment. Figure 4.8 shows final round valuations by treatment. Within treatments, the valuations and choices produce similar proportions favouring each bet. However, there are large differences between treatments. In treatments with feedback, a much larger proportion chose the P-bet and valued it highest. Under the hypothesis that experience of a repeated market causes behaviour to converge on the optimal, we should not expect any differences between treatments. Table 4.8 shows the results of testing a series of hypotheses concerning the differences in choices and valuations between treatments. In 6 of the 7 cases, we can reject the null hypothesis of no difference between treatments at the 0.1 percent level.

Table 4.7: Within Subject Choice-Valuation Preference Reversals

	Reversals	Standard Reversals	Proportion	P-value
Round 1 Buying	118	78	0.66	0.001
Round 1 Selling	123	87	0.71	0.000
Round 10 Buying	101	43	0.43	0.163
Round 10 Selling	106	51	0.48	0.771

A preference reversal is said to occur if in a given round the lottery that a subject values highest is not the same as the one they chose at the end of the experiment. A standard reversal is when the $-bet is valued highest but the P-bet is chosen. The null hypothesis that standard and non-standard reversals are equally likely is tested against the two sided alternative that they are not equally likely. The P-values are the probabilities of observing the obtained or a more extreme result under the null hypothesis.

Table 4.8: Testing for Differences in Choice and Final Round Valuations across Treatments

Hypothesis that the following is equal across treatments	Test	P value
Prob[Choose $-bet]	Pearson's chi-squared	0.000
Prob[Value $-bet highest in buying]	Pearson's chi-squared	0.000
Prob[Value $-bet highest in selling]	Pearson's chi-squared	0.000
Mean bid buying $-bet	Analysis of Variance	0.001
Mean bid buying P-bet	Analysis of Variance	0.000
Mean bid selling $-bet	Analysis of Variance	0.071
Mean bid selling P-bet	Analysis of Variance	0.000

The p values are the probability of observing the obtained statistic under the null hypothesis of no differences between treatments.

Figure 4.10: Preference Reversals without Feedback and with Feedback

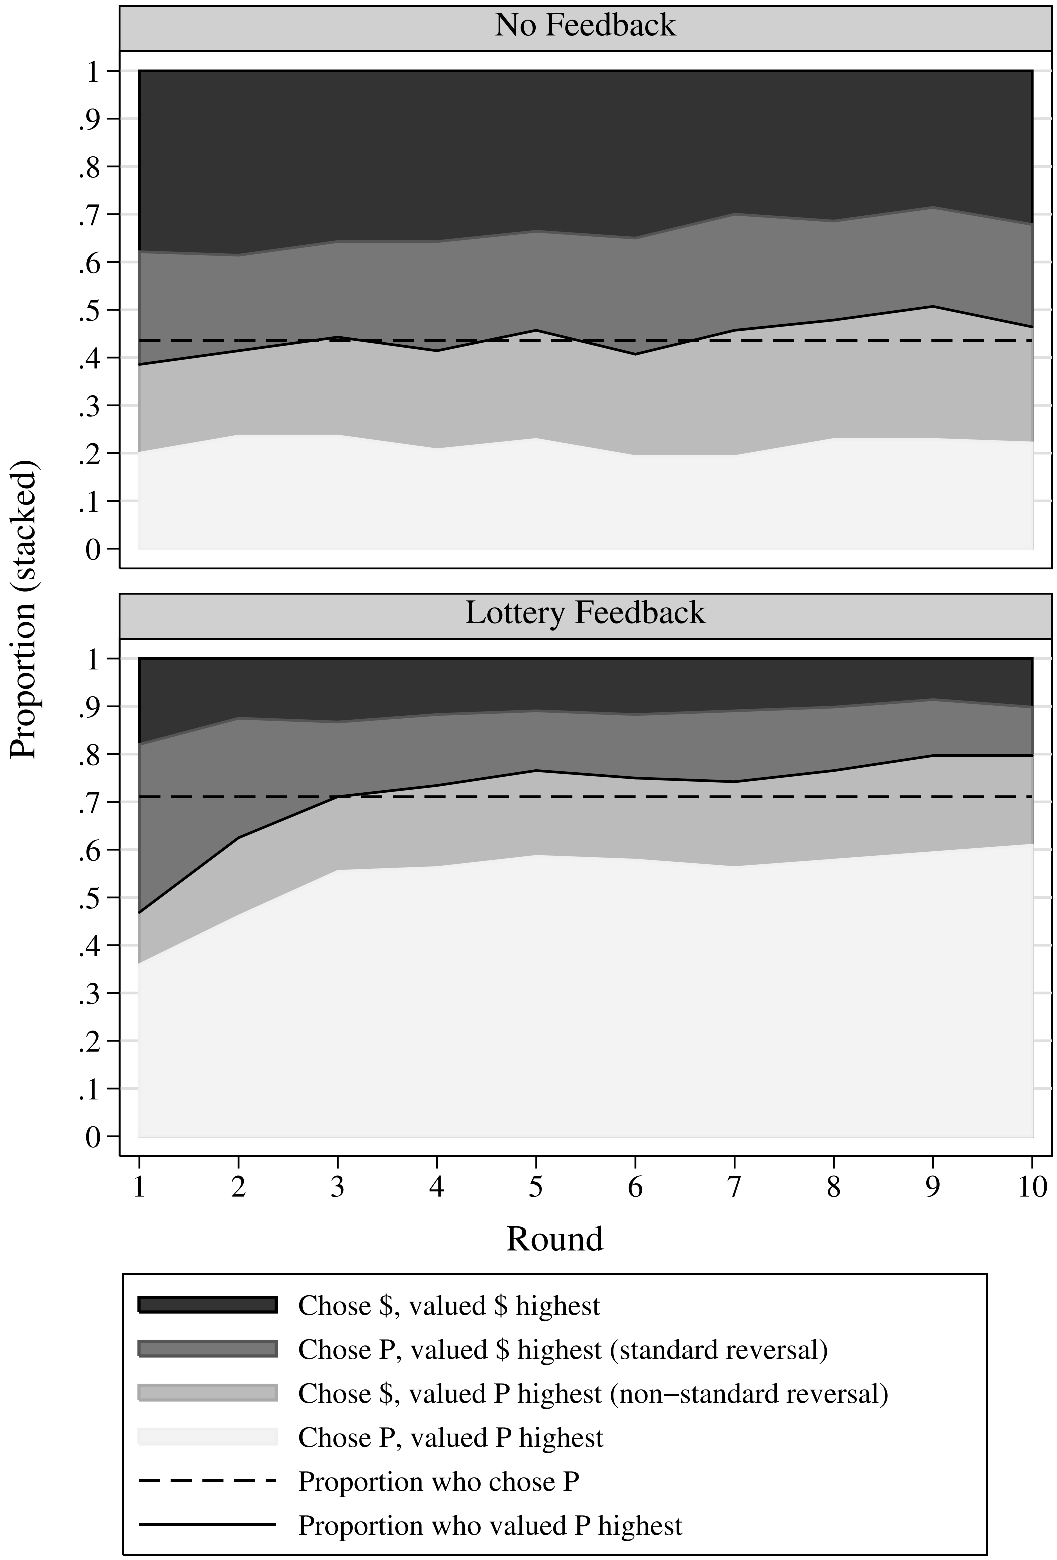

4.5.5 Attitude to Risk and Probability Weighting

Result 4.6 *Over weighting of small probabilities decays in the repeated markets with lottery feedback but not those without it.*

Support. The P and $ bets used in the experiment had the same expected value. A risk averse expected utility maximiser should prefer the P-bet; a risk loving one should prefer the $-bet. Table 4.9 shows implied preferences over the P and $ bets. In round 1 the majority preferred the $-bet. By round 10, the majority in the treatments without feedback still prefer the $-bet but the majority in the treatments with feedback have switched to preferring the P-bet. A possible explanation is that subjects have a tendency to over weigh small probabilities. The $-bet pays out a relatively large amount with a small probability. Subjects over weigh this probability and hence overvalue the $-bet relative to the P-bet. Those in the feedback treatments repeatedly experience the resolution of the lotteries and this experience reduces the tendency to over weigh small probabilities. Those in the no feedback treatments do not have this experience so continue to overvalue the $-bet.

Table 4.9: First and Final Round P-Bet/$-Bet Preferences

Treatment	Prefer $	Prefer P	Indif-ferent	p given risk averse	p given risk loving
Round 1					
All	167	79	22	0.000	1.000
Unknown odds	88	40	8	0.000	1.000
Known odds	79	39	14	0.000	1.000
Round 10					
All	112	135	21	0.937	0.081
Unknown odds	48	17	7	0.000	1.000
Known odds	30	31	7	0.601	0.500
Unknown odds & feedback	16	43	5	1.000	0.000
Known odds & feedback	18	44	2	1.000	0.001

For each subject in each auction round, implied preferences were calculated by comparing their bids in selling auctions for the two bets. The table reports the number of subjects with the specified implied preferences over the P and $ bets. Let *diff = value of P-bet - value of $-bet*. The p values were calculated by running one sided sign tests with the null hypothesis that selling bids for the two lotteries are equal. The *p | risk averse* column shows the probability of observing the obtained results if the median of *diff* is more than or equal to 0. We would expect this if subjects were risk averse. The *p | risk loving* column shows the probability of observing the obtained results if the median of *diff* is less than or equal to 0. We would expect this if subjects were risk loving.

4.6 Discussion

The experiment reproduced several of the results found in previous studies and provides some new insights into how repeated markets affect the occurrence of anomalies. As numerous other studies have found, a disparity between buying and selling prices was observed. In a repeated market where the lottery was not resolved after each round, the disparity decayed and after 5 rounds had been completely eliminated. The first set of novel findings relate to how anomalies decay or persist in repeated markets. In a repeated market where the lottery was resolved each round, the disparity persisted. The reverse pattern emerged with probability weighting effects. The initial anomaly (over valuing the $-bet) persisted in the repeated market where the lottery was not resolved but was eliminated in the treatments where the lottery was resolved each round. Feedback in the form of resolving the lottery was needed to eliminate one anomaly but causes another one to persist.

The experiment also investigated preference reversals. In line with previous studies, initially subjects tended to value the $-bet highest but (at the end of the experiment) choose the P-bet. As subjects completed more auction rounds, they tended to switch from valuing the $-bet highest to valuing the P-bet highest. By the final round, the number of subjects who valued the P-bet highest actually exceeded the number who chose it when faced by a straight choice between the two lotteries. The second set of novel findings is that although the repeated market considerably reduced within treatment preference reversals, it also produced striking between treatment preference reversals. Subjects in treatments with feedback were considerably more likely to value the P-bet highest and choose it over the $-bet. Subjects in treatments with unknown odds valued lotteries lower when buying than subjects in corresponding treatments with known odds. What is more, these effects increased with the number of rounds completed. As a consequence, which treatment a subject was randomly assigned to considerably influenced their final valuations of the two lotteries and their choice between them.

These two sets of novel results provide a compelling reason to reject the hypothesis that exposure to a repeated laboratory market is sufficient to cause behaviour to converge on that predicted by the standard theory. Furthermore, the results have implications for the Discovered Preference Hypothesis. Even though the hypothesis need not assert that ten rounds of a repeated laboratory market are sufficient to eliminate all anomalies, it still does not envisage that experience might actually reinforce an anomaly, as found by the experiment for the case of the WTA/WTP gap.

How might market experience alter behaviour, then, if it is not through the discovery of preferences? One explanation is that people use a caution heuristic and the level of cautiousness exhibited is influenced by their recent experience. As a person's level of cautiousness increases, they become less willing to trade, less willing to take risks, and more averse to uncertainty. The experiments findings can be interpreted as follows. As subjects repeat the auction, they become more familiar with it and less cautious (more willing to trade so less likely to set buying prices below selling prices). In the treatments with feedback, subjects experienced playing the lottery and losing (receiving a payout of zero) and this increased their level of cautiousness. The effect of this increase in cautiousness was that they became less willing to trade (so the gap between buying and selling prices persisted) and less willing to take risks (so they were more likely to choose the P-bet and value it highest).

Appendix I

Example of Calculating Expected Value Conditional on an Observed Signal

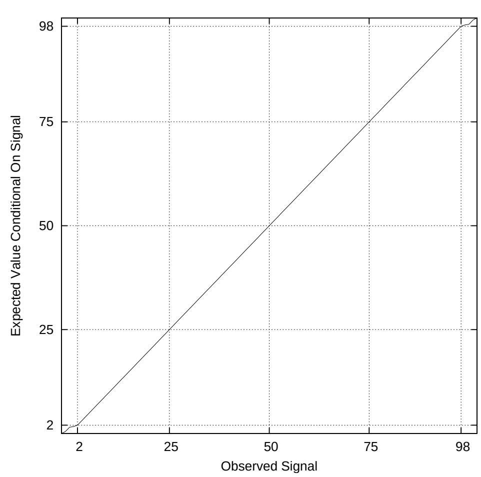

Value V is an integer uniformly distributed on $[0,100]$, so $Pr(V=v)=1/101$. Player i observes one of four possible signals, $X_i \in \{V-2, V-1, V+1, V+2\}$. $Pr(X_i = x_i|V=v) = 1/4$ for each of these signals and $Pr(X_i = x_i|V=v) = 0$ for all other values of X_i. For a given signal, $Pr(X_i = x_i) = \frac{1}{101}\sum_{v=0}^{100} Pr(X_i = x_i|V=v)$. The figure shows $E[V|X_i = x_i]$ against x_i. $E[V|X_i = x_i] = \frac{1}{Pr(X_i=x_i)}\sum_{v=0}^{100} V \times Pr(X_i = x_i|V=v) \times Pr(V=v)$. This simplifies to give $E[V|X_i = x_i] = \sum V \times Pr(X_i = x_i|V =$

$v)/\sum Pr(X_i = x_i|V = v)$. The following is an example of how $E[V|X_i = x_i]$ is calculated. $E[V|X_i = 0] = \sum V \times Pr(X_i = 0|V = v)/\sum Pr(X_i = 0|V = v)$. A signal of 0 could only occur if $V = 1$ or $V = 2$, hence $\sum V \times Pr(X_i = 0|V = v) = 1 \times 1/4 + 2 \times 1/4 = 3/4$. Similarly, $\sum Pr(X_i = 0|V = v) = 1/4 + 1/4 = 2/4$. Therefore $E[v|X_i = 0] = 3/2$. Values of $E[V|X_i = x_i]$ for other values of x_i are calculated using the same approach. When x_i is in the range $[2, 98]$ then $E[V|X_i = x_i] = x_i$; when x is in the range $[-2, 1]$ then $E[V|X_i = x_i] > x_i$; when x is in the range $[99, 102]$ then $E[V|X_i = x_i] < x_i$.

Appendix J

Interdependent Values in Vickrey Buying and Selling Auctions

This appendix concerns Vickrey auctions where bidders have only partial information about the value of the item auctioned and different bidders have different information. In particular, each bidder receives a private signal and the value of the item to the bidder is (a) an increasing function of the signal they observed and (b) a non decreasing function of each signal observed by other bidders. Assumptions are made concerning the symmetry of valuations and signals. I derive symmetric equilibrium strategies for k^{th} price Vickrey buying and selling auctions. The derivations are an extension of the one for a second-price buying auction presented by Krishna (2002, p87) and I use similar assumptions and setup.

The Model

1. There are N risk neutral bidders.

2. Each bidder, i observes a private signal $X_i \in [0, \omega]$. The joint density function of the signals is a symmetric function of its arguments.

3. The signals $\mathbf{X} = X_1, X_2, \ldots, X_N$ are positively affiliated.[1]

4. The value of the item to bidder i is a function of all bidders' signals, $V_i = v_i(X_1, X_2, \ldots, X_N)$. v_i is non decreasing in all its arguments and increasing in X_i. Further, valuations are symmetric in the sense that $v_i(\mathbf{X}) = u(X_i, \mathbf{X}_{-i})$ and u is symmetric in the last $N-1$ components. What this means is that among the other bidders, it does not matter who observed which signal.

5. Let $Y_1, Y_2, ..., Y_{n-1}$ be the largest, second largest, ..., smallest from among $X_m \in \mathbf{X}_{-i}$. Note, $\mathbf{X}_{-i}$ is the set of signals observed by the other bidders. Define the function

   $v_j(x,y) = E[V|X_1 = x_1, Y_j = y]$

   Assume that v is an increasing function of x and y for all j.

Buying Auctions

In a buying auction the kth highest bid is the price p and all bidders who bid more than p buy at price p.

Proposition J.1 *Symmetric equilibrium strategies in a kth-price buying auction are given by:*

$$\beta_k(x) = v_{k-1}(x,x)$$

Proof.

Suppose bidders $2, .., N$ follow the strategy of bidding β_k. Bidder 1 bids b: he trades (buys) if and only if $b > \beta_k(Y_j)$ where $j = k-1$. (Given that bidder 1 trades, the kth highest bid will be submitted by the bidder who observes signal Y_j). For example, in a second price auction he trades if $b > \beta_k(Y_1)$. If he does not trade,

[1] See Krishna (2002, p269) or Milgrom and Weber (1982) for details of the mathematical properties of affiliated random variables.

his payoff is zero. If he trades, his expected payoff is the difference between the expected value of the item and the kth bid. Hence, bidder 1's expected payoff on trading given $X_1 = x$ and $Y_j = y$, is $v_{k-1}(x,y) - \beta_k(y)$. Now, let $g_j(\cdot|x)$ be the density of Y_j conditional on $X_1 = x$. His expected payoff when his signal is x and he bids an amount b can be found by integrating with respect to y on the interval where $b > \beta_k(y)$ (that is where $\beta_k^{-1}(b) > y$)

$$\begin{aligned}\Pi(b,x,k) &= \int_0^{\beta_k^{-1}(b)} [v_{k-1}(x,y) - \beta_k(y)]\ g_{k-1}(y|x)dy \\ &= \int_0^{\beta_k^{-1}(b)} [v_{k-1}(x,y) - v_{k-1}(y,y)\ ]\ g_{k-1}(y|x)dy\end{aligned}$$

Since v is an increasing function of its first argument $[v_{k-1}(x,y) - v_{k-1}(y,y)] > 0$ if and only if $x > y$. Hence Π is maximised when $\beta_k^{-1}(b) = x$, i.e. when $b = \beta_k(x) = v_{k-1}(x,x)$. ■

Selling Auctions

In a selling auction the kth highest bid is the price p and all bidders who bid less than p sell at price p.

Proposition J.2 *Symmetric equilibrium strategies in a kth-price selling auction are given by:*

$$\beta_k(x) = v_k(x,x)$$

Proof. As above, suppose bidders $2,..,N$ follow the strategy of bidding β_k. Bidder 1 bids b: he trades (sells) if and only if $b < \beta_k(Y_k)$. Given that bidder 1 trades, the kth highest bid will be submitted by the bidder who observes signal Y_k. Notice this differs from the buying auction: when buying, bidder 1 trades if and only if his bid one of the $(k-1)$ highest bids, so the bid that determines the price will of the $(k-1)^{th}$ highest bid among the other bids; when selling, bidder 1 trades if and only if his bid is one of the $(N-k)$ lowest bids, so the bid that determines the price will be the k^{th} highest bid among the other bids. When bidder 1 trades, his payoff will be the price p minus the opportunity cost of selling the item. His expected payoff when his signal is x and he bids an amount b is

$$\begin{aligned}
\Pi(b,x,k) &= \int_{\beta_k^{-1}(b)}^{\omega} [\beta_k(y) - v_k(x,y)]\ g_k(y|x)dy \\
&= \int_{\beta_k^{-1}(b)}^{\omega} [v_k(y,y) - v_k(x,y)]\ g_k(y|x)dy
\end{aligned}$$

Since v is an increasing function of its first argument $[v_k(y,y) - v_k(x,y)] > 0$ if and only if $x < y$. Hence Π is maximised when $\beta_k^{-1}(b) = x$, i.e. when $b = \beta_k(x) = v_k(x,x)$. ■

Comparision

Propositions J.1 and J.2 imply that a bidder's optimal bid b given signal x will be the same in a buying and selling auction if k^{th} price selling and $(1+k)^{th}$ price buying auctions are used.

Appendix K

Instructions and Forms

The instructions and forms used in the experiment are shown on the following pages.

K.1 Instructions

[Where the instructions differ between treatments the text is highlighted as shown here. Which treatments saw the text is specified in square brackets.]

Instructions

Introduction

You are about to participate in an experiment investigating how people make decisions in markets. During the session, please do not talk or communicate with any of the other participants. If you have a question, please raise your hand and I will come to your desk to answer it.

Payment

You will be paid in cash at the end of the experiment. The amount you receive will depend on the decisions you and other participants make and the outcome of random events. During the course of the experiment you will complete a number of tasks. In each task you will have the chance to gain or lose 'credits'. At the end of the experiment one of these tasks will be selected at random and you will be paid according to what happened in this task. You will receive £1 for every 100 credits you earned in the selected task.

Outline

The experiment is divided into five parts. Parts one to four involve buying and selling lotteries. There are 10 rounds in each of these parts. Part five involves a simple choice between lotteries.

In the parts involving buying, each round the experimenter will endow you with some credits (for the purposes of the explanation, call this amount of credits x). You will have the chance to buy a lottery from the experimenter. In parts involving selling, each round the experimenter will endow you with one lottery and you will have the chance to sell it to the experimenter. The price of the lottery will be determined by an auction (how the auction works is described in detail later). At the end of each round you will be told the result of the auction [Feedback] and how much the lottery paid out. The diagrams below show how much you would earn from the round in different cases. In a buying round (figure 1), you start with x credits. If you buy the lottery from the experimenter, p credits (the price of the lottery determined by the auction) are subtracted from the x credits you started with and you also receive however many credits the lottery pays out. If you don't buy, you keep all of the x credits. In a selling round (figure 2), you start with one lottery. If you sell the lottery, you receive the p credits (the price determined by the auction). If you don't sell, you receive however many credits the lottery pays out.

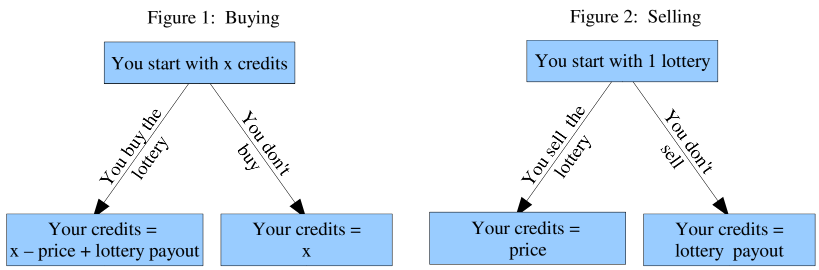

How do the lotteries work?

The lotteries' payout will be determined by the colour of a ball drawn from an urn. The urn will contain 100 balls: [known odds] 81 are white and 19 are orange. [unknown odds] some are orange and some are white. The number of credits the lottery pays out for different colours will change from task to task. During each task you will be told how much the lottery pays out on screen.

[No feedback]At the end of the experiment, (as mentioned above) a task will be selected at random by the computer to be the task you get paid for. After this task has been selected, the experimenter will draw a ball from the urn to determine how much the lottery pays out in the selected task.
[Feedback] After each task, a ball will be drawn from the urn by the experimenter to determine how much the lottery pays out. The ball will be put back in the urn so that the number of white and orange balls isn't changed.

How do the auctions work?

As stated above, during the experiment you will participate in a series of auctions to buy and sell lotteries. You and three other participants will be bidding in the auction. It will be the same set of people bidding every auction in the experiment. After the auction, you won't be told the value of other participants' bids and they won't be told the value of your bid. No one will be told who the other participants in their group are. However, you and the other participants will be told the price which the auction produced [Feedback] and the outcome of the lottery before starting the next round.

How does the buying auction work?

At the start of the auction the experimenter will endow you with some credits. You and the three other participants will then bid to buy a lottery from the experimenter. Two lotteries will be available to buy. Each participant can only buy one lottery per auction. The price and who buys will be determined as follows.

(1) How are bids entered? Each participant will be prompted to type a bid into the box shown on the screenshot below. If you read the text at the bottom of the screen you can see that what it is essentially asking you to do is to state the maximum amount that you are willing to buy the lottery for.

[A screen shot of the relevant enter bid screen appeared here.]

(2) How is the price determined? The computer will record the bids made by each of the participants and order them from lowest to highest. Suppose the bids were as given in the table below. (The numbers are just examples and contain no significance beyond this.)

Participant	*Bids*	*Buys Lottery*	*Credits from task*
Y	351	-	x
W	**367**	-	x
Z	560	Yes	x – 367 + lottery payout
X	685	Yes	x – 367 + lottery payout

,
In all buying auctions, the second lowest bid determines the price. So in this case the price would be 367.

(3) Who buys the lotteries and how much do they pay?
In **buying** auctions the participants with the **two highest bids** buy the lottery. They pay the price, not the amount they bid. So, in the example above Z and X buy the lottery. They each pay 367 (the price) and receive whatever the lottery pays out.

The selling auction

The selling auctions are similar to the buying auctions but have some differences. The experimenter endows you with a lottery instead of some credits. Two of the group members sell their lottery to the experimenter and the other two keep the lottery and receive whatever it pays out.

(1) How are bids entered? Each participant will be prompted to type a bid into the box shown on the screenshot below. If you read the text at the bottom of the screen you can see that what it is essentially asking you to do is to state the minimum amount that you are willing to sell the lottery for.

[A screen shot of the relevant enter bid screen appeared here.]

(2) How is the price determined? The computer will record the bids made by each of the participants and order them from lowest to highest. Suppose the bids are the same as before.

Participant	*Bid*	*Sells lottery*	*Credits from task*
Y	351	Yes	560
W	367	Yes	560
Z	**<u>560</u>**		lottery payout
X	685		lottery payout

In all selling auctions, the second highest (not the second lowest) bid determines the price. So in this case the price would be 560.

(3) Who sells the lotteries and how much do they receive?
In **<u>selling</u>** auctions the participants with the **<u>two lowest bids</u>** sell the lottery. They receive the price, not the amount they bid. So, in the example above Y and W sell the lottery. They each receive 560 (the price). The others, X and Z keep the lottery and receive whatever it pays out.

How you will be told the results of the auction [feedback] and the outcome of the lottery

Once you and the other participants in your group have submitted your bids, you will be told the result of the auction. [Feedback] When all the groups have completed the auction, the experimenter will draw a ball from the urn to determine how much the lottery pays out. Then the experimenter will enter the result into his computer and it will appear on your screen. The screenshot below shows an example of what you might be shown. (The numbers are just examples and contain no significance beyond this.)

[A screen shot of the relevant auction and lottery outcome screen appeared here.]

When you click continue you will begin the next task.

Remember that any one of the tasks could be selected to be the one that determines how much you get paid, so you should complete all the tasks carefully.

K.2 Comprehension Test

Your computer number: …………

Questions

Please answer the following questions. They serve as a test of your understanding of how the auctions work and how the payoffs are determined. Feel free to ask questions if there is anything you don't understand. The experimenter will collect and check your answers before starting the experiment.

1. It is a buying task. Each participant has been allocated 1000 credits.

Participant	*Bid*	*Buys Lottery*	*Credits from task*
M	100	□	…………
N	200	□	…………
P	350	□	…………
Q	620	□	…………

(a) Circle the bid in the 'Bid' column that determines the price.
(b) In the 'Buys Lottery' column indicate which two participants buy the lottery
(c) How much do they pay?........
(d) Suppose the lottery pays out 1000 credits. Write values in the 'Credits from task' column for all four participants.

2. It is a selling task. Each participant has been allocated one lottery.

Participant	*Bid*	*Sells Lottery*	*Credits from task*
M	200	□	…………
N	340	□	…………
P	900	□	…………
Q	1000	□	…………

(a) Circle the bid in the 'Bid' column that determines the price.
(b) Tick the boxes in the 'Sells Lottery' column indicate which two participants sell the lottery
(c) How much do they receive?........
(d) Suppose the lottery pays out 0 credits. Write values in the 'Credits from task' column for all four participants.

K.3 Debrief Questions

[Question 5, highlighted, was only asked in treatments with unknown odds]

(1) How did you decide how much to bid in the auctions?

……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………

(2) Did you change your bidding strategy or use the same strategy for the whole experiment? Why?

……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………

(3) Did you use different strategies for buying and selling auctions? Why?

……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………

(4) For the last task in the experiment, you were asked to choose between two lotteries. What influenced your decision?

……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………

(5) The Urn contained 100 balls. Some were white and some were orange. How many do you estimate were orange?

……………………………………………………………………………………
…..

(6) Any other comments?

……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………
……………………………………………………………………………………

Appendix L

Screenshots

L.1 Enter Bid

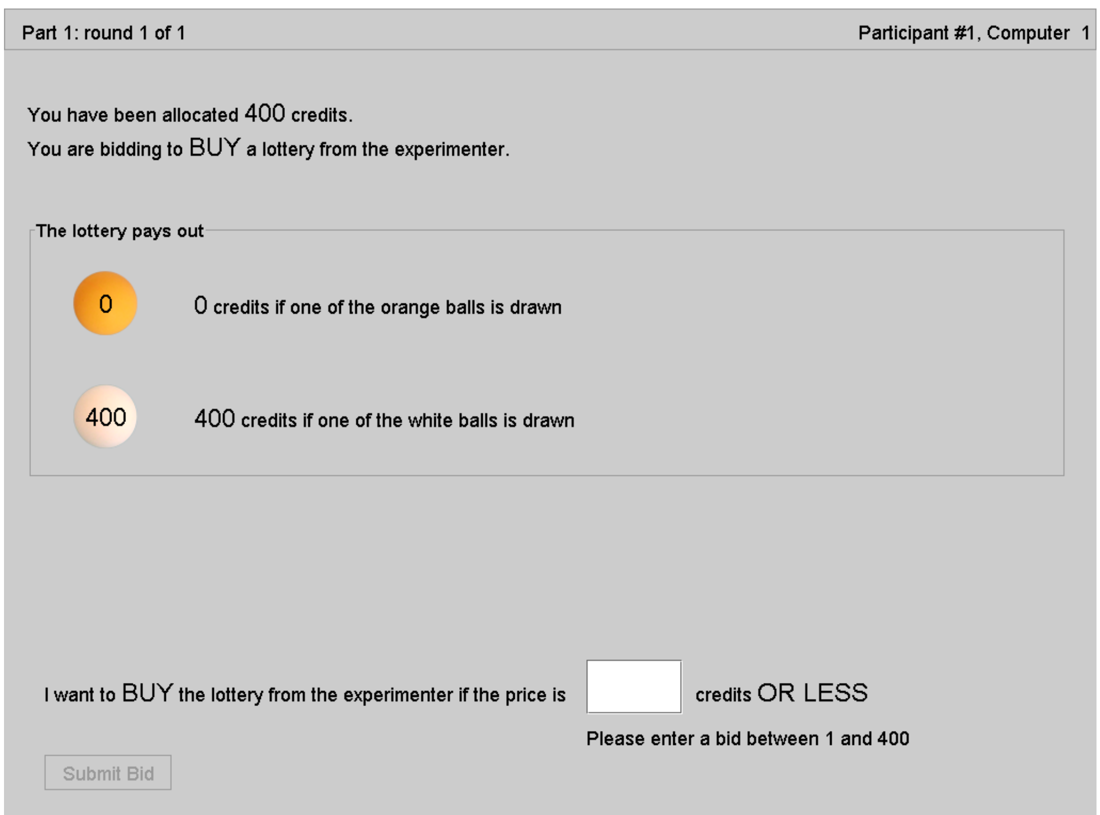

L.2 Auction Outcome

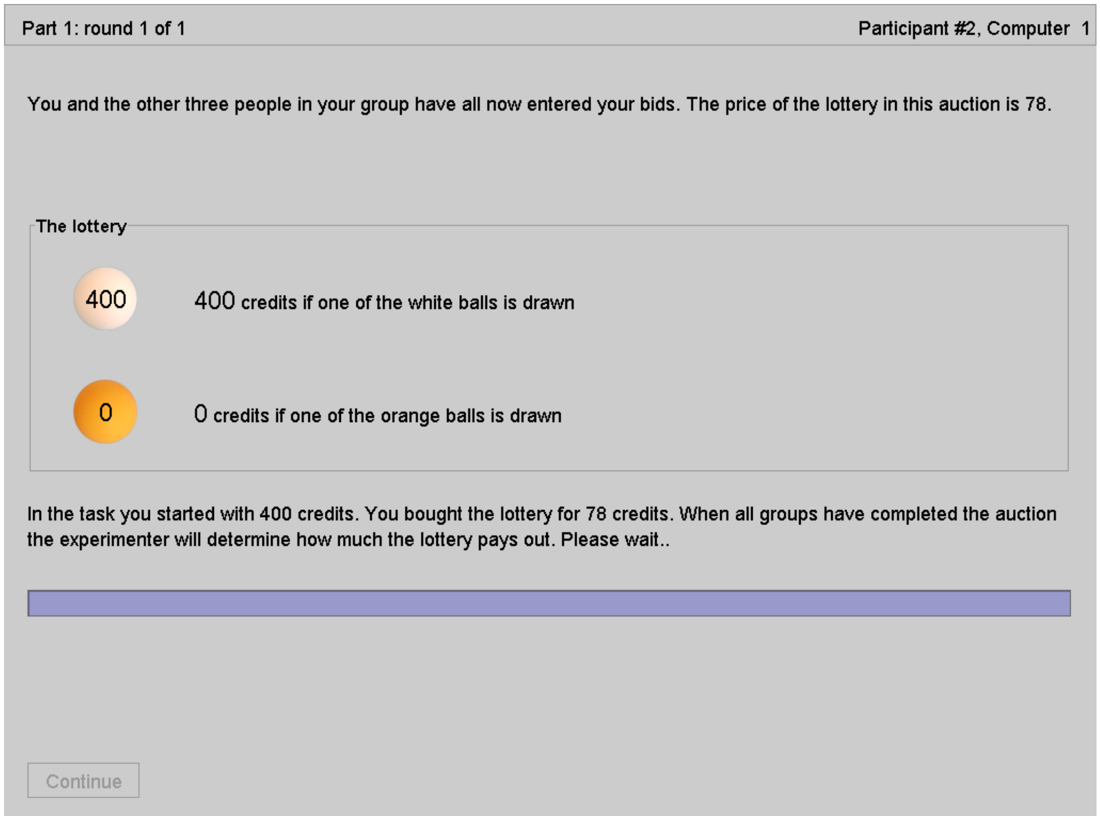

L.3 Lottery Resolution

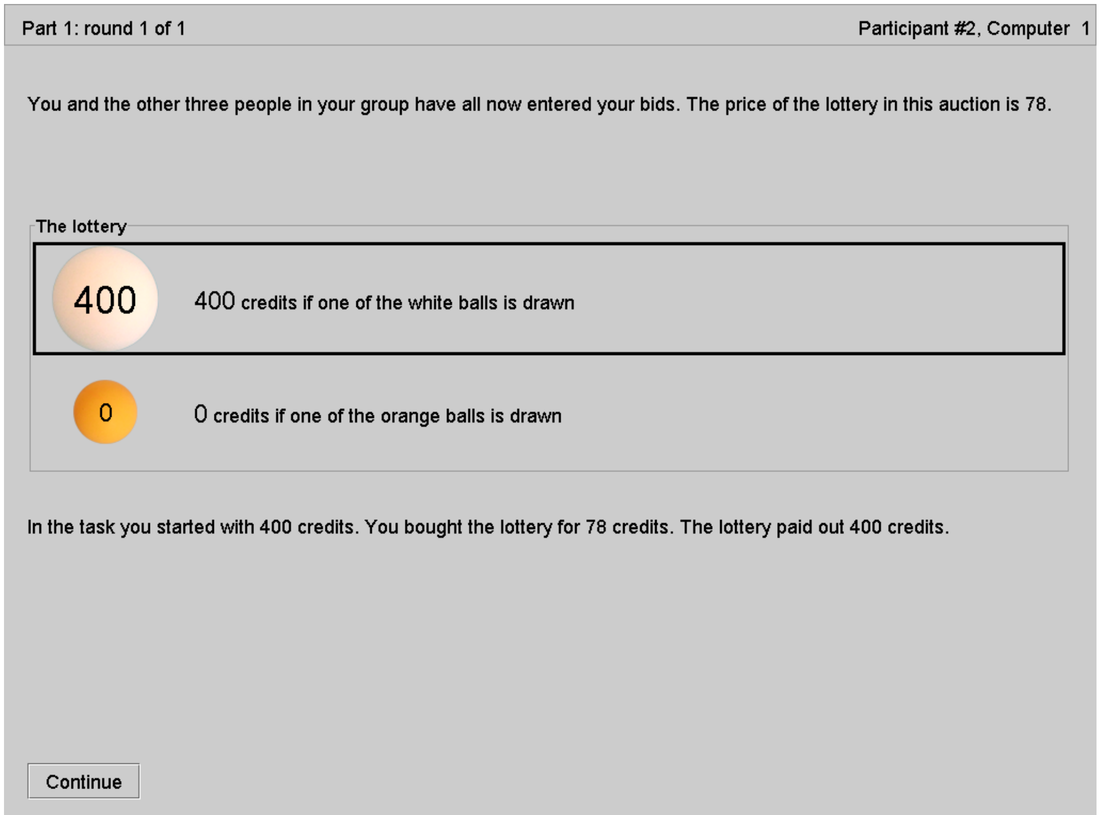

L.4 Choice Task

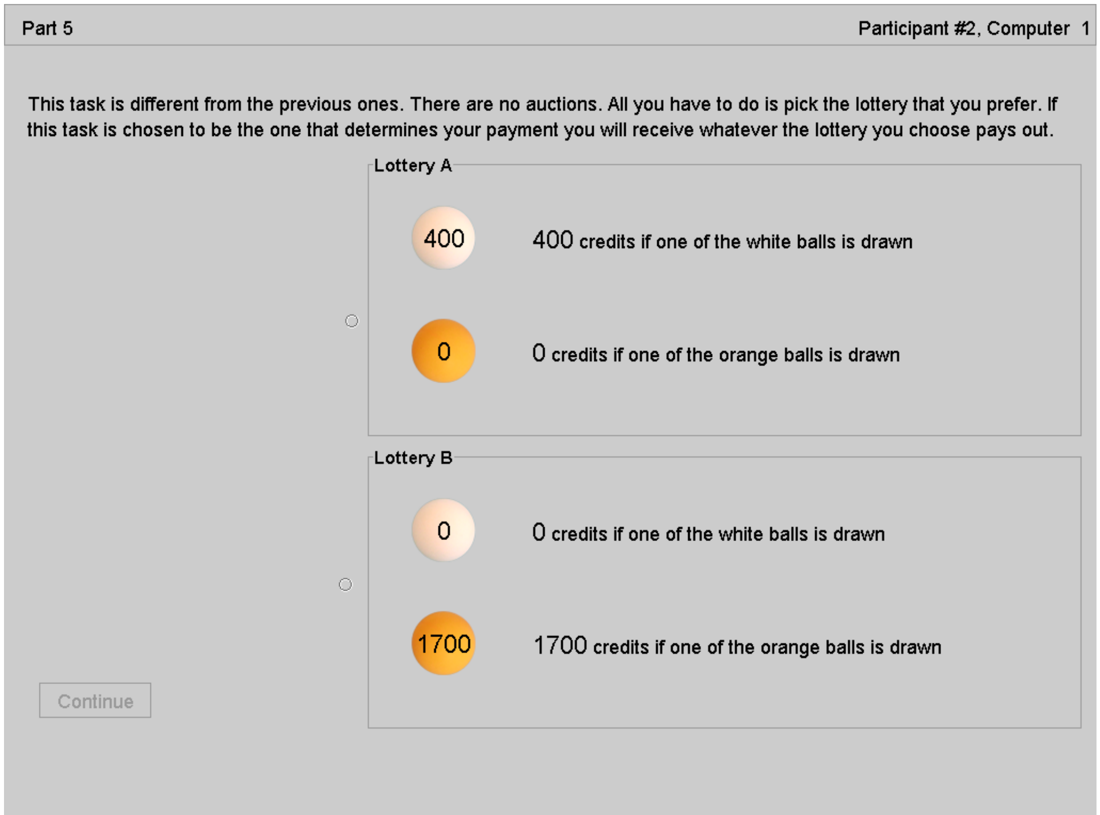

Chapter 5

Concluding Remarks

There is a large literature on loss aversion and the endowment effect. Early studies involving hypothetical valuations of environmental goods found that the minimum amount people were willing to accept in compensation for giving up an entitlement typically exceeded the amount they were willing to pay for the same entitlement. Subsequent controlled and incentivised experiments and some field studies produced similar results. These results are inconsistent with the standard theory of choice used in Economics. The result have been interpreted as evidence that the preferences which determine choice are reference dependent. In particular, losses relative to the reference point carry more weight than corresponding gains. Kahneman and Tversky (1979) called this effect loss aversion.

Is loss aversion a stable feature of preferences? The evidence is mixed. There is evidence that market experience eliminates the endowment effect. Some, but not all, studies that have investigated WTA/WTP gaps in repeated markets have found that while a gap is observed in early rounds, it decays and is eliminated in later rounds. Furthermore, studies with subjects with different levels of experience of trading in naturally occurring markets have found that those with more market experience tend not to exhibit the endowment effect.

Several theories of how experience affects behaviour have been suggested. Plott (1996) proposed the 'Discovered Preference Hypothesis' and argues that rational behaviour emerges from a process of reflecting, experience and practice. Plott and Zeiler (2005) argue that (a) the WTA/WTP gap observed in many studies is not due to loss aversion being a feature of people's preferences. Instead, they argue (b) the gap is due to experiment subjects' misconceptions about what it is in their best interest to do and (c) when misconceptions are controlled for, the gap disappears. Again, however, the evidence is mixed and while many studies find that experience influences behaviour, it is not clear its influence is to eliminate anomalies for the standard theory. For example, Braga et al. (forthcoming) investigate preference reversals in a repeated market and find that market experience eliminates one anomaly but creates a new one. As things stand, we know that experience influences behaviour but it is not clear what aspects of experience matter and what the mechanisms are.

This thesis investigated how experience affects loss aversion using three carefully designed experiments. The 'Crisps and Lemonade' (CL) experiment, chapter 2, investigated how experience of consuming, owning and choosing goods affects willingness to exchange one good for another. The 'Spinning Arrow' (SA) experiment, chapter 3, investigated the WTA/WTP gap in repeated markets with symmetric and asymmetric information. In two of the four treatments the market regime switched between symmetric and asymmetric information after 10 of the 20 rounds were completed. The experiment investigated the direct effect of asymmetric information on the gap and the effect of experiencing asymmetric information on later behaviour. The 'Drawing Balls from Urn' (DBU) experiment, chapter 4, consisted of repeated markets for lotteries. It investigated how lottery feedback and the presentation of odds affected the evolution of bids in buying and selling auctions.

The design of the experiments followed several patterns. (a) A common task preceded by different experiences. This allowed between subject comparisons of the effect of different types of experience. This was used in the CL and SA experiments. (b) Repetition of a task. This allowed within subject comparisons of behaviour in early and later trials. This was used in the SA and DBU experiments. (c) (i) Repetition with varying feedback. This allowed between subject comparisons of the effect of different aspects of experience. (ii) Repetition with identical feedback but varying initial information. This allowed between subject comparisons of the effect of the same feedback when it carries different levels of new information. These last two patterns were used in the DBU experiment.

Each of the three experiments reproduced results reported in the existing literature in addition to their novel findings. The CL experiment found a statistically significant endowment effect. Subjects tended to stick with the goods they were endowed with. In the SA experiment, under symmetric information there was a statistically significant WTA/WTP gap that decayed in later rounds. Similarly, the DBU experiment results include a WTA/WTP gap that decays with repetition. Moreover, the experiment reproduces the recent finding of Braga et al. that in a repeated market for P and $ bets, the tendency in early rounds to value the $ bet highest but choose the P bet is replaced in later rounds by the reverse tendency. Reproducing the results of previous studies has two implications. First it increases our confidence that the results of the earlier studies were not due to chance. Second, it shows that the earlier results are repeatable and robust to minor changes in experimental procedures. Third, it suggests the way I ran experiments was comparable to other studies. That is, factors such as the clarity of the instructions and the strength of the incentives were not different enough, if at all, to produce markedly different behaviour among subjects.

The three experiments also produced a set of novel results. In the CL experiment, the proportion of subjects who swapped an item from their endowment for an alternative was considerably higher than in previous studies but still below the rate predicted by the standard theory. One explanation is that loss aversion is weaker when a person owns several identical items and is faced with giving up one than when a person is faced with giving up their only instance of an item. Another explanation is that when a person takes ownership of an item, all items of that type become more desirable to the person. Hence, when they are faced with swapping an item they own for an alternative, they will be more likely to swap if they own an item of the same type as the alternative. The experiment also found that when a subject acquired their endowment in two steps, they were less willing to swap items than if they had acquired the same endowment in one step. Furthermore, among subjects who acquired their endowment in two steps, those with experience of choosing between the goods prior to owning either of them were less likely to exhibit the endowment effect than those with just experience of owning one of the goods.

The SA experiment found that the WTA/WTP gap decayed under symmetric information but persisted under asymmetric information. The standard theory predicts a gap under asymmetric information but not under symmetric information. The results suggest that incentives play a role in changing behaviour. There was also evidence of spillover effects. When the regime switched between symmetric and asymmetric information, uninformed subjects did not immediately adjust their bidding strategy. Furthermore, when the regime switched from asymmetric to symmetric information, the previously informed subjects traded more intensely than the previously uninformed subjects, even though all subjects faced the same decision problem.

The DBU experiment found that different feedback about the consequences of choices has different effects on behaviour. When the lotteries subjects were trading were resolved after each round, the WTA/WTP gap persisted; when the lotteries were resolved at the end of the experiment, the gap decayed. Conversely, the over weighting of small probabilities decayed in treatments with lottery feedback but not in those without.

Several common themes in the results of the experiments were identified. In the CL experiment, subjects' attitudes to risk were assessed using two surveys. The results suggest risk attitudes that are inconsistent with an expected utility of wealth model but could be described by reference dependent theory and loss aversion. Furthermore, subjects who identified themselves as more loss averse in the attitude to risk survey were less likely to swap goods when given the option. This suggests loss aversion is the common cause of both the endowment effect and aversion to 50-50 lose X, win $X+Y$ bets. In the SA experiment, the magnitude of costly errors was greater in auctions to buy lotteries than auctions to sell lotteries. This can be interpreted as loss aversion or a caution heuristic having two effects: causing reluctance to trade and causing reluctance to play lotteries. In buying, these two effects act in the same direction; in selling they act in opposite directions. A related effect was observed in the DBU experiment. There was evidence of ambiguity aversion in buying but not selling. Again, this can be interpreted as ambiguity causing a reluctance to trade and causing a reluctance to play lotteries.

Another finding repeated across experiments was that choices were sensitive to experimental procedure, sometimes in unexpected ways. In the CL experiment, the proportion of subjects willing to swap an item from their endowment was lower if the endowment had been acquired in two steps. In the DBU experiment, the proportion of subjects who, at the end of the experiment, chose the dollar bet over the p bet was higher in treatments where the lottery was resolved after each auction round.

There were also some inconsistencies between the results of the experiments. In the SA experiment lotteries were resolved after each auction round. Under symmetric information, the WTA/WTP gap decayed. In the DBU experiment, in treatments where the lotteries were not resolved after each round, the gap also decayed. In treatments where the lotteries were resolved, however, the gap persisted. The puzzle is why lottery feedback caused the gap to persist in the DBU experiment but not in the SA one. Perhaps differing behaviour was caused by differences in experiment design. The DBU experiment differed from the SA one in the following ways. Subjects traded lotteries with two not four outcomes. A physical random device, an urn of balls, not a computer simulated one was used. Subjects were paid based on one randomly selected task not based on earnings accumulated over all tasks. A within subject not a between subject design was used. That is, each subject completed both buying and selling auctions not one or the other. Finally, the feedback subjects received was based just on the last round not on all rounds completed. In the SA experiment, before starting the next round subjects were told how much they earned from the round and then how much they had accumulated over all previous rounds. Which of these differences accounts for the differences in observed behaviour between the experiments is open to speculation. One plausible explanation is that the type of feedback caused the differences. Perhaps when subjects are told the consequences of their choices in one round in isolation, losses and gains relative to their endowment at the start of the round are salient. So they exhibit the endowment effect in the form of a WTA/WTP gap. In contrast, perhaps when subjects are also shown the consequences of the choices they have made in all rounds so far for accumulated earnings, the relative salience of the loss or gain they made in the round just completed decreases and that of the accumulated earnings increases. Since, for accumulated earnings equal sized losses and gains cancel each other out, if subjects focus on the effect of their behaviour on accumulated earnings, they will be less likely to weigh losses and gains differently.

Several broad conclusions can be drawn from the results of the three experiments. First, experience does change behaviour attributed to loss aversion. In the CL experiment, swap rates varied depending on experience subjects received before making the swap decision. In the SA and DBU experiments, in some but not all treatments, the WTA/WTP gap decayed. What explains the impact of experience on loss aversion? Plott and Zeiler suggest the endowment effect is due to subjects' misconceptions. On this view, the effect of experience could be to correct these misconceptions. This view, however, is not supported by the experimental evidence. In the CL experiment, the set of controls for misconceptions was constant across treatments but the observed endowment effect varied between treatments. Likewise, in the DBU experiment, the set of controls was constant across treatments but the WTA/WTP gap decayed in some treatments but not others. Another explanation is Plott's 'Discovered Preference Hypothesis'. Again, however, this is not supported by the experimental evidence since it cannot explain the results of the DBU experiment.

If the mechanism by which experience influences behaviour is not the reduction of misconceptions or the discovery of preferences, what is it? The results of the SA experiment suggest that heuristics play a role in determining behaviour and experience might change which heuristic gets used. On this view, the endowment effect can be interpreted as the result of a caution heuristic and changes in behaviour, interpreted as people changing their level of cautiousness based on previous trading success.

The following are possible extensions of the experiments reported in this thesis. First, investigating why the swap rates in the CL experiment were higher than in earlier studies. Is the endowment effect weaker when giving up one of several identical items? Is it weaker when a person already owns an item like the one they are gaining? Why does acquiring the endowment in two steps strengthen the endowment effect? Second, investigating how people with experience of asymmetric information behave under symmetric information. Will a difference in cautiousness be present in individual decision making problems? Third, further investigation of how lottery feedback affects the WTA/WTP gap. What accounts for the different results from the SA and DBU experiments?

The findings reported in this thesis have implications for which theory of choice is selected to analyse economic problems. The results suggest loss aversion should be taken seriously. First, each of the experiments found evidence of loss aversion. Second, loss aversion persisted in some cases where we might have expected it not to. Notably, the choosers treatment of the CL experiment and in the lottery feedback treatments of the DBU experiment. Second, the results are not consistent with Plott's 'Discovered Preference Hypothesis'. This is most striking in the DBU experiment: with lottery feedback, probability weighting effects decay but the WTA/WTP gap persists; without lottery feedback, probability weighting effects decay but the WTA/WTP gap persists. This implies using the standard theory on the assumption that experience will eliminate any anomalies is not a defensible position. Third, the results suggest that heuristic driven behaviour can in part account for (a) the presence of anomalies and (b) changes in behaviour in repeated tasks. Developing models incorporating heuristics may be a fruitful line of future research.

The broader implications include some for the debate on the justification of paternalism. Thaler and Sunstein (2003) argue that the common assumption among economists that people choose what is in their best interest is false. Further, how choices are presented influences the option people choose. Consequently, people can be made better or worse off by varying how choices are presented. The results of the three experiments show that how choices are presented and what happened before the choice both affect behaviour. This increases the scope of the type of paternalism that Thaler and Sunstein propose to include the type of feedback people receive on past choices.

Finally, if peoples' choices do not follow their best interest and further if their best interest is independent of their current endowment, then people who exhibit the endowment effect are likely to miss out on advantageous trades. That is, the endowment effect in the form of a WTA/WTP gap will prevent the potential gains from trade being fully exploited. If the caution heuristic interpretation is right, then policy interventions that reduce cautiousness could increase the number of mutually advantageous trades that are executed, and so increase welfare. The results of the SA experiment suggest that cautiousness is determined by previous trading success, which is assessed using an availability heuristic. Two policy interventions that could reduce cautiousness are (a) reducing cases of asymmetric information and (b) where the success of a trade depends on the resolution of a risk, giving people feedback on the aggregate success of their choices not individual successes and failures.

Bibliography

Adamowicz, W. L., V. Bhardwaj, and B. Macnab: 1993, 'Experiments on the Difference between Willingness to Pay and Willingness to Accept'. *Land Economics* **69**(4), 416–427.

Akerlof, G. A.: 1970, 'The Market for 'Lemons': Quality Uncertainty and the Market Mechanism'. *The Quarterly Journal of Economics* **84**(3), 488–500.

Andreoni, J. A. and J. H. Miller: 1993, 'Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence'. *Economic Journal* **103**(418), 570–85.

Arlen, J., M. Spitzer, and E. Talley: 2002, 'Endowment Effects within Corporate Agency Relationships'. *The Journal of Legal Studies* **31**(1), 1–37.

Banford, N., J. Knetsch, and G. Mauser: 1979, 'Feasibility Judgements and Alternative Measures of Benefits and Costs'. *Journal of Business Administration* **11**(1), 25–35.

Bateman, I., D. Kahneman, A. Munro, C. Starmer, and R. Sugden: 2005, 'Testing Competing Models of Loss Aversion: an Adversarial Collaboration'. *Journal of Public Economics* **89**(8), 1561–1580.

Bateman, I., A. Munro, B. Rhodes, C. Starmer, and R. Sugden: 1997, 'A Test of the

Theory of Reference-Dependent Preferences'. *Quarterly Journal of Economics* **112**(2), 479–505.

Becker, G., M. DeGroot, and J. Marschak: 1964, 'Measuring Utility by a Single-Response Sequential Method'. *Behavioral Science* **9**(3), 226–32.

Beggs, A. and K. Graddy: 2005, 'Testing for Reference Dependence: An Application to the Art Market'. Economics Series Working Papers 228, University of Oxford, Department of Economics.

Benartzi, S. and R. H. Thaler: 1995, 'Myopic Loss Aversion and the Equity Premium Puzzle'. *The Quarterly Journal of Economics* **110**(1), 73–92.

Binmore, K.: 1999, 'Why Experiment in Economics?'. *Economic Journal* **109**(453), F16–24.

Birnbaum, M., G. Coffey, B. Mellers, and R. Weiss: 1992, 'Utility Measurement: Configural-Weight Theory and the Judge's Point of View'. *Journal of Experimental Psychology: Human Perception and Performance* **18**(2), 331–46.

Bishop, R. C. and T. A. Heberlein: 1979, 'Measuring Values of Extramarket Goods: Are Indirect Measures Biased?'. *American Journal of Agricultural Economics* **61**(5), 926–930.

Bleaney, M. and S. Humphrey: 2006, 'An Experimental Test of Generalized Ambiguity Aversion using Lottery Pricing Tasks'. *Theory and Decision* **60**(2), 257–282.

Bowman, D., D. Minehart, and M. Rabin: 1999, 'Loss Aversion in a Consumption-Savings Model'. *Journal of Economic Behavior & Organization* **38**(2), 155–178.

Boyce, R. R., T. C. Brown, G. H. McClelland, G. L. Peterson, and W. D. Schulze:

1992, 'An Experimental Examination of Intrinsic Values as a Source of the WTA-WTP Disparity'. *The American Economic Review* **82**(5), 1366–1373.

Braga, J., S. Humphrey, and C. Starmer: 2008, 'Market Experience Eliminates Some Anomalies - and Creates New Ones'. *European Economic Review* (forthcoming).

Braga, J. and C. Starmer: 2005, 'Preference Anomalies, Preference Elicitation and the Discovered Preference Hypothesis'. *Environmental and Resource Economics* **32**(1), 55–89.

Brehm, J. W.: 1956, 'Postdecision Changes in the Desirability of Alternatives'. *The Journal of Abnormal and Social Psychology* **52**(3), 384–389.

Brookshire, D. S. and D. L. Coursey: 1987, 'Measuring the Value of a Public Good: An Empirical Comparison of Elicitation Procedures'. *The American Economic Review* **77**(4), 554–566.

Brookshire, D. S., A. Randall, and J. R. Stoll: 1980, 'Valuing Increments and Decrements in Natural Resource Service Flows'. *American Journal of Agricultural Economics* **62**(3), 478–488.

Brown, T. C. and R. Gregory: 1999, 'Why the WTA-WTP Disparity Matters'. *Ecological Economics* **28**(3), 323–335.

Camerer, C., L. Babcock, G. Loewenstein, and R. Thaler: 1997, 'Labor Supply of New York City Cabdrivers: One Day At A Time'. *Quarterly Journal of Economics* **112**(2), 407–441.

Camerer, C. F.: 1998, 'Prospect Theory in the Wild: Evidence From the Field'. Working Papers 1037, California Institute of Technology, Division of the Humanities and Social Sciences.

Camerer, C. F. and R. M. Hogarth: 1999, 'The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework'. *Journal of Risk and Uncertainty* **19**(1-3), 7–42.

Carmon, Z. and D. Ariely: 2000, 'Focusing on the Forgone: How Value Can Appear So Different to Buyers and Sellers'. *Journal of Consumer Research: An Interdisciplinary Quarterly* **27**(3), 360–70.

Casey, J. T.: 1995, 'Predicting Buyer-Seller Pricing Disparities'. *Management Science* **41**(6), 979–999.

Chen, M. K.: 2008, 'Rationalization and Cognitive Dissonance: Do Choices Affect or Reflect Preferences?'. Cowles Foundation Discussion Papers 1669, Cowles Foundation, Yale University.

Chen, M. K., V. Lakshminarayanan, and L. R. Santos: 2006, 'How Basic Are Behavioral Biases? Evidence from Capuchin Monkey Trading Behavior'. *Journal of Political Economy* **114**(3), 517–537.

Chou, Y. K.: 2002, 'Testing Alternative Models of Labour Supply: Evidence from Taxi Drivers in Singapore.'. *Singapore Economic Review* **47**(1), p17 –.

Coursey, D. L., J. L. Hovis, and W. D. Schulze: 1987, 'The Disparity Between Willingness to Accept and Willingness to Pay Measures of Value'. *The Quarterly Journal of Economics* **102**(3), 679–690.

Cox, J. C. and V. Sadiraj: 2006, 'Small-and Large-Stakes Risk Aversion: Implications of Concavity Calibration for Decision Theory'. *Games and Economic Behavior* **56**(1), 45–60.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. G. Wagner: 2005, 'Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey'. IZA Discussion Papers 1730, Institute for the Study of Labor (IZA).

Dupont, D. Y. and G. S. Lee: 2002, 'The Endowment Effect, Status Quo Bias and Loss Aversion: Rational Alternative Explanation'. *Journal of Risk and Uncertainty* **25**(1), 87–101.

Eisenberger, R. and M. Weber: 1995, 'Willingness-to-Pay and Willingness-to-Accept for Risky and Ambiguous Lotteries'. *Journal of Risk and Uncertainty* **10**(3), 223–33.

Farber, H.: 2008, 'Reference-Dependent Preferences and Labor Supply: The Case of New York City Taxi Drivers'. *American Economic Review* **98**(3), 1069–1082.

Festinger, L.: 1957, *A Theory of Cognitive Dissonance*. Stanford University Press.

Franciosi, R., P. Kujal, R. Michelitsch, V. Smith, and G. Deng: 1996, 'Experimental Tests of the Endowment Effect'. *Journal of Economic Behavior & Organization* **30**(2), 213–226.

Friedman, M.: 1953, 'The Methodology of Positive Economics'. In: *Essays in Positive Economics*. University of Chicago Press.

Genesove, D. and C. Mayer: 2001, 'Loss Aversion and Seller Behavior: Evidence from the Housing Market'. *Quarterly Journal of Economics* **116**(4), 1233–1260.

Gilovich, T., D. Griffin, and D. Kahneman: 2002, *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press.

Goette, L., D. Huffman, and E. Fehr: 2004, 'Loss Aversion and Labor Supply'. *Journal of the European Economic Association* **2**(2-3), 216–228.

Greiner, B.: 2004, 'The Online Recruitment System ORSEE 2.0 - A Guide for the Organization of Experiments in Economics'. Working Paper Series in Economics 10, University of Cologne, Department of Economics.

Grether, D. M. and J. C. Cox: 1996, 'The Preference Reversal Phenomenon: Response Mode, Markets and Incentives'. *Economic Theory* **7**(3), 381–405.

Hammack, J. and G. Brown: 1974, *Waterfowl and Wetlands: Toward Bioeconomic Analysis*. Baltimore, MD: Johns Hopkins University Press.

Hardie, B. G. S., E. J. Johnson, and P. S. Fader: 1993, 'Modeling Loss Aversion and Reference Dependence Effects on Brand Choice'. *Marketing Science* **12**(4), 378–394.

Harless, D. W.: 1989, 'More Laboratory Evidence on the Disparity between Willingness to Pay and Compensation Demanded'. *Journal of Economic Behavior & Organization* **11**(3), 359–379.

Heifetz, A. and E. Segev: 2004, 'The Evolutionary Role of Toughness in Bargaining'. *Games and Economic Behavior* **49**(1), 117–134.

Hey, J.: 2005, 'Why We Should Not Be Silent about Noise'. *Experimental Economics* **8**(4), 325–345.

Holt, C. A. and S. K. Laury: 2002, 'Risk Aversion and Incentive Effects'. *American Economic Review* **92**(5), 1644–1655.

Holt, C. A. and S. K. Laury: 2005, 'Risk Aversion and Incentive Effects: New Data without Order Effects'. *American Economic Review* **95**(3), 902–904.

Horowitz, J. K. and K. E. McConnell: 2002, 'A Review of WTA/WTP Studies'. *Journal of Environmental Economics and Management* **44**(3), 426–447.

Humphrey, S.: 2006, 'Does Learning Diminish Violations of Independence, Coalescing and Monotonicity?'. *Theory and Decision* **61**(2), 93–128.

Kachelmeier, S. J. and M. Shehata: 1992, 'Examining Risk Preferences under High Monetary Incentives: Experimental Evidence from the People's Republic of China'. *American Economic Review* **82**(5), 1120–41.

Kahneman, D.: 2003, 'Maps of Bounded Rationality: Psychology for Behavioral Economics'. *American Economic Review* **93**(5), 1449–1475.

Kahneman, D., J. L. Knetsch, and R. H. Thaler: 1990, 'Experimental Tests of the Endowment Effect and the Coase Theorem'. *The Journal of Political Economy* **98**(6), 1325–1348.

Kahneman, D. and J. Renshon: 2007, 'Why Hawks Win'. *Foreign Policy* **158**, 34–38.

Kahneman, D., P. Slovic, and A. Tversky: 1982, *Judgment under Uncertainty: Heuristics and Biases*. Cambridge University Press.

Kahneman, D. and A. Tversky: 1973, 'Availability: A Heuristic for Judging Frequency and Probability'. *Cognitive Psychology* **5**, 207–232.

Kahneman, D. and A. Tversky: 1979, 'Prospect Theory: An Analysis of Decision under Risk'. *Econometrica* **47**(2), 263–91.

Kahneman, D. and A. Tversky: 1982, 'The Psychology of Preferences'. *Scientific American* **246**(1), 160–173.

Knetsch, J., F.-F. Tang, and R. Thaler: 2001, 'The Endowment Effect and Repeated Market Trials: Is the Vickrey Auction Demand Revealing?'. *Experimental Economics* **4**(3), 257–269.

Knetsch, J. L.: 1989, 'The Endowment Effect and Evidence of Nonreversible Indifference Curves'. *American Economic Review* **79**(5), 1277–84.

Knetsch, J. L. and J. A. Sinden: 1984, 'Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value'. *The Quarterly Journal of Economics* **99**(3), 507–21.

Knez, M. and V. Smith: 1987, *Laboratory Experimentation in Economics: Six Points of View*, Chapt. Hypothetical Valuations and Preference Reversals in the Context of Asset Trading, pp. 131–154. Cambridge University Press.

Knez, P., V. L. Smith, and A. W. Williams: 1985, 'Individual Rationality, Market Rationality, and Value Estimation'. *The American Economic Review* **75**(2), 397–402.

Koszegi, B. and M. Rabin: 2006, 'A Model of Reference-Dependent Preferences'. *Quarterly Journal of Economics* **121**(4), 1133–1165.

Krishna, V.: 2002, *Auction Theory*. Academic Press.

Lindsay, L.: 2004, 'Investigating the Robustness of the Endowment Effect and Its Relation to Loss Aversion'. Master's thesis, University of Nottingham.

List, J. A.: 2003, 'Does Market Experience Eliminate Market Anomalies?'. *The Quarterly Journal of Economics* **118**(1), 41–71.

List, J. A.: 2004a, 'Neoclassical Theory Versus Prospect Theory: Evidence from the Marketplace'. *Econometrica* **72**(2), 615–625.

List, J. A.: 2004b, 'Substitutability, Experience, and the Value Disparity: Evidence from the Marketplace'. *Journal of Environmental Economics and Management* **47**(3), 486–509.

List, J. A. and J. F. Shogren: 1999, 'Price Information and Bidding Behavior in Repeated Second-Price Auctions'. *American Journal of Agricultural Economics* **81**(4), 942–949.

Loomes, G.: 2005, 'Modelling the Stochastic Component of Behaviour in Experiments: Some Issues for the Interpretation of Data'. *Experimental Economics* **8**(4), 301–323.

Loomes, G., C. Starmer, and R. Sugden: 2003, 'Do Anomalies Disappear in Repeated Markets?'. *Economic Journal* **113**(486), C153–C153.

Loomes, G., C. Starmer, and R. Sugden: 2007, 'Reversals and Disparities between Willingness to Pay and Willingness to Accept in Repeated Markets'. (2007-10).

Loomes, G. and R. Sugden: 1982, 'Regret Theory: An Alternative Theory of Rational Choice under Uncertainty'. *Economic Journal* **92**(368), 805–24.

Loomes, G. and R. Sugden: 1995, 'Incorporating a Stochastic Element into Decision Theories'. *European Economic Review* **39**(3-4), 641–648.

MacDonald, H. and J. Bowker: 1994, 'The Endowment Effect and WTA: A Quasi-Experimental Test'. *Journal of Agricultural and Applied Economics* **26**(2), 545–551.

Marshall, J. D., J. L. Knetsch, and J. A. Sinden: 1986, 'Agents' Evaluations and the Disparity in Measures of Economic Loss'. *Journal of Economic Behavior & Organization* **7**(2), 115–127.

Marx, L. M.: 1999, 'Adaptive Learning and Iterated Weak Dominance'. *Games and Economic Behavior* **26**(2), 253–278.

Mellers, B. A., S. jie Chang, M. H. Birnbaum, and L. D. Ordonez: 1992, ‘Preferences, Prices, and Ratings in Risky Decision Making’. *Journal of experimental psychology. Human perception and performance* **18**(2), 347–361.

Milgrom, P. R. and R. J. Weber: 1982, ‘A Theory of Auctions and Competitive Bidding’. *Econometrica* **50**(5), 1089–1122.

Morrison, G. C.: 1997, ‘Willingness to Pay and Willingness to Accept: Some Evidence of an Endowment Effect’. *Applied Economics* **29**(4), 411–17.

Neilson, W., M. McKee, and R. P. Berrens: 2008, ‘Value and Outcome Uncertainty as Explanations for the WTA vs WTP Disparity: Theory and Experimental Evidence’. Working Papers 08-07, Department of Economics, Appalachian State University.

Novemsky, N. and D. Kahneman: 2005, ‘The Boundaries of Loss Aversion’. *Journal of marketing research* **42**(2), 119–128.

Peters, E., P. Slovic, and R. Gregory: 2003, ‘The Role of Affect in the WTA/WTP Disparity’. *Journal of Behavioral Decision Making* **16**(4), 309–330.

Plott, C.: 1996, ‘Rational Individual Behavior in Markets and Social Choice Processes: The Discovered Preference Hypothesis’. In: M. P. K. Arrow, E. Colombatto and C. Schmidt (eds.): *Rational Foundations of Economic Behavior*. Macmillan, pp. 225–50.

Plott, C. R. and K. Zeiler: 2005, ‘The Willingness to Pay-Willingness to Accept Gap, the "Endowment Effect," Subject Misconceptions, and Experimental Procedures for Eliciting Valuations’. *American Economic Review* **95**(3), 530–545.

Rabin, M.: 2000, ‘Risk Aversion and Expected-utility Theory: A Calibration Theorem’. *Econometrica* **68**(5), 1281–1281.

Rasiel, E. B., K. P. Weinfurt, and K. A. Schulman: 2005, 'Can Prospect Theory Explain Risk-Seeking Behavior by Terminally Ill Patients?'. *Medical Decision Making* **25**(6), 609–613.

Rowe, R. D., R. C. D'Arge, and D. S. Brookshire: 1980, 'An Experiment on the Economic Value of Visibility'. *Journal of Environmental Economics and Management* **7**(1), 1–19.

Sayman, S. and A. Onculer: 2005, 'Effects of Study Design Characteristics on the WTA-WTP Disparity: A Meta Analytical Framework'. *Journal of Economic Psychology* **26**(2), 289–312.

Schmidt, U., C. Starmer, and R. Sugden: 2008, 'Third-Generation Prospect Theory'. *Journal of Risk and Uncertainty* **36**(3), 203–223.

Schmidt, U. and S. Traub: 2006, 'An Experimental Investigation of the Disparity between WTA and WTP for Lotteries'. Economics working papers 2006,09, Christian-Albrechts-University of Kiel, Department of Economics.

Shefrin, H. and D. Caldwell: 2001, 'Determinants of the Magnitude of Willingness to Accept Relative to Willingness to Pay'. *Journal of Behavioral Decision Making* **14**(2), 87–105.

Shogren, J. F., S. Cho, C. Koo, J. List, C. Park, P. Polo, and R. Wilhelmi: 2001, 'Auction Mechanisms and the Measurement of WTP and WTA'. *Resource and Energy Economics* **23**(2), 97–109.

Shogren, J. F., S. Y. Shin, D. J. Hayes, and J. B. Kliebenstein: 1994, 'Resolving Differences in Willingness to Pay and Willingness to Accept'. *The American Economic Review* **84**(1), 255–270.

Singh, H.: 1991, 'The Disparity between Willingness to Pay and Compensation Demanded: Another Look at Laboratory Evidence'. *Economics Letters* **35**(3), 263–266.

Sloman, S. A.: 2002, 'Two Systems of Reasoning'. In: T. Gilovich, D. Griffin, and D. Kahneman (eds.): *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge University Press, Cambridge, pp. 379–396.

Smith, V. L.: 1980, 'Experiments with a Decentralized Mechanism for Public Good Decisions'. *The American Economic Review* **70**(4), 584–599.

Sugden, R.: 2003, 'Reference-Dependent Subjective Expected Utility'. *Journal of Economic Theory* **111**(2), 172–191.

Thaler, R. H.: 1980, 'Toward a Positive Theory of Consumer Choice'. *Journal of Economic Behavior and Organisation* **1**, 39–60.

Thaler, R. H. and C. R. Sunstein: 2003, 'Libertarian Paternalism'. *The American Economic Review* **93**(2), 175–179.

Tversky, A. and D. Kahneman: 1974, 'Judgment under Uncertainty: Heuristics and Biases'. *Science* **185**(4157), 1124–1131.

Tversky, A. and D. Kahneman: 1991, 'Loss Aversion in Riskless Choice: A Reference-Dependent Model'. *The Quarterly Journal of Economics* **106**(4), 1039–61.

Tversky, A. and D. Kahneman: 1992, 'Advances in Prospect Theory: Cumulative Representation of Uncertainty'. *Journal of Risk and Uncertainty* **5**(4), 297–323.

Vickrey, W.: 1961, 'Counterspeculation, Auctions, and Competitive Sealed Tenders'. *The Journal of Finance* **16**(1), 8–37.

von Neumann, J. and O. Morgenstern: 1947, *Theory of Games and Economic Behavior.* Princeton, NJ: Princeton.

Willig, R. D.: 1976, 'Consumer's Surplus Without Apology'. *The American Economic Review* **66**(4), 589–597.