How do agents facing a consistent loss scenario adapt their aversion to risk, with respect to prospect theory and the reflection effect, and how can this be applied to traders in a derivative market?

Submitted by *Jack William Harrison*

to the University of Exeter

as a dissertation towards the degree of

Master of Science by advanced study in

Behavioural Economics and Finance

in September 2018

Word Count: 10945

Acknowledgements

I would like to thank my friends and family for their immeasurable mental, emotional, and financial help towards my completion of this dissertation. As well as this I would like to thank the University of Exeter Business School staff, the lecturers throughout the years who have prepared me for this dissertation, and Professor Dieter Balkenborg for his insightful advice and support, without which this dissertation would have been impossible.

How do agents facing a consistent loss scenario adapt their aversion to risk, with respect to prospect theory and the reflection effect, and how can this be applied to traders in a derivative market?

Abstract: An economic experiment was conducted to examine if exposure to repeated losses affected a participants' subsequent level of risk aversion. This experiment involved a small number of participants and acted as a pilot study and proof of concept for further research. An extensive literature review is included to provide context and a motivation for this research.

Table of Contents:

Introduction……..……....4

Literature review…….....6

i. **Prospect Theory ……..……6**
ii. **Risk and Loss Aversion……..…….8**
iii. **Gambling……..…….…….15**

Benefits of experimental versus empirical data……...……18

i. **Experimental data……...……...18**
ii. **Empirical data……...…….20**

Methodology…….…….…….……25

Ethical Concerns……30

Results…….……31

i. **Presentation of results……31**
ii. **Analysis of results: regressions…….……37**
iii. **Analysis of results: non-parametric tests……...……..41**

Conclusion…...……...….43

References…….……..44

Appendices……..……46

Introduction

Our research question aims to address inconsistencies in human behaviour, namely, how our behaviour changes when we have been exposed to repeated losses, and if this makes us more or less prone to risk aversion. Classical economics would predict that economics agents exhibit consistent levels of risk aversion, however, newer research by behavioural economics would indicate this is not the case and that previous results of risky decisions would influence subsequent behaviour. Anecdotal evidence would provide support the latter, if one thinks of gamblers experiencing heavy losses, they rarely quit, they instead increase their stakes and continue in an effort to recoup their losses.

The potential cause of this change in behaviour can be linked to a previously researched theory, the reflection effect, a part of Kahneman and Tversky's famous 'prospect theory'. This states that when outcomes are framed as a loss, with respect to some reference point, participants become more risk loving, and when outcomes are framed as a gain, participants become more risk averse. This leads to an asymmetric utility function, where losses (outcomes below the reference point) exhibit a convex function, whereas gains (outcomes above the reference point) exhibit a concave function.

This is an important question to answer, as there are many examples, throughout history, of individuals being unable to maintain a consistent level of risk aversion, which has led to them taking larger and larger risks to compensate for previous losses.

One of which would be chronic gamblers; problem gamblers tend to be unable to regulate their risk, and take on greater and greater risks, incurring large losses and debts, as they try to chase back their original amount of money. Understanding this behaviour could allow us to introduce better government regulation to protect at-risk gamblers, and prevent them from incurring unsustainable losses, and accruing large debt.

Another example would be that of Nick Leeson; a former derivatives trader, also known as the ‘rogue trader’ whose actions led to the collapse of Barings Bank in 1995, which at the time was the world’s second oldest merchant bank. After experiencing initial success from administering risky, dangerous and unauthorised futures trades, he then began to lose larger and larger amounts of money, which he hid in an error account, this led to Leeson taking on bigger and bigger risks to try and regain his initial amount, and he ended up taking on enough losses to cause Barings Bank to collapse. Again, understanding the potential for this kind of behaviour could allow for preventative regulation to be introduced.

To answer this question, we will run an experiment to test how risk aversion changes after repeated losses, the experiment will roughly take the form of an investment simulation, participants will choose between two investment decisions, one of which is a ‘winner’ and one of which is a ‘loser’, the ‘winner’ will double their stake, and the ‘loser’ will lose their stake. If participants experience a run of four ‘losers’ they will be able to change their stake, the percentage they choose to invest will serve as a gauge of their level of risk aversion, and whether or not this has changed due to the run of losses. As a result of our budget constraint of £100, we will use time as a means to increase the significance, and the magnitude of emotion, associated with a loss, by making each decision take thirty seconds.

The null hypothesis

A run of consecutive losses will have no effect on participants’ subsequent levels of risk aversion.

The alternate hypothesis

A run of consecutive losses will lead to a change in participants’ subsequent level of risk aversion.

We will then conduct a significance test on the difference in the percentage remaining of the endowment staked at the beginning of the task and during the task, after the run of four consecutive losses. If we reject our null hypothesis that there will be no change in percentage staked after the repeated losses, we can state there is a probable link between losses and varying risk aversion and would suggest implementation of policy to protect at risk gamblers, such as maximum stake limits, or even a complete ban from gambling after a run of bad bets, and for derivative traders, a potential temporary ban if investments repeatedly turn bad.

Literature Review

i. Prospect Theory

The following are several papers on the topic of Prospect Theory, and its different applications and implications.

The most prominent of papers written regarding risk aversion was titled: 'Prospect Theory: An analysis of decision under risk" (Kahneman and Tversky, 1979). They wished to develop a more psychologically accurate model of decision making, compared to expected utility theory, which is a normative model, and assumes economic agents have infinite calculating capability, have access to all information and have purely outcome-based preferences.

The model is divided into two phases: during the initial phase, called editing, potential outcomes are ordered due to a developed heuristic. Primarily, whether or not outcomes are perceived as a loss or a gain relative to a set reference point. The editing phase aims to eliminate any potential framing effects, and can be broken down to coding, combination, segregation, cancellation, simplification and detection of dominance.

In the second phase, evaluation, participants behave as they would under the expected utility theory and choose the outcome to which they have assigned the highest utility.

Kahneman and Tversky wrote an equation for the evaluation phase in its simplest form:

$$V = \sum_{i=1}^{n} \pi\left(p_i\right) v(x_i)$$

where V is the expected utility of making a decision, x_i is the potential outcome of a decision, p_i is the probability of that outcome, v is a function that assigns a value to an outcome, and π is probability weighting function that captures the idea that people tend to overreact to small probability events and underreact to large probability events.

To better understand the application of this theory we can imagine two scenarios, in the first of which you start with an endowment of £1000 and gives you a choice between two options:

A. You have a 50% chance of winning an additional £1000
B. You gain an additional £500 with certainty

In the second scenario, you have an initial endowment of £2000, and are given the choice between the following two options:

C. You have a 50% chance of losing £1000
D. You lose £500 with certainty

In A and C, the payoffs are identical; 50% chance of £1000, 50% chance of £2000, as are B and D, 100% chance of £1500. Kahneman and Tversky found the proportion of people who chose the gamble increased when the question was framed as a loss, as in the second scenario, this provided evidence for their theory.

The Neural Basis for Prospect Theory, (Tom et al, 2007), investigated neural correlates of loss aversion while participants were asked whether to accept or reject 50/50 offers of gaining or losing money. They found ‘A broad set of areas showed increasing activity as potential gains increased. Potential losses were represented by decreasing activity in several of these same gain-sensitive areas. Finally, individual differences in behavioural loss aversion were predicted by a measure of neural loss aversion in several regions.’ This paper found a link between changes in neural activity and increased sensitivity to losses over gains.

Thirty Years of Prospect Theory in Economics, (Barberis, N., 2009), was a critique of prospect theory. In this he argued that even thirty years after the publication of the famous theory there were few aspects of it being applied in modern economic thought, and that the only field where prospect theory has been extensively applied is finance and insurance, two areas where risk plays a crucial role. However, he also argues that there is a trend of prospect

theory extending into other areas of economics, such as consumption choice, industrial organisation and contract theory.

He states that the most significant implication of prospect theory is reference dependence, as well as its statements on loss aversion. However, the writings on diminishing sensitivity to losses or gains are of little significance Barberis argues, as they are too difficult to apply to modern economic analysis.

ii. Risk and Loss Aversion

The following are several papers on the topic of risk and loss aversion, they all contain implications and research that exploration of our question will build upon.

Risk Aversion and Incentive Effects, (Holt and Laury, 2002), tried to provide a model for risk aversion, using lottery choice from a field experiment (Binswanger, H. P., 1980), which found that most farmers exhibit a significant amount of risk aversion that tend to increase as payoffs increase. They also used information regarding overbidding at auctions, (Harrison, G. W., 1989), to help their model, and a follow up experiment where there were increased payoffs and consequences for overbidding, (Smith and Walker, 1999).

Risk Aversion in the Large and Small, (Pratt, 1975) found that as potential losses and gains increased participants became more risk averse, this would provide support for prospect theory, as it suggests people are more sensitive to potential risk as losses increase.

Risk Breeds Risk Aversion, (He, T. S. and Hong, F., 2017), in this paper, the research question was does exposure to more or less risky scenarios affect a subject's subsequent risk-taking behaviour. This had several policy implications they listed as, for example, will an investor who has just experience a period of volatile market condition be more likely to invest in less risky assets such as government bonds or invest in more risky assets, such as small cap stocks, understanding this behaviour will allow us to predict future behaviour in

investors. Another application of their research would be on the effect of a social welfare safety net, as this effectively reduces risk for people in a society, and therefore may cause more risk-loving behaviour, which is the opposite of the result the government would want.

Their null hypothesis was that the exposure to risky scenarios will have no effect on subsequent risk-taking behaviour, against which they tested their alternate hypothesis that exposure to risky scenarios will have some effect on subsequent risk-taking behaviour.

Their experimental design was as follows: participants complete twelve rounds of the game, within each round they faced twelve decisions, a choice between lottery and a sure figure, the lottery was of the form: fifty percent chance of X and fifty percent chance of four hundred minus X, within the round the lottery remained constant and the value of Z, the sure figure increased. If a participant was risk neutral they would pick the lottery between decisions one and nine, at decision ten, they would bet apathetic between them, and for decisions eleven and twelve they would pick the lottery. The twelve rounds were divided into six each, where in the first six, they could face either low, medium, or high levels of risk, and then in the last six, they would only face medium risk. The frequency of risk neutrality was calculated using a chi squared test, and 92.7% overall, with 92.4% for the high group, 92.5% for the medium group and 93.2% for the low group.

Their results showed participants in the low group took fewer safe options, and participants in the high group took a greater amount of safer option, which supports their alternate hypothesis. They could assume this was due to the different choices they were exposed to; however, the experimenters were worried it could be due to the participants being on autopilot, and just responding without analysing their decision appropriately. To test for these, they analysed their reaction times, and found that their reaction time did significantly shorten, so there is a chance they were not being completely analytical. There was also a

concern that participants in the medium group, who were exposed to the same six rounds twice, would make different decisions in each of the identical rounds, but they found that these participants reacted the same 72% of the time.

Those in the high group, on average, switched to the safe option after six decisions in the final six rounds, those in the medium group, switched to the safe option after eight decisions, and those in the low group, after nine decisions. This would indicate exposure to risk increases risk aversion.

The authors attributed this change to three factors; firstly, evolution, humans have evolved, in high risk scenarios, to prioritise survival over maximising expected return, secondly, stress, exposure to risk induces a feeling of stress, which participants may choose to lessen by picking a safe option. Finally, they stated the main effect was the environment, participants learn from their environment, and pick safer choices after being exposed to risk.

It is important to note with this paper, the participants were not aware of the outcome of the lotteries they picked until after the experiment had ended, so the change in behaviour was purely to do with the risk and not the return.

They also concluded that the stakes involved in the risk taking were so minimal that they may have had an insignificant effect, this, as well as the small amount of time over which the decisions were made could mean that the real-world effects of this theory could be of a much greater magnitude.

Our research question could build on the findings of this paper, as they were looking at how the risk of a loss affects risk aversion, where as we would be looking at the effect of a guaranteed loss.

The Effect of Myopia and Loss on Risk Aversion, (Thaler and Kahneman, 1997), the finding from which we could build on to answer our research question. They defined myopic loss

aversion, as a combination of greater sensitivity to losses than gains and a tendency to evaluate outcomes frequently. They tested for two implications of this theory, firstly, that investors who display myopic loss aversion will be more willing to take on risk if they evaluate their outcomes less often, i.e. have access to less information, and secondly, if all payoff are increased so that player never incur losses, they take on less risk.

They found that in a task in which investors can learn from their decisions, then both of their predictions which they tested for are supported. Investors who get the most frequent feedback (therefore the most information), take on the least risk and therefore earn the least money. They also found the provision of aggregated data, as well as a restriction of opportunities to change decisions reduced the effect of myopic loss aversion. They concluded that only additional research on this affect could help to determine the right mix of information to give to employees, as their main motivation, was to understand how to maximise returns for pension investments. They found participants who were able to change their investments after feedback often did so to reduce risk and therefore return.

Study of our research question could help to build on this, as our research would uncover the exact effect of repeated losses on risk aversion and would therefore allow us to know what type of information to provide inexperienced investors to allow them to maximise their returns.

Time-varying Risk Aversion, (Guiso, Sapienza, and Zingales, 2013), analysed the effect of the global financial crisis on investors' willingness to take on risky assets. They discovered that both a qualitative and quantitative measure of risk aversion, had increased substantially since the crash. They aimed to test whether this was due to a fear factor, i.e. investing became scarier after the crash, or if it was due to some other factor, such as averseness to volatility. To test if their first theory for this increase in risk aversion was correct, they conducted a lab

experiment, where they tested whether watching a horror film affected participants levels of certainty, they found watching this film, therefore experiencing fear, led to reduced levels of certainty, they therefore concluded their first theory for an increase in risk aversion was likely to be correct and they then used the fear based model to test trading behaviours.

To test their theory, they first had to create a way to measure changes in risk aversion, this is particularly pertinent to our research question, there were two different existing approaches: the first relied on a revealed preference strategy, the second involved direct elicitations of risk attitudes from decisions in the laboratory or in survey. The revealed preferences method relies on the following equation:

$$\theta_i = \frac{r}{\alpha_i \sigma^2},$$

where θ measures relative risk aversion for investor i, α measures proportion of portfolio in risky assets, and σ^2, the variance of the risky asset. This approach is easy to apply but has several shortcomings, for example, it requires strong assumptions about beliefs, mainly, that all investors use the same historical distribution of returns, and secondly, levels of risk aversion can only be computed if investors have a positive amount invested in the risky assets market, which may not always be the case if they are highly risk averse. The main shortcoming of this approach is that to test time series changes in risk aversion we would need investors' portfolios to update instantaneously.

To overcome these shortcomings, experimenters used the second method to measure risk aversion, by using a survey, where they asked questions like the following: "Which of the following statements comes closest to the amount of financial risk that you are willing to take when you make your financial investment? 1) Take substantial financial risks expecting to earn substantial returns; 2) Take above average financial risks expecting to earn above

average returns; 3) Take average financial risks expecting to earn average returns; 4) Not willing to take any financial risks."

They state that it is broadly believed that risk aversion fluctuates with the business cycle, but these fluctuations tend to be larger than the change in wealth, this, they conclude, is due to the fear factor. Their repeated questionnaire, found that since the 2008 crisis, levels of risk aversion have increased substantially, this does not vary across differences such as wealth or habits, but most importantly this increase is present even amongst those who did not possess any risky assets. Their paper opens up the possibility that emotions play an important role on economic behaviour and can have long term, relevant effects.

This paper concerns how loss affects risk aversion, similar to our research question, however, the loss in this paper is public loss, not personal, that we wish to test, in fact some of those questioned did not lose anything in the crash as they had no risky investments, but even they experienced increased levels of risk aversion. Our paper could build on this research, to examine what role emotions play on our economic behaviour, but on a smaller scale to that pursued in this paper.

Wealth Fluctuations Generate Time-varying Risk Aversion: Micro-evidence on Individuals' Asset Allocation, (Brunner Meier, M., and Nagel, S., 2006) investigated the relationship between a households' total wealth and the proportion of their assets allocated to risky investments, they hypothesised that a reduction in total wealth would lead to a decrease in the proportion allocated to riskier investments. They used information from PSID (panel study of income dynamics) by the University of Michigan, which is the longest running longitudinal household survey in the world. They hypothesised that 'Persistent habits, consumption commitments, and subsistence levels can generate time-varying risk aversion with the consequence that when the level of liquid wealth changes, the proportion a household invests

in risky assets should also change in the same direction.' They instead found evidence to the contrary, that the share of liquid assets that households place in risky investments is completely independent of total wealth. Instead, they find the major driver behind portfolio allocation is inertia, households rebalance very slowly following capital gains or losses.

One explanation they provide for this inertia is that changing portfolio allocation could be seen as too costly, which, if this is the reason, is more due to the cost of paying any attention to the portfolio as opposed to a fixed cost in the form of explicit transaction costs, because households react exactly the same following a small wealth change or a large wealth change. This does not completely explain why there is no link between household wealth and risk aversion, as even when later changing portfolio allocation, controlling for the cost of paying attention to the portfolio, there is no link between the two factors.

In this study, they conclude no link between losses, described here as reductions in household wealth, and risk aversion, proportion of portfolio in risky assets, however in this study they are allowing for there to be a cost for adjusting risk, where as in our study for our research question we would assume that changing levels of risk incurs no cost and can be done freely and instantaneously. As well as this, there research is reliant on people self-reporting their investments, which the authors acknowledge, they may not be completely honest about.

iii. Gambling

Now we will look at academic papers regarding changes in gamblers' behaviour and how this can be related to our research question.

Risk Dishabituation, (Demaree et al, 2012) found that in repeated gambling, risk is reduced following surprising results such as low-probability wins or losses. Their original hypothesis stated that in path-dependent risk taking, like playing a slot machine, the wager on one trial may be affected by the outcome of the preceding trial. This was based on previous research

that showed that risk taking preferences in experiments was dependent on previous outcomes. The paper aimed to test two previous theories; ‘the house money effect’, which suggests risk-taking increases after a win, and the ‘break even effect’ which states that risk-taking increases after a loss, as well as the ‘mood maintenance hypothesis’, which supports the notion that risk-taking decreases after a win and increases after a loss.

To test this, they created three simple lotteries, where the chance of winning was either 13%, 50% or 87%, they found evidence for the ‘break even effect’ and ‘the house money effect’ in the 13% and 87% tests respectively. These tests produced evidence for their hypothesis, as risk aversion increased after the unlikely event occurred.

Study of our research question would build on this, as we would be testing the effect of a loss on risk aversion, whatever the chances of that loss were, and we would expect their risk aversion to decrease, which would support the ‘break even effect’.

The Role of Cognitive Bias and Skill in Fruit Machine Gambling, (Griffiths, M., 1994), detailed six cognitive biases that the author believed all chronic gamblers possessed. The first of which came from previous research, (Langer, 1975), into the illusion of control, which is defined as an expectancy of success higher than the objective probability would warrant, this was down to several factors: choice, i.e. picking your own lottery numbers, familiarity, i.e. having a favourite fruit machine, involvement, i.e. throwing your own dice in crabs, and competition, i.e. competing with a friend at roulette, all of which have no effect on the likelihood of any outcome.

The second bias, flexible attributions, came from previous research, (Gilovich, 1983), which states that gamblers attribute losses to bad luck, and wins to skill. This research showed that most gamblers viewed losses as close wins, and attribute fluke events to them, but could not

do the same for wins, this is also a result of hindsight bias, whereby gamblers are not surprised by the results of gamblers when they win as they had predicted the outcome.

The third bias, representativeness, again came from previous research (Kahneman and Tversky, 1971), and occurs where people expect to find a representative relationship between samples drawn from the population and the population itself. For example, when participants have seen a long string of coin flips where the results have all been tails, they assume heads becomes more likely, this is also called 'the gambler's fallacy', where gamblers believe wins become more likely as a string of losses increases.

The fourth bias, availability bias, supported by research, (Kahneman and Tversky, 1973), states that when a person is evaluating the probability of a chance event they based it on the relative ease with which they can get the information, i.e. winners are more public about their wins than losers about their losses.

The fifth bias, illusory correlations, supported by research, (Henslin, 1967), occurs where people believe two variables are connected when they are not, Henslin's research found craps players rolled the dice lighter to try and get a low number and harder to get a high number.

The sixth and final bias, fixation on absolute frequency, refers to when people measure success based on absolute frequency of wins as opposed to relative frequency of wins. Heavy gamblers do win more often than non-gamblers, but they also lose significantly more.

These biases show the extent to which gamblers will willingly distort reality to justify their risk taking, our study would build on this research by seeing how gamblers show cognitive biases when they are losing.

Our research will build on the work of the aforementioned studies and help to provide greater depth and understanding of the data within these papers.

Benefits of experimental versus empirical data

When answering any research question in the field of economics, there are two avenues available, pursuing an answer through the study of empirical data, or creating a simulation in a laboratory, and observing the results. The following is a justification for the latter approach with this research question, along with the potential positives and negatives with either approach.

i. Experimental Data

The advantages of using a laboratory experiment for this research question can be broadly split into three categories: ability to isolate cause and effect, precise control of variables and easy to replicate. With regard to our research question, isolating cause and effect accurately is essential, an experiment to answer the question would have to create a loss scenario that can be easily repeated, such as a random event like flipping a coin, from this repeated loss we would be able to see at what point, if at all, a participant's behaviour changes, the cause here being our repeated random event and the effect being the change in behaviour.

The precise control of variables is required to accurately isolate the aforementioned cause and effect. By controlling external variables, we can make sure the dependent variable, in our case, the level of risk aversion, is only affected by the independent variable, the number of consecutive losses the participant has experienced. However, even in a laboratory setting, it is impossible to completely control every variable, as participants may be experiencing issues outside of the laboratory that could affect their risk aversion, such as ill health, or financial or relationship troubles, though, the precisely controlled nature of a laboratory setting would reduce the magnitude of these effects.

The ability to easily replicate an experiment, allows us to gather large amounts of data, with identical conditions, so we can detect anomalous results significantly easier, and remove

them from our results, this would help reduce the effect of the aforesaid potential external effects, as well as this, it allows us to see if repeated losses have a consistent effect on risk aversion, as for some participants, it may lead to increased risk aversion, i.e. to cut losses, or to reduce risk aversion, i.e. to chase losses.

The disadvantages of a laboratory experiment can also be divided into three categories: artificiality, too tightly controlled behaviour, and demand characteristics. For our research question, artificiality is a large issue, as it is difficult to recreate a significant loss, as participants are unlikely to willingly participate in an experiment where they could lose large amounts of money, and furthermore, the length of time exposed to these losses would be very small compared to a trader who has been losing money for months, or a gambler who has lost everything over the course of a year.

Tightly controlling variables to isolate cause and effect, may restrict participants' behaviour, and may lead them to behaving in a different manner to how they would in a real-world scenario, this could lead us to an incorrect conclusion, participants' may have time varying risk aversion, which could be hard to detect in a restrictive laboratory experiment.

There is also a risk that as participants become aware of what the experiment is testing, in this case, changes to risk aversion, they change their behaviour, in our experiment they may try to be more consistent with their levels of risk.

Moreover, due to ethical concerns surrounding deception, it is hard to create an experiment in which participants repeatedly experience losses, the only way to circumvent this issue is to create a game in which there are many repetitions, and then observe when repeated losses naturally occur, a downside of this is that due to budget constraints, to implement this, the potential financial losses and gains involved in each game would have to be very small, this

negative can be lessened by increasing the time for each decision to occur so participants are in the loss frame of mind for longer.

ii. Empirical Data

Dutch Psychologist, Adriaan de Groot, outlined five steps that all empirical research must follow, which he called the empirical cycle, firstly, observation of a phenomenon and an inquiry into its causes. For our research question, this is the change in a person's aversion to risk. Secondly, there is induction, the formulation of hypotheses, a generalised explanation for the phenomenon, our null hypotheses would be that repeated losses would have no effect risk aversion, against our alternate hypothesis, that there would be some effect.

Thirdly, deduction, the formulation of experiments that will test the hypotheses, we have two main ways to complete this step to help with our research question, conduct an experiment in a laboratory, which has the aforementioned advantages and disadvantages, or we can analyse data from real world scenarios that closely recreate our situation of repeated losses, examples of which we will discuss shortly, along with their benefits and shortcomings. Fourthly, testing, where we conduct the previously constructed experiments, or analyse the data from our previously stated examples. Finally, evaluation, we interpret the results and construct a theory, an abductive argument, where we present the results of the experiment, as the most reasonable explanation of our phenomenon, with our research question, we would see whether or not repeated exposure to loss has affected our subjects' risk aversion. (A De Groot, 1956).

To help answer our research question, we could observe changes in practice of investment firms, before and after periods of particularly poor returns, and see whether or not aversion to risk changes. A particularly negative period could be the 2008 financial crisis and the subsequent recession. To help answer our research question, we could build on the findings of

the paper Time-varying Risk Aversion (Guiso et al, 2013), through isolating the exact point where behaviour began to change, and whether this was a result of the repeated losses, or due to reduced funds, or even potentially due to new regulation on the market. However, due to the nature of empirical research, it would be difficult to isolate which of these caused the change in behaviour. A way to assess change in investment activity, could be to see how leveraged firms are with their investments, to see how diversified their portfolios, to see the ratings of the bonds they take on, all of these could be indicators for change in attitude to risk. Due to the global nature of the financial crisis, we may observe slightly different effects than if we just observed a firm that was unique in its losses, the crisis may have changed the entire attitude within the financial sector towards risk, and therefore might not be fully representative of individual losses.

Another economic agent that could be observed to answer our research question is chronic gamblers, problem gamblers tend to be unable to regulate their risk, and often take on larger and larger risks, incurring huge losses, as they try to chase back their initial sum of money. The controlled nature of a casino or a bookmaker would allow us to see when behaviour changes and in what form this change occurs, i.e. through increased stakes, or through longer odds, however chronic problem gamblers may not be completely representative of the population, they may be more risk loving than the average economic agent, which is why they would first be attracted to gambling. Furthermore, they are unlikely to provide accurate records of their losses, due to sensitivity around sharing the information surrounding these losses, or they may simply be unaware of how much they have lost. To counteract this problem, we could begin to observe gamblers, and keep our own records, however, there could be many issues with this approach, firstly, demand characteristics, once they are aware of the purpose of the experiment, they may take on greater risks to appease the experimenter, and there would be ethical concerns with observing someone lose significant amounts of

money and not working to prevent it, especially if there is a chance the participants are increasing their risk due to being observed.

An economic agent that could be observed to help answer our research question is asylum seekers from war zones, similar to the paper, Violent Conflict and Behaviour, (Voors et al, 2012) which focused on the relationship between risk aversion and living in a war zone, they found that exposure to war reduced risk aversion as people had to take large risks to escape. We would interview asylum seekers about their previous investment decisions, or previous attitudes to risk and to see how these attitudes and behaviours have changed, the effect of war would definitely be significant enough time and cost wise and would provide a clear event window to see if behaviour changes. However, there would be a multitude of issues with this approach, namely, that previous countries may have societies entirely incomparable to the United Kingdom, where behaviour is either restricted due to a repressive government, or cultural attitudes towards risk are different, this would make any comparisons between before and after the event irrelevant. For example, a large number of refugees in the West are from majority Muslim countries, where accruing interest on a loan is considered haram, and therefore their attitudes to investing would inevitably be different, furthermore, most refugees are likely to be on the lower end of the income scale, and therefore unlikely to be investing their limited income.

Another slightly more light-hearted methodology could be to analyse the decisions of participants on ITV's The Chase, as participants in this show compete in an initial 'cash builder' round, where they must answer as many general knowledge questions in a minute as they can, earning £1000 for each correct answer, if we categorise each incorrect answer as a loss, then we can observe the participants with many incorrect answers in the next round, where they must make a risk-based decision. They choose between three offers, one low-risk and low-reward, one medium-risk and medium-reward, and one high-risk and high-reward.

Our hypothesis would predict that those with many incorrect answers in the first round would always pick the low risk and reward offer, but this is not always the case. An analysis of the decisions of participants could help answer our research question, however there are many other factors involved that could influence their decision, so we would not be able to make any definitive conclusions from this data.

We could also observe individual case studies, for example, we could analyse the behaviour of Nick Leeson, the former derivatives rogue trader, whose actions led to the collapse of Barings Bank in nineteen ninety-five, after experiencing initial success from administering risky, unauthorised trades, he then began to lose significant amounts of money, which he hid in an error account, this led to Leeson taking on greater and greater risks to try and regain his initial losses, and he ended up taking on enough losses to cause the bank to collapse. In this example, it is easy to isolate the cause and effect of Leeson's actions, the initial loss was the cause, and the effect was the subsequent reduction in risk aversion, this would support the theory that gamblers become more risk loving as they experience loss. However, analysing the behaviour of one individual is not enough to be able to draw any significant conclusions, as his behaviour may have been anomalous, as well as this, due to the nature of his profession, he is likely to have an unusual attitude to risk, which would make him unrepresentative of a typical economic agent. (*Rogue Trader*, N Leeson, 1996).

These examples of potential empirical research all share similar advantages and disadvantages, their main advantages being the same as the disadvantages for conducting a laboratory experiment, namely: no issue of artificiality, for the simple reason that the scenarios are not artificial. Also, there is no issue of demand characteristics, as most of the data will be gathered post event-window, so the participants cannot be affected by the purpose of the experiment as they have already made their decisions. Finally, as participants are being observed in their normal environment, their behaviour is not restricted in any way,

as it might be in a laboratory, therefore we are better able to see what exactly causes a change in behaviour in an uncontrolled environment.

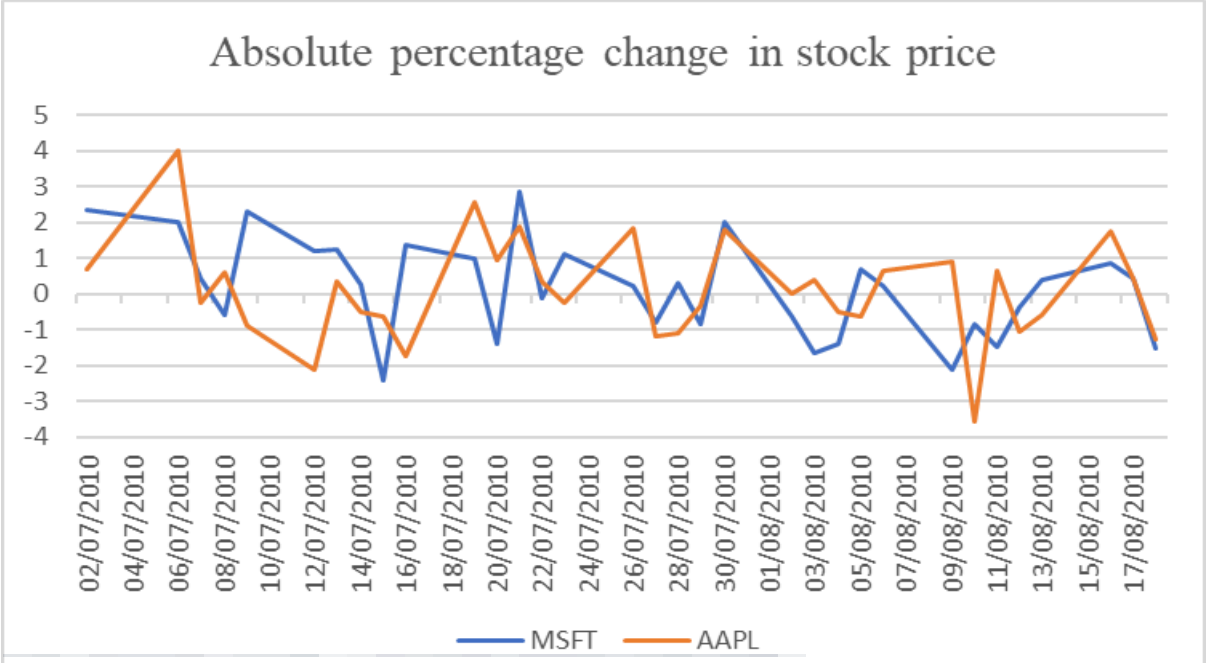
The main disadvantages of this form of empirical research are threefold, namely: very hard to isolate cause and effect, a change in risk aversion may be coincidental or may be due to some other unknown cause, this issue would be lessened in a laboratory. As well as this, it is hard to precisely control variables, but as we previously mentioned, this can be an advantage as this allows participants to behave realistically. Furthermore, for certain examples, such as that of Nick Leeson, it is difficult or even impossible to replicate data, and it is unwise to draw conclusions from a single data point, so these examples would function better as demonstrations of our theory in action. However, for some of our examples, such as those observing asylum seekers, or gamblers, there are a large number of potential data points, but each case will be unique and therefore it would be very difficult to isolate a specific individual cause. Therefore, we should conclude that the most effective way to answer our research question is through experimental data.

Methodology

To test our research question, we shall run several sessions outside a laboratory in controlled conditions, these locations include the university library. The participants will complete a task which will involve several investment decisions. The participants will come from as varied background as possible to prevent any potential heuristics or implicit biases affecting decision making, there will be multiple sessions in the laboratory.

The budget of this experiment is £100, provided by the University of Exeter Business School, £8.97 was spent on materials for the experiment, such as counters, and the remaining £91.03 will be spent on paying participants.

The participants will be told they are analysts at an investment firm, InvestCo, and they must make decisions on how to invest their money. They shall each be given an endowment of five pounds, throughout the experiment, they will use coloured counters, each worth twenty pence, to represent this endowment, at the end of the experiment they will exchange these for real money. They will face thirty-three decisions total, all these decisions will be between two options, Microsoft stock and Apple stock, between the second of July 2010, and the eighteenth of August 2010, this data was gathered from Yahoo Finance, if participants correctly choose the stock which performed better on that day, they will double their investment, if they choose incorrectly, they will lose their stake.



The above graph shows the percentage change in stock prices across time periods, this was calculated by dividing the closing value of the stock on day t+1 by the closing value of the stock on day t.

For the first decision, participants will be able to invest anywhere between zero and five pounds, in increments of twenty pence, the stock that performed best will be revealed and participants will receive pay outs accordingly. They will enter the next stage of the task, where they will have to make the same decision as before, however this time they must stake twenty pence, to increase the significance of a loss, each decision will last thirty seconds. Participants will then receive their pay out, however, if they experience a loss, they will receive a red counter, the task then repeats. If a participant loses again, they will receive another red counter, however, if they win, they lose their initial red counter. If they accrue four red counters, through four consecutive losses, they are told “your manager at InvestCo has decided to give you a chance to make more money due to your recent run of bad form, and there are now no limits on your level of investment”, and then they are allowed to set their stake anywhere between zero and the remaining amount of their endowment. if they get this decision wrong, they are again, allowed to stake as much as they would like. If at any point, their funds reach zero they are out of the task.

As the task involves a choice between two stocks, there is roughly a 50% chance of the participant choosing a bad investment, this means a run of four bad investments, should happen roughly one in sixteen times, and therefore twice in our experiment, however as this can continue if they answer wrong again, this should happen on average more than twice in our experiment.

The extent to which the percentage of their total endowment that they stake changes after a run of four losses will serve as a metric to see how levels of risk aversion have changed.

This experiment was chosen to maximise the possibility of exposure to repeated losses, subject to our budget constraint, and to maximise the number of potential participants. Due to the university's code of ethics, participants cannot be deceived, therefore losses have to occur naturally, so we created a design with a 50% chance of a loss, to remove any possible effect of a gambling heuristic developing, participants will choose between two stocks that's relative performance will be random, and unknown to the participant, but will create a frame of mind of an investor, as opposed to a gambler.

We will conduct several econometric tests to test our research question. The first will be to construct a regression equation, of the form:

$$r_L = \alpha + \beta r_I + \varepsilon,$$

where r_L is the level of risk after a loss (average percentage of remaining endowment staked at each decision), α is a constant, β is a coefficient that determines the effect of loss on risk aversion, r_I is the initial level of risk (percentage staked in the initial decision), and ε is an error term. We will test the significance of the coefficient β using a two-sided t-test, our hypotheses will be:

Null hypothesis: Levels of risk aversion will not be affected by repeated losses and therefore $\beta=0$.

Alternate hypothesis: Levels of risk aversion will be affected by repeated losses and therefore $\beta\neq0$.

For our regression to be valid, three caveats must be satisfied, those being:

1. The variables must be independent
2. The data must not be heteroskedastic
3. The variables must not be serially correlated

The first caveat can be tested with a Chi Squared Test, to test if our two variables are related; we will test the null hypothesis that they are not related, against the alternate hypothesis, they are related.

The second caveat will be tested with a Breusch Pagan Test, to test if our data is heteroskedastic, we will test the null hypothesis that the data is homoscedastic, against the alternate hypothesis that the data is heteroscedastic.

For the final caveat, as our data is not time series, it cannot be autocorrelated, therefore we do not need to worry about this caveat.

From calculating our regression, we will also calculate three other variables, the correlation coefficient, r, the residual sum of squares, RSS, and the coefficient of determination, R^2, from these figures we will be able to infer information about our data.

We will then run a multiple regression, taking into account the effect of the number of decisions faced, of the form:

$$r_L = \alpha + \beta_1 r_I + \beta_2 x + \varepsilon,$$

we will then conduct significance tests on the two coefficients.

Following these regression tests, we will conduct a non-parametric test, as we cannot guarantee our data is normally distributed.

We will use a Spearman's Rank Correlation Coefficient test to examine whether there is a relationship between the change in levels of risk aversion before and after consecutive losses and the number of critical decisions faced.

We will use this formula to calculate the coefficient from our data: $R_s = 1 - \left(\frac{6\sum d^2}{n^3 - n}\right)$,

R_s, will show us how correlated these two variables are, with 1 indicating absolute positive correlation, -1 indicating absolute negative correlation. The absolute value of the coefficient, as well as the direction, will indicate how strong the correlation is and whether it is positive or negative.

Ethical Considerations

There are minimal ethical concerns with this experimental design, as participants will not be using their own money, and all the investment decisions will be observed in a controlled environment. Furthermore, all participants will be able to give informed consent, meaning they are not a child, or disabled in a manner that means they cannot give informed content. All data will be kept confidential and will be deleted after the completion of this dissertation.

A potential ethical concern is that of triggering a gambling addiction in a participant, due to the gambling nature of the task, to prevent this, participants will be screened for any previous problems, and asked if they have ever suffered from a gambling addiction, if they have, they will not be allowed to complete the task to prevent any chance of regression.

Results

i. Presentation of results

Due to our budget constraint, we were only able to gather results of 17 participants, however our results serve as a proof of concept for further research with a larger budget to build on. Due to the small number of data points, what follows is primarily a description of the data

along with statements regarding potential trends in the data, that could be explained via existing research, or could be developed by further research.

The data here has been calculated and displayed using Microsoft Excel, unless otherwise stated.

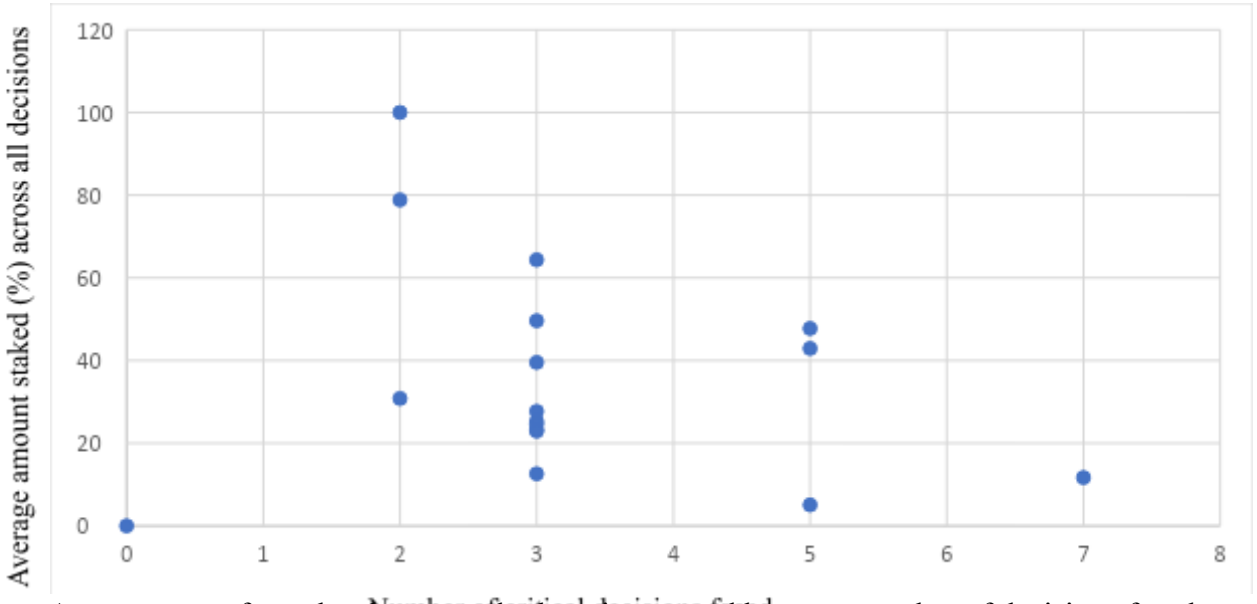
The following the percentage each participant staked at each decision, as well as how many decisions they faced, and the amount they received at the end of the experiment:

	Initial stake	1st	2nd	3rd	4th	5th	6th	7th	Final Amount
1	8	57.6923	100						16.6
2	16								6.6
3	8	12.5	9.52381	15.7895					5.2
4	100	19.2308	42.3729						7
5	4	4.7619	5	4.7619	5.26316	5.55556			4
6	12	23.5294	26.3158	20					4.4
7	40	7.14286	38.4615	37.5					2
8	100	100	100						0
9	40	42.8571	25	42.3077	28.5714	100			7.8
10	40	18.1818	62.5	68.1818					7.4
11	100	4.08163	5.76923	66.0377					3.6
12	60	38.4615	41.6667	38.4615					3.6
13	20	4.16667	8.69565	9.52381	21.7391	16.6667	13.3333	7.69231	2.4
14	40	31.25	18.1818	22.2222	100				0
15	40	21.4286	27.2727	25					1.6
16	20	55.5556	37.5	100					6.8
17	28	31.0345	21.0526	16.6667					10

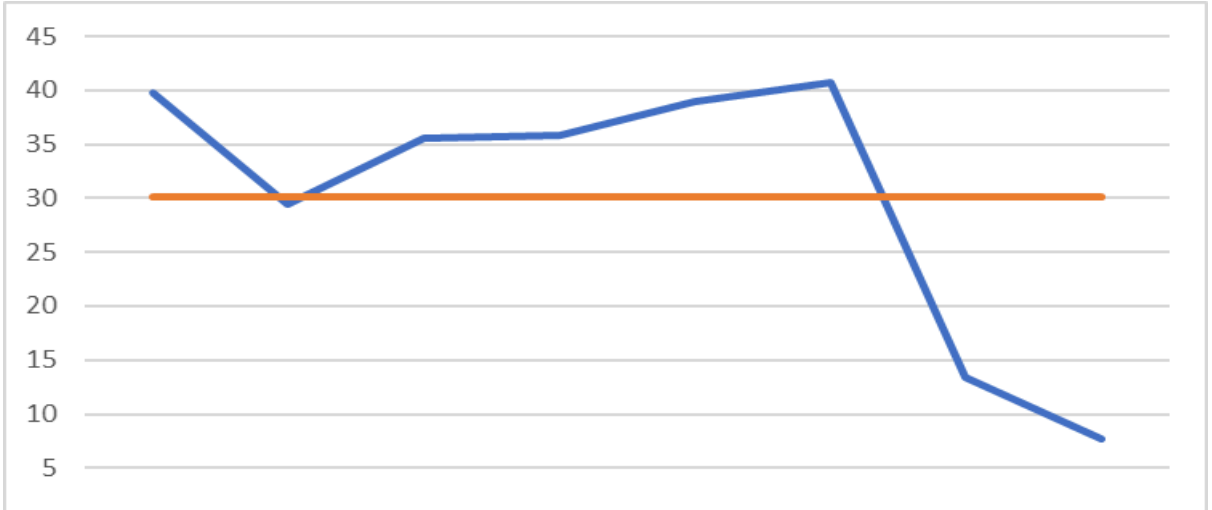
Using the aforementioned experimental design, we gathered data on 17 participants. They came from a varied background, with an age range from 18 to 68, with 12 being male, and five being female, with seven being students, and ten being non-students. Within the category of student, three were from the business school, two were from the college of life and environmental science, one was from the college of Engineering, Mathematics and Physical Sciences, and one was from the college of Humanities. Six participants were non-UK natives. The diversity of the background of our participants should have reduced the effect of any implicit biases or heuristics.

The average amount each participant received was £5.24, which was slightly above the expected £5.00, this meant that only 17 participants were involved, as opposed to the 18 predicted, there was £2.03 remaining of our budget at the conclusion of our experiment.

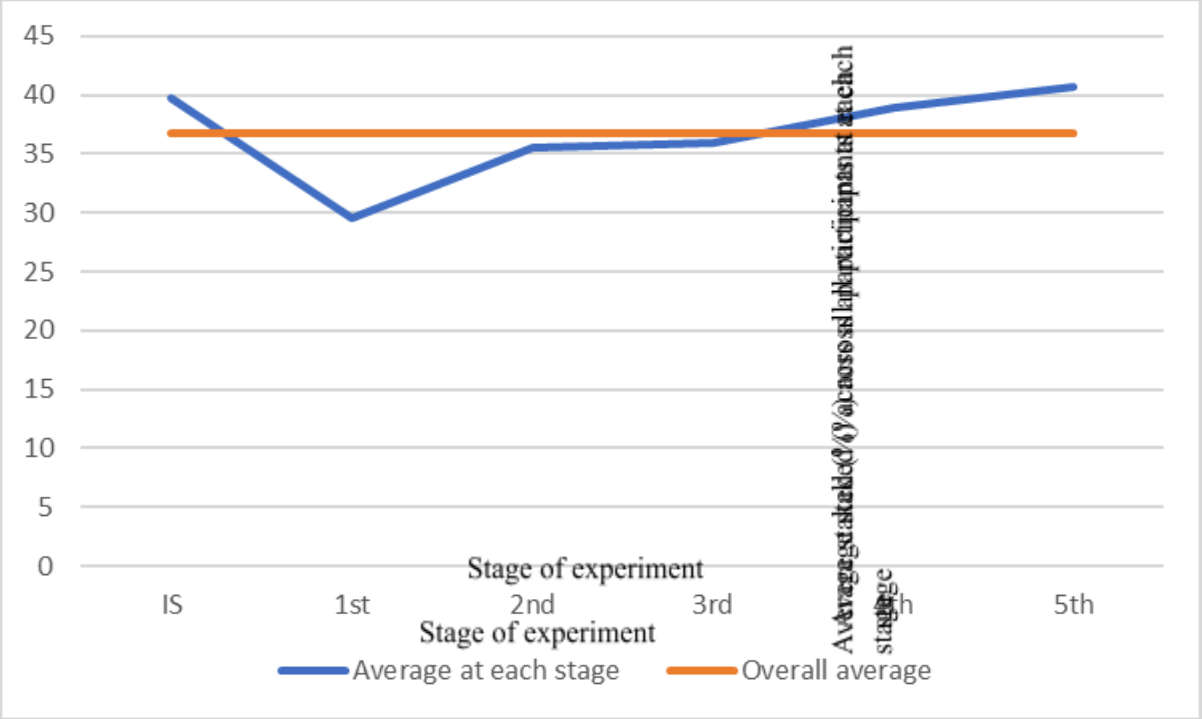
The average number of open ended decisions each participant faced was 3.05882, which was broadly in line with our prediction, that with 32 even decision, a run of four losses would occur twice, with a 50% chance after this of facing another decision, i.e. five consecutive losses, and then diminishing chances of six or more consecutive losses, so just above three was our prediction. Of this, one participant faced seven critical decisions, three faced five decisions, one faced four decisions, nine faced three decisions, three faced two decisions, and one faced no decisions, this is excluding the initial open-ended decision.



As we can see from the above graph, there is no trend between number of decisions faced, displayed here on the x-axis, and the average stake, displayed here on the y-axis.



The average initial stake was 39.764%, the average for the first stake was 29.4922%, for the second stake, 35.582%, for the third, 35.881%, for the fourth, 38.893%, the fifth, 40.741%, the sixth, 13.333%, and the seventh was 7.692%. The last two were only faced by one participant. The above graph shows the average staked at each stage versus the overall average, 30.1725%, as we can see, for the first five decisions, it is primarily above the average, indicating a reduction in risk aversion with increasing losses. However, if we remove the last two decisions, as they were faced by only one participant, we produce the following graph, where the average is now 36.7257%.

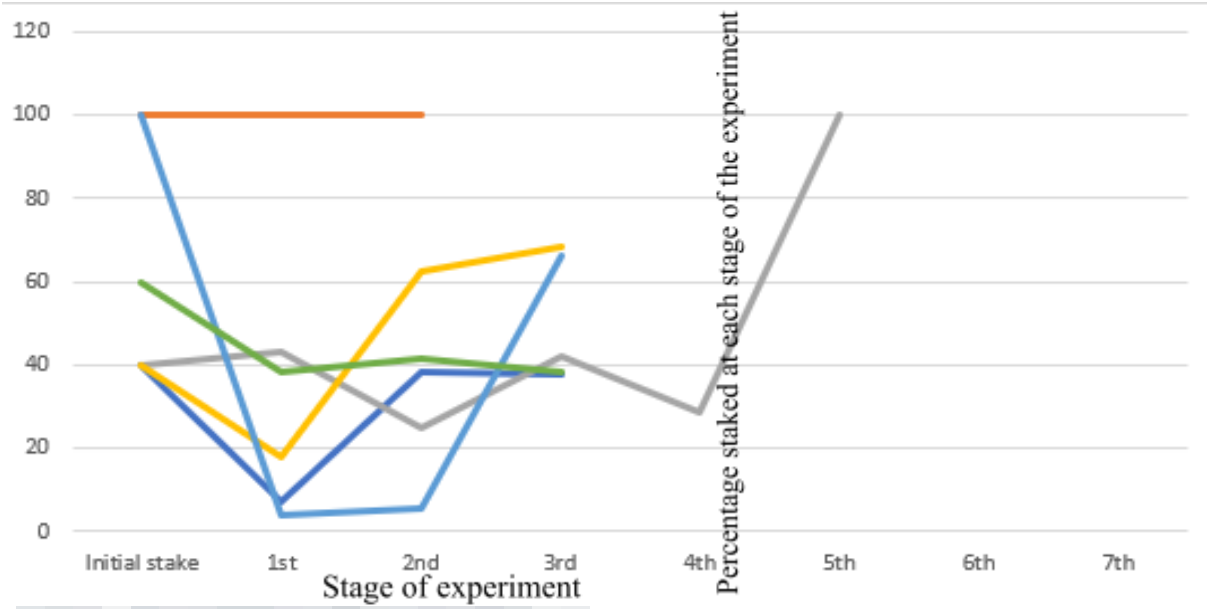


As we can see, the initial stake is above the average, which as it occurs before any losses occur, indicates that participants become more risk averse after the initial losses, but as the game progresses, around the 4th decision, their level of risk aversion decreases. However, the difference is minimal at these different stages, and due to the small number of data points, meaning one particularly risk averse, or risk loving individual could have a large influence on the average, it is hard to infer any significant conclusions from this, but future research could explore this avenue further. This would seem to support our null hypothesis that there would

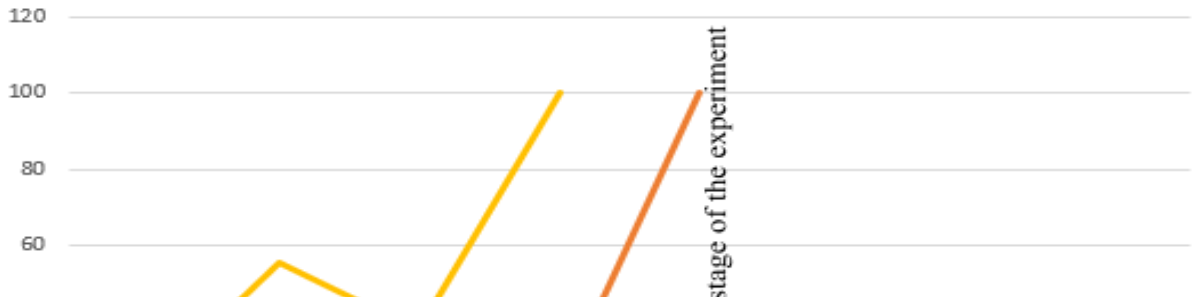
be no significant change in percentage staked, as the differences are small, however this will be tested shortly using a significance test.

We can see from the following graphs, which have been split into three to make the information easier to interpret, how each participants behaviour changed as the experiment progressed.

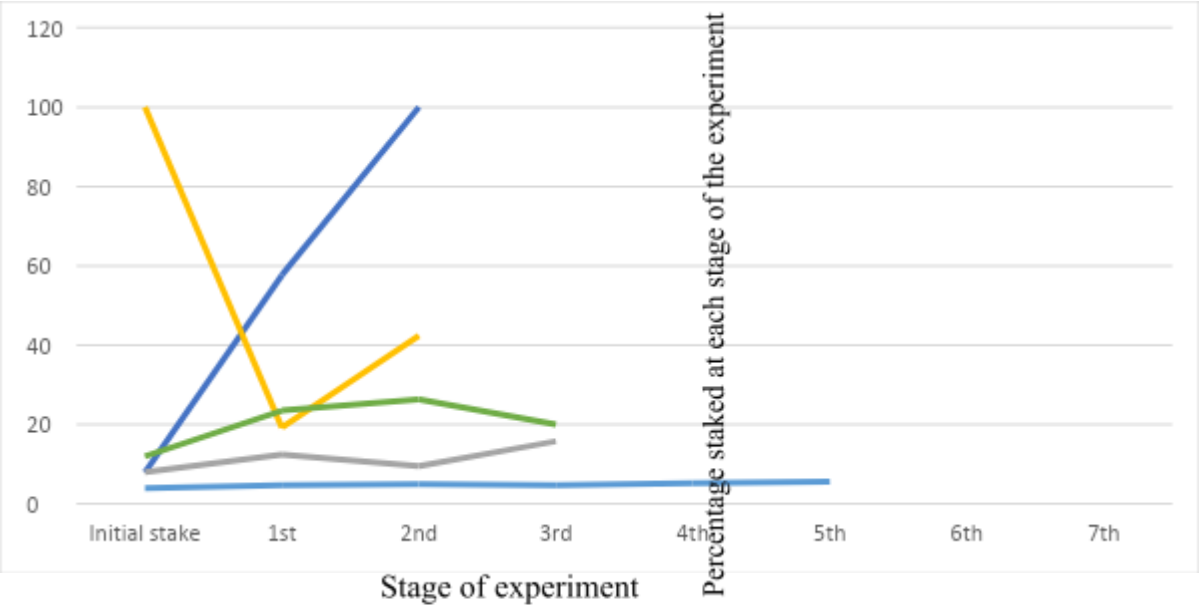
The 1st 6:



The 2nd 6:



The final 5:



As we can see from these graphs, participants did not display consistent tactics, with some reducing risk as the experiment progressed, some increasing risk, and some not changing tactics at all.

ii. Analysis of results: regression

To test our null hypothesis, we shall construct a regression from our data and then conduct a significance test on the effect of the consecutive losses.

However, before we can construct an estimation of our regression, we must test that our data does not violate the three caveats required for a regression, these being:

1. The variables must be independent
2. The data must not be heteroskedastic
3. The variables must not be serially correlated

We can test the first caveat by conducting a chi squared test, we will use a chi square statistic to test the null hypothesis that the two variables are not related. Using the Excel function CHISQ.TEST and taking our range of observed values as the average percentages staked after consecutive losses, and our range of expected values as the initial stake, we obtain a Chi Square value of 6.94×10^{-178}, which is below our critical value of 7.962, which is taken from the Chi Square table, with a probability of 0.95, the left hand tail, and 16 degrees of freedom, so we fail to reject our null hypothesis that the two variables are not related.

To test our second caveat, we will use the Breusch-Pagan test to test our null hypothesis that the variables are homoscedastic, against the alternate hypothesis that the variables are heteroscedastic. To calculate the test statistic for the Breusch-Pagan test we use the formula nkR^2, where n is the sample size, k is the degrees of freedom, and R^2 is the coefficient of

determination calculated from our regression. This is 17x16x0.1194, which equals 32.4768, which is greater than our critical value of 26.3, which is taken from the Chi Square table, with a probability of 0.05, the right-hand tail, and 16 degrees of freedom, so we therefore fail to reject our null hypothesis that our variables are homoscedastic.

As previously mentioned the third caveat is not relevant to our data, so therefore we can ignore it.

As all conditions have been met we can now construct our regression.

As we are estimating the regression, the expected value of the error term is zero. We calculated the following regression using the webtool *xuru*, and it as follows:

$$\hat{r}_L = 24.51 + 0.2815\hat{r}_I,$$

as $E[r_L]=\hat{r}_L$, $E[r_I]=\hat{r}_I$, $E[\alpha]=24.51$, $E[\beta]=0.2815$, and $E[\varepsilon]=0$.

Also calculated from our regression, were the correlation coefficient, r, which was equal to 0.3455, the residual sum of squares, rss, which was equal to 9848, and the coefficient of determination, R^2, which was equal to 0.1194.

To test our original null hypothesis, using this regression, we must create a new set of null and alternate hypotheses, H_0: $\beta=0$, i.e. r_I has no effect on r_L, and H_1: $\beta\neq0$, i.e. r_I has some effect on r_L.

To test this, we can conduct a two-sided t-test, with 95% confidence and 16 degrees of freedom, this gives us a critical value of 2.12. However, firstly we can construct confidence intervals for $\hat{\beta}$, of the form: $\hat{\beta}-2.12\sigma\leq\hat{\beta}\leq\hat{\beta}+2.12\sigma$, however we do not know the standard deviation, σ, so instead we use the standard error. Using the Excel function STDEV.S we calculate the standard error of r_I to be 32.449, we create a confidence interval of $-68.5104\leq\beta\leq69.0734$, 95% of values of β should fall within this interval, assuming our results

are normally distributed. As we only have 17 results, the central limit theorem does not hold, so we cannot guarantee the results are normally distributed, we shall, however, continue as if the theorem does hold.

To proceed with our t-test, we need to calculate the test statistic, using the formula: $\frac{\hat{\beta}-\beta^*}{s.e(\hat{\beta})}$, where β*, is the value of β in our null hypothesis, in this case, 0. This leads us to (0.2815-0)/32.449=0.008675, which is less than our critical value of 2.12, so we fail to reject our null hypothesis, so we must conclude from this test with 95% probability that r_I has no effect on r_L.

We can interpret different results from our information from our regression, such as the value of r, or the Pearson correlation coefficient, which is equal to 0.3455, indicates a weak positive relationship between r_L and r_I. Our value of rss, the residual sum of squares, is 9848, which is large compared to the average values of our data points. Another value we calculated was R^2, the coefficient of determination, which was relatively small at 0.1194, which means that 11.94% of changes in r_L were attributable to r_I.

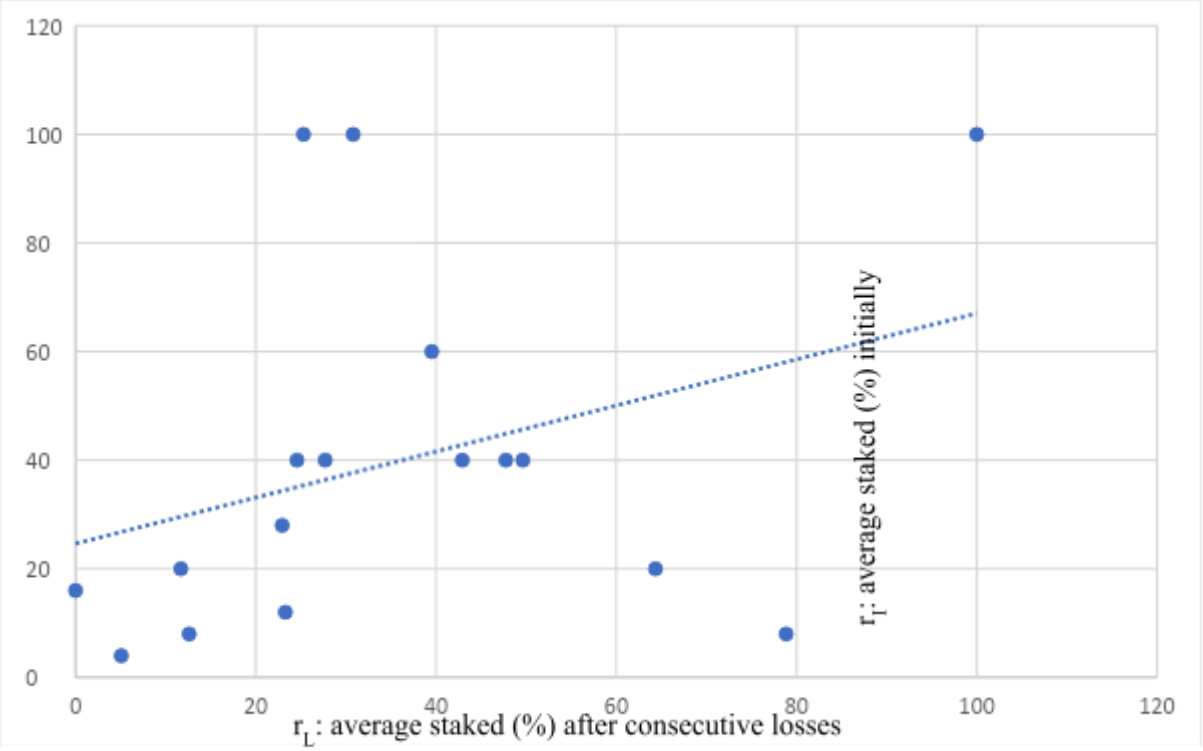


Figure 8: A scatter diagram with rI on the y-axis, and rL on the x-axis, with a line of best fit

We can additionally run a multiple regression, to include the number of decisions each participant faced, to produce the following equation, produced using the multiple regression calculator on *socscistatistics.com*:

$$\hat{r}_L = 32.1362 + 0.2603\hat{r}_I - 2.137\hat{x},$$

where $\hat{r}_L$, is an estimate of the average percentage staked after repeated losses, $\hat{r}_I$, is an estimate of the initial percentage staked, and $\hat{x}$, is an estimate of the number of decisions each participant faced, excluding the initial decision. For the latter variable, an increased number of critical decisions reduces overall levels of risk taking, as the regression produces a negative coefficient.

Using Excel we can calculate the standard error of our variables, s.e($\hat{r}_I$)=32.449, and s.e($\hat{x}$)=1.5098, using this we can test our null and alternate hypotheses that these variables have a significant effect on our dependent variable, r_L.

Our first set of hypotheses are as follows: H_0: $\beta_1=0$, and H_1: $\beta_1\neq0$, where β_1 is the coefficient of regression of r_I, to test these we can use the following equation: $\frac{\hat{\beta}-\beta^*}{s.e(\hat{\beta})}$, where β^* is the value of β in the null hypothesis, in this case 0. (0.2603/32.449) =0.008022, which is lower than our critical value of 2.12, so we fail to reject our null hypothesis that r_I has no effect on r_L.

Our second set of hypotheses are as follows: H_0: $\beta_2=0$, and H_1: $\beta_2\neq0$, where β_2 is the coefficient of regression of x, we will use the aforementioned equation to test this hypothesis. (-2.137/1.5098) = -1.4154, which is less than our critical value of 2.12, in this instance the negative sign is irrelevant as it just indicates direction, so we fail to reject our null hypothesis that x has no effect on r_L.

iii. Analysis of results: non-parametric tests

Due to our limited number of data points we cannot guarantee our data is normally distributed, therefore we can also conduct a non-parametric test. The test we will conduct is the Spearman's Rank Correlation Coefficient, which we will use to test the strength of the link between the difference in initial levels of risk aversion and subsequent average risk aversion, r_I and r_L, and the number of critical decisions faced. The null hypothesis of this test will be there is no link between the number of decisions faced, and the difference in initial and subsequent risk aversion, tested against our alternate hypothesis that the difference will change with increasing numbers of critical decisions.

Participant	Difference between rL and rI	Rank difference	no of decisions	Rank no.	Difference	d^2
1	0	17	2	3	14	196
2	1.068504595	16	0	17	-1	1
3	2.913510101	15	3	9	6	36
4	4.604427736	14	2	3	11	121
5	5.082072999	13	5	15.5	-2.5	6.25
6	7.747252747	12	3	9	3	9
7	8.311776216	11	3	9	2	4
8	9.621212121	10	2	3	7	49
9	11.28173375	9	5	15.5	-6.5	42.25
10	12.2985348	8	3	9	-1	1
11	15.43290043	7	3	9	-2	4
12	16	6	3	9	-3	9
13	20.47008547	5	7	1	4	16
14	44.35185185	4	4	14	-10	100
15	69.19817471	3	3	9	-6	36
16	70.84615385	2	3	9	-7	49
17	74.70380024	1	3	9	-8	64

Table 2: Spearman's Rank Correlation Coefficient Data Table

The difference between r_L and r_I is absolute, so negative signs are ignored. The formula for the Spearman's Rank Correlation Coefficient is as follows:

$$R_s = 1 - \left(\frac{6\sum d^2}{n^3 - n}\right),$$

using our data, we produce the following equation:

$$R_s = 1 - \left(\frac{6x743.5}{17^3 - 17}\right) = 1 - \frac{4461}{4896} = 0.08841,$$

this figure, is below the critical value of 0.521, taking from the Spearman's Rank Correlation Coefficient table, with a 5% significance level and 15 degrees of freedom, we therefore fail to reject the null hypothesis, and conclude that any correlation between the two variables is coincidental.

Conclusion

Our initial prediction was that a run of losses would lead to a change in levels of risk aversion, however the econometric analysis of our results was ultimately inconclusive. This was primarily due to the limited number of participants, as well as this, the small financial incentives for each section of the game, twenty pence, were insignificant, and participants were unlikely to have reacted in the same manner as if these amounts were larger, as would likely be the case in most scenarios. However, our research did provide a basis for further research, which could improve our findings, by increasing the relative amount each participant stakes in the game, this could be achieved either via increasing the absolute amount of cash involved, or performing the experiment in an area with lower average incomes, such as rural India, however in this scenario, researchers would have to control for any cultural differences in attitudes to risk.

The literature reviewed provided many different conclusions to how risk would be affected by losses, with some previous research indicating an increase in risk taking, and some indicating a decrease was to be expected, however, from our own research we were unable to find evidence either way.

Our initial alternate hypothesis was that the run of losses would lead to some change in risk aversion, however, due to the reasons mentioned previously we were unable to draw any significant conclusions. Hopefully further research will be done on this question as an understanding of time-varying risk aversion is essential in protecting economic agents, or their clients and employers, from catastrophic psychological foibles.

References

Brunnermeier, M. and Nagel, S. (2005). Do Wealth Fluctuations Generate Time-Varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation. *SSRN Electronic Journal*.

De Groot, A. (1956). The Meaning of "Significance" for Different Types of Research.

Demaree, H., Burns, K., DeDonno, M., Agarwala, E. and Everhart, D. (2012). Risk dishabituation: In repeated gambling, risk is reduced following low-probability "surprising" events (wins or losses). *Emotion*, 12(3), pp.495-502.

Gilovich, T. (1983). Biased evaluation and persistence in gambling. *Journal of Personality and Social Psychology*, 44(6), pp.1110-1126.

Griffiths, M. (1994). The role of cognitive bias and skill in fruit machine gambling. *British Journal of Psychology*, 85(3), pp.351-369.

Guiso, L., Sapienza, P. and Zingales, L. (2013). Time Varying Risk Aversion. *SSRN Electronic Journal*.

He, T. and Hong, F. (2017). Risk breeds risk aversion. *Experimental Economics*.

Henslin, J. (1967). Craps and Magic. *American Journal of Sociology*, 73(3), pp.316-330.

Holt, C. and Laury, S. (2002). Risk Aversion and Incentive Effects. *SSRN Electronic Journal*.

Kahneman, D. and Tversky, A. (1979). Prospect Theory: An analysis of decision under risk

Langer, J. (1975). The Illusion of Control.

Leeson, N. and Whitley, E. (n.d.). *Rogue trader*.

Pratt, J. (1976). Risk Aversion in the Small and in the Large. *Econometrica*, 44(2), p.420.

Thaler, R., Tversky, A., Kahneman, D. and Schwartz, A. (1997). The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test. *The Quarterly Journal of Economics*, 112(2), pp.647-661.

Tom, S., Fox, C., Trepel, C. and Poldrack, R. (2007). The Neural Basis of Loss Aversion in Decision-Making Under Risk. *Science*, 315(5811), pp.515-518.

Tversky, A. and Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), pp.105-110.

Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), pp.207-232.

University of Michigan (1968) Panel Study of Income Dynamics

Voors, M., Nillesen, E., Verwimp, P., Bulte, E., Lensink, R. and Soest, D. (2012). Violent Conflict and Behavior: A Field Experiment in Burundi. *American Economic Review*, 102(2), pp.941-964.

Appendices

Stock market data: AAPL vs MSFT (green indicates this stock performed best on that day):

Date	MSFT	AAPL
02/07/2010	2.364	0.684
06/07/2010	2.015	4.038
07/07/2010	0.453	-0.224
08/07/2010	-0.574	0.593
09/07/2010	2.307	-0.897
12/07/2010	1.208	-2.134
13/07/2010	1.234	0.369
14/07/2010	0.275	-0.506
15/07/2010	-2.43	-0.616
16/07/2010	1.366	-1.729
19/07/2010	0.991	2.569
20/07/2010	-1.413	0.933
21/07/2010	2.866	1.88
22/07/2010	-0.116	0.355
23/07/2010	1.124	-0.254
26/07/2010	0.23	1.851
27/07/2010	-0.803	-1.181
28/07/2010	0.308	-1.092
29/07/2010	-0.845	-0.333
30/07/2010	2.014	1.788
02/08/2010	-0.646	0.031
03/08/2010	-1.644	0.401
04/08/2010	-1.399	-0.487
05/08/2010	0.709	-0.615
06/08/2010	0.235	0.638
09/08/2010	-2.109	0.894
10/08/2010	-0.838	-3.554
11/08/2010	-1.488	0.64
12/08/2010	-0.368	-1.068
13/08/2010	0.41	-0.586
16/08/2010	0.857	1.749
17/08/2010	0.445	0.437
18/08/2010	-1.531	-1.261

Results including amount staked at each decision, endowment at time of decision (light blue), percentage staked (dark blue) and final amount (red); presented in landscape:

