Learning and loss aversion
Evidence from a Financial Betting Market

Tomás Ó Briain

A thesis submitted for the degree of

Doctor of Philosophy

University of Edinburgh

2015

Declaration

This is to certify that the work contained within has been composed by me. No part of this thesis has been submitted for any other degree or professional qualification.

The work presented in Chapters 5 and 6 forms the basis of co-authored working papers with Dr. Peter Moles and Dr. Ufuk Güçbilmez. Chapter 5 is based on a working paper entitled Reinforcement learning and Overconfidence. Chapter 6 is based on a working paper entitled Learning Theories and the Disposition Effect. The work presented in Chapter 7 is based on a working paper for which I am the single author.

Tomás Ó Briain

Dedication

Ar son mo ṁáṫair.

Ar ḋeis Dé go raiḃ a hanam dílis.

Acknowledgements

I would like to express my profound gratitude to my supervisors, Dr. Peter Moles and Dr. Ufuk Güçbilmez, for their support, guidance and Jobian patience.

Abstract

This research is motivated by a number of open questions in the behavioural finance literature. Firstly, if investors do not learn in a rational Bayesian manner but rather suffer from biases set out in the naïve reinforcement hypothesis, rationality assumptions in individual preference models may not hold. I use a unique longitudinal dataset comprising in excess of 1.5 million fixed-odds financial bets, where bettors perform identical, consecutive decisions which mimic financial choices made in a laboratory, but the use of their own funds departs from the artificiality of an experiment. I present evidence of unwarranted overconfidence generated by reinforcement learning in both real and simulated markets.

Secondly, Kahneman and Tversky (1979) state that losses loom larger than gains. I examine whether the disposition to avoid losses is driving behaviour in the losing domain in the dataset and conclude that there is little evidence of loss aversion. I differentiate between betting on Financial Markets, in which agents may perceive an internal locus of control, and betting on the simulated market, where results are uncorrelated and in which the emotions of regret and disappointment may not loom as large.

Finally, Odean (1998) provides evidence that investors readily realise paper gains by selling their winning stocks, yet hold on to their losing stocks too long. This loss aversion is consistent with Kahneman and Tversky (1979) prospect theory, however, how long would the investor hold on to a stock that is losing value on a day-to-day basis? Conversely, would an investor rush to sell a stock that has yielded positive returns in each month during the past year? I test the interaction between learning and loss aversion in a financial betting experiment in which two treatment groups are subjected to consecutive gains or losses.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Summary

This thesis is an empirical examination of several topics in the realm of behavioural finance and economic decision making. In short, it concerns decision making under risk. Its theoretical framework includes the Expected Utility Hypothesis, Subjective Utility Theory, Prospect Theory and Learning Theory and the research questions are motivated by empirical evidence in the fields of behavioural finance, behavioural economics, psychology and management science.

The methodological approach is primarily quantitative, however some qualitative results are discussed briefly in the final empirical chapter. Chapters 1 to 4 serve to introduce the motivation, data, institutional setting and relevant literature surrounding this research. The empirical body of the thesis is contained in Chapters 5 to 7. Chapter 8 concludes.

1.2 Motivation

Normative, descriptive and prescriptive theories of decision making under risk have undergone significant revisions in the literature. The expected value rule was the predominant descriptive theory of rational choice at one point, but did not explain behaviour in lotteries, gambling or insurance. Expected utility theory allowed for choices in the presence of risk seeking and risk aversion, and thereafter, the Bernouli (1738) utility concave function posited that agents maximize their utility and not expected value or wealth. Kahneman and Tversky (1979) proposed a utility curve with an antipodal shape in each of the domains of gains and losses, however empirical observations such as the House Money Effect and Break-Even effect (Thaler and Johnson, 1990) are at odds with this. Moreover, Barron and Erev (2003) argue that prospect theory may not be generalizable to settings with multiple decisions and feedback. Empirical work aimed at informing further refinement of models of rational choice is therefore warranted.

While we have good models for understanding how asset prices should be determined, it is still a real puzzle as to whether individuals actually behave in the way theory suggests. Economics assumes agents form rational expectations and update their prior beliefs as new information is revealed, with a view to converging on the equilibrium or true level of a variable. If the market consists only of rational agents, this assumption is valid. However, if market participants are biased or 'heuristic' learners, it may not be. In that sense, an understanding of the conditions under which investors can endeavour to 'learn' their way out of biases, or in which no such learning can take place, is crucial. The motivation behind this thesis is to try and shed light on this issue.

Learning is central to some of the most fundamental issues in finance and economics. For example, we also know that there is considerable heterogeneity among agents in financial markets with regard to skill. If both uninformed and informed agents exist, this begs the question of how the informed cohort have acquired their skill (Seru, Shumway, and Stoffman, 2010). Are agents endowed with an innate ability which can be discovered after a certain amount of practical experience? Or is it the case that ability can be improved with further experience? In that sense, the focus of this thesis is on the effect prior outcomes have on future behaviour.

1.3 Data

Thaler and Ziemba (1988) state that using stocks, which have no defined termination point or expiry date, makes testing for market efficiency and rationality difficult. They argue that although a stock's price is a function of discounted future dividends, it also depends on the price another investor may pay for it in the future. According to Thaler and Ziemba (1988), wagering markets concern contingent claims that both expire at a known point in the future and also give quick feedback, characteristics which facilitate expedient learning and give them a better chance of being efficient. Sauer (1998) suggests that betting markets are essentially simple financial markets, where the scope of the pricing problem is reduced, and as such, can offer clear insights into pricing issues which are much more involved in conventional financial markets.

Vaughan Williams (2005) notes that while betting markets share some characteristics with conventional financial markets such as a large number of heterogeneous participants, public access to a large data set and the potential for insider trading, they have unique features which make them attractive for empirical studies. He argues that the absence of infinitive expected future outcomes and a clear expiry point enable a clearer and more productive learning process in betting markets as opposed to stock markets.

Hirshleifer (1966) states that gambling is not a wealth-orientated activity, but rather a pleasure-orientated one. However Fong (2014) states that retail investors in financial markets not only invest in high-beta, low book-to-market ratio, volatile stocks, but also have a preference for lottery-type stocks and may trade actively to gamble. While bettors may derive utility from gambling, Barber and Odean (2000) and Grinblatt and Keloharju (2009) show that individual investors in financial markets may derive utility in the same way from financial markets.

The research questions in this thesis concern learning by individuals. Brav and Heaton (2002) note that learning in experiments requires immediate outcomes while Russell and Thaler (1985) state that without well-structured timely feedback, learning may be negligible. In that respect, the short duration between decision and feedback in betting markets makes them attractive as a setting in which to test hypotheses. I therefore examine the trading history of more than 30,000 individual financial bettors. With this dataset, we can observe how much each agent is betting, the evolution of successive wins and losses throughout time, and the individuals' subsequent behaviour in the faces of losses and wins. Agents will be observed making consecutive financial decisions and we will test hypotheses on what is driving behaviour, all motivated by prior theoretical work and empirical observations from experimental work and from the field.

I will also present laboratory evidence which supports the suggestion that reinforcement learning interacts with other salient biases such as loss aversion, manifested as the disposition effect. A trading interface was coded in z-Tree (Fischbacher, 2007) and subjects were presented with five rounds of betting and the choice to quit after the second and subsequent rounds. I suggest that reinforcement learning and Bayesian learning affect behaviour in a systematic manner and that the presence of strong reinforcement has a salient affect on financial biases.

1.4 Structure of thesis

Chapter 2 introduces the financial fixed odds betting market and provides an overview of the mechanism of the betting proposition offered by online bookmakers in the industry. Chapter 3 presents an introductory literature review with broad scope across the topic. Each of the three main empirical chapters contain a focused literature review targeting the motivation of the hypotheses being tested, however there is no such restriction in the initial literature review chapter.

Chapter 4 sets the scene for the empirical chapters which follow. From the offset, there is a description of the betting offering, how prices which are being offered are being determined, which financial assets underlie the betting propositions and who is making the market. There is an overview of the competitive environment for bookmakers and a brief discussion of relevant regulatory issues. As the title suggests, this chapter also provides an overview of the dataset under analysis, however, by defining the betting proposition and introducing the competitive environment, legal framework and regulatory setting in which the market exists, this chapter should provide the reader with an adequate overview with which to approach the subsequent empirical chapters.

Chapter 5 is the first of three empirical chapters and examines the effect learning has on individual behaviour. Using the literature, I motivate hypotheses on the learning behaviour of individual investors. In any environment in which there are variables to be predicted, individuals are expected to update their priors as more of the underlying, unobserved distribution is revealed. One of the key assumptions in asset pricing is that agents are rational and that changes in their behaviour is a result of rational Bayesian updating. However, we know that individuals are also subject to reinforcement learning, such that they are biased learners. This setting provides a unique test. I observe behaviour on real and simulated financial markets. I hypothesise that stake size changes on real financial markets are correlated with past returns. In turn, I hypothesise that behaviour in the simulated market is independent of previous outcomes. I find that strong positive reinforcement prompts bettors to increase their stake sizes and that such stake size increases decrease their wealth with each subsequent bet. Strong negative reinforcement has an effect on attrition, in that unsuccessful bettors exit the market. However, those that remain do not temper their stake sizes significantly. In addition, after controlling for the path of wins and losses, bettors have a tendency to increase the amount they wager with each subsequent bet. As such, learning manifests itself in exit from the market in the losing domain, rather than a rational Bayesian decrease in stake size.

In Chapter 6, having presented evidence on the saliency of reinforcement learning in the previous chapter, I expand on this to analyse whether there is an interaction between reinforcement learning and loss aversion. I motivate the analysis with a literature review which incorporates a wide range of topics from theoretical work on expected utility theory to empirical evidence on observations predicted by prospect theory. In effect, this chapter advances the results presented earlier by combining the analysis with predictions from prospect theory. I hypothesise that agents who have been subject to a series of consecutive winning or losing 'streaks', have the propensity to behave in a manner inconsistent with behaviour predicted by prospect theory, and in particular in the losing domain. The analysis presented in this natural experiment setting, using the same data as the first, motivates the final of the three main empirical chapters: the introduction of evidence to support the hypotheses using data from a focused and controlled laboratory experiment.

Chapter 7 introduces evidence from a bespoke financial betting experiment and synthesises the previous two chapters on learning and loss aversion. I outline the motivation and experimental design for the experiment which was administered at the University of Edinburgh. Subjects were given detailed instructions, shown a trading interface with a financial chart and their trading history, and encouraged to place bets on a simulated market. In five rounds of betting in total, they were given the choice whether to continue betting or book gains and cut losses. The experiment involved real money and the expected gains for participants were substantial at £25 per subject, with a maximum potential payout of £50 for each subject. I also gathered both demographic and behavioural metrics from subjects in a subsequent questionnaire. I hypothesise that agents exposed to a treatment of consecutive gains and losses would behave in a manner inconsistent with the loss aversion component of prospect theory, manifested by the disposition effect.

Chapter 8 summarises the thesis, provides a synopsis of the results and discusses areas for further research.

1.5 Summary of results

The overall results of the three main empirical chapters are as follows.

In Chapter 5, I show that individual investors exhibit behaviour consistent with Bayesian updating in financial markets, however behaviour consistent with naive reinforcement learning in simulated markets. They show that rates of attrition in financial markets are consistent with that predicted by the literature: agents are more likely to continue to trade in the face of losses in financial markets than they are in a simulated market, manifested by lower attrition rates in financial markets.

The second empirical project is outlined in Chapter 6. I first test whether behaviour in the winning and losing domains of the dataset is driven by predictions made in prospect theory. The disposition effect suggests that agents in the losing domain should exhibit risk seeking, while those in the positive domain should be risk averse. I find that behaviour differs according to market setting and risk profile. In aggregate, there is weak support for the existence of a disposition effect. However, when I isolate agents who have experienced successive wins and losses, there is evidence of behaviour consistent with prospect theory. I show that agents in the reinforcement learning cohort do not follow the predictions of the disposition effect. I further show that agents betting on financial markets exhibit a disposition effect, while those betting on the simulated market do not. These results motivate the final empirical chapter where I seek further support for the hypothesis in the form of a controlled test.

Chapter 7 outlines a targeted experimental design in which I test whether agents behave in accordance with predictions outlined in prospect theory. I find some support for the hypothesis that agents subject to consecutive wins and losses (analogous to stock investors having a series of positive or negative monthly returns) behave in a manner inconsistent with the disposition effect. In effect, agents may not have a higher propensity to sell 'winning' stocks than 'losing' stocks: in this case, it may be reinforcement learning which is prompting agents in the winning domain to hold on to their positions and those in the losing domain to cut their losses.

1.6 Conclusion

We will return to the results in further detail in Chapter 8. However, before proceeding, it may be instructive to define some of the concepts and betting market nomenclature which will be referred to in the literature review chapter. Thus, the next chapter provides a broad overview of the fixed-odds betting market and a definition of the exact nature of the betting proposition being offered by bookmakers in the industry.

Chapter 2

Fixed-odds betting market

2.1 Introduction

This chapter introduces the fixed-odds betting market and outlines some of the mechanics behind the betting propositions in this industry. A precise definition of the products being offered by bookmakers is presented along with accompanying payoff diagrams to explain the varying odds prices being offered for betting. The motivation for this chapter is to introduce some of the terminology which will make the literature review and subsequent chapters more accessible.

2.2 Key Concepts

A fixed-odds financial bet is analogous to a vanilla 'Higher/Lower' binary, 'cash-or-nothing', bet option or digital option with a specified strike and maturity, where the option pays a notional 100 if the option expires 'in-the-money' and zero if it does not. They are cash settled, but with a unique discontinuous payoff pattern: a graph of the payoff presents a jump with regard to the strike price. Figure 2.1 presents a payoff diagram for a vanilla 'higher' and 'lower' financial bet.

In option parlance, these are European-style options, as the underlying must be above the strike at maturity, regardless of whether it is hit beforehand (Sebehela, 2011). When traded as OTC products in traditional financial markets, such options are used to hedge against 'jump risk' and are relatively easy to trade as they require only a sense of the direction of price movement in the underlying rather than a view on both the direction and the magnitude of price movement.

This style of bet pays a fixed amount (determined by the stake wagered and the price-odds associated with the bet) on expiry if the level of the underlying is higher (or lower) than the strike price chosen, according to the direction of

Figure 2.1

Payoff & P&L Diagrams

This figures demonstrates the payoff to a financial fixed-odds bet. These figures were adapted from payoff diagrams presented in Raw (2008).

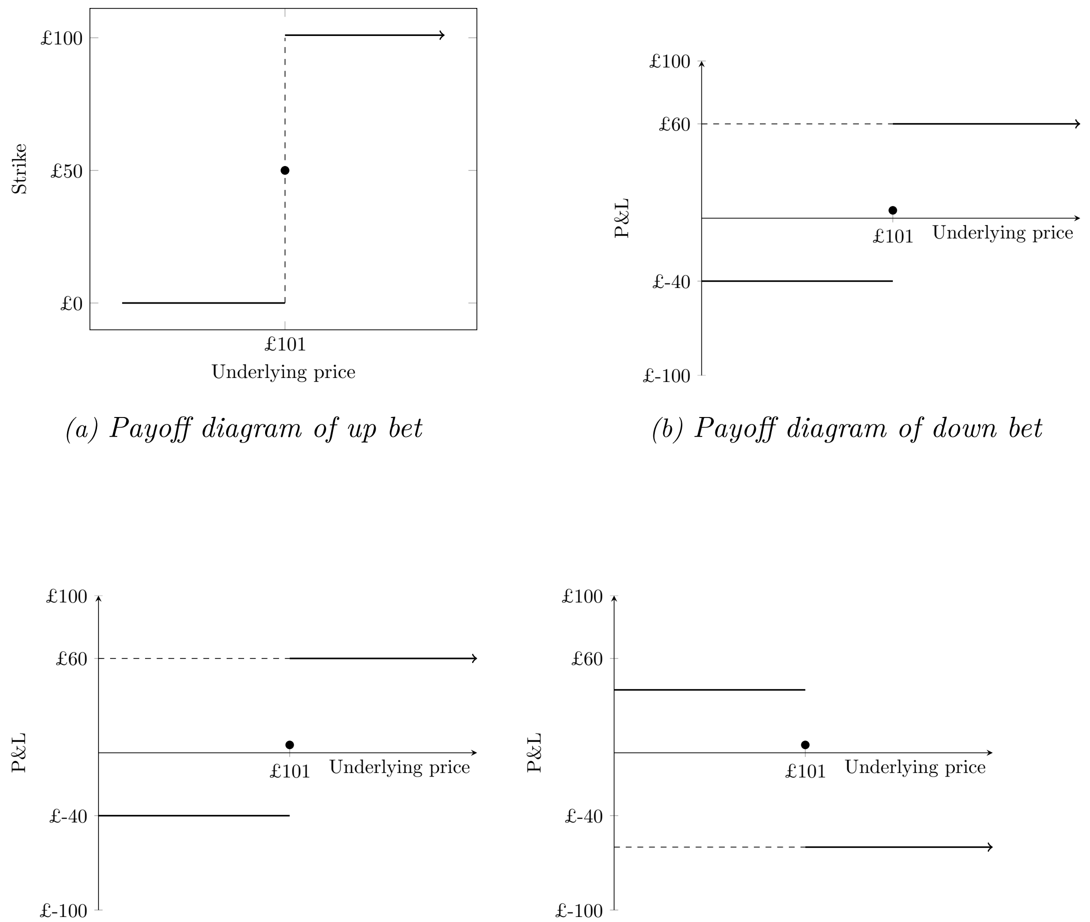

(a) Payoff diagram of up bet

(b) Payoff diagram of down bet

(c) P&L diagram of up bet

(d) P&L diagram of down bet

the trade. It is a predetermined payout regardless of how deep in-the-money the option is and it pays out only if it is at least at-the-money at maturity. Otherwise, it expires worthless regardless of how much time it may have been in-the-money during its life (RBC Capital Markets, 2001).

As both the risk and reward are pre-determined, they are widely use in the sports betting industry (Kotze and Joseph, 2009). A similar style of bet is quoted by a subset of bookmakers where price levels of the underlying are fixed, and the price-odds associated with each level are updated in real-time. However, that style of bet is not the focus of this study as I did not have access to such data.

The usual treatment of binary options in the recent literature refers to products traded in a continuous double-auction on betting exchanges. Since it is traders who are making prices (although some liquidity may be provided by bookmakers), trading directly affects prices as bets are initiated. In effect, there is a limit-order book of fixed-odds bets displayed and matched at the exchange. If an exchange-member initiates a trade, the trade is executed against resting orders in the order book. If an investors places a trade via a third-party, the third-party will usually execute a corresponding trade at the exchange for the same product and amount. This is not the case with fixed-odds financial bets as bookmakers are laying all bets. In effect, the bet is not hedged, nor is any corresponding order placed at the exchange. The customer is not paying a wider bid-ask spread to transact as they do not have a seat at the exchange, for example, they are paying the bookmaker's overround in order to transact a trade and are betting against the bookmaker alone. Thus, in contrast with financial markets, where investors need to 'beat the market', customers in betting markets must endeavour to 'beat the bookie'.

When making markets on sporting events, bookmakers manipulate prices to balance the amount wagered on each side of the bet, however this is frequently not the case with fixed-odds betting and bookmakers are effectively running stock market 'bucket shops'[1]. Willimas, Wood, and Parke (2012) state that while sports bookmakers offering wagers on financial markets is a relatively new development, such betting on financial markets with an entity external to financial exchanges is very similar to 'bucketeering'.

European bookmakers[2] such as Ladbrokes, William Hill, Paddy Power and Victor Chandler offer fixed-odds bets on a variety of financial products, with time horizons from expiries on a rolling one-minute basis to ad-hoc bets with expiries of longer than one year. Bets may be placed on stock indices such as the FTSE, DAX, CAC, Eurostoxx50, NIKKEI, Dow Jones Industrial Average (DJIA) and S&P 500, commodities such as WTI futures, Brent Crude futures and all major currency pairs.

2.3 Mechanics of financial betting

The data for this study comes a transactional database of customer bets at an online bookmaker. The company offers bets on sporting events such as soccer,

[1] A bucket shop traditionally refers to a business which deals in derivatives of financial assets. However, when a trade is placed by a customer, no subsequent transaction is committed to the exchange by the proprietor. The proprietor is laying customer bets, rather than placing their trades at the exchange as an intermediary. The market is described in detail by Lefevre (1923) in 'Reminiscences of a Stock Operator'.

[2] At the time of writing, this style of financial bet is regulated as a derivative in the U.S. and Australia, where such products are typically only offered by spread-betting companies and other regulated entities.

hockey, rugby, but also bets on political outcomes, novelty bets and any ad-hoc stories of public interest for which a bet could be designed. Part of the company's offering involves bets on financial markets.

Bets are offered on stock indices, commodities, currency pairs and a random-number generated market. At times, the company has also offered lottery-type bets on Financial markets, such as a bet on the final digit of the settlement price of the FTSE 1000, for example, or a bet on whether the price of a barrel of oil would above a certain level before the end of the year. The most popular bets, however, are standard fixed-odds bets on stock indices.

The company offers bets on sporting outcomes on a round-the-clock basis. It acquired a betting company in Australia in 2010, and as a result has a risk management function in operation throughout the night. There are dedicated IT teams which monitor the performance and uptime of the company's website and as a result of all the this, betting can take place 24 hours a day. In Europe, there are no sporting events until the late morning, and the company's retail outlets do not open until approximately 10am. Thereafter, there are continuous betting markets offered throughout the day until the late evening, unless there are exceptional international sporting events such as the World Cup or Olympic Games which have events late at night and over the weekend. The day with the most throughput on the website is Saturday when the most high-profile sporting events take place.

During the night, bets are offered on US sports and a dedicated team of both sport traders and risk management officers are in place in both the office in Europe and in Australia. Betting continues on US sports until the early morning CET time. There is a gap of one or two hours in the early morning in Europe when there is little or no betting activity.

As regards financial betting, the European exchanges open at 7am and there is an underlying price with which to calculate betting odds throughout the day until 4/5pm. Thereafter, the only exchanges open are the US exchanges such as the CBOT, CME, NYSE and NASDAQ, which trade until 9PM CET. There is a gap of two or three hours until the Asian exchanges open, typically at midnight or 1am CET. In terms of volume of betting, from midnight the focus is on betting on the NIKKEI. From 8am most of the bets are on the FTSE. In the late afternoon, betting switches to the US stock indices such as the DOW and S&P. There are no markets to bet on between 9pm and midnight. In addition, there can be no betting on financial markets when the exchange is closed, which precludes betting on public holidays, weekends and other market holiday.

To fill in the gaps on such days, the company developed a virtual market. Virtual market type events are popular in retail betting, with most betting companies offering bets on virtual horse-racing or virtual greyhound racing on a continuous basis throughout the day. Customers tend to bet on these events when they are present in one of the retail outlets and there is a break in sporting events for a number of minutes. The virtual financial market allows the company to offer financial-type bets all weekend, after the US markets have closed, on Christmas Day, New Year's Eve and on other market holidays. There are a significant number of bets in the dataset for this study transacted on Christmas Day and New Year's Eve.

The virtual financial market is open for betting 24 hours a day, 7 days a week. As regards the proposition, customers have an opportunity to bet on the level of the virtual index at any time and bets are settled once every minute or every second minute. In effect, a customer can predict the level of the market at any minute (on the minute) during the day, or every 'odd' minute during the day. In the financial markets, there are markets in which the customers bet on the level of an index, commodity or currency at the end of the day, or in some of the more ad-hoc markets on the price of a barrel of oil, the price of oil a number of months in the future.

On the company's trading interface, the customers see a price changing in real-time with the title 'Virtual Financial Market'. The level of the price could be anything from 10,000 to 50,000 for example, and the price changes every second, very similar to the price action one would expect from a stock index. Every minute, on the minute, a display is shown with a message such as 'Betting Open'. The price will then change in real-time and a chart with a time-series of previous price changes will be shown. During this period, customers can interact with the customer-facing interface and place a bet. Bets are generally in the form of betting that the price, at expiry time, will be higher or lower than a specific level. In a one-minute market, for example, each expiry time is one minute hence.

If a customer decides not to place a bet but merely to observe what has happened and try to learn how the market functions, they will see the following. Each minute, a message is displayed indicated that bets will be taken. The price of the index will then change in real-time for 50 seconds. With 10 seconds to go before the next minute, a message will be displayed saying 'No More Bets'. Thereafter, the price will continue to change in real-time, however the betting interface will be disabled so that the customer cannot place a bet. On the minute, the current level of the index is the settlement price. This is displayed at the bottom of the screen, with a message such as 'Settlement Price @ 12.08: 10000.67'. The 'Betting Open' message will then be displayed again, indicating that the customer may place a bet on the next expiry (12:09, for example) and the whole process continues again.

In the case of a market where the settlement price is struck only against the daily settlement price (closing price) of the actual underlying stock index or commodity, a customer will see the following. There will be no market displayed for the FTSE before 8am, for example. Shortly after 8am, a market will be visible in the main trading interface, indicating that bets may be placed on the FTSE. This market may be called 'FTSE 100: Daily Closing Level', for example. If the customer clicks into that market, they will see the level of the market offered for betting by the bookmaker change in real-time throughout the day, from approximately 8am until 4.30pm. There are also a number of other market levels displayed.

2.4 Odds prices and margin

Since this is a fixed-odds offering, the bookmaker offers bets above or below the current price at a price of 5/6. This indicates that if a customer clicks the button that says '5/6 Above' with the current level of the index, the customer is placing a bet that the FTSE will settle at a level above the current price, and vice versa. The bet will then be displayed in the customer's trading interface throughout the day until shortly after 4.30pm, at which time it will be settled.

In addition to betting that the settlement price of the index will be above or below the current price, the bookmaker offers a number of other options. The current price level, which updates in real-time in both the '5/6 Above' and '5/6 Below' buttons on the interface is essentially a strike price. As the odds are fixed and do not change dynamically to indicate changes in the underlying, the strike prices changes dynamically. In effect, the bookmaker is taking the view that the current price is the best indicator of the settlement price of the index.

If the underlying market level rises by ten ticks, the bookmaker will no longer be offering bets at the level ten ticks lower, but rather at the new strike price 10 ticks higher than the previous. In this way, the strike price changes while the odds remain fixed. As the bookmaker is taking the view that the current price is the best indicator of the price of the index in the future, they are effectively suggesting that there is a 50/50 chance of the price going up or down. In a fair bet, the outcome of this bet to both the customer and bookmaker would be zero. In this case, the fixed-odds offered should be even money, or 1/1, indicating that if the customer chooses 'Above', bets $10 and the price does ultimately go up, the customer would win another $10. If not, and the price subsequently decreases, the customer would lose his/her \$10. This situation describes a fair bet and is essentially a coin toss.

However then bookmaker only pays out 5/6 for each of the outcomes, which means that if the customer bet \$10, price goes up and the customer was correct in their prediction, the bookmaker would return the original \$10 to the customer and pay them \$8.30. If the customer was wrong, the bookmaker keeps the \$10. This overround means that it is not a fair bet for the customer and that the bookmaker will always collect margin of 8% on this bet. I set aside arguments about whether the bet is correctly priced at 5/6 above and below until later, as I will compare and contrast the situation where the bookmaker is pricing bets on financial markets where an underlying exists and also on the virtual market where the price is random.

Figure 2.2 shows a time series with the full range of bet prices outlined. In addition to the odds of 5/6, the bookmaker also offers three more bet prices to bet on: 1/1, 5/1 and 1/20. The strike for these bet prices is slightly further away from the current price than the 5/6 strike. For example, as we saw above, if the current price is 15000, the bookmaker will be offering bets of 'Above' and 'Below' at 5/6 with the strike prices of the bet set at 15000. However, in addition, the bookmaker offers a bet at 5/1 'Above', at a strike which is much higher than the current price, 15200, for example. If a customer bets on a 5/1 bet, they are expecting a significant jump in the level of the underlying market. In this case, for the bet to be profitable, the underlying price would have to rise by 200 ticks before the market expiry so that the settlement level would be higher than the strike.

A 1/20 bet is offered at the same strike as the 5/1 bet, but customers are betting than the level of the market will stay below this level rather than rise above it. The format of the bet is not a 'one touch', for example, and the effective price of the option may drift in and out of moneyness until expiry. The important thing for the customer is that the price must be below the strike at expiry: in the meantime, the price can drift higher than the strike but the bet will not be settled as a losing bet and the customer will not lose their stake. The last fixed-odds price is 1/1, and this is offered at a strike which is between the 5/6 and 5/1 strikes. The setup of this bet is very similar.

As shown earlier, some of the bet prices are symmetric, will others are quite skewed. Some of the betting propositions are a bet on volatility, offering customers the opportunity to buy volatility in the case of a 5/1 bet and sell volatility with a 1/20 bet for example. In practice, a bookmaker will not offer all betting outcomes at all times, and if there is a scheduled market-moving economic figure

Figure 2.2

Bet prices

This figure presents the range of bet prices offered by the bookmaker, assuming a current market price of 240.00. The strike prices for 5/1 Above and 1/20 Below and 5/1 Below and 1/20 Above bets are identical. In each figure the wide dashed line indicates the 'no more bets' period, the solid line to the right indicates the settlement/expiry time and the dashed horizontals lines indicate the strike prices for both bullish and bearish bets. The probabilities inferred from the bet prices indicate the probability (given a forecast of volatility) of the bet being 'in-the-money'.

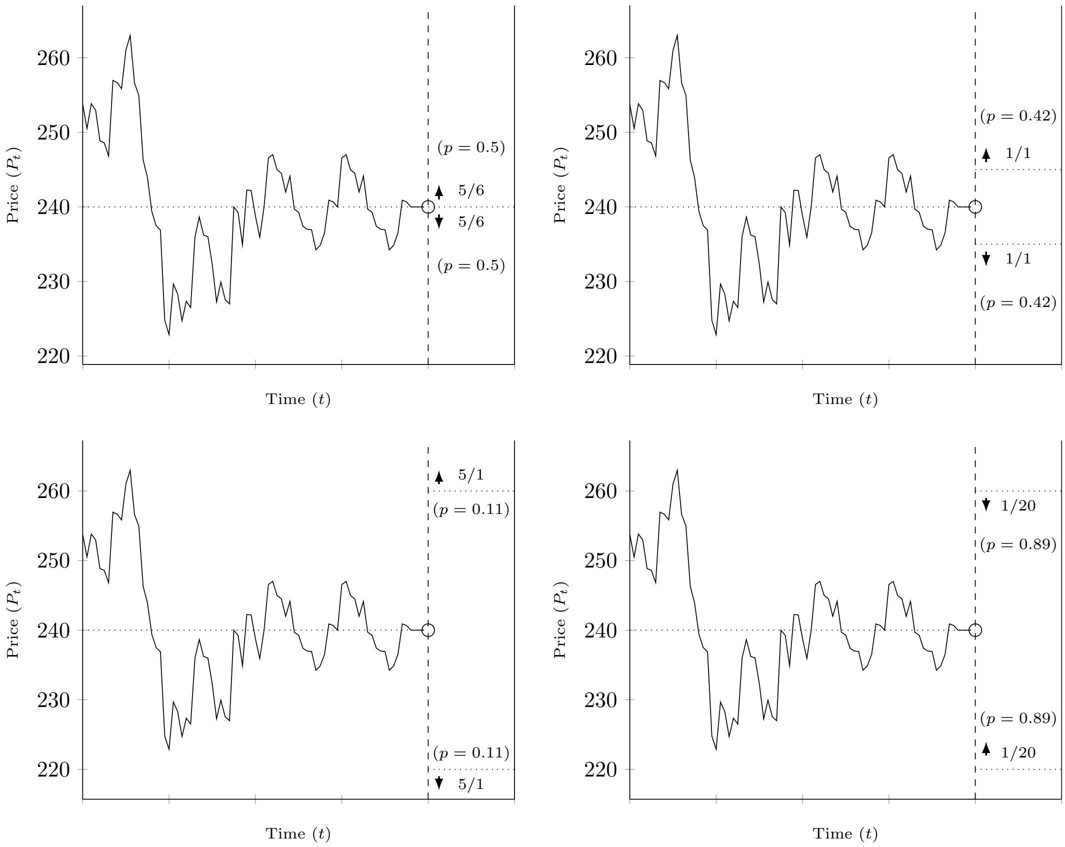

to be released, the 5/1 bets will be disabled in the betting interface. If these bets were available over scheduled announcements, customers could place a bet on both the 'Above' and 'Below' outcomes, and if they expected significant volatility, could profit at a rate of five times their stake on whichever direction the price went after the announcement which would offer profit even with the loss of the stake on the other side of this bet.

2.5 Information set

As regards relevant information on the betting interface, there are no analyst opinions or other fundamental information available to the customer as would be offered by a direct-access broker or spread-betting company, for example.

Some bookmakers have components to the interface which offer some indication of whether the bulk of customers are betting on 'Above' or 'Below' bets for a particular index, which some customers may interpret as the market being 'Bullish' or 'Bearish' for the day. Apart from current strike price, the specified fixed-odds associated with the strike, a simple chart showing a time series of past prices, a settlement price table displayed all previous settlement prices for the day and the name of the underlying product upon which the market is based, the customers do not have access to any further information. The amount of information presented to bettors varies from bookmaker to bookmaker. A screenshot of the trading interface for the four main players in this industry is presented in Figure 2.3.

This is in strict contrast to the information offered to customers on the sporting interface of the website, where for each horse or football team, for example, the customers can browse statistics on scores, times, history and past outcomes upon which to base an opinion. There is a component to the website which offers a 'Trading Academy', however the focus is mostly on the functionality of the trading interface as does not offer any opinion on market efficiency or microstructure. Moreover, bookmakers are generally not regulated by the SEC or FSA or any other financial regulatory, but rather by offshore gaming regulators such as the Isle of Man Gaming Commission. Bookmakers do not have permission, as a result, to offer financial advice to customers betting on the financial market component of the website. As regards their prior trading history, customers can view a full history of their previous bets, which provides variables such as the market bet on, the strike price, the associated fixed-odds, the size of the stake the customer placed, the outcome, and the size of the customer's win or loss.

2.6 Bets vs trading data

It is important to reiterate the difference between a financial bet and a financial trade. A financial bet is a single atomic transaction with a predefined expiry date. It is closer in definition to an option rather than an equity trade, for example, in that corresponding purchases and sales do not have to be identified in the data set i.e the profit and loss for each round-trip transaction is available in the settlement details of each bet. In effect, as the holding time is fixed, we know the investment horizon of each bet. No further matching of trades in a dataset is necessary as a result. As regards trading costs, although there is no bid/ask spread, however there is an implied spread in the overround calculated for each bet. For example, 'midpoint' or 'over/under' fixed-odds bets are priced at '5/6', rather than even money. Some of the similarities and differences between financial and financial

Figure 2.3

Information available to bettors

The following figures display screenshots from the trading interface of a number of bookmakers who offer bets on financial markets. Although there is considerable detail in terms of the specification of bets (i.e. payoff, expiry time, stake sizes etc), there is little in the way of news or analytics except for a time-series of very recent quoted prices.

(a) Ladbrokes

(b) William Hill

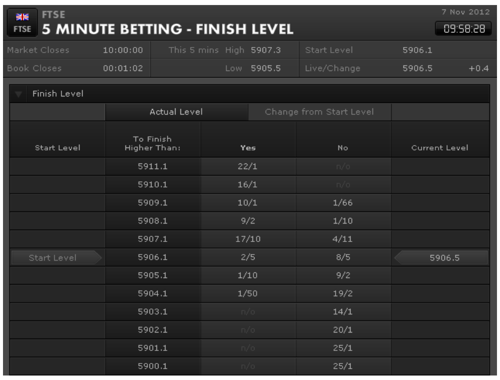

Start Level	To Finish Higher Than:	Yes	No	Current Level
	5911.1	22/1	n/o	
	5910.1	16/1	n/o	
	5909.1	10/1	1/66	
	5908.1	9/2	1/10	
	5907.1	17/10	4/11	
Start Level	5906.1	2/5	8/5	5906.5
	5905.1	1/10	9/2	
	5904.1	1/50	19/2	
	5903.1	n/o	14/1	
	5902.1	n/o	20/1	
	5901.1	n/o	25/1	
	5900.1	n/o	25/1	

(c) Bet365

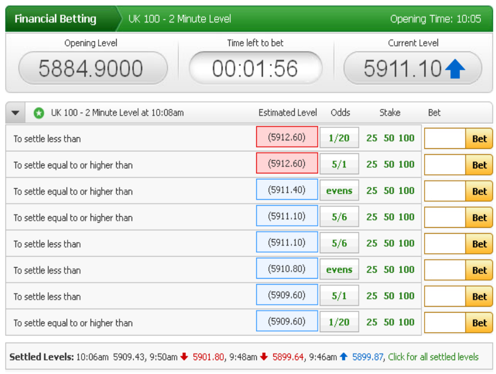

	Estimated Level	Odds	Stake	Bet
To settle less than	(5912.60)	1/20	25 50 100	Bet
To settle equal to or higher than	(5912.60)	5/1	25 50 100	Bet
To settle equal to or higher than	(5911.40)	evens	25 50 100	Bet
To settle equal to or higher than	(5911.10)	5/6	25 50 100	Bet
To settle less than	(5911.10)	5/6	25 50 100	Bet
To settle less than	(5910.80)	evens	25 50 100	Bet
To settle less than	(5909.60)	5/1	25 50 100	Bet
To settle equal to or higher than	(5909.60)	1/20	25 50 100	Bet

(d) Paddy Power

Table 2.1

Fixed-odds financial betting markets vs Financial markets

This table compares and contrasts FFO (Financial fixed-odds) betting markets with financial markets, focusing on issues such as valuation, transaction costs, trading venues, taxation and regulation.

Topic	Fixed-Odds Financial Betting Markets	Financial Markets
Characteristics	Derivative of a financial market underlying.	Financial product or derivative of financial product.
Valuation	Specified holding period. Expected cash-flows known ex-ante. In some case, subjective probabilities for some outcomes may be inferred.	Unknown for common stocks. Specified ex-ante for bonds and short-term interest rates. Some derivatives such as digital, binary or cash-or-nothing options may have characteristics similar to betting products.
Regulation	Unregulated. Regulated by offshore regulators such as the IOM Gambling Supervision Commission. Customers have recourse to IBAS (The Independent Betting Adjudication Service), but membership is voluntary and decisions are not binding.	Regulated by the FSA, CFC or SES (or equivalent).
Taxation	No income or capital gains taxes on betting or spread-betting winnings. Spread betting generally refers to retail futures trading with a broker.	Capital Gains or Income Tax applied but varies from jurisdiction to jurisdiction.
Trading location	In bookmakers' retail outlets, via online trading platforms on websites, over the phone or through an interface which allows trading on mobile (WAP etc), tablet and other devices.	Directly by exchange-members on open-outcry & electronic exchanges. Indirectly via brokers, retail trading platforms, financial intermediaries and other counterparties for a fee.
Slippage	Bet prices are marginated to include the bookmakers overround/vig/take. The subjective probability of an outcome is the marginated bet price minus the overround.	The bid/ask spread if traded directly at the exchange. If traded via an intermediary, either a wider bid/ask spread, management fee, variable transaction fee or other.

betting markets are outlined in Table 2.1.

There are similarities between financial trading and betting markets. Betting market propositions have a well-defined start and expiry time, potential payoffs known in advance, relatively short-duration between action and feedback, market-makers in the form of 'bookies', market analysts in the form of 'tipsters', centralised exchanges and low-barriers to entry for individuals. Thaler and Ziemba (1988) argue that betting markets are better suited to testing for market efficiency and rationality than stock markets, as the quick, repeated feedback bettor receive lends itself to expedient learning.

Levitt (2004) states that both financial and betting markets share characteristics of a zero-sum game where agents with heterogeneous beliefs trade on uncertain outcomes which are resolved over time. Betting market prices react quickly to new information and are reliable forecasts of the true probabilities of outcomes in aggregate (Wolfers and Zitzewitz, 2006b). There is the possibility of insider trading, the existence of a subset of informed agents, and opportunities for arbitrage, all of which exist in traditional markets, however the assets being traded in betting markets are less complex products, and hence more trivial to price (Vaughan Williams, 2005).

For the purposes of this research, the well-defined, regular feedback received by agents in betting markets is an advantage, as is the absence of uncertain future cash flows and dividends for the purposes of asset pricing. The betting propositions concern simple economic propositions as both implied probabilities and expiry times are known in advance, and full-information on actual and counterfactual outcomes is given as feedback to agents as they trade.

2.7 Gambling as entertainment

While financial betting markets are offered by traditional bookmakers, it is not trivial to determine whether such products are pure gambles similar to sporting outcomes and casino games, or whether they are analogous to ultra-short-term, high-frequency financial derivatives. Certainly, as regards utility from gambling, there is a distinction between the utility bettors may derive from participating in an evening at a luxurious casino, as opposed to interacting with the trading interface of a bookmaker on a laptop at home.

There is evidence that some individuals participate in the stock market for entertainment or treat it as gambling. Fong (2014) considers portfolios with high-beta stocks, extremely volatile stocks, stocks with low book-to-market ratios and stocks with high recent maximum daily returns as essentially gambling accounts consisting of lottery-type stocks. An important issue to address at this point is whether participation in financial betting markets constitutes investment, speculation or gambling. Increasingly, the lines between investment, speculation and gambling have been blurred, especially with the proliferation of derivative markets on diverse underlying products, and the increased participation in such markets by individual investors.

Bay, Sjödin, and McGoun (2011) state that there is a consensus on the definitions of investment and gambling: investment is a productive activity which has benefits for the economy; gambling is undertaken for entertainment. However, they acknowledge that the definition of speculation poses a problem, in that bookmakers have confounded the distinction between financial and non-financial betting by offering betting propositions on everything from derivatives on financial indices to bets on currency pairs, from wagers on sporting outcomes to bets on political and monetary policy outcomes.

Willimas et al. (2012) argue that there has always been a similarity between gambling and investment or speculation in financial markets and that bookmakers offering betting on financial indices, currencies and commodities have reduced the distinction between the two. They also state that speculation on financial markets meets the common definition of gambling, however the negative-expected value of gambling differentiates it from traditional investment, which in most cases does not have negative-expected value. Nevertheless, they also note that 'day-trading', short-selling of stocks and betting on the movements of financial products with entities external to the exchanges has challenged this distinction.

Derivative markets may have a societal benefit in terms of the transfer of risk from one market participant to another, however many participants transact such securities in order to profit from the hedging of risks they are not in fact exposed to at all. In this respect, Bay et al. (2011) argue that while hedging is proposed as the main function of derivative markets, it is merely an alibi, masking the absence of a distinction between gambling and hedging.

While addressing whether or not financial betting is investment or speculation, we must also consider the opposite: whether trading itself constitutes gambling. In fact, there is ample empirical evidence that, for some market participants at least, this is indeed the case. Anderson (2008) analysed a detailed dataset from a Swedish online broker and found strong evidence of gambling behaviour. They also find the same characteristics among high-volume traders as compulsive gamblers. They are more likely to be younger, male, less wealthy and have a lower level of education. The refute the hypothesis that the lowest staking investors are 'gambling with peanuts': those who trade the highest stakes are those with the highest turnover in their portfolios.

Grinblatt and Keloharju (2009) offer more empirical evidence on the link between utility derived from gambling and utility derived from trading. They analyse overconfidence and sensation-seeking in a dataset consisting of all household investors in Finland from 1995 through 2002. They conclude that agents who had received a large number of speeding tickets (their proxy for sensation-seeking) traded more frequently. Dorn and Sengmueller (2009) find that investors who reported enjoying trading or gambling had twice the rate of turnover in their portfolios as other investors. They suggest that such non-monetary 'entertainment' received from investing may explain some of the excess-trading puzzle.

2.8 Conclusion

Prior to embarking on a review of the relevant literature in the next chapter, a brief note on the specification of bets in financial fixed-odds markets and some initial background was required.

Having introduced the concept of financial fixed-odds betting, defined the nature of the betting proposition being offered and given some indication of the bettors' information set, I now move on to a review of relevant literature in the next chapter.

Chapter 3

Literature Review

3.1 Introduction

This chapter provides a broad overview of relevant literature to this research. Each of the subsequent results chapters contains a focused literature review with direct relevance to the hypotheses being tested. The present chapter, however, has no such restriction and this broader scope allows the introduction of evidence from a wide range of topics in behavioural finance.

3.2 Betting Markets

Prices in gambling markets are determined either by the wagering of participants in a parimutuel call-auction setting, by investors via a betting exchange, or by bookmakers in a continuous auction. The availability of real-time data from various online bookmakers and betting exchanges offer a rich ground for investigating issues such as market microstructure, asset pricing and behavioural finance. While the corollary between sports gambling, prediction betting and traditional markets has been made frequently in the literature, the opinion on whether conclusions drawn from one field may generalise directly to the other is divided. Thaler and Ziemba (1988) and Vaughan-Williams (1999) state that wagering markets are more suited for studies of market efficiency than financial markets, whereas Tetlock (2004) argues that the characteristics of both these markets and their participants do not conform to the theoretical assumptions underlying market efficiency.

Studies of gambling markets have been concerned with four styles of markets: parimutuel markets, prediction markets, fixed-odds betting markets and spread-betting. Evidence of inefficiencies in parimutuel markets have frequently been studied[1], where the price-odds for bets in these markets are not determined entirely by the bookmaker but are directly affected by the relative weight of wagering on either side of the bet. Prediction markets[2] have also been the subject of some discourse with contradictory arguments on whether their prices can be interpreted as the probability of outcomes[3]. The rationality of investors and tests of market efficiency have been undertaken in fixed-odds markets in racetrack and sports betting (Egon Franck and Nuesch (2010), Marshall (2009) and Hausch and William (1995)), with index or spread-betting in sports has also been studied (Gilbert W. Bassett (1981)).

Durham, Hertzel, and Martin (2005) state that betting markets share many characteristics of traditional financial markets such as large volume, liquidity and many sources of publicly available information. They argue that betting markets constitute a 'halfway-house' between the laboratory and financial markets. In comparison with a laboratory setting, they feature large volumes of investors actual wealth being wagered. In contrast with financial markets, there is a well-defined terminal value associated with each bet. They also mention three other characteristics of financial markets which are present in betting markets: the presence of experts marketing analysis, bookmakers who act as a substitute for market-makers and large numbers of arbitrageurs ready to take advantage of short-term price discrepancies.

Avery and Chevalier (1999) state that betting markets over advantages over financial markets in that there are regular settlement intervals with which to compare the evolution of valuations before expiry with fundamental value. Regular settlement also gives participants in betting markets the opportunity to evaluate the relative success of their transactions, which is also absent from financial markets. Avery and Chevalier (1999) compares the existence of stylised facts such as the hot-hand bias, gamblers fallacy and home-team bias with their analogues in financial markets: momentum, belief in mean-reversion and the equity home-bias puzzle.

Levitt (2004) argues that in both betting and financial markets, agents with heterogeneous beliefs and access to varying sources of information place large amounts of their own wealth at stake and that uncertainty surrounding valuation is resolved over time. He does however point out some crucial differences. While

[1]See Lange and Economides (2005), Thaler and Ziemba (1988), Qiu (2007) and Gjerde (1994).

[2]Prediction markets are generally contingent on a binary outcome related to politics, sporting events or project success in large companies. The fact the many of them are exchange-traded and feature large volumes lends them to be used as estimates of the subjective probability of an outcome.

[3]See Wolfers and Zitzewitz (2004), Manski (2004) and Wolfers and Zitzewitz (2006a).

market markets at financial exchanges match buyers and sellers, bookmakers maintain large outright positions with respect to certain events[1]. While casino games offer a clear edge to bookmakers, this is not necessarily the case with sporting events where it is possible for the bookmaker to generate losses in the long-term. In terms of microstructure, there are also less price updates in betting markets, especially the greater the duration to expiry.

Despite the proliferation of empirical studies on diverse gambling markets, there is a paucity of empirical evidence on the market efficiency or behaviour of individuals in fixed-odds betting or handicap markets. The focus in this chapter is on motivating a rationale for using financial betting markets as opposed to an alternate market setting to test the hypotheses presented later on in the thesis. Each empirical chapter contains a focused literature review and motivation, where issues of particular relevance to the chapter are outlined and gaps in the literature are highlighted. As a result, this particular chapter is intended as an exploratory overview of literature on this subject, without restricting the analysis to topics in later chapters. While the papers presented in each results chapter have their origins in the behavioural finance literature for the most part, there is no such constraint on this section. As a result, this chapter should serve to give a wide ranging introduction to literature on financial betting in a general sense.

The follow sections are structured around three main topics. Firstly, there has been considerable debate about the efficiency of betting and financial markets. In particular, the debate revolves around different trading mechanisms such as at betting exchanges or with bookmakers. This topic arises quite naturally from the the study of any nascent financial market as market fragmentation takes places.

Secondly, the literature on informed trading is particularly relevant to fixed-odds wagering markets, as the bookmaker is confronted by particular issues of adverse selection. Madhavan (2002) states that given the impossibility of identifying informed investors in a market-making setting, prices adjust in the direction of money flow. However, as the bookmaker acts in a market-making capacity and has full information on individual traders' identities, the adverse selection problem may not be as severe as that encountered by market-makers dealing with unidentified traders. Therefore, this necessitates a different treatment than that given to traditional financial markets. Finally, the issue of pricing, and whether bookmakers can efficiently process and utilise the set of current information, as to consistently outperform both the large body of uninformed traders and, in aggregate, the subset of informed traders.

[1]Levitt (2004) observes two-thirds of bets falling on only one side of sporting outcomes in approximately half of his sample, contrary to the premise that bookmakers adjust the 'line' to attract an equal amount of wagers on either side of a bet.

In the following three sections, the literature concerning financial markets, prediction markets and wagering markets are covered. Of particular interest are studies involving bookmaker prices, those undertaken using the prices of binary options from continuous double-auction markets and also longitudinal studies of customer behaviour. The main issues surrounding tests of market efficiency, asset pricing and trader performance are assessed. A number of potential methodologies for the present study are noted and the dataset under analysis is compared and contrasted with those used in the literature on wagering markets. It is shown that, while there is limited high-frequency studies concerning betting markets in general, there is almost no examination of the topic of financial fixed-odds betting.

3.3 Market Efficiency

Fixed-odds financial markets exhibit characteristics of traditional financial products. They are a derivative of an underlying product, the current value of which (to professional market agents at least) is in the public domain at all times. In contrast with financial securities, they have a defined expiry time, at which point the value of the bet converges to the underlying value[1]. These similarities between sports gambling, prediction betting and traditional markets have been noted in the literature, with some authors stating that wagering markets are more suited for studies of market efficiency than financial markets[2].

Egon Franck and Nuesch (2010) compare the prices of eight bookmakers with those of the betting exchange Betfair for 5478 European football league matches, and document the superiority of the predictive power of the latter. They run a univariate probit regression on the outcome of an event with the probabilities of the different outcomes from each source and also develop a simple betting rule capable of generating abnormal profits based on their results. A possible weakness in the present study is that the price dataset from the bookmakers was recorded from an odds comparison website and not directly from the source of the data. Furthermore, the data was collected as a snapshot on the day prior to events taking place. Any effect of short-term trading by informed traders is omitted from the study as a result.

Assuming that market microstructural effects on the price of the underlying are controlled for, in markets where causality flows from price to trade only (trading by market participants does not affect prices), the source of inefficien-

[1]Just before expiry, the probability that the bet will pay $Stake + (Stake \; * \; OddsPrice)$ approaches unity.

[2]See Thaler and Ziemba (1988) and Vaughan-Williams (1999).

cies lies with the market maker (the bookmaker). It is the price input into the bookmaker's pricing model that is affected by latency, or the bookmaker's pricing model is simply not an adequate predictor of the price of the underlying at expiry. In any event, since the price is being re-quoted on a continuous basis, the more robust way to test for market efficiency is to test using real-time data.

To this end, there is a paucity of high-frequency market efficiency tests in the wagering market literature. A notable except is Croxson and Reade (2007), who analyse efficiency in Betfair soccer prices with particular regard to goal arrival using second-by-second data. This is a particular case where the study of betting markets offer a more robust treatment than financial markets. While it can be argued that at least some market participants are privy to private information before it is released into the public domain, the same is not the case for market-moving events such as goals during sporting games. Croxson and Reade (2007) find these markets to be semi-strong efficient.

Although traded in a different market structure, person-to-person binary option markets share some characteristics with fixed-odds financial bets. Tetlock (2004) uses data on both sports and financial binary option markets collected on a 30 minute frequency from the Tradesports.com website, performing non-parametric tests of market efficiency, and finds that financial betting contracts exhibit less mispricing than sports markets. He suggests that conclusions drawn from wagering markets do not generalise to financial markets. However, since data was polled at a 30 minute frequency, there may be issues of data quality in the study, whereas a more robust approach would be to revisit this study using tick data.

Vlastakis, Dotsis and Markellos (2006) examine the predictability of football match outcomes using the closing odds[1] quoted by six bookmakers and find evidence of significant and long-lived mispricing. Tetlock (2008) revisits the issue, with an analysis of whether liquidity improves the predictive power of binary option prices, and finds that liquidity does not mitigate the mispricing of these contracts on the Tradesports exchange. As the data is collected at 30 minute intervals, the same questions concerning data quality arise.

Some of the literature has examined efficiency with regard to arbitrage opportunities available to informed traders by trading betting exchanges against individual bookmakers. Marshall (2009) examined how quickly arbitrage opportunities in sports betting markets converge to efficient levels using aggregate data from a number of online bookmakers. The analysis involves a dataset of prices rather than bets. His paper suggests studying wagering markets as a resolution

[1] Closing odds are generally quoted when a 'no more bets' period has been reached.

to an outstanding weakness in tests of market efficiency using financial markets: that of constraints on short-selling. A regression model using the duration in minutes as a dependent variable and variables to proxy for profit, event outcome, regulatory environment, public or private company status, and sport is proposed. Standard errors in this model are adjusted using White (1980). The model is subjected to robustness tests with control variables for bookmaker credit risk and membership of the Independent Betting Adjudication Service (IBAS). While skewness in the dependent variable is mentioned, this is examined by way of taking the natural log of the variable before re-estimating the regression. In contrast to Tetlock, Marshall (2009) finds evidence of long-lived mispricing with the duration of mispricing rising with the complexity of the bet involved.

As mentioned by Marshall (2009), the short-selling issue is a significant one in the literature, as agents who attempt to exploit arbitrage opportunities may be forced to 'pay up' and close their positions at a loss by noise traders who move prices even further from efficient price levels. This is not a reason for concern in wagering markets for two reasons. Firstly, there are no margin calls, as the traders liability is limited to the stake of the bet. Secondly, since financial fixed-odds markets are generally in a continuous auction setting rather than a double-auction, noise traders cannot directly affect the prices of bets. A potential weakness in Marshall's paper is that his arguments are based on the premise that arbitrage opportunities are removed by the activities of informed traders and the fact that such arbitrage opportunities may exist in spite of trading by such agents is only mentioned in passing.

In reality, the fact the bookmakers frequently hold prices which are at inefficient levels for marketing purposes may have a more significant affect on the results than acknowledged. It is of great value to the bookmaker to be seen as offering the best price of all competing bookmakers for a particular outcome on odds comparison sites such as Oddschecker.com. There may also be issues of data quality in this study, as although the median duration of arbitrage opportunities is reported as 15 minutes, there may be latency associated with prices which are not captured at source directly from the bookmaker.

In summary, a considerable body of the empirical studies involving financial market data has focused on price-related components of the microstructure such as the bid-ask spread, liquidity, traded volume and actual prices. While informed trading and trading times has been analysed by Menkhoff and Schmeling (2010), for example, no such study has been undertaken with prediction or wagering markets. Since the dataset in the present study contains a large database of bookmaker bets, as in Levitt (2004), we are in a position to examine this issue in greater detail. In particular, an analysis of whether informed traders choose to trade at certain pattern times before expiry, may shed light on their activities. For example, an initial investigation of the dataset for this study has shown that the bookmaker does not attract an equal wagered amount on both sides. Moreover, there exists a small group of individual or collusionary customers who can earn abnormal profits from the bookmaker[1].

3.4 Asset Pricing & Forecast Accuracy

The literature on bookmaker pricing has focused, for the most part, on the UK racing and sports market. Evidence of pricing anomalies such as the favourite-longshot bias and the gamblers 'hot-hand' fallacy are frequently addressed issues (see Woodland and Woodland (1994), Hodges, Tompkins, and Ziemba (2003) and Smith and Vaughan-Williams (2010)). While the price-odds of fixed-odds financial bets offered by bookmakers are static, the strike or 'handicap' levels are requoted, for the most part, in real-time. At any point before maturity of a bet, the combination of price-odds and handicap level represent the bookmaker's forecast of the future path of the underlying. For example, the 'Over/Under 5/6' bet is the bookmakers expectation of the current price of the underlying at expiry, adjusted for the overround. The price-odds and handicap level of the '5/1' and '1/20' bets represent the bookmakers estimation of short-term volatility.[2] In effect, a binary, digital or 'cash-or-nothing' option is being offered with a strike equivalent to the handicap level and a payoff equal to the stake times the price-odds if the handicap level is breached at expiry, and zero if it is not. Prices are quoted in the form of dynamic price odds for a given outcome, or static odds for a dynamic outcome. Since such prices can be interpreted as probabilities, the forecast accuracy of price-odds is a frequently addressed issue in the literature with regard to sporting and prediction market events.

Wolfers and Zitzewitz (2004) describe the market design of prediction markets and estimate the consensus forecasts of economic indicators. They describe how a series of contracts on certain prediction market can constitute a probability distribution of market expectations of their future outcomes. They compare the standard errors of forecasts with the standard errors of reported economic

[1]Chapter 4 presents some statistics on the betting behaviour of the top (and bottom) decile customers, who are analysed according to their 'skill flag' i.e. 0 for the worst customers, 2 for the most skilled and 0 for the rest.

[2]I present an overview of the betting proposition offered by the data provider for this thesis in the next chapter, and will return to the various odds priced being quoted in real-time by the bookmaker.

indicators, using the Census Bureau's estimation of the standard error associated with each individual statistic, and show that traders of these contracts are over-estimating their values. The interpretation of a series of bets on either side of a point estimate for future values of a financial fixed-odds bet as a probability distribution of estimates, may yield interesting results.

Wolfers and Zitzewitz (2006a) examine the efficiency of prediction markets on political events, sporting outcomes and synthetic binary equity index options and conclude that prediction market prices approximate mean beliefs in the outcome of events. Their paper is a direct response to Manski (2004), who suggested that there was scant evidence that prediction market prices were good estimators of future probabilities. They propose an equilibrium model where demand from traders is motivated by any deviation of prices from aggregate beliefs and where beliefs are heterogeneous. The model is then developed to account for changes in beliefs, where any significant bias in prices is offset by noise-traders searching for 'action' or arbitrageurs taking advantage of price inefficiencies. This development of the theory, especially in the approach the authors take as regards the categorisation of informed and noise traders, has specific applications to fixed-odds betting markets.

Of particular relevance to this study are bookmakers estimates of short-term volatility and the previous literature that has examined the pricing of binary options on financial products at betting exchanges. Using data on every trade on the Dow Jones Industrial Average (DJIA) via Tradesports from June 2003 to August 2005, Zitzewitz (2006) examines the pricing of binary option prices, adjusting standard errors for heteroskedasticity and return correlations on expiry, and finds that the implied volatility of binary options on both the intraday and closing values of the index can be used to predict realised short-term volatility. Average returns to expiry are measured, and market efficiency is tested by examining whether price predicts returns. To test whether investors were overstating volatility, a probit regression of binary option payouts on a factor adjusting implied volatility by investors estimation of volatility is constructed. Only a subset of option expiries showed significant results for this test, however he finds evidence of a favourite-longshot bias (consistent with traders overstating volatility) for these expiries and shows that the prices of these binary options add predictive power to a model of high frequency volatility.

Using the Zitzewitz (2006) approach, backing out the implied volatility from bookmakers quoted fixed-odds prices may be possible, if the model of volatility does not change from one period to the next. For example, with 28800 seconds in a daily trading session for the FTSE 100 Index, and all prices for the underlying and front-month future price in the public domain, assuming no exogenous inputs to a bookmaker's pricing model, the dataset a bookmaker uses to price markets is available. It should be possible, with some trial and error, to develop a model of volatility to exactly mirror every price change of a bookmaker publishing a range of price levels and associated price-odds during the trading session. However, we must leave such analysis to further research, as the dataset being provided for this study does not contain prices, by rather is a database of customer trades.

Egon Franck and Nuesch (2010) examine the prediction accuracy of bookmaker prices with that of the betting exchange Betfair for the same markets. They propose a univariate probit model to compare implied probabilities from eight bookmakers and the betting exchange to the actual outcome of the event, and use four goodness-of-fit measures to analyse forecast accuracy. One of the goodness-of-fit measures used in the Brier Score[1](which measures the mean squared difference between the event outcome and forecast outcomes) and McKelvey and Zovoina's R2[2], which may be applicable to fixed-odds financial price quotes if adapted to account for the range of price-odds associated with each betting expiry. The regression is then re-run with a ratio measure of the difference in probability forecast by each bookmaker and that of the betting exchange. The inclusion of this variable improves the forecasting accuracy of the bookmakers' prices alone, and indicates that there is some extra informational component in the Betfair implied probabilities.

In summary, the literature has heretofore focused on prediction markets, the quoted price-odds of bookmakers for sporting events and binary option pricing at betting exchanges. With the growing popularity of person-to-person betting exchanges, market design is also a frequently addressed issue. The pricing of financial products by bookmakers in a continuous auction is not addressed, nor is financial fixed-odds betting in particular. While asset pricing is examined with reference to prices only, there is valuable information in betting by arbitrageurs and informed traders that has not yet been researched.

3.5 Informed Trading

According to Glosten and Milgrom (1985), market-makers can mitigate issues of adverse selection via the bid-ask spread. Bookmakers who operate in a pari-mutuel setting where there is no bid-ask spread, or those who strategically change their quoted odds to attract equal wagered amounts on both sides of the mar-

[1]See Brier (1950).
[2]See Kelvey and Zovoina (1975).

ket, earn the 'overround' as compensation for this. A strand of the literature models the behaviour of individual investors and notes characteristics of investor irrationality. Barber and Odean (2000) show that trading of individual investors exhibits characteristics such as overconfidence and overtrading which lead to net losses. For the bookmaker to continue as a going concern, the number of uninformed traders must be sufficiently large for their loses to offset the winnings of informed traders. Moreover, bookmakers move to identify informed trading and either limit the size of stakes accepted from such traders, using their trades as information, or close accounts completely.

This is done by limiting the amount a particular customer can wager on a specific event, or at an aggregate level from event to event. In fact, bookmakers impose different levels of risk limits accordingly to the informational asymmetry in the market and will not lay economically significant bets on markets in which they have no information advantage. However, if the bookmaker has employed experts who are skilled at consistently pricing those markets correctly and monitoring the betting patterns of informed traders, more liberal risk levels will be set. As such, the question of whose bets provide information is paramount to the bookmaker.

Previous studies involving informed trading have focused on market microstructure issues such as the bid-ask spread, using data on prices rather than individual trades. Marshall (2009) investigates the activities of sports arbitrageurs using a dataset of prices from a company that identifies arbitrage opportunities. While he finds evidence of short-run mispricing and draws the conclusion that it is the activities of arbitrageurs that is re-aligning prices with market values, it may be insightful to investigate the actual trading activities of these agents, rather than drawing conclusions based on the price changes of bookmakers' odds.

Using a dataset covering six years of individual investor trading at a Dutch discount broker, Rob Bauer and Eichholtz (2008) identify a subset a investors whose return performance is consistently better than the rest. They rank traders into decile portfolios during an initial selected period and calculate decile returns during a second performance evaluation period, readjusting weights to control for survivorship bias. A t-test on the performance difference between the lowest and the highest return portfolios indicates persistent outperformance by the highest decile.

Menkhoff and Schmeling (2010) perform a multivariate cross-sectional analysis of informed and uninformed trader characteristics using data from the foreign exchange markets and show that informed traders use medium-sized orders (consistent with the concept of 'stealth-trading'), trade large volumes and trade early in the session. Their methodology examines the price-impact of trades, but since trading does not affect the fixed-odds prices of bookmakers in the present study, a different approach is needed.

Levitt (2004) analyses a dataset of bookmaker prices and customer bets during an NFL season and concludes that bookmakers change their quoted odds infrequently and do not attempt to balance the amount wagered on each side of the market. He attributes this to the bookmaker's superior skills in setting prices and the fact that bookmakers are aware of the bounds at which informed traders will enter the market if prices are distorted to attract betting and account for this. This study of bookmaker and customer activity is insightful, but as pointed out by Levitt (2004), there are weaknesses with this dataset. It consists of a series of virtual bets in a win-picking contest rather than actual customer bets and there is a high rate of attrition in the sample. A revisiting of this issue with a dataset of actual customer bets may yield significantly different results.

There is also the issue of the relationship between informed traders, noise traders and the bookmaker. According to Glosten and Milgrom (1985), there exists a situation where the adverse selection problems of a market maker are so great, that the market inevitably fails. While it is the case bookmakers are reliant on uninformed traders to offset losses from trading with informed traders, informed traders are reliant on uninformed traders themselves. If noise traders are not gaining utility from their expected negative profit (including the bookmaker's overround), there is no scope for the market to continue.

3.6 Conclusion

This literature review has examined issues concerning market efficiency, asset pricing and trader performance. Despite numerous treatments of wagering markets in previous studies, there is a paucity of empirical studies using financial fixed-odds markets. Consistent with the literature on traditional financial markets, issues of market microstructure have been examined with prices rather than trades. Moreover, the literature on market efficiency in wagering has, for the most part, been undertaken with low-frequency data.

A potential limitation of the use of fixed-odds betting markets to test for market efficiency is the issue of rationality. In the Fama (1970) model, investors are assumed to be risk-averse and utility maximizers, however, numerous studies of wagering markets and lottery-type stocks have shown this not to be the case. If a certain group of traders are betting for reasons other than profit-maximisation, assumptions of rationality cannot be made. However, studies have shown that the same applies to financial markets when individual traded volume is small[1]. It may be that investors gamble when investing small amounts in stocks, and are only rational when they perceive the amounts invested to be of economic significance.

Having presented a broad overview of some of the relevant literature in the realm of behavioural finance, I now proceed to a detailed description of the dataset for this research. The next chapter will outline in detail the dataset selected to address the hypotheses and provides an exploratory analysis of the data. Summary statistics on betting frequency, stake size, the study period, bet prices and underlying products are introduced. I compare and contrast conventional financial panel data with transactional betting data while also addressing some of the limitations of the dataset.

[1] See Statman (2001), Anderson (2008).

Chapter 4

Data

4.1 Introduction

The dataset in this paper has been provided by an online bookmaker. It contains the details of all bets placed on the company's financial fixed-odds betting product via the web platform of the bookmaker, as well as those placed over the phone and via mobile devices over a four-year period. When bets are struck, an entry into a table in the company's database is created with the details of each bet. Bets spanning a number of years are stored in a current database and thereafter archived off-site to a database of historical bets containing all bets struck with the bookmaker since the company started operating.

In order to collect the data for this research, a database analyst ran an SQL query on the current database to return all bets which were placed on the company's financial betting offering. No bets from the historical database were included, therefore the data spans the period 2008 to 2012 and comprises 1,692,252 bets placed on the financial market product.

A number of variables were then encrypted, namely: *Username*, *Customer ID* and *IP Address*. While a table of customer-specific information which could be linked to each bet was available, this was not provided by the company. All usernames and customer-specific details are obfuscated, however individual accounts can be identified and tracked throughout the entire period. As there are no customer-specific details provided, no demographics on individual customers are available. A data licence was drafted by the company specifying the conditions under which the data was being provided for research and the Research Support Department at the University of Edinburgh provided a non-disclosure agreement in turn The data was provided on a memory stick in raw, comma-delimited, *.csv* format. The data was then imported to SAS and STATA for analysis.

The next sections provide an overview of the raw bets provided by the bookmaker and an exploration of the data. Section 4.4 presents summary statistics on stake size by product, bet price, bet channel and bet type. Section 4.5 explores the distribution of bets (and the daily number of bets) by year and by day of the week and Section 4.6 provides some detail on the average duration between bets. Section 4.7 presents the distribution of customer stake and profit in USD size by year, by bet price, the distribution of profit on exit and the profit for each bettor cycle[1]. Section 4.5 concludes.

4.2 Data Exploration

Table 4.1 presents details of the variables in the raw dataset supplied by the provider of the data. There are 30 variables associated with each row in the dataset. The variables in each row describe characteristics of bets such as the timestamp, market, expiry time, stake, market/handicap level, settlement level, result, winnings and refund (if applicable). All timestamps in the dataset are to the second.

To comply with data protection legislation, the company supplying the data obfuscated a number of variables. The unique bet identifier, event identifier, username and bet ip address have been encoded, meaning that the value of the original variable has been passed through an encoding algorithm (AES, Blowfish, SHA-3, MD5 or RSA, for example). The original value of the data cannot therefore be inferred or decrypted. The variable *currency* was also obfuscated. In the case of *currency*, there were two unique values in the dataset: *cur1* and *cur1*.

With the knowledge that the company accepts bets only in EUR and GBP, and by analysing the minimum and maximum stake sizes, it was possible to infer which currency related to which value. Therefore, after importing the data to Stata, the value of stake, winnings and refund were converted to USD from *cur1* and *cur2* using historical FX rates for each date in the dataset. While all usernames are obfuscated, each bettor in the dataset has a unique encrypted value, therefore individual accounts can be identified and tracked throughout the entire period. Two variables, index and completed, are variables which were added by the database analyst who compiled the data and do not from part of the original data.

There are no customer-specific details provided, therefore no demographics on individual customers are available. However, during data exploration, to account for heterogeneity between bettors, a series of customer activity metrics have been calculated and a second source table created. A list of these calculated variables

[1]After a lapse of 90 days between bets, bettors was attributed to a further ‘cycle’ of betting.

Table 4.1

Dataset Variables

This table presents a list of the variables in the raw data for this study along with a brief description, variable type, length and the Stata format. Individual customer identities have been obfuscated and as such, no demographics on bettor characteristics are available. However, individual account IDs can be identified and tracked throughout the dataset. Each row in the table constitutes a unique bet.

Panel A: Raw data file

#	Name	Renamed	Description	Type	Stata Format
1	index	N/A	Bet row identifier	long	%12.0g
2	field1	EVENTTYPE	Product	str54	%54s
3	field2	EVENTTYPEID	Product id	int	%8.0g
4	field4	EVENT	Product type	str45	%45s
5	field5	STARTTIME	Date market offered	str19	%19s
6	field6	SUSPENDAT	No more bets' timestamp	str19	%19s
7	field7	ESTSTARTTIME	Market starting date	str19	%19s
8	field8	MARKET	Market	str33	%33s
9	field9	MARKETID	Market id	long	%12.0g
10	field10	MKTHCAPMAKEUP	Settlement price	double	%10.0g
11	field11	SELNID	Bet price id	long	%12.0g
12	field12	SELECTION	Bet price id	str72	%72s
13	field13	PERCENTAGEMAXBET	Percentage of customer's max stake	double	%10.0g
14	field14	STAKEFACTOR	Customer limit	double	%10.0g
15	field16	BETDATE	Timestamp	str19	%19s
16	field17	SETTLEDAT	Settlement timestamp	str19	%19s
17	field18	SETTLEDDATE	Settlement date	str19	%19s
18	field19	BETTYPE	Bet type	str4	%9s
19	field24	DECIMALBETPRICE	Bet odds	double	%10.0g
20	field25	BETHCAPVALUE	Strike price	double	%10.0g
21	field28	BETCHANNEL	Channel	str1	%9s
22	completed	N/A	DBA flag during obfuscation	byte	%8.0g
23	field3enc	EVENTID	Encoded product type id	str21	%21s
24	field15enc	BETID	Encoded bet identifier	str27	%27s
25	field20obf	CURRENCY	Obfuscated currency	str4	%9s
26	field21obf	STAKE	Obfuscated stake	double	%10.0g
27	field22obf	REFUND	Obfuscated refund	double	%10.0g
28	field23obf	WINNINGS	Obfuscated winnings	double	%10.0g
29	field27enc	USERNAME	Encoded username	str50	%50s
30	field29enc	BETIPADDRESS	Encoded bet ip address	str45	%45s

is presented in Table 4.2. While all the analysis has been done using the original, formatted dataset, the calculated variables presented in Table 4.2 have been used as part of the analysis in the current chapter.

4.3 Singles, doubles and accumulators

Financial bets are not just transacted in isolation, but may be combined with bets on other products i.e. sports bets, novelty bets or political bets. Such bets are known as doubles, trebles or accumulators, and the outcome of such bets depend on the performance of each component. Accumulators can be used as a method of masking informed trading by certain groups of customers and an examination of the activities of customers trading accumulators may yield some insights into

Table 4.2

Dataset Variables

This table presents the calculated variables in the customer table. All the variables are derived from the rows of raw data in the main bet table using the obfuscated customer IDs to aggregate transactions according to bettors. Each row in this table refers to a unique bettor.

Panel A: Aggregate Customer Table					
#	**Name**	**Description**	**Type**	**Length**	**SAS Format**
1	USERNAME	Obfuscated bettor username	Char	50	$100.
4	AVGSTAKEFACTOR	Average 'Stake Factor' applied to account	Num	8	
6	FIRSTSTAKEFACTOR	Initial stake factor of account in dataset	Num	8	6.2
7	LASTSTAKEFACTOR	Final stake factor of account in dataset	Num	8	6.2
8	FIRSTBETPRICE	Initial bet price of account in dataset	Num	8	6.2
9	FIRSTBETTYPE	Initial type of first bet of account in dataset	Char	3	
10	FIRSTEVENTTYPE	Initial market of first bet of account in dataset	Char	24	
11	FACTORCHANGES	Number of stake factor changes	Num	8	
12	BETCOUNT	Number of bets	Num	8	
13	TRADINGDAYS	Number of trading days	Num	8	COMMA15.
14	FIRSTBETDATE	First timestamp customer appears in dataset	Num	8	DATETIME.
15	LASTBETDATE	First timestamp customer appears in dataset	Num	8	DATETIME.
18	ACCOUNTDURATION	Days between first and last bet in dataset	Num	8	BEST8.
19	PRODUCTCOUNT	Number of products bet on	Num	8	
20	BETPRICECOUNT	Number of prices	Num	8	
21	BETTYPECOUNT	Number of bet types	Num	8	
22	FTSEBETCOUNT	Number of bets on the FTSE	Num	8	
23	DOWBETCOUNT	Number of bets on the DOW	Num	8	
24	VIRTUALBETCOUNT	Number of bets on the Virtual Market	Num	8	
25	FINANCIALSBETCOUNT	Number of bets on the Financial Markets	Num	8	
26	FIVETOSIXBETCOUNT	Number of 5/6 bets	Num	8	
27	ONETOTWENTYBETCOUNT	Number of 1/20 bets	Num	8	
28	ONETOONEBETCOUNT	Number of 1/1 bets	Num	8	
29	FIVETOONEBETCOUNT	Number of 5/1 bets	Num	8	
31	MAXSTAKE	Highest stake placed	Num	8	DOLLAR11.2
32	AVGSTAKE	Average stake placed	Num	8	DOLLAR11.2
33	SUMSTAKE	Total amount bet	Num	8	DOLLAR11.2
34	VARSTAKE	VAR (Stake)	Num	8	DOLLAR11.2
35	SDSTAKE	STDEV (Stake)	Num	8	
39	AVGPROFIT	Average profit per bet	Num	8	DOLLAR11.2
40	SUMPROFIT	Total profit.	Num	8	DOLLAR11.2
41	VARPROFIT	VAR(Profit)	Num	8	
42	SDPROFIT	STDEV(Profit)	Num	8	DOLLAR11.2
46	AVGMARGIN	Average return per bet	Num	8	PERCENT8.2
47	TOTALMARGIN	Total return	Num	8	PERCENT8.2
53	AVGBETDURATION	Average holding period	Num	8	

the risk management of betting customers. While a separate study of financial bets as selections in multi-leg bets would certainly be informative, it is not the focus of this study. Moreover, multi-leg bets, doubles, trebles and accumulators are rare in the dataset.

4.4 Financial betting products

Table 4.3 presents stake sizes in USD by financial betting product, by bet type, by bet price and by betting channel. Panel A provides a breakdown of total stake (turnover) according to betting product. The three products with the most betting activity are the FTSE 100 Index, Dow Jones Industrial Index and the Virtual Market, a simulated market where the underlying price is generated by a pseudo-random number generator. The median bet size for bets on the FTSE is \$17.08, with bet sizes ranging from \$0.13 to \$930.07 in the first and 99th percentiles, respectively. Bets are offered for trading in stock indices, commodities, currency pairs and a stock lottery (betting on the last digit of a stock index's settlement price).

Table 4.3

Statistics

This table presents stake sizes by financial product, by bet type, by bet price and by bet channel. All stakes are in USD.

	# Bets	Stake	Mean	Std	P1	P50	P99
All	1,692,252	113,440,049	67.03	286.24	0.13	14.12	871.04
Panel A: Stake by Product							
FTSE 100 Index	901,018	67,647,131	75.08	331.14	0.13	17.08	930.07
DJIA	435,820	31,395,899	72.04	277.59	0.13	15.04	986.54
Virtual Market	260,324	9,408,003	36.14	125.85	0.13	6.36	590.00
Other	93,068	4,980,012	53.51	139.24	0.13	16.81	635.62
Stock Lotteries	2,022	9,004	4.45	11.13	0.11	1.26	61.68
Panel B: Stake by Bet Type							
Single	1,686,283	113,301,724	67.19	286.71	0.13	14.18	874.13
Double	3,451	94,139	27.28	79.83	0.03	4.62	399.12
Treble	1,185	30,107	25.41	69.12	0.02	4.94	340.52
Accumulator	1,151	12,433	10.80	31.95	0.02	2.41	154.25
Trixie	63	820	13.02	38.47	0.02	0.05	206.42
Lucky 15/31/63	39	331	8.48	9.72	0.05	5.38	48.92
Yankee	35	232	6.64	5.38	0.03	4.66	21.67
Patent	21	153	7.30	7.61	0.39	3.03	30.49
Canadian	12	58	4.87	3.30	0.07	6.13	9.19
Heinz/Super Heinz	12	51	4.22	3.64	1.13	3.56	14.41
Panel C: Stake by Bet Price							
5/6	1,232,659	89,203,919	72.37	312.31	0.15	16.44	893.17
1/20	283,800	19,180,825	67.59	237.28	0.12	10.76	985.02
Evens	113,932	4,263,550	37.42	134.41	0.13	6.43	550.61
5/1	59,839	782,750	13.08	34.64	0.12	2.66	155.05
8/1	2,022	9,004	4.45	11.13	0.11	1.26	61.68
Panel D: Stake by Bet Channel							
Internet	1,658,571	108,856,609	65.63	284.17	0.13	13.76	854.43
Mobile	26,266	2,486,303	94.66	227.79	0.30	30.45	1,115.17
Phone	7,415	2,097,136	282.82	635.35	1.53	72.24	3,177.04

It is not surprising that the FTSE would have the highest turnover among customers of a bookmaker based in the UK and Ireland. It is curious, however, that the product with the highest turnover after the DOW and FTSE is the Virtual Market, a Random-Number-Generated (RNG) market. This is a time-series generated using a proprietary implementation of a Mersenne Twister RNG and has been certified as suitably random by the Isle of Man Gaming Commission, the regulator under whom the bookmaker operates. A stock lottery bet on the last digit of the closing price of each financial product (with betting odds of 8/1) was also offered for betting by this bookmaker, but was discontinued and therefore does not constitute a large number of bets in this data set.

Panel B shows the distribution by bet type, with the majority of activity in single[1] bets, while Panel C presents the distribution according to bet price[2]. Panel D indicates the breakdown according to source, with the focus of betting on bets transacted via the company's online web interface. While the volume of bets transacted over the phone is clearly much lower, the median bet size is much higher at \$72.24 as opposed to \$13.76 for internet bets.

To proceed with an exploratory data analysis, a number of histograms are presented in Figure 4.1. As we proceed, the key points to take from the summary statistics presented earlier is that the mean and median number of bets across all product types in aggregate is quite small at 4.9 and 1, respectively. The furthest point from the origin in each case does not show the full extent of the data as the x-axes have been censored for illustrative purposes, but it is clear that the median amounts in each case are quite small. As such, these are mainly individuals and not institutional or professional bettors. Panel A shows the distribution of the total amount bet and total profit for customers. As expected, total profit is centred just below zero. Panel B provides detail on the total number of bets placed and the number of different financial instruments bet on by each account in the dataset, while Panel C show the distribution of both account tenure in days (in the period covered by the dataset) and trading days.

4.5 Study Period

The data set contains 1,692,252 bets in total over a four year period. The first bet in the dataset is on 17 April 2008. There are 105,563 bets in 2008, 457,899 bets in 2009, 490,203 bets in 2010, 453,856 bets in 2011 and 184,731 bets in 2012. The last bet in the dataset is 6 June 2012.

Table 4.4 presents the distribution of bets across weekdays. Prior to the introduction of the virtual financial market, there were no bets transacted on the

[1] A 'double' is a bet which contains two 'legs' (i.e. one on the outcome of a football match and another on the result of a horse race), with the bet only paying off if both legs are successful.

[2] As this is a fixed-odds betting proposition, there are only four discrete bet prices offered for financial bets: 5/6, 1/20, Evens and 5/1. The odds of 8/1 refer to the price of the stock lottery product.

Figure 4.1

Histograms of Betting Variables

The following histograms treat each individual bettor as an observation. Figures (a) and (b) show the distribution of the total bet amount and total net profit for each customer. As expected, total profit is centered just below zero. The furthest point from the origin in each case does not show the full extent of the data as the x-axes have been censored for illustrative purposes, but it is clear that the amounts under discussion are not of institutional size. As such, these are mainly individuals and not institutional or professional bettors. Figures (c) and (d) provide detail on the total number of bets placed and the number of different financial instruments bet on by each account in the dataset, while Figures (e) and (f) show the distribution of both account tenure (during the period covered by the dataset) and trading days. Account tenure and trading days have been aggregated in groups of ten and five, respectively.

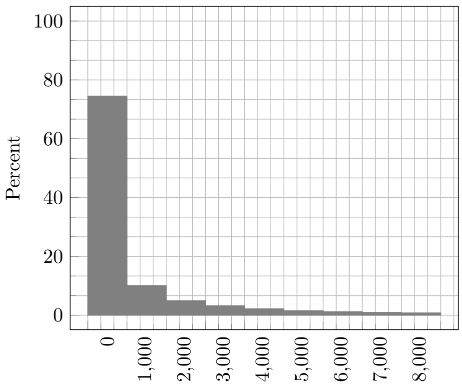

(a) Total bet amount

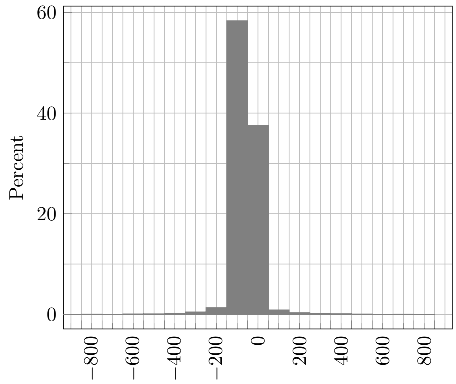

(b) Total net profit

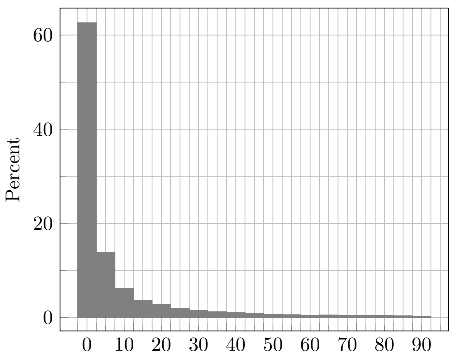

(c) Number of bets

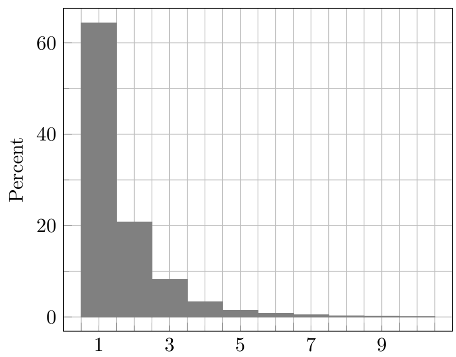

(d) Number of products

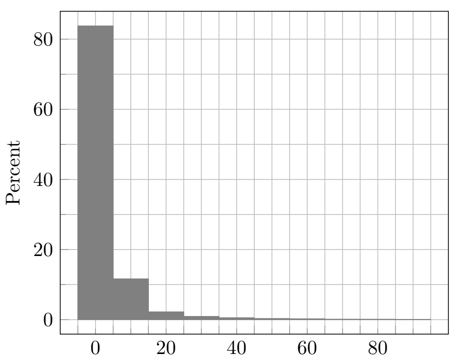

(e) Account Duration

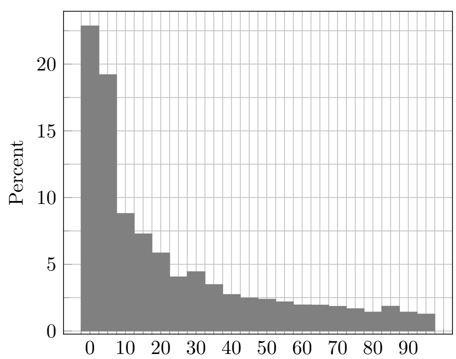

(f) Trading Days

weekend as the financial exchanges are closed. However, with the introduction of a betting market with an underlying generated by a pseudo-random number generator in 2010, a significant number of bets were placed over the weekend in the latter part of the dataset.

Table 4.4

Weekly distribution of betting

This table presents the distribution of betting across days of the week before and after the introduction of betting on the virtual market. There were no bets transacted over the weekend prior to that. Throughout the dataset, the volume and mean stake sizes are lower at the weekend. All stake sizes are in USD.

Day	N	Mean	P1	P25	Median	P75	P99
Stake by day of the week (Pre 2010)							
1	72	1460.88	400.00	941.00	1444.00	1928.50	2,791.00
2	77	1590.60	466.00	975.00	1672.00	2122.00	3,068.00
3	78	1401.31	354.00	725.00	1490.00	2004.00	3,313.00
4	75	1456.92	282.00	785.00	1632.00	2034.00	3,557.00
5	70	1398.03	6.00	692.00	1567.50	1879.00	2,922.00
Total	372	1462.61	332.00	807.00	1590.50	1953.00	2,922.00
Stake by day of the week (Post 2010)							
0	86	513.70	35.00	371.00	472.00	614.00	1,215.00
1	126	1541.51	352.00	1171.00	1511.00	1887.00	3,022.00
2	126	1713.00	646.00	1324.00	1652.00	2028.00	2,820.00
3	126	1668.17	810.00	1319.00	1581.50	1931.00	3,179.00
4	123	1696.53	853.00	1402.00	1619.00	1949.00	2,809.00
5	124	1652.73	765.00	1322.50	1580.50	1880.50	2,748.00
6	86	590.03	128.00	452.00	552.50	710.00	1,722.00
Total	797	1416.30	223.00	966.00	1458.00	1843.00	2,846.00

4.6 Feedback frequency

Table 4.5 presents the distribution of the number of bets per account each year, the number of bets per day, the time-lapse between bets and the time duration between bet placement and settlement. The average number of yearly bets per account has been decreasing since 2008, with a significant drop from almost 90 bets per account per year to 27.49 bets per account per year in 2011, the last full year in the dataset. The median number of bets also decreased during the period covered by the dataset, from 4 per year to 2.

As regards the average number of bets per day, this also decreased from 2009 to 2012, with a drop of almost a quarter during this period. The time between bets has increased, with an increase in the mean ‘rest period’ between bets from 1.39 days (2000.51 minutes) to 3.02 days (4342.38 minutes).

Table 4.5

Betting frequency & Feedback

This table presents the distribution by year of the number of bets per account, the number of bets per day transacted with the bookmaker, the distribution of time between bets and time duration between bet placement and bet settlement. The mean (median) bettor places 47 (3) bets, however the median is 4 (2) for accounts opened in 2008-9 (2011-12). On average, there are in excess of 1,300 bets/day, with lower activity in 2008 and a fall in volume after 2010. The median timespan between action and feedback in the entire sample is 2.62 minutes (0.04 hours).

Day	N	Mean	P1	P25	Median	P75	P99
Bets per account by year							
2008	4,848	88.91	1.00	1.00	4.00	16.00	1,475.00
2009	8,235	62.13	1.00	1.00	4.00	19.00	873.00
2010	8,864	46.08	1.00	1.00	3.00	11.00	562.00
2011	9,400	27.49	1.00	1.00	2.00	8.00	445.00
2012	4,568	18.11	1.00	1.00	2.00	6.00	300.00
Total	35,915	47.12	1.00	1.00	3.00	11.00	672.00
Daily # bets by year							
2008	187	564.51	2.00	475.00	563.00	651.00	1,014.00
2009	259	1767.95	445.00	1480.00	1772.00	2110.00	3,068.00
2010	290	1690.36	147.00	1217.00	1691.50	2202.00	3,179.00
2011	361	1257.22	199.00	718.00	1335.00	1664.00	2,503.00
2012	146	1265.28	296.00	644.00	1406.50	1710.00	2,435.00
Total	1,243	1361.43	199.00	709.00	1402.00	1844.00	2,843.00
Time between bets (minutes)							
2008	100,715	2233.86	0.25	9.35	22.35	131.75	51,883.25
2009	449,664	2000.51	0.00	2.55	5.08	14.48	38,863.68
2010	481,339	2738.72	0.02	1.53	3.27	8.27	41,907.32
2011	444,456	3692.54	0.03	1.23	2.50	6.83	57,215.57
2012	180,163	4342.38	0.08	1.23	2.48	6.05	58,644.53
Total	1,656,337	2937.99	0.03	1.70	3.92	11.35	46,076.75
Time to feedback (minutes)							
2008	105,563	165.21	10.97	14.97	18.93	44.85	465.35
2009	457,899	17.98	1.37	2.57	4.07	5.93	345.50
2010	490,203	16.25	0.95	1.77	2.43	4.18	291.13
2011	453,856	12.29	0.63	1.17	1.80	3.32	265.68
2012	184,731	7.73	0.62	1.00	1.47	2.65	169.63
Total	1,692,252	24.02	0.68	1.65	2.62	5.48	333.33

Nevertheless, customers have been betting on increasing shorter expiries with each subsequent year. The mean time between bet execution and feedback (i.e. bet settlement) reduced from 17.98 minutes to 7.73 minutes in 2012[1].

[1]The statistics for 2008 are driven by an outlier market which was a bet on the price of oil at the end of the year, with a time duration between bet execution and settlement spanning a period of months.

4.7 Customer Betting Activity

While the evolution of betting could be analysed on a bet-by-bet basis as each bettor enters the dataset, there are a number of issues of concern with approach. Firstly, there is no way to say with certainty that the dataset contains each bettor's first bet with this bookmaker. The dataset spans a number of years after the product was offered for betting by the bookmaker and is essentially left-censored. I therefore use 90 days as a cut-off point at the start and end of the dataset, excluding any bettor who placed a bet in the first 90 days covered by the dataset. Of the 35,915 bettors in the dataset, 31,050 were not present in the first or last three months of the dataset and those accounts placed 922,732 of the 1,692,252 bets in the dataset.

I also used 90 days as a time interval to create 'cycles' for each bettor, reactivating any bettor who has been dormant for more than this time period as a new bettor-observation. A new cycle is defined as trading activity by a bettor after a break of 90 days from the previous bet. In terms of the distribution of cycles, 82% of bettors only have a single cycle of betting, 13% have two cycles, a total of 82% of bets are in the first cycle and 94% within the first two cycles. As a result, we can conclude that the vast majority of betting happens in at most two bursts of trading activity, with a break of 90 days between cycles. The combined effect of excluding bettors that were present at the start and end of the dataset, and also recoding accounts with time intervals longer than 90 days, has increased the total number of bettor-observations from 35,628 to 44,041. A new cycle of betting was attributed to a customer after there were a lapse of 90 days between the last bet.

Table 4.6 presents summary statistics on a number of betting variables with a focus on the highest and lowest percentiles of profit across bettors. *Skill* is equal to 0 if a bettor is in the bottom 1% in terms of profit after their last bet, 3 if in the top 1%, and 1 otherwise. The worst bettors have on average almost two cycles in total. They also bet more frequently on average that those with the highest ability. The lowest percentile bettors in terms of skill (i.e. profit) execute 784.96 bets, whereas the top 1% of skilled bettors on transact and average of 47.12 bets. The worst performers also have less time duration between bets.

This is consistent with evidence in the literature that those with the least ability trade more Barber and Odean (2000). However, the mean stake size of 133.41 in the top percentile of bettors is considerably larger the average stake size of 93.79 placed by the lowest skilled bettors.

Table 4.7 presents the distribution of stake size by year, profit by year, profit

Table 4.6

Stats by skill flag

This table presents the number of bets, the number of cycles, summary stats on stake size and the length of inactive periods between between bets for three levels of bettor skill, calculated on the basis of the highest and lowest percentiles of profit on exit. The variable *Skill* is given a value of zero (two) for the lowest(highest) percentiles of bettor in terms of P&L after their final bet, and a value of 1 otherwise. Summary stats for stake size are in USD.

Skill	N	Mean	P1	P25	Median	P75	P99
			# Bets by Skill flag				
0	1,378	784.96	3.00	67.00	192.00	565.00	11,437.00
1	34,471	17.62	1.00	1.00	3.00	9.00	244.00
2	66	47.12	1.00	4.00	15.00	45.00	642.00
Total	35,915	47.12	1.00	1.00	3.00	11.00	672.00
			#Cycles by Skill				
0	1,378	1.82	1.00	1.00	1.00	2.00	6.00
1	34,471	1.23	1.00	1.00	1.00	1.00	4.00
2	66	1.20	1.00	1.00	1.00	1.00	3.00
Total	35,915	1.25	1.00	1.00	1.00	1.00	4.00
			Stake by Skill				
0	1,081,676	93.79	0.16	7.83	27.57	68.76	1,135.20
1	607,466	19.06	0.12	1.27	4.76	15.04	238.47
2	3,110	133.41	0.32	4.77	47.97	122.76	1,303.47
Total	1,692,252	67.03	0.13	3.56	14.12	48.68	871.04
			Time between bets by skill				
0	1,080,298	740.75	0.13	1.58	3.67	10.07	8,816.88
1	572,995	7085.59	0.00	1.90	4.28	15.52	203,614.28
2	3,044	1990.84	0.00	0.80	3.30	13.79	31,707.15
Total	1,656,337	2937.99	0.03	1.70	3.92	11.35	46,076.75

by bet price, profit on exit and profit by bettor cycle. The average stake size is $67, with a median of $14, and decline in the average stake in subsequent years.

4.8 Limitations

The dataset has a number of limitations. Firstly, no demographics are available for bettors. I cannot present summary statistics on age, location, gender, level of education, financial literacy or market experience.

Secondly, as the dataset does not include all bets on the company's financial product since its inception, we cannot be sure that we are capturing each customers initial betting experience. However, 30,000 accounts did not place a bet in the first and last months of the dataset. Only 67 accounts in total placed a bet in the first or last three months. I therefore excluded any customer who placed a bet in the first three months in the dataset and considered the remaining bettors

Table 4.7

Customers

This table presents the distribution of stake size by year, profit by year, profit by bet price, profit on exit and profit by bettor cycle. All stake sizes and profit are in USD.

Skill Flag	N	Mean	P1	P25	Median	P75	P99
			Stake by year (USD)				
2008	105,563	108.38	0.18	6.31	19.17	83.40	1,408.04
2009	457,899	81.04	0.13	3.60	14.19	53.05	1,037.22
2010	490,203	61.26	0.12	3.75	15.16	49.06	886.75
2011	453,856	52.63	0.13	3.28	12.91	42.79	654.76
2012	184,731	59.43	0.13	2.53	10.12	38.06	834.71
Total	1,692,252	67.03	0.13	3.56	14.12	48.68	871.04
			Profit per account by year (USD)				
2008	105,563	-64.86	-1670.70	-39.83	-0.44	15.83	691.06
2009	457,899	-36.45	-1070.71	-22.27	0.15	9.82	462.22
2010	490,203	-28.18	-838.11	-20.01	0.14	11.27	392.66
2011	453,856	-22.37	-679.20	-13.54	0.31	7.26	280.24
2012	184,731	-23.15	-844.09	-10.13	0.27	6.43	350.99
Total	1,692,252	-30.60	-946.77	-17.57	0.19	9.44	413.30
			Profit by price (USD)				
1/20	283,800	-6.11	-172.37	0.03	0.41	1.75	48.34
5/1	59,839	-12.52	-268.54	-16.10	-3.59	-1.57	267.05
5/6	1,232,659	-37.53	-1099.85	-30.38	0.29	15.52	476.19
8/1	2,022	-4.01	-123.37	-4.88	-2.36	-0.61	75.68
Evens	113,932	-26.65	-625.87	-16.33	-2.25	3.88	259.41
Total	1,692,252	-30.60	-946.77	-17.57	0.19	9.44	413.30
			Profit on exit (USD)				
2008	4,848	-4379.06	-78255.84	-257.92	-34.27	-3.11	155.46
2009	8,235	-2055.68	-20946.10	-142.19	-22.79	-2.62	88.30
2010	8,864	-970.68	-12167.21	-59.69	-10.39	-1.40	160.49
2011	9,400	-423.14	-5118.30	-37.39	-6.95	-1.28	35.96
2012	4,568	-228.69	-3459.41	-25.67	-5.07	-1.07	39.36
Total	35,915	-1441.86	-16587.27	-72.63	-12.23	-1.60	86.79
			Profit by cycle (USD)				
1	29,447	-1050.12	-8449.19	-46.44	-7.56	-1.08	88.36
2	4,634	-2414.51	-40183.87	-255.11	-45.20	-7.71	88.41
3	1,269	-4612.10	-130783.99	-763.96	-147.03	-29.74	48.83
4	395	-7249.31	-207209.84	-1843.14	-333.25	-60.32	243.68
5	118	-5096.25	-54014.49	-2511.21	-604.48	-113.23	44.55
6	41	-7774.58	-154000.70	-4580.60	-505.58	-273.50	28.24
7	9	-2963.95	-10492.35	-5926.44	-506.90	-193.69	-6.79
8	2	-4808.46	-9611.52	-9611.52	-4808.46	-5.41	-5.41
Total	35,915	-1441.86	-16587.27	-72.63	-12.23	-1.60	86.79

as essentially newly opened accounts. I used the same approach to deal with the latter part of the dataset and considered any bettor who had not had a bet in the last three months as having exited the dataset.

Thirdly, a full time-series of the prices offered by the bookmaker throughout the dataset was not available. As a result, we cannot examine issues of pricing and how they may have related to subsequently transacted bets. Nevertheless, as the focus of this research is on the behaviour of individuals, rather than bookmakers, the panel dataset is sufficient in order to address the hypotheses in later chapters.

4.9 Conclusion

In the preceding chapters, the motivation for the research, the suitability of the dataset, the strengths and weaknesses of the setting and a broad literature review were described. This chapter gave a broad overview of the mechanics, chronology and specification of the betting offering in this setting. The competitive environment between the main players in this industry were introduced and a note on the regulatory environment was presented.

At this point, we have sufficient background information and motivation to introduce the first of three results chapters, beginning with the first, on the overconfidence of individual investors.

Chapter 5

Reinforcement learning and overconfidence

5.1 Introduction

This chapter examines how reinforcement learning induces overconfidence. I analyse how successive wins and losses can cause biased learning, resulting in investors incorrectly estimating their abilities. In addition, I examine overconfidence in two distinct market settings: one where ability can result in superior returns and a second where returns are uncorrelated and skill is futile. This novel market segmentation allows the testing of predictions from the literature in a setting where confidence could be warranted and in a second setting where they are unwarranted.

I find that strong positive reinforcement prompts bettors to increase their stake sizes and that such stake size increases decrease their wealth with each subsequent bet. Strong negative reinforcement has an effect on attrition, in that unsuccessful bettors exit the market. However, those that remain do not temper their stake sizes significantly. In addition, after controlling for the path of wins and losses, bettors have a tendency to increase the amount they wager with each subsequent bet. As such, learning manifests itself in exit from the market in the losing domain, rather than a rational Bayesian decrease in stake size. I perform this analysis in a setting which has distinguishing features, including its similarity to a controlled laboratory experiment, the existence of a simulated market and the expedient feedback received after each transaction.

The rest of this section outlines the motivation for the chapter and introduces some relevant concepts from the literature, while a focused literature review is presented in Section 5.2. The research questions are introduced in Section 5.3, and are addressed in Section 5.4 when the results are presented. Section 5.5

concludes.

Individual preference models assume that agents are rational while empirical research in the area of behavioural finance has suggested otherwise. The possibility of irrational agents in a competitive market is accounted for with the following proposals: (a) irrational agents execute trades randomly and their net effect is negligible (Kyle, 1985), (b) irrespective of trading by irrational agents, a subset of informed arbitrageurs ensure that prices are efficient (Shleifer and Vishny, 1997) or (c) prices approach equilibrium as agents 'learn by trading' (Seru et al., 2010). This dictates that if investors are not rational from time to time, they can be relied upon to update their sources of private information (or the weights they apply to various sources of information) and learn in a Bayesian manner to be rational. The premise is that any market inefficiencies that are caused by such biases are eventually 'traded out'. However, Brav and Heaton (2002) note that the empirical research on convergence to rational expectations equilibrium has demonstrated that this will not just 'happen', even if agents have the possibility of learning their way to it. Various forms of non-Bayesian learning have been posited as explanations for asset pricing anomalies (Kaustia and Knüpfer, 2008).

A consequence of such biased learning is overconfidence (Gervais and Odean, 2001). Rather than updating their beliefs about their ability in a purely Bayesian manner, investors may overestimate their ability, attributing success to their own superior knowledge. In the literature, it has manifested itself in a three-fold pattern of behaviour: *overplacement* (the tendency of agents to rank themselves higher than their peers), *overprecision* (indicating overconfidence in the accuracy of one's predictions) and *overestimation* (unwarranted confidence in abilities)[1]. It has also been shown to affect equilibrium outcomes in almost every market setting (Benoit et al., 2014). In fact, in financial markets, overconfidence has been shown to lead to overtrading (Odean, 1998b), which in turn reduces investor wealth (Barber and Odean, 2000).

Overconfidence, however, is quite distinct from sentiment, or excessive optimism or pessimism, although there is a correlation between the two. Optimism refers to the overestimation of the frequency of positive outcomes and the underestimation the frequency of negative outcomes. Barone-Adesi, Mancini, and Shefrin (2012) state that excessive optimism occurs when mean returns are overestimated, and overconfidence occurs when the volatility of such returns is underestimated, such that optimism leads to disappointment with mean returns and surprise at their subsequent standard deviation. Although distinct, these two biases have

[1]See Benoit, Dubra, and Moore (2008), Benoit, Dubra, and Moore (2014) and Benoît and Dubra (2011) for a debate on the prevalence of these three traits.

been shown to be correlated (Shefrin, 2008). A further distinction can be made between the objective and subjective or dispositional components of optimism. The latter refers to a general disposition towards unknown events in the future, while the former refers to the actual expected return which is a product of such personality traits (Cervellatia, Pattitoniac, and Saviolib, 2013).

There is at times, however, no clear distinction in the empirical literature between optimism and overconfidence. After giving a definitition of both terms, Forbes (2009:141) states: 'Since the outcome of these two differentiable biases is often the same behaviour, I simply note the difference before proceeding to ignore it'. Trevelyan (2008) argues that optimism is a personality trait that is relatively stable, wherein one can be overconfident in one's abilities in a certain task, but underconfident about one's abilities in another. This is consistent with evidence in Gervais, Heaton, and Odean (2002) that individuals' overconfidence varies with task difficulty. However, empirical studies which use gender as a proxy for overconfidence are at odds with this. Fabre and François-Heude (2009) states that overconfidence and optimism are closely related, likely to appear jointly and correlated, but that each bias is distinct theoretically and empirically. They argue that optimism refers to a belief in a positive outcome, independent of the ability an individual can harness to effect the outcome.

Malmendier and Tate (2005a) note that in the psychology literature, confidence is used to refer to biases in self-assessment, whereas optimism commonly describes beliefs about the likelihood of exogenous events in the future. They choose the term 'confidence' in order to make a distinction between overoptimism about mean returns which occurs as a result of overconfidence about ability, and the general positive disposition that results in 'optimism'.

Malmendier and Taylor (2015) state that terminology about biases borrowed from psychology can be quite broad, and that different theoretical approaches and empirical predications can result, depending on whether one is referring to overoptimism (overestimating the frequency or magnitude of positive outcomes beyond one's control), overconfidence (overestimating outcomes under one's control, or at least perceived to be) or overprecision (biased beliefs about volatility).

In this chapter, I define overconfidence as the overestimation of one's abilities in the case of negative outcomes on both financial and simulated markets. I suggest how the current analysis could be extended to disentangle overconfidence from overoptimism in the case of negative returns in the simulated market in the conclusion.

Empirical testing of such behavioural finance theory is complicated by a paucity of transactional panel data. As an empirical analysis of investor behaviour ideally entails a study of panel-data involving the same investors over a period of time, any analysis of investor behaviour in the stock market is only possible where such information is disclosed in a fully transparent manner to the regulator. Although scarce, a number of such datasets are available, a notable one of which is data collated by the Finnish Central Securities Depository which includes the details of shareholdings and financial transactions for all investors in Finland[1].

A natural response to this lack of available data is to turn to a laboratory setting as is common in the experimental economics literature. However, respondents in such experiments may not always be sampled randomly, as many have students as respondents due to their proximity to the location of the lab on campus. In addition, the costs associated with providing respondents with adequate consideration as to make the contingent claims being traded economically significant are prohibitive. An alternative approach used in previous literature on informed trading has used prices to draw conclusions on the activities of informed traders, however with the dataset used in this study, behaviour can be observed directly.

By using a longitudinal dataset comprising in excess of 1.5 million individual-level fixed-odds financial bets[2] we have a natural-experimental setting with which to test hypotheses. The sample includes transactions from more than 30,000 customers from an online bookmaker on major stock indices, currency pairs, commodities and also on a simulated market (Virtual Market), and is a similar dataset to that exploited by academics performing empirical tests of behavioural finance theories with brokerage data.

In this setting, bettors are performing identical, consecutive decisions which mimic financial choices made in a laboratory, but the use of their own funds departs from the artificiality of an experiment. Also, in contrast with learning in other markets, for example IPO markets,[3] not only is this a clean experiment (i.e. with no 'hot' and 'cold' IPOs or issue-specific characteristics) but there is also a relatively short time between action and response which should facilitate more expedient learning. Indeed, Brav and Heaton (2002) note that learning

[1]See Grinblatt, Keloharju, and Ikäheimo (2008), Kaustia and Knüpfer (2008), Keloharju, Knüpfer, and Torstila (2008), Grinblatt and Keloharju (2009), Seru et al. (2010), Shive (2010), Kaustia (2010), Linnainmaa (2010), Grinblatt, Keloharju, and Linnainmaa (2011), Linnainmaa (2011), Keloharju, Knüpfer, and Linnainmaa (2012) and Linnainmaa and Saar (2012).

[2]The bets have a specific strike, maturity and terminal value, and given their discontinuous payoff functions, are analogous to ultra short-term digital, binary or 'cash-or-nothing' options. This specification makes them trivial to price as (similar to short-term interest rates and fixed-income securities) they expire at a pre-determined price and at a pre-determined time in the future.

[3]See Kaustia and Knüpfer (2008).

in experiments requires immediate outcomes while Russell and Thaler (1985) state that without well-structured timely feedback, learning may be negligible. This novel data set allows the assessment of how behaviour changes according to different learning outcome paths and how biased learning induces overconfidence.

Crucially, it also facilitates distinguishing between luck and smart behaviour. The dataset includes bets on a financial betting market in which the underlying is a stock index, commodity or currency pair (i.e. Dow, S&P, FTSE, Crude, Brent, EUR/USD, GBP/USD). In this market, the existence of informed traders is possible as the betting proposition is analogous to a binary or cash-or-nothing option. However, this setting is unique in the following respect. The dataset also contains bets on a simulated financial market (Virtual Market) where the underlying time-series is created by a random-number generator, where returns are uncorrelated and success is due to luck alone. In that respect, we can examine behaviour in a setting where the confidence of informed agents could be warranted[1] and a setting where confidence is unwarranted. This delineation between the characteristics of both market settings affords a rare opportunity to distinguish between luck and skill[2].

The main finding in this chapter is that strong positive (negative) reinforcement induces overconfidence as proxied by increased (reduced) stake sizes. In addition, bettors on a simulated market which is driven purely by a random-number generator seem to increase and decrease their stake sizes in a path-dependent manner, disregarding the fact that returns are uncorrelated. Such behaviour is evidence of an emotional response to random reinforcement. Also, the most salient learning effect is in the initial stages and the same pattern of reinforcement has a lesser effect when experience in later rounds of betting. Moreover, even when the path of results is controlled for, there is a tendency to increase stake sizes with each subsequent bet. This effect is stronger for the simulated market, which should not be the case if agents are rational.

The next section provides an overview of relevant literature and motivates the hypotheses outlined in a subsequent section.

[1] See Song, Jang, Hanssens, and Suh (2014) for a discussion of the positive effects of both managerial overconfidence and overconfidence as a mitigating factor in biases by individual investors.

[2] This study also constitutes the first analysis of the financial fixed-odds betting market, and in doing so, sheds light not only on a heretofore opaque market setting but also on the activities of relatively recent entrants into the market-making sphere: traditional sports bookmakers.

5.2 Literature Review

The literature defines overconfidence as the overestimation of one's knowledge, the precision of private information, or the underestimation of the volatility of unknown processes (Skala, 2008). Much of the empirical work in this domain has focused on corporate finance. Malmendier and Tate (2005a) show that firms with overconfident CEOs (based on a personal portfolio characteristics) display suboptimal investment behaviour.[1] Gervais, Heaton, and Odean (2011) (and Gervais, Heaton, and Odean (2003)) study managerial overconfidence in the context of compensation and subsequent investment decisions, suggesting that overconfident managers tend to accept more convex compensation structures which are associated with very high levels of risk. Overconfidence has also been suggested as one of the causes of SSW[2] bubbles in asset prices. (Scheinkman and Xiong, 2003). Broihanne, Merli, and Roger (2014) interviewed finance professionals and noted that they exhibited overconfidence in terms of miscalibration of probabilities but not in terms of a better than average effect. As regards individual investor behaviour, Barber and Odean (2001) showed that overconfident investors trade too much.

However, accurate measures of either overplacement, overprecision or overestimation in transactional data are elusive (Yeoh and Wood, 2011). Prior literature has relied on proxies for overconfidence, such as sophistication (Song et al., 2014) or have elicited responses to questions designed to screen for overconfidence such as miscalibration, the better than average effect, illusion of control and unrealistic optimism (Glaser and Weber, 2007).

Gender has been frequently used as a proxy for overconfidence, however evidence has been mixed and there is a lively debate in the literature on whether gender alone accounts for the heterogeneity in overconfidence. Barber and Odean (2001) focus on gender differences in trading volume, showing that men trade more excessively than women and that this excess trading is wealth-reducing. Bengtsson, Persson, and Willenhag (2005) also find support for the premise that overconfidence is gender specific. In contrast, however, Deaves, Lüders, and Luo (2008) find little evidence that there are differences in gender in an experimental setting. Biais, Hilton, Mazurier, and Pouget (2005) correlated miscalibration with returns in an experimental asset market, noting that miscalibration reduced trading performance as overconfident traders were particularly subject to the winners curse. However, in terms of gender, they also found no significant difference

[1]See also Malmendier and Tate (2005b), Malmendier and Tate (2008), Brown and Sarma (2007) and Goel and Thakor (2008).

[2]Smith, Suchanek, and Williams (1988).

in miscalibration between men and women.

Barber and Odean (1999) suggest overconfidence as the overriding factor causing the high levels of trading evident in financial markets. Grinblatt and Keloharju (2009) derive an overconfidence measure from a psychometric test given to Finnish males upon entry into military service. They find no correlation between the overconfidence measure and turnover in a panel dataset of household investors, however the overconfident individuals do trade more often. Statman, Thorley, and Vorkink (2006) indicate that investors become more confident after positive returns and less overconfident after negative returns. They show that trading volume is correlated with past returns and interpret this as support for the overconfidence hypothesis.

Glaser and Weber (2007) test whether overconfident investor trade more, using a questionnaire to create overconfidence scores (in terms of miscalibration, volatility estimates and the better than average effect). They find no correlation between miscalibration and trading volume, however those that thought they were above average in terms of skill or performance transacted more. Gervais and Odean (2001) state that overconfidence does not make for successful trading, but successful trading can cause overconfidence. Indeed, Merli (2014) tests for overconfidence in a large brokerage dataset, concluding that assets being bought by overconfident investors underperform the assets they are selling.

Thus, the consensus in the early behavioural finance literature is that overconfidence is not a good thing. More recently, however, research has been suggesting that overconfidence may not be an overwhelmingly negative trait.

Hirshleifer (2001) states that the theory suggests that only moderate overconfidence should result in successful investment decisions and empirical work may be overstating poor performance by extremely overconfident investors. Moreover, Benoit et al. (2008) contend that many studies examining overconfidence do not indicate *true* overconfidence. In a follow-up paper in this debate, Benoît and Dubra (2011) show that people ranking themselves as better than average on simple tasks can be explained by rational Bayesian updating. In addition, Ko and Huang (2007) suggest that some biases and overconfidence may make the market more efficient and show that moderate overconfidence can improve market pricing by generating excess information acquisition by overconfident investors.

However, overconfidence alone does not convincingly explain observed behaviour. Mahani and Bernhardt (2007) show that past experience is correlated with trading intensity in the future. If overconfidence were a fixed characteristic, however, returns in the past would be uncorrelated with future decisions which is not observed in empirical studies (Linnainmaa, 2011). Mahani and Bernhardt

(2007) state that standard risk-seeking or overconfidence models in which learning is not central cannot explain empirical market irregularities. Gervais et al. (2002) state that agents tend to be underconfident when completing relatively easy tasks, but overconfident when addressed moderate to extreme difficulty tasks. They argue that one of the reasons we should expect overconfidence in managers is that capital budgeting decisions do not lend themselves well to learning. One of their arguments is that such decisions are not encountered on a frequent basis, may be fundamentally different with each iteration and are not associated with clear, expedient feedback. Menkhoff, Schmidt, and Brozynski (2006) present mixed results on the relationship between experience and overconfidence. When strictly defined as miscalibration, overconfidence is reduced with experience, however for certain tasks experience had the opposite effect. They stress the importance of learning in the process in which overconfidence is fostered.

In addition to these suggested causes, there is evidence in the literature that learning is key to the formation of biases. Gervais and Odean (2001) develop a model in which traders are subject to biased learning which generates overconfidence. However, overconfidence in this model is dynamic, increasing from the offset as investors wrongly attribute too much of their success to their own skill, but decreasing as more information is revealed about their true ability. Hoffmann and Post (2013) show that more confident investors rely on reinforcement learning, extrapolating from their own return experiences into future transactions. In addition, Mahani and Bernhardt (2007) argue that investors learn about their ability by observing their performance and that the data is inconsistent with explanations which fail to incorporate this naive learning. Gervais and Odean (2001) show that biases in learning can cause overconfidence when investors incorrectly infer ability from their successes. In their model, investors learn about their ability as they trade. Those who have had positive returns naïvely update their beliefs about their ability, becoming overconfident as their attribute success to their own skill.

Hirshleifer and Luo (2001) show that overconfident traders persist in markets, in spite of their underestimation of risk and overestimation of the expected return from their investments. However, they acknowledge that their model does not incorporate changes in overconfidence over time, nor do they allow for 'learning about ability'[1]. Camerer and Lovallo (1999) show that when agents are subjected to an experimental treatment where their payoffs depend on their own abilities, they overestimate their own skill and participate at a higher intensity. They pro-

[1]Learning about ability suggests explorative investment by novice traders in order to infer information about their relative prowess in trading. Those that infer that they are in the low-skill cohort should rationally exit the market.

pose that overconfident agents would not only be relatively insensitive to risk, but may actually prefer riskier assets as they believe they can 'beat the odds'. Daniel, Hirshleifer, and Subrahmanyam (2004) propose a theory based on overconfidence and biased self-attribution as factors which cause investors to overweight private signals and underweight public information.

Thus, it seems as though there may be an interaction between reinforcement learning and overconfidence. This is not the first study to note this. Song et al. (2014) examine purchasing decisions (i.e. whether to participate, what to purchase and how much) following negative and positive returns in the ELN market, showing that negative returns are associated with attrition and lower levels of participation.[1] However, they provide empirical evidence that proxies for overconfidence can attenuate the reinforcement learning effect. While this is in contrast with studies of individual investors showing that overconfidence is detrimental (Barber and Odean, 2000), it is consistent with arguments in favour of wealth-creating managerial overconfidence. In their sample, men and online channel investors (linked with overconfidence in the literature) take more risk after experiencing negative returns.

The crux of the matter is whether individuals learn rationally. While a central tenet of asset pricing models, the rationality assumption has been subject to scrutiny. De Bondt and Thaler (1985) find evidence of systematic price reversals for extreme-return stocks consistent with the overreaction hypothesis, in contrast to the rational Bayesian response to new information. De Bondt and Thaler (1987) re-iterate that the strategy tested in that paper was motivated by the premise the investors are poor Bayesian decision makers. De Bondt and Thaler (1990) further articulate the research of Kahneman and Tversky (1973) which states that people overweight salient information and underweight less salient information when making predictions. They also emphasize that behavioural explanations for anomalous stylised facts observed in financial markets should be taken seriously. If investors do not learn in a rational Bayesian fashion and instead suffer from a similar bias to that set out in the naïve reinforcement hypothesis, the rationality assumption may not hold. To that effect, as suggested by Barberis and Thaler (2003, p. 1118), the 'continued empirical scrutiny of assumed behaviour is essential to validating the claims of behavioural finance theorists'.

There is also empirical support for the hypothesis that investors are subject to reinforcement learning. In a study of individual investors at a large discount brokerage, Strahilevitz, Odean, and Barber (2011) identify patterns in trading by

[1] According to Mahani and Bernhardt (2007), agents must be overconfident in order to generate high initial losses and high attrition rates.

individuals which are affected by emotions i.e. investors repeating actions which resulted in a profit, while avoiding actions which resulted in a loss. Di Guisto, Brown, and Maughan (2013) show there is persistence in bettor returns and that betting improves with more experience, however the effect is due to attrition by less skilled bettors. In this case, bettors learn about their ability rather than learning by doing.

Kaustia and Knüpfer (2008) examine the relationship between returns on previous IPO subscriptions and the likelihood of subsequent participation in further IPOs. They conclude that personally experienced returns are an important determinant of future activity and that this is consistent with reinforcement learning theory. The results of the three tests performed by Kaustia and Knüpfer (2008) indicate that individual investors are affected by personally experienced performance and are more likely to participate, more likely to participate sooner and more likely to participate at a higher intensity if they have experienced positive returns. They state that these results are consistent with reinforcement learning, however they do note a number of alternative reasons for this behaviour, including further unobserved differences between investors, portfolio re-balancing, wealth effects, expectations of preferential treatment by investment banks and the existence of a 'hot issue' market during the sample time period. In outlining their contribution to the literature, they hint at implications for the IPO and asset pricing literature, the role of sentiment in economic decision making and empirical tests of the reinforcement learning hypothesis.

Chiang, Hirshleifer, and Qian (2011) expand the Kaustia and Knüpfer (2008) study by examining whether investors improve their ability by rational learning or whether their performance deteriorates due to reinforcement learning. They also contend that the Kaustia and Knüpfer (2008) results are also consistent with rational Bayesian learning, as those investors who experience positive returns will tend to participate more often than those who have experienced negative returns. While Kaustia and Knüpfer (2008) analyse whether investors participate move if they have experienced positive returns, Chiang et al. (2011) examine what effect this continued participation actually has on returns. Their dataset includes details on IPO subscriptions in the Taiwanese market and in contrast to that of Kaustia and Knüpfer (2008), includes data on both individual and institutional investors. They also expand on the analysis of Kaustia and Knüpfer (2008) by differentiating in all of their tests between individual and institutional investors. They indicate that individual investors (but not institutional investors) are subject to naïve reinforcement learning, as evidenced by their deteriorating returns as they gain experience and in contrast with the rational Bayesian learning hypothesis.

Seru et al. (2010) investigate how individual investors are affected by two distinct types of learning: learning about their own abilities and learning by trading. They conclude that investor performance improves with experience, however they highlight that attrition due to investors learning about their lack of ability may be an important factor. The authors stress that the performance of investors who remain active should improve, not just that the performance of the sample in aggregate improves over time. In approaching the problem, they estimate a simple learning model which looks for evidence of learning in the sample in aggregate, assuming that the attrition from the sample is random and that all investors are homogeneous. The overall result of the Seru et al. (2010) paper is that the correlation of both performance and disposition with investor experience and survival rates suggests that investors learn by trading. The authors state that without having controlling for endogenous attrition and individual heterogeneity, however, the literature overestimates the effect of experience on learning. They suggest that the existing literature on learning, stating that authors overestimate how quickly investors become better at trading because they ignore attrition from the sample of investors who have 'learned about their ability'. In fact, learning by trading happens slowly, indicating the possibility of persistent market inefficiencies while investors are in this 'learning phase'.

Benos (1998) indicates that the persistence of overconfidence over time and learning are still open questions in the literature and that the investigation of information markets may be fruitful in this respect. Fixed-odds betting markets offer a quasi-experimental setting in which to perform such empirical work. According to Chiang et al. (2011), it is important to examine in what contexts individuals can learn their way out of cognitive biases and in what contexts learning actually exacerbates bias. Hirshleifer (2001) note that in the presence of deferred or inconclusive feedback, investors have the propensity to be more overconfident.

This setting addresses these two issues. Firstly, there is regular feedback, which should allow for more expedient (biased) learning. Secondly, the dataset incorporates bets where the underlying is a financial product but also one in which the underlying is based on a random-number generator. This allows the testing of hypotheses in a setting in which rational Bayesian learning is possible and confidence is warranted, and a setting in which positive returns are purely due to happenstance. In the Virtual Market, I argue, there can be no Bayesian learning, only naive reinforcement learning. Moreover, any overconfidence we note in this setting will be unwarranted.

Finally, rather than selecting a proxy for overconfidence, I use a particular feature of the data in order to classify agents as overconfident. I analyse stake size changes, classifying those who are initially successful and increase their stake sizes as confident, and those who decrease them or quit as underconfident. Initially, we could expect some agents to place explorative stake sizes in order to 'dip their toes the water'. However, an increase in stake size in subsequent bets not only places more capital at risk, but also increase the variance of the bet and is an indicator of confidence.[1] While it may be consistent with rational Bayesian updating for those who bet on the Financial markets and are successful to increase their stake sizes, the same cannot be said for the simulated market. As outcomes are independent and driven by a pseudo-random number generator, I argue that stake size increases for successful agents in the simulated market is indicative of overconfidence.

The papers discussed in this section investigate how past performance affects future behaviour. They find evidence of two types of learning: 'learning about ability' and 'learning by doing'. Some of the evidence uncovered could be argued to be consistent with both the rational Bayesian learning hypothesis and naive reinforcement learning. However useful techniques proposed in the Seru et al. (2010) paper suggest a way to disentangle both types and account for the learning that takes place as lower-ability agents naturally leave the sample. The following section motivates the main hypotheses.

5.3 Hypothesis Development

In this section, as the set of bettors is heterogeneous, I proceed by dividing the sample into three distinct subsets: (i) those who show a preference for betting only on Financial Markets, (ii) those who show a preference for betting on the Virtual Market and (iii) those who bet on both. The initial analysis will focus on the first two groups, however, all bettors will be included in later analysis. I differentiate between the financial market setting and the simulated market, establishing boundaries for what type of learning can take place and what type cannot. In order for the results to have broader implications, I first establish that the first treatment in this setting is indeed synonymous with conventional financial markets. I motivate the three main hypotheses with prior literature and learning theories, and posit predictions for expected behaviour on this basis.

[1] Statman et al. (2006) use the magnitude of trading volume as a proxy in a similar way.

Treatment A: Financial markets

While betting on the financial markets, agents cannot strictly observe each other's behaviour, but herding and information cascades are possible as bettors aggregate signals from various sources (including their own private signals). While rational learning can manifest itself in a number of ways in this setting[1], we could expect that agents who have experience positive returns will increase their bet sizes, while those who have experienced negative returns will exit the market.

While we could expect to see losing bettors in this market as a result, those who are rational may infer that they are likely to be unskilled and leave the dataset. This type of behaviour is called 'learning about ability' by Seru et al. (2010). If we see changes in stake size in Financial Markets, this could be consistent with either rational or reinforcement learning. Investors could be becoming confident by increasing their bet sizes, however this may be completely rational as there is an element of skill involved in this setting.

Treatment B: Simulated market

We cannot therefore disentangle rational Bayesian stake size changes from naive reinforcement learning in the Financial market setting, as behaviour could be driven by either effect. In the simulated market, however, there is no such ambiguity. Again, as agents cannot observe each other's behaviour, there can be no rational belief-learning. There is also full information about foregone payoffs, as bettors know after each losing bet how much they would have won had the counterfactual taken place. As market prices are based on a random-number generator, there can be no private information and outcomes are uncorrelated. As a result, there is no herding or information cascades, no tax-loss selling and no observation or imitation, only private signals and random outcomes. As it is not possible to 'learn about ability' or 'learn by trading', therefore, if we see changes in stake size in the simulated market, this can be consistent only with reinforcement learning. Changing bet sizes as a result of feedback from the Virtual Market therefore constitutes an 'emotional response to random rewards'. In that sense, any increase in stake sizes for successful traders in this setting is evidence of overconfidence which is both irrational and unwarranted.

Reinforcement learning theory, or the 'law of effect', dictates that agents will repeat behaviour that has been associated with positive feedback and avoid behaviour that has resulted in negative feedback. It dictates that agents should

[1] Bettors could be merely repeating actions which have resulted in positive feedback and avoiding behaviour which has resulted in negative feedback or could be rationally increasing and decreasing their stake sizes in a Bayesian manner (Chiang et al., 2011).

stick to given choices as long as they generate rewards, otherwise they should switch (Roth and Erev, 1995). Rational learning, however, incorporates both private signals and public information, updating beliefs about payoffs accordingly. For example, Bayesian learning refers to weighing both ‘experienced’ and ‘observed’ outcomes equally whereas reinforcement learning over-weighs ‘experienced’ outcomes. In contrast with a pure ‘stay/switch’ reinforcement model, Bayesian belief-learners rationally learn from experience (Camerer and Ho, 1999). I argue that agents who have been subject to successive strings of winning or losing bets have been subject to strong reinforcement cues, while those who have had mixed results have not.

In the model presented in Gervais and Odean (2001), traders are initially unaware about their true ability. As they trade, they overweight their own investing prowess and ‘learn’ to be overconfident. Gervais and Odean (2001) suggest that agents assess their abilities by observing their successes and failures rather than by rational introspection and that this biased process leads to overconfidence. Overconfidence peaks in the initial stages of a trader’s career in their model since he or she is inexperienced and successful, but is tempered as more experience is gained.

It is a well-documented stylised fact that only a very small fraction of individual traders are successful (Barber and Odean, 2000). Therefore, after a certain number of trades or period of time, traders who have had negative returns will infer that they may be part of the majority, stop trading and leave the sample. Indeed, Seru et al. (2010) show that learning about ability is more important than learning by doing and investors with poor performance are more likely to cease trading. I suggest that those who are successful during this initial period may infer that they do indeed have ability and will continue to trade. I argue that those who are successful and continue to trade, increasing their bet stake sizes, are becoming more confident. Those who have negative returns in the initial phase and quit are exhibiting underconfidence.

At the offset, investors do not know their ability. If it is variable, it can be improved with experience (Seru et al., 2010). If it is a constant, the level of which agents can only discover by trading, by transacting on the financial market, agents can ascertain their level of skill after a certain number of bets. Those who have experienced positive returns may increase their bet sizes, attributing success to their own skill. It is at this point that learning can be biased. If this success continues, the increase in stake size was rational and a result of Bayesian learning i.e. placing the appropriate weights on all information, including one’s own skill. However, if subsequent returns are negative, such stake size increases were a result of overconfidence due to naive reinforcement learning, with too much weight being attributed to one's own skill. In the financial market, we have no ex-ante way to disentangle these two explanations. However, the presence of the virtual market setting allows this.

Positive returns in the simulated market setting can come about as a result of happenstance. As there is no skill involved, attributing success to one's trading prowess and increasing stake sizes as a result of such overconfidence is irrational and constitutes overconfidence as a result of naive reinforcement learning. If trading on the Virtual Market is essentially an emotional response to random outcomes, then trading on the financial market types should be motivated by one of two types of learning: learning about ability or learning about trading. Rational learning theory, therefore, predicts changes in bet sizes in financial markets but no path-dependent behaviour in the simulated market.

Overconfidence, as commonly defined in the literature, is a wealth-reducing characteristic. If overconfidence is indeed defined as a bias, it is necessary to show that ex-post, such behaviour ultimately results in negative returns. Barber and Odean (2000), for example, show that overconfident investors (as proxied by gender) trade too much and experience negative returns as a result. While there are no demographics in this data set and hence no such proxy, I propose a novel way in which to disentangle confidence from overconfidence. The virtual market setting offers such a possibility.

Increased stake sizes for winners on the Financial markets is consistent with rational Bayesian updating. Increased stake sizes on the simulated market, however, can only be consistent with naive reinforcement learning. Increased stake sizes for winners on the Financial markets is consistent with rational Bayesian updating and is in effect 'warranted confidence'. Increased stake sizes for winners betting on the simulated markets is 'unwarranted confidence'. Thus, the magnitude of stake sizes changes in each setting allows the comparison of the level of confidence in each setting. I therefore analyse the levels of confidence in a setting in which there is an element of skill, and one in which results are driven purely by chance. While we cannot say, ex-ante, that increased stake sizes on the financial markets are consistent with overconfidence, I suggest that such behaviour in the simulated market would be indicative of overconfidence.

I test for the existence of such behaviour with the following two complementary hypotheses.

> **H1.** *A string of wins/losses will lead to an increase/decrease in the stake size in the Financial market*

> **H2.** *A string of wins/losses will not affect the stake size in the simulated market.*

In addition, I perform a second split in the data. We can examine the magnitude of stake size changes for agents who have experienced successive wins and losses in comparison with those who have had mixed experiences of wins and losses. I suggest that strong positive reinforcement cues, defined as successive winning bets, will induce higher levels of overconfidence than mixed experiences of wins and losses. We therefore expect that those agents with the highest numbers of successive wins (losses) to increase (decrease) their stake sizes and remain in (exit from) the market. Thus, the length of the string of results is crucial. I suggest that those who have been subject to strong positive reinforcement will become more overconfident, increasing their stake sizes as a result. Equally, those who have experienced successive negative returns should decrease their stake sizes or exit the market entirely. We could therefore expect three and four losses or wins in a row to have a much stronger effect on behaviour than a single loss or gain. in effect, the longer the string of successive negative or positive feedback, the more salient the effect on the stake size we expect.

Thus, the final hypothesis is as follows.

> **H3.** *The magnitude of change in the stake size is positively related to the length of the string of wins/losses.*

In summary, learning theory predicts path-dependent behaviour in the financial markets but reinforcement learning only in the simulated market. If we see path dependent behaviour in the virtual market setting, it is consistent with irrational overconfidence. Having motivated the hypotheses, I now present the methodological approach.

5.3.1 Methodology

In order to proceed with hypothesis testing, I first undertake a univariate analysis. Initially, I look at behaviour for all 44,041 bettors in aggregate and test the first two hypotheses that reinforcement learning induces confidence among bettors[1] To address the hypotheses, I divide the sample into three subsets: (i) 30,662 bettors who show a preference for betting only on Financial Markets, (ii) 11,389 who

[1]From the entire sample of bettors, I excluded any bettor who had transacted on the stock lottery product i.e. betting on the last digit of the closing price of the FTSE index. This is a standalone product marketed on a different part of the website and therefore constituted a separate undertaking from financial or simulated market betting. A total of 50 bettors (with a combined total of 154 bets) were excluded.

show a preference for betting on the simulated market and (iii) the 1,990 bettors who executed bets on both.

If agents are rational, there should be no path dependent change in stake size for those who bet on the simulated market. As returns are due to luck alone, we should see neither evidence of overconfidence by winners nor underconfidence by losers. Those who bet on Financial markets and are successful during this initial period may infer that they have ability, and will continue to trade. I suggest that such agents become confident, attributing success to their own skill, in effect applying the appropriate weight to their own signals using Bayesian updating. Those who are unsuccessful and do not exit from the sample (i.e. do not learn about ability), and those who change their bet sizes based on random rewards may have been subject to naïve reinforcement learning

To examine changes in bettor behaviour across different groups of bettors, I adopt a similar approach to Strahilevitz et al. (2011), using decision trees to present the evolution of mean stake over the first five bets[1]. I first standardised all bet sizes, by dividing by the first bet. Stake size changes were then calculated out to the fifth bet. Therefore, for each customer, initial bet sizes have been standardized to one, with subsequent bets expressed as a proportion of the initial bet. The intuition here is that groups of customers who are not changing their behaviour should have a standardized mean bet size of one after a certain number of bets, while those who do change their behaviour will have bet sizes greater than or less than one.

5.4 Results

Figure 5.1 shows the number of bettors playing at each round of betting, the mean stake change for survivors and an attrition-adjusted mean, the number of winners and losers and the subsequent attrition rates at each node for the aggregate group. It is clear that bettors are not exiting the market randomly resulting in endogenous attrition from the sample which may introduce bias and reduce power (Goodman and Blum, 1996). In addition, agents with the worst performance are likely those first to exit the market before we can analyse the effect their behaviour has on the whole cohort of bettors. In light of this, it would be insufficient to report mean stake sizes changes without reference to the reduction in observations with each subsequent round of betting or any adjustment to the mean in each

[1]The reason I focus on the first five bets is that the literature predicts that the most salient learning experience is the initial one. In addition, the median number of bets placed in the sample is three. Although it is not possible to present results in tree format for a large number of bets, I do incorporate further rounds of betting later in the analysis.

Figure 5.1

Evolution of betting: All bettors

This sample contains all bettors, irrespective of bet price (N=44,041). The initial stake size for each bettor has been standardised to one and subsequent stake size changes calculated as a proportion of the first bet. Stake size changes have then been winsorized at 1%. The number of bettors participating at each round, the mean stake size change, an attrition-adjusted mean (including bettors who exited the market at the previous bet) and the number of winners and losers are presented. The rate of attrition from the previous bet is displayed in parentheses above each node. A standard one-sample *t-test* with $H_0 : \mu_0 = 1$ is performed on the survivor-only mean and the attrition-adjusted mean, with asterix at standard significance levels.

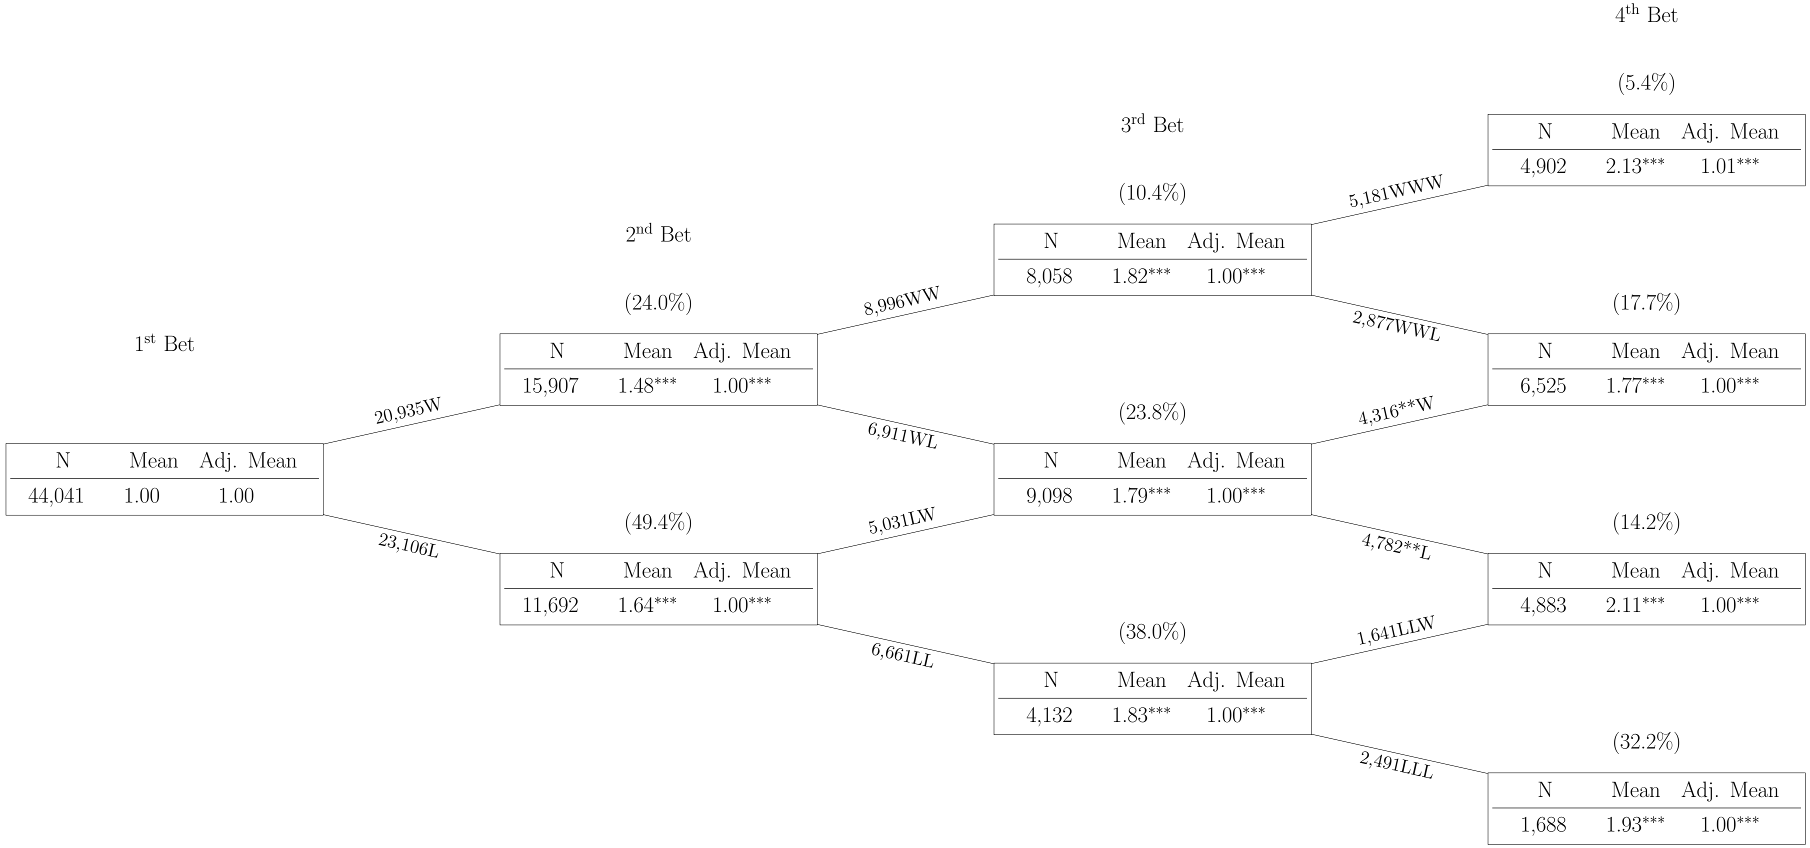

cohort to account for the reduced number of bettors. Seru et al. (2010) state that survivorship bias may give the appearance of learning when in fact, no learning is taking place. To address the issue of endogenous attrition, the attrition-adjusted mean is calculated on the basis that bettors who exited the market are included for a further bet with a stake size of zero.

As can be seen from the attrition rates in parentheses above each node, agents in the losing domain exhibit higher attrition rates than those in the winning domain. When comparing attrition rates between the upper and lower domains, attrition rates are higher at each round of betting. In addition, at any given round, there is a monotonic increase in attrition rates from top to bottom, indicating that the higher the instances of successive negative feedback, the higher the attrition rate. As regards stake sizes, however, there is an increase in stake sizes with each subsequent bet, irrespective of domain. In particular, there is no decrease in stake size beyond the level of the first bet in the losing domain.

In Figures 5.2 and 5.3, I have restricted membership to each of the two main groups of financial and simulated bettors. I again standardized each bettor's initial stake size to one and examine the subsequent changes in stake size at each round of betting with a one-sample t-test , however, I do present both parametric and non-parametric tests in a later analysis. At each node, I first show statistics for survivors and the attrition-adjusted sample: the level of attrition in parentheses, the count of bettors in each group, the standardized mean stake (set to one for the initial bet), the attrition-adjusted mean, asterisks indicating the results of a one-sample mean test testing for a statistically significant difference in mean stake size from one, and also the count of losers and winners at that round. Asterisks on the mean tests indicate significance at the standard levels.

As regards attrition rates, there is a difference between behaviour in the financial market and the simulated market. Attrition rates in both the losing and winning domains of the simulated market are higher with each subsequent bet than the financial market. As regards stake size changes, there does not seem to be a marked difference between the two settings as of the fourth bet. However, when the attrition adjusted mean stake sizes are compared between the two settings, the cumulative effect of higher attrition rates at each node in the simulated market results in lower mean stake sizes when attrition is adjusted for.

The decision trees presented in Figures 5.1, 5.2 and 5.3 examine the evolution of betting for all bettors across the first four bets. However, we are mainly interested in agents who have been subject to strong positive of negative reinforcement. Thus, I focus only on agents who have had consecutive strings of wins and losses. Path dependent behaviour is likely to be most salient for those bet-

Figure 5.2

Evolution of betting: Financial Markets

This sample is restricted to those who bet only on Financial markets in the first four bets, irrespective of bet price (N=24,624). The initial stake size for each bettor has been standardised to one and subsequent stake size changes calculated as a proportion of the first bet at each round of betting. Stake size changes have then been winsorized at 1%. At each node, the number of bettors participating at that round, the mean stake size change, an attrition-adjusted mean (including bettors who exited the market at the previous bet) and the number of winners and losers are presented. The rate of attrition from the previous bet is displayed in parentheses above each node. A standard one-sample *t-test* with $H_0 : \mu_0 = 1$ is performed on the survivor-only mean and the attrition-adjusted mean, with asterix at standard significance levels.

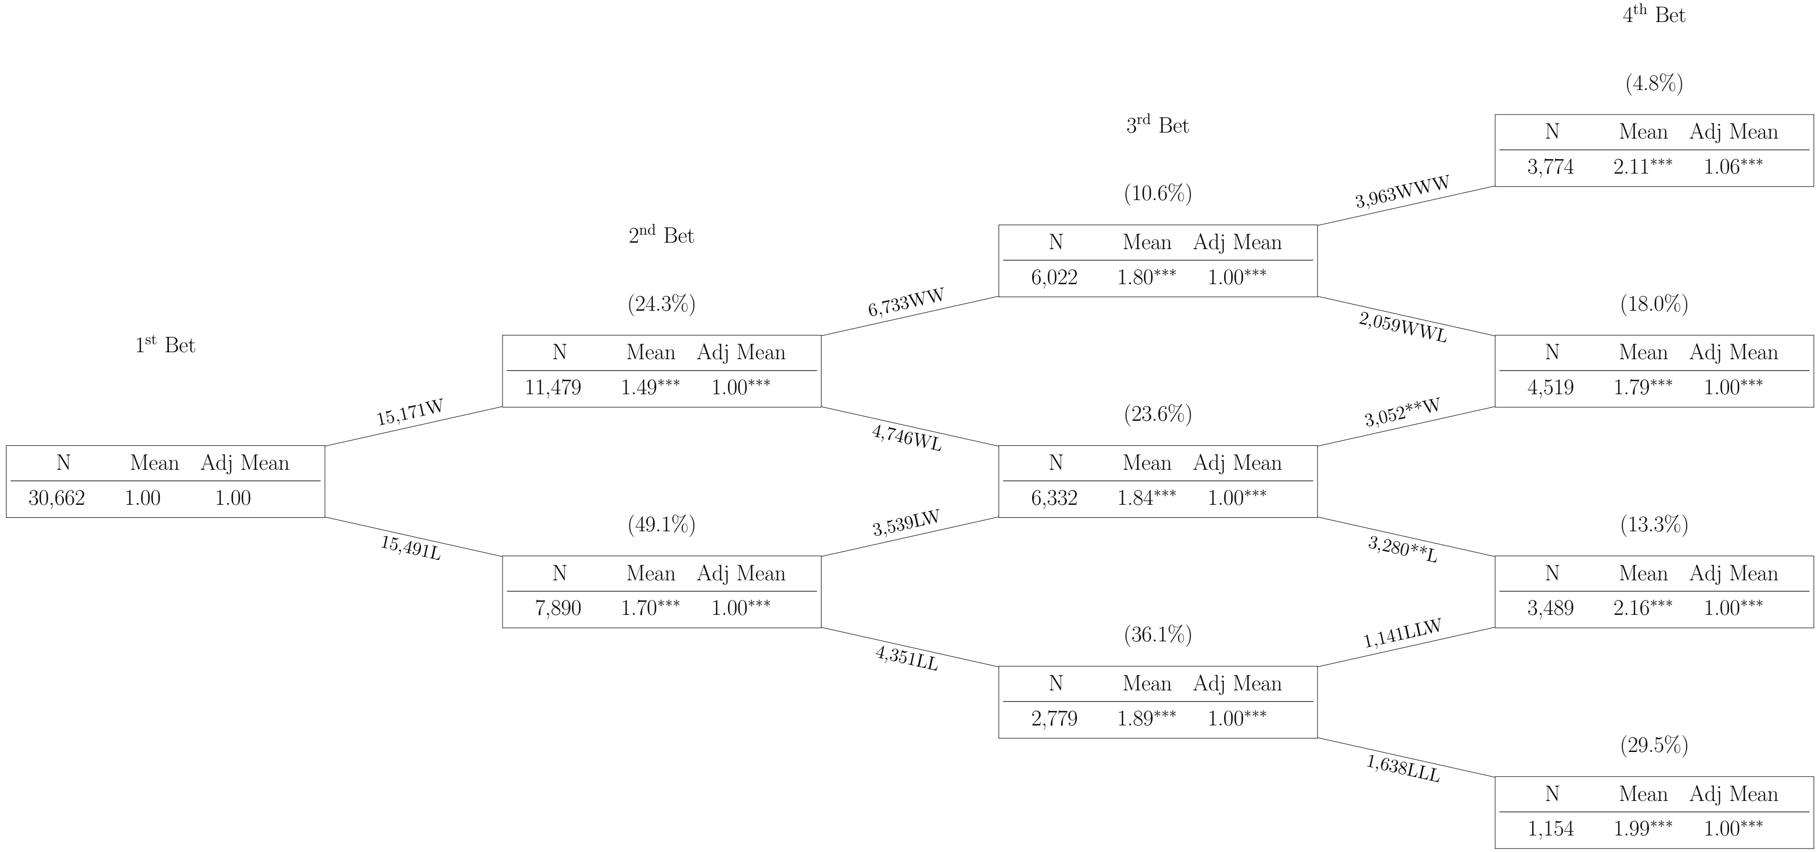

Figure 5.3

Evolution of betting: Simulated Market

This sample is restricted to those who bet only on the simulated market in the first four bets, irrespective of bet price (N=8,982). The initial stake size for each bettor has been standardised to one and subsequent stake size changes calculated as a proportion of the first bet at each round of betting. Stake size changes have then been winsorized at 1%. At each node, the number of bettors participating at that round, the mean stake size change, an attrition-adjusted mean (including bettors who exited the market at the previous bet) and the number of winners and losers are presented. The rate of attrition from the previous bet is displayed in parentheses above each node. A standard one-sample *t-test* with $H_0 : \mu_0 = 1$ is performed on the survivor-only mean and the attrition-adjusted mean, with asterix at standard significance levels.

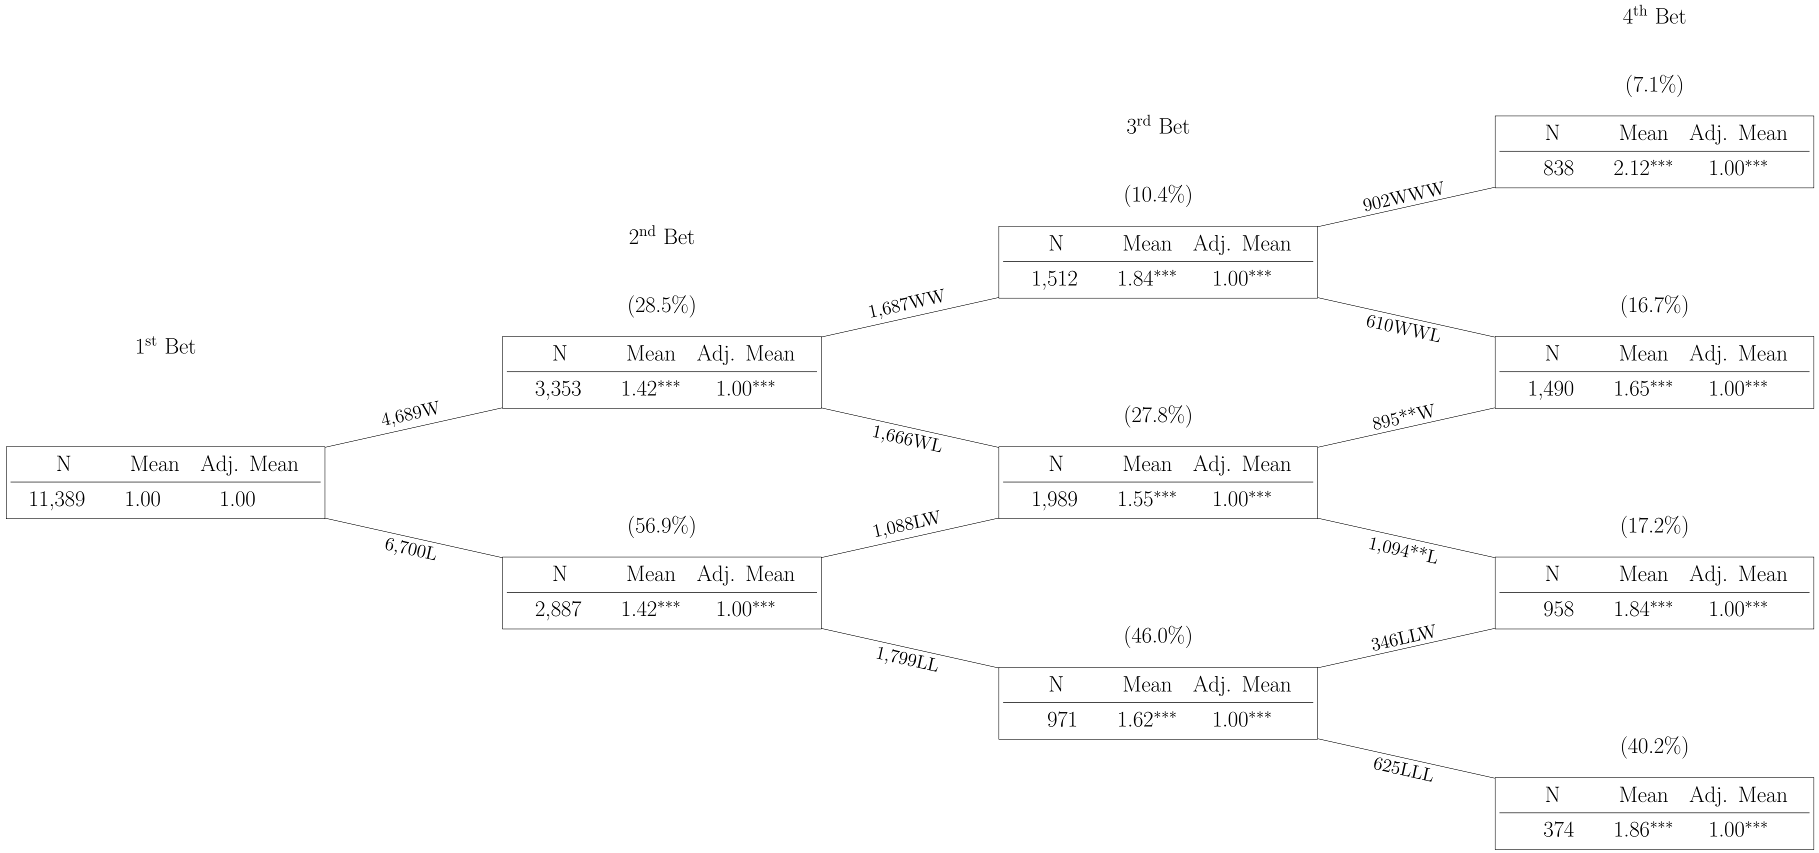

tors who have experienced a series of wins or losses, rather than those who have had mixed results. Therefore, I collate the results presented in Figures 5.1, 5.2 and 5.3 in Table 5.1, examining only on bettors in the pure losing and winnings domains i.e. those with strings of consecutive wins and losses. I also extend the analysis to include the fifth bet, include the median in the analysis and perform non-parametric tests as well as the one-sample t-tests presented earlier. Tabulating the results in this manner will allow the presentation of evidence to address the first two hypotheses.

5.4.1 Univariate results

The first hypothesis suggests that in aggregate, consecutive wins will lead to overconfidence and higher stake sizes, while consecutive losses will lead to lower stake sizes. Panel A of Table 5.1 shows standardised stake size changes and attrition for all bettors across the first five bets for losses and wins. For wins, there has been a monotonic increase in both mean and median stake size at each round from the first to the fifth bet and all of the mean stake sizes changes are significantly different from one. For losses, however, mean stake sizes have also increased. The results in the domain of wins offer some support for the first hypothesis that strings of wins induces confidence, but the domain of losses does not. Where we do see clear evidence of underconfidence in the domain of losses is in the rates of attrition.

After a single losing bet, 49.4% of bettors exit the market as opposed to an attrition rate of only 24% for winners. The levels of attrition stay high in the losing domain across the first five bets, but drop off significantly for winners with each round of betting. In terms of attrition, a very high number of bettors subject to strong negative reinforcement exit the market, but when we analyse mean stake sizes at each node, those who remain do not significantly decrease their stake sizes[1]. The results may constitute an empirical test of the reinforcement learning hypothesis as in Kaustia and Knüpfer (2008), but we must exercise some caution in drawing conclusions from the mean at each node alone, however, as there is no change in median stake sizes for losers, while the median bet size in the winning domain more than doubles over the course of the first five bets.

The attrition rates we observe are consistent with prior studies such as Chiang et al. (2011), who show that individual investors are less likely to bid in future

[1]While I can only account for voluntary attrition and must leave inconclusive any explanations based on involuntary attrition (i.e. lack of funds to continue, technical issues with the trading platform, external reasons for stopping betting), it is evident that the rates of attrition are much higher after losing bets, as expected.

IPOs given low returns in previous auctions. However, Chiang et al. (2011) do not shed light on the behaviour of those who have experienced low returns and continue regardless, which this result does. We therefore find support for the first hypothesis in the winning domain in terms of stake sizes and attrition rates, but underconfidence only manifests itself in terms of attrition rates in the losing domain. While there is no change in the median stake for losers, as a cohort, the mean stake size for losing bettors who do not exit the market does not decrease. This trade-off between attrition and bet sizing is similar to the result observed in Di Guisto et al. (2013), where learning in the losing domain manifests itself in exit from the market, rather than a rational Bayesian reduction in stake size.

I now address the second hypothesis. Using the methodology outlined above, the null hypotheses suggests that mean standardized stake sizes in the simulated market should be equal to one and path-independent. Mean standardized stake sizes for financial markets are expected to be different from one, indicating either reinforcement learning or rational learning, and path-dependent.

Panel B of Table 5.1 shows stake size changes and attrition rates for bets on the Financial Market and the simulated market. By the fifth bet, successful Financial market bettors have increased their stake sizes by a factor of 2.45, however mean stake sizes in the losing domain have doubled. The picture in the winning domain is quite clear, with monotonic increases in both mean and median stake size changes with each subsequent bet. In the losing domain, there is no change in the median bet size however. As regards simulated market bettors, contrary to expectations, the mean and median standardised stake changes have also increased in the simulated market, both in the winning and losing domains. We therefore must reject hypothesis H2 with regard to simulated market bettors.

Clearly, increases in stake size in the winning domain in the financial market setting may be consistent with rational Bayesian updating of stake size, therefore we cannot conclude that such behaviour is evidence of reinforcement learning induced overconfidence. However, placing more stake at risk in the simulated market is evidence of reinforcement learning induced overconfidence.

It may be that some bettors do not differentiate between a real financial market and one where the underlying process is random. There is a precedent for such behaviour in the literature. Barberis, Shleifer, and Vishny (1998) state that the mistaken belief that random sequences exhibit patterns is a manifestation of the representativeness heuristic (Kahneman and Tversky, 1973). Bloomfield and Hales (2002) also suggest that agents failing to differentiate between random walks and real data has a strong psychological foundation.

As regards the third hypothesis, therefore, we observe a similar trade-off be-

Table 5.1

Strong positive/negative reinforcement

This sample includes bettors who have had successive strings of winning or losing bets, irrespective of bet price and reflects the upper and lower domains of the decision trees presented in Figures 5.1, 5.2 and 5.3. Panel A shows the results for winners and losers in aggregate, Panel B presents statistics for Financial Market bettors while Panel C displays results for simulated market bettors. At each round of betting, I present the mean (with a standard one-sample *t-test* with $H_0 : \mu_0 = 1$), 25th percentile, median, 75th percentile, the results of a Rank Sum test on the raw (non-standardized) mean difference in stake, the count of bettors and the attrition rate. Initial stakes have been standardised to unity, with stake sizes changes calculated at each round of betting and winsorized at 1%.

	Losses							*Wins*						
Bet#	Mean	25th	Median	75th	WMW	Count	Attrition	Mean	25th	Median	75th	WMW	Count	Attrition
							Panel A : All							
2	1.64***	0.98	1.00	1.50	-11.28***	11,692	49.4	1.48***	0.99	1.00	1.20	-6.03***	15,907	24.0
3	1.83***	0.75	1.00	2.00	-7.42***	4,132	38.0	1.82***	1.00	1.00	1.69	-13.39***	8,058	10.4
4	1.93***	0.75	1.00	2.00	-4.85***	1,688	32.2	2.13***	1.00	1.01	2.00	-17.48***	4,902	5.4
5	2.08***	0.63	1.00	2.00	-3.95***	805	24.4	2.51***	1.00	1.21	2.04	-19.27***	3,298	3.3
							Panel B : Financial Market							
2	1.70***	0.98	1.00	1.65	-10.49***	7,890	49.1	1.49***	0.98	1.00	1.25	-5.27***	11,479	24.3
3	1.89***	0.71	1.00	2.00	-8.30***	2,779	36.1	1.80***	0.99	1.00	1.67	-10.61***	6,022	10.6
4	1.99***	0.72	1.00	2.00	-6.01***	1,154	29.5	2.11***	1.00	1.06	2.00	-14.27***	3,774	4.8
5	2.14***	0.63	1.00	2.00	-4.62***	558	20.7	2.45***	1.00	1.21	2.00	-15.44***	2,604	3.2
							Panel C : Simulated Market							
2	1.42***	1.00	1.00	1.11	-4.79***	2,887	56.9	1.42***	1.00	1.00	1.07	-3.86***	3,353	28.5
3	1.62***	0.99	1.00	1.86	-1.02	971	46.0	1.84***	1.00	1.00	1.83	-7.51***	1,512	10.4
4	1.86***	0.99	1.00	1.99	0.21	374	40.2	2.12***	1.00	1.00	2.00	-7.91***	838	7.1
5	1.98***	0.60	1.00	1.99	0.30	174	32.3	2.71***	1.00	1.20	2.50	-9.02***	514	3.6

tween stake size and attrition for those who decide to keep playing or who exit. It seems as though a string of wins increases the stake size irrespective of market setting. When bettors encounter consecutive losing bets, the high rates of attrition indicate that they exit the market. Rates of attrition in the simulated market are higher for losers and remain high with each subsequent bet. However, those that remain increase their stake sizes even in the losing domain, albeit at a lower rate than winners.

There is lower attrition in the Financial Market setting in the losing domain, however this behaviour is not surprising, as bettors may perceive an internal locus of control in the Financial market setting. Pastor and Veronesi (2009) state that persistent trading in the face of losses can be consistent with rational learning. In effect, losses are the cost of 'learning about ability' before attrition (on the basis of perceived poor ability) or further 'learning by trading' (in the event of success). Seru et al. (2010) also show that learning about ability is more evident than learning by doing and that investors with poor performance are more likely to cease trading. Rational learning theory, therefore, predicts that bettors in the financial market will exhibit behaviour consistent with 'learning about ability' as agents trade in the face of persistent losses in order to update their beliefs about subjective ability.

In summary, I find evidence that positive feedback results in higher participation, but negative feedback does not induce lower participation in all bettors in the losing domain. Furthermore, I do not find the differences one would expect when differentiating between market type. However, the result in the simulated market is consistent with reinforcement learning fostering overconfidence. In the next section, I focus on stake sizes changes, attrition rates and introduce lifetime profitability as a measure of ex-post overconfidence. I also expand the sample under analysis in order to analyse the effect that longer successive strings of positive or negative feedback affect attrition and stake size changes, and in turn confidence.

5.4.2 Expanded analysis

Up to this point, I have restricted the analysis to the first five bets, as the literature predicts that overconfidence peaks early and is tempered by experience. Now, I relax this restriction and analyse behaviour out to further rounds of betting. I analyse attrition rates and stake size changes and also introduce some summary stats on lifetime profit for bettors as an ex-post measure of overconfidence. This analysis will allow testing of the third hypothesis on the power of reinforcement with each subsequent bet.

Table 5.2 shows the percentage of bettors who quit given a certain path of wins and losses. Although the columns and rows are numbered 0 to 10, the analysis runs from the first bet up to and including the twentieth bet, thus the cell in the lower right hand corner of the matrix represents a path of ten wins and ten losses. The first bet is represented in the first anti-diagonal from (1,0) to (0,1) comprising the 'L' and 'W' nodes, and indicates that 49.40% (24.02%) of bettors who experienced a loss (win) on the first bet quit[1]. After 5 (10) losses or wins, the percentages who quit are 25.73% (19.35%) and 3.04% (0.94%), respectively. It is clear from this table, therefore, that attrition rates for losers remain relatively high with each subsequent bet, whereas attrition levels for winners fall significantly. However, comparisons between bettors with differing numbers of wins and losses can be made on a row-by-row basis, column-by-column basis, on both sides of the diagonal and along the diagonal.

Firstly, comparing the cells to the left of the diagonal with those to the right, bettors with larger numbers of wins than losses exhibit lower attrition rates. The closer the cell to the bottom left corner of the table, the lower the attrition rate. Comparing across numbers of losses along the top row, attrition rates are relatively high initially, but taper off with each subsequent losing bet. We can also examine point differences between numbers of wins at losses by comparing the first column to the first row. For example, three wins in a row prompts 5.39% of bettors to quit, whereas the corresponding figures for losses is 32.24%. Finally, along the diagonal, it seems that attrition rates are falling with each subsequent bet, even when the number of losses and wins are identical. For example, five wins and five losses is associated with an attrition rate of 6.21%, which drops to 4.08% in the case of ten wins and ten losses. In short, attrition rates are high initially, but bettors are less likely to quit with each subsequent round of betting. This result is consistent with that of Chiang et al. (2011) in that positive feedback results in higher participation, while string negative reinforcement results in exit.

Table 5.3 examines attrition rates according to the number of wins and losses experienced by bettors, but allows a comparison across bettors with differing numbers of consecutive wins and losses in a row. There was no reference to rounds of betting in the previous table, however we can compare across the columns for each round of betting here.

The first column in this table shows the aggregate attrition rate at each round of betting, and includes all winning and losing bettors who left the market. This column shows the unconditional or benchmark attrition rate. Thereafter, the first

[1]The percentages in the first five anti-diagonals correspond to the attrition rates for the first five bets presented in the earlier analysis.

Table 5.2

Path-dependency: Attrition

This table presents the percentage of bettors who quit given a certain path of wins and losses. The first bet is represented in the first anti-diagonal from (1,0) to (0,1), indicating that 49.40% (24.02%) of bettors who experienced a loss (win) on the first bet quit. After 5 (10) losses or wins, the percentages who quit are 25.73% (19.35%) and 3.04% (0.94%), respectively.

	Number of Losses										
# Wins	0	1	2	3	4	5	6	7	8	9	10
0		49.40	37.97	32.24	24.41	25.73	20.72	18.46	17.14	6.82	19.35
	44,041	*23,106*	*6,661*	*2,491*	*1,065*	*517*	*251*	*130*	*70*	*44*	*31*
1	24.02	23.82	20.07	19.07	18.27	16.07	17.00	13.62	13.30	10.96	14.85
	20,935	*11,942*	*6,423*	*3,472*	*1,888*	*1,039*	*594*	*367*	*233*	*146*	*101*
2	10.43	12.78	13.38	13.47	11.19	13.88	15.51	13.17	8.47	12.07	9.46
	8,996	*7,193*	*5,306*	*3,586*	*2,315*	*1,484*	*877*	*562*	*366*	*232*	*148*
3	5.39	9.19	9.41	9.75	9.87	10.08	9.04	10.21	11.09	8.99	7.98
	5,181	*4,745*	*4,145*	*3,312*	*2,452*	*1,716*	*1,128*	*754*	*541*	*367*	*263*
4	3.28	7.65	6.99	7.56	8.19	8.15	9.65	7.62	8.56	8.20	3.55
	3,410	*3,178*	*3,061*	*2,803*	*2,282*	*1,803*	*1,388*	*932*	*724*	*488*	*366*
5	3.04	6.37	6.12	5.07	6.50	6.21	6.97	7.64	6.40	7.32	6.89
	2,504	*2,197*	*2,224*	*2,287*	*2,046*	*1,675*	*1,406*	*1,047*	*812*	*615*	*450*
6	1.88	5.55	4.55	4.62	4.82	4.35	6.06	6.05	6.39	6.06	7.82
	1,973	*1,622*	*1,648*	*1,775*	*1,742*	*1,610*	*1,403*	*1,124*	*892*	*710*	*537*
7	1.64	5.47	4.59	3.47	4.34	4.33	4.88	4.55	5.00	5.94	4.11
	1,647	*1,279*	*1,241*	*1,355*	*1,406*	*1,385*	*1,291*	*1,120*	*940*	*741*	*584*
8	1.44	4.52	3.89	3.61	4.02	2.90	4.76	3.87	4.80	5.04	5.14
	1,392	*1,040*	*952*	*1,025*	*1,094*	*1,171*	*1,155*	*1,059*	*958*	*793*	*642*
9	1.41	3.18	4.36	2.19	2.06	2.55	3.52	2.34	3.61	4.05	4.32
	1,207	*911*	*757*	*775*	*827*	*942*	*995*	*984*	*914*	*814*	*718*
10	0.94	4.25	3.71	5.05	2.73	2.30	1.76	3.08	2.29	2.37	4.08
	1,066	*824*	*647*	*653*	*659*	*740*	*795*	*843*	*829*	*760*	*711*

entry in each column for losses (wins) corresponds to a cell in the topmost row (leftmost column) of Table 5.2 for the first five bets. However, each subsequent row presents the attrition rate if such a series of wins or losses was experienced in later rounds of betting. For example, after the first round of betting, 49.40% of bettors who experienced zero wins and a single loss up to that point (having placed one bet) did not place a second bet. However, after two bets, only 26.96% of bettors who experienced the same result in the last two bets (of a total of two bets) exited the market.

Comparing the unconditional attrition rate in the first column to the attrition rate for losers and winnings at each round of betting also yields interesting results. The unconditional rate of attrition at the fourth round is 15.45%, for example. However 32.24% of agents that experienced zero wins and three losses did not

Table 5.3

Strength of reinforcement: Attrition

The first column shows the percentage of bettors who quit at each round of betting, irrespective of their prior results. Thereafter, all figures in columns are percentages of bettors who quit given a certain number of wins and losses.

		# of Losses					# of Wins				
Bet#	Attrition	0W1L	0W2L	0W3L	0W4L	0W5L	1W0L	2W0L	3W0L	4W0L	5W0L
2	37.33	49.40					24.02				
3	22.87	26.96	37.97				19.50	10.43			
4	15.45	17.63	24.40	32.24			12.33	9.26	5.39		
5	12.11	14.78	20.53	21.76	24.41		9.57	7.63	5.44	3.28	
6	10.56	13.08	17.40	20.02	21.29	25.73	8.09	6.05	3.65	3.67	3.04
7	8.43	10.30	14.73	16.46	15.62	23.27	6.41	4.31	4.09	3.50	3.32
8	8.01	11.42	12.17	14.02	17.85	22.60	6.15	4.76	2.56	2.74	3.90
9	7.04	8.94	11.07	15.74	19.42	21.28	5.57	3.89	2.25	3.36	2.45
10	6.26	8.50	9.48	11.42	15.54	18.12	4.84	3.39	2.92	1.85	3.09

place a fourth bet, compared with 5.39% of agents experiencing three wins and no losing bets. The contrast between winners and losers and the unconditional rates of attrition become more pronounced with each round of betting, and by the tenth bet, 11.42% of agents that experienced zero wins and three losses (after their last winning bet) exit the market, and only 2.92% of agents with three consecutive winning bets (following a loss) quit, compared with an unconditional rate of attrition at the tenth round of 6.26%.

It seems that the effect of negative reinforcement is declining with each round of betting, which is an important result. Kaustia and Knüpfer (2008) also presented evidence consistent with investors avoiding repeating actions that had resulted in negative feedback, however I show that this effect is reduced with the passage of time. In addition, the same schedule and type of reinforcement does not have the same effect if it is experienced later rather than in the early stages. Reinforcement seems to generate overconfidence (and underconfidence in losing bettors) at a lower rate if it is not experienced at the start. When we focus on stake sizes and profit in the coming tables, the picture with regard to stake size and the number of strings and wins and losses becomes even clearer. We will find further support for the last hypothesis in the analysis that follows.

Panel A of Table 5.4 presents median stake size changes at each round of betting for all bettors across the first twenty bets. Such strings of wins and losses should facilitate addressing the third hypothesis. All stake sizes are standardised by dividing by the initial stake size, resulting in an initial standardised stake size of one and subsequent standardised stake sizes in terms of the first stake. At this point, I dispense with the survivor-only analysis in order to account for the amount of attrition at each round. Analysing median stake size changes for

Table 5.4

Path-dependency: Stake size changes

Panel A presents the median subsequent stake size change by bettors with a given number of wins and losses. Bettors who exit the market are shown placing a zero stake size in the next bet. The counts in each cell are identical to those in Table 5.2. Panel B shows the effect certain paths of wins and losses have on median stake sizes changes according to when such paths occur. All figures are median standardised stake size changes.

Panel A

	# of Losses										
# Wins	0	1	2	3	4	5	6	7	8	9	10
0	1.00	0.07	0.50	0.80	0.99	0.99	1.00	1.00	0.99	1.01	1.00
1	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
2	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
3	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
4	1.20	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
5	1.27	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
6	1.34	1.02	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
7	1.41	1.20	1.00	1.00	1.00	1.00	1.00	1.01	1.00	1.01	1.00
8	1.48	1.31	1.11	1.01	1.01	1.00	1.00	1.01	1.00	1.00	1.00
9	1.55	1.50	1.11	1.07	1.25	1.00	1.01	1.01	1.01	1.01	1.01
10	1.64	1.69	1.49	1.20	1.46	1.01	1.01	1.03	1.04	1.02	1.01

Panel B

		# of Losses					# of Wins				
Bet#	Stake	0W1L	0W2L	0W3L	0W4L	0W5L	1W0L	2W0L	3W0L	4W0L	5W0L
2	1.00	0.07					1.00				
3	0.80	0.99	0.50				1.00	1.00			
4	1.00	1.00	1.00	0.80			1.00	1.00	1.00		
5	1.00	1.00	1.00	1.00	0.99		1.00	1.00	1.00	1.20	
6	1.00	1.00	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.17	1.27
7	1.00	1.00	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.03	1.75
8	1.00	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.00	1.00	1.49
9	1.00	1.00	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.02	1.15
10	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.01	1.02

survivors only does not give a complete picture of what is happening in the losing domain.

In order to allow attrition to have an effect on each round of betting, I include dropouts from the previous bet which while facilitate an analysis of stake size changes for the entire cohort at each bet outcome. By doing this, we can analyse the joint behaviour of those who exited the market in addition to those who choose to continue. At each round of betting, bettors who quit at the previous bet are therefore represented as having a zero stake size. However, dropouts are carried on no further than the next bet. The median in the first five cells in the top row and the leftmost column are therefore equivalent to the median stake sizes presented in Table 5.1, but are adjusted for attrition at the previous bet. The counts in each cell are identical to those in Table 5.2 and are omitted.

The value in the cell (0,0) refers to the median initial bet size (standardised to one), the topmost row to the domain of losses and the leftmost column to the domain of wins. The outcome of the first bet is represented in the first anti-diagonal from (1,0) to (0,1), indicating that the median standardised stake change for bettors who experienced a loss (win) on the first bet was \$0.00 (\$1.00)[1]. The bottom left corner indicates that bettors who have experienced ten wins in a row increase their stake by a factor of 1.64. Each subsequent bet in the first column is associated with an increase in the median bet size. Across the first row, there is a sudden initial drop in bet size, however this slowly increases back to the size of the first bet with each subsequent loss, however does not exceed the value of the first bet.

I find some support for the third hypothesis here. The longer the string of wins, the higher the increase in stake size. However, this is clearly not the case in the losing domain. Chiang et al. (2011) found that as individual investors gain more experience, their ability decreases and they become more aggressive in terms of bid size. There may be similar behaviour here. Negative reinforcement has a strong initial effect on confidence, but with time, the effect of negative feedback diminishes, which is an interesting result. The behaviour we observe here may be driven by bettors' diminishing willingness to exit as they keep placing bets and committing capital, irrespective of the reinforcement schedule to which they are subjected.

Panel B of Table 5.4 shows the median standardised stake sizes changes associated with different paths of wins and losses. I analyse whether it is a bettor's initial experience that shapes future behaviour and in turn, whether later reinforcement is as powerful as initial reinforcement. All statistics are attrition-adjusted using dropouts at the previous bet.

The first column presents the median stake size change at each bet, irrespective of result, and is the unconditional or benchmark stake size placed by bettors. Subsequent columns show the effect of different strings of losses and wins as the occur at different rounds of betting. For example, the '0W2L' column indicates that a string of two losses has a large effect on stake sizes if this occurs after two bets, but has less effect if this occurs over three bets. In the domain of wins, a single win has the same effect on the median stake size irrespective of where it occurs and is not associated with any change in stake size. Although a string of five successive wins and no losses does have an effect on the median stake size, there is no monotonic pattern with each subsequent bet and the effect of such losses fades with each subsequent bet.

[1]The values in the first five anti-diagonals correspond to the median stake size in the first five bets presented in the earlier analysis. The zero median stake in the (1,0) indicates that the modal decision was to quit.

I find further evidence to support the third hypothesis here. In the domain of losses, it seems that the initial experience is quite salient, and later reinforcement is not as powerful as reinforcement that happens in the early stages. Increases in stake size are positively related to the length of the string of wins or losses.

We now turn to an analysis of lifetime profitability of bettors as an ex-post measure of overconfidence. Panel A of Table 5.5 outlines the median profit on exit by bettors experiencing differing paths of wins and losses. Firstly, across the first row, it is clear that the longer the string of losses, the lower the median profit of the lifetime of the account. However, the same is true of the first column. Given that there is a negative expected value from bets offered by a bookmaker, this is not surprising. As stated in Statman (2001) among others, stock trading and lottery buying are a negative sum game. Each additional win in a string of wins reduces profit, however, not to the same extent as in the losing domain. The median stake size for bettors experiencing five (ten) wins in a row in the first five bets is -\$11.64 (\$17.99).

Panel B of Table 5.5 outlines the effect strings of wins and losses have on median lifetime profit at each round of betting, according to when such consecutive strings occur. The '0W3L' columns indicates that if a bettor experiences a string of three losses in the first three bets, this is associated with a median profit of \$16.59. However, if three losses are experienced at any point beyond that, the median profit for such bettors is reduced. For winners, profit is also reduced if consecutive strings of wins are experienced in the later stages of betting, however the median profit amounts are not as negative. Interestingly, if a single win occurs any time after the second bet, it is associated with a monotonic decrease in median profit for bettors in the domain of wins.

The surprising results in the winning domain in the last two tables with regard to profit may not be counter-intuitive, given that there is a negative expected value in betting on the virtual market. We also know from Odean (1998a) that trading reduces individual traders' wealth. In addition, Gervais and Odean (2001) predict that overconfident traders increase their trading volume and ultimately lower their expected profits.

Taken together, the results of this expanded analysis offer support for the third hypothesis that the magnitude of stake size changes is related to the length of the string of wins, however the evidence in support of the same behaviour in the domain of losses was weaker. We did see evidence in the form of monotonically higher attrition rates in the domain of losses, however median stake sizes did not decrease monotonically with each subsequent loss. There was an initial shock to stake sizes and confidence after negative feedback, but after ten bets, stake sizes

Table 5.5

Lifetime profitability

Panel A shows median profit for bettors with certain paths of wins and losses. The analysis includes all bets up to and including the twentieth bet. The counts in each cell are identical to those in Table 5.2. Panel B shows the effect that certain strings of wins and losses have on median profit on exit at each round of betting. Each column concerns a certain number of wins and losses. Each row indicates at which point in the evolution of betting such a series of wins and losses was experienced. All figures are median profit amounts in USD.

Panel A

	# of Losses										
# Wins	0	1	2	3	4	5	6	7	8	9	10
0		-12.49	-27.48	-46.17	-65.35	-79.59	-101.90	-118.02	-141.89	-124.14	-128.52
1	-5.28	-23.24	-44.56	-61.50	-78.86	-98.29	-110.20	-122.90	-126.31	-118.28	-152.98
2	-23.78	-34.81	-56.51	-77.32	-92.18	-106.88	-120.04	-131.75	-138.63	-124.79	-176.59
3	-39.18	-49.93	-68.02	-96.10	-119.12	-135.14	-152.60	-155.16	-195.91	-172.27	-172.27
4	-50.65	-63.06	-83.85	-110.45	-124.83	-146.12	-176.67	-172.79	-206.85	-226.09	-229.12
5	-57.12	-74.49	-91.46	-127.86	-151.29	-164.47	-193.24	-200.06	-211.84	-211.00	-199.04
6	-62.93	-82.13	-109.62	-132.91	-161.62	-182.55	-197.18	-207.83	-239.03	-275.22	-296.16
7	-68.22	-87.29	-125.69	-159.64	-182.18	-210.89	-251.42	-256.88	-251.98	-276.27	-351.82
8	-71.65	-80.16	-119.62	-164.51	-213.63	-228.76	-272.98	-296.16	-296.74	-361.54	-374.71
9	-77.76	-87.93	-111.56	-190.66	-240.34	-248.10	-309.81	-333.66	-304.38	-359.84	-382.69
10	-79.72	-102.90	-98.83	-156.69	-234.54	-255.46	-310.93	-366.58	-331.93	-366.94	-383.64

Panel B

		# of Losses					# of Wins				
Bet#	Profit	0W1L	0W2L	0W3L	0W4L	0W5L	1W0L	2W0L	3W0L	4W0L	5W0L
2	-9.76	-12.49					-5.28				
3	-9.76	-19.44	-27.48				-28.28	-23.78			
4	-24.99	-35.12	-39.13	-46.17			-42.36	-47.55	-39.18		
5	-40.11	-50.74	-48.96	-59.82	-65.35		-59.31	-62.21	-64.41	-50.65	
6	-55.40	-67.67	-61.64	-64.22	-74.25	-79.59	-80.82	-80.48	-77.52	-71.51	-57.12
7	-68.75	-80.65	-79.31	-72.86	-83.80	-80.50	-98.25	-109.54	-103.68	-87.46	-75.31
8	-84.27	-102.93	-96.75	-103.92	-88.95	-80.54	-104.81	-112.57	-135.22	-111.33	-103.29
9	-98.91	-101.60	-127.45	-103.43	-104.37	-87.90	-130.83	-138.71	-138.00	-136.24	-143.69
10	-112.90	-128.24	-120.61	-138.14	-103.76	-130.22	-142.71	-159.68	-179.17	-138.00	-120.86

increased to the level of the first bet.

We observe mixed results depending on whether we analyse mean/median stake sizes or attrition at each node. Grinblatt and Keloharju (2009) examined sensation-seeking and overconfidence and also found mixed results depending on the use of the decision to continue, volume and turnover as the dependent variable. Up to this point, however, I have not controlled for bet price, which may provoke different behaviour in the domains of wins and losses.

It is clear that strings of successive strings of losses causes bettors to exit the market, which could be interpreted as underconfidence by such bettors. However, those that remain do not significantly reduce their stake sizes. Strings of winning bets cause much less bettors to exit and those that remain do increase their stake sizes. Moreover, as we saw in the last two tables, such bettors experience reduced profit with each subsequent winning bet, consistent with overconfidence.

In the following section, I continue the analysis in a multivariate setting where I can control for bet price as a proxy for risk preferences. I also incorporate more of the sample by extending the analysis in a panel regression and incorporate all bets by agents in the sample.

5.4.3 Multivariate results

In the previous section, I presented evidence in the form of path-dependent behaviour in terms of attrition, stake size and profitability. In Section 5.4.1, we also saw some evidence supporting the hypotheses in the form of stake size changes and attrition rates, however the behaviour in both the simulated market and the losing domains in both settings was not as we had expected. Even when returns were uncorrelated, bettors exhibited path-dependent behaviour. Moreover, losers did not significantly reduce their stake sizes and the net effect on mean stake sizes could be attributed to the high rates of attrition in the losing cohort. The previous analysis pooled all bet types and did not take account of differing bet prices and hence risk preferences, however. In this section, I turn to a multivariate analysis of stake size changes in the winning and losing domains for further support and use bet price as a control for risk.

Similar to the previous analysis, I first pool all bettors irrespective of betting product, and then reintroduce the two main groups from the previous analysis: those who bet on financial markets and those who bet on the simulated market. The reason for the reintroduction of this split across categories is that we can differentiate between the types of learning that can take place in both settings. In addition, I suggest that agents may perceive an internal locus of control in a financial market setting, a phenomenon which may not be present for those agents betting on the simulated market. In effect, losses may loom large for bettors who feel that their own decisions (in the face of the available information set) have resulted in poor returns. For those who understand that they are betting on uncorrelated outcomes, losses may not have the same saliency.

Initially, I focus on behaviour in the first few bets of a bettor's account as the number of bets per account is quite low, however the literature predicts that overconfidence will peak in the early stages and peter out with experience in any event (Menkhoff et al., 2006)[1]. I then introduce a panel regression which includes all bettors and all bets.

I first turn to the second hypothesis concerning behaviour subsequent to strings of wins and losses in the simulated market and financial markets. The

[1] I also perform further robustness tests with expanded rounds of betting in a later section.

intuition behind this section is as follows. If there is no change in behaviour as a result of strings of wins and losses, there should be no subsequent change in stake sizes for those who have experienced such strong reinforcement. Therefore, if I regress the first stake on each subsequent stake size for those who have been subject positive or negative reinforcement, and the coefficient on the first stake is not significant, I cannot reject the null hypothesis that wins and losses have no effect. In addition, as there is asymmetric attrition from the dataset, I account for this by estimating the regression model for survivors only and also on an attrition-adjusted sample. The attrition-adjusted sample includes those who exited the market at the previous round of betting. I therefore perform this test and split the dataset according to wins and losses, and also differentiate between bets on the Financial Market and simulated market.

As returns on the Financial markets are not uncorrelated, 'learning by doing' is possible and will manifest itself in both continued participation (lower attrition) and overconfidence (larger bet sizes) by bettors in the domain of wins. We still do not expect to find evidence of stake size changes in the simulated market as a result of strings of wins and losses. For the simulated market bettors, if the first bet is an unbiased estimator of each subsequent bet, a simple regression model with the current bet at each round as dependent and the first bet as independent should yield a constant of zero and a slope of one. If this is the case, it indicates no path-dependent behaviour in the simulated market group and therefore no effect by reinforcement on behaviour.

However, given the results presented earlier, we have reason to doubt that this will be the case in the losing domains of either setting. In the prior analysis, bettors with strings of losing bets did exit the market, but those that remained did not reduce their stake sizes. It may be that these bettors were indeed underconfident, however the threat of looming losses seemed to be more salient for these bettors than those that exited[1].

The section therefore formalizes the earlier tests presented in Table 5.1. I also control for risk preferences by including bet price as an independent variable. I firstly estimate a model with only a single independent variable and thereafter add further controls and interaction terms to examine bettor behaviour. The initial model is specified as follows:

$$ln(S_{it}) = \beta_0 + \beta_1 * ln(S_i(1)) + e_i \tag{5.1}$$

where $ln(S_{it})$ is a bettors current stake and $ln(S_i(1))$ is a bettors initial stake.

[1]This is a topic explored in further detail in the next chapter.

The expectation for the simulated market is $\beta_0 = 0$ and $\beta_1 = 1$ if $S_i(1)$ is an unbiased forecast of S_{it}. I first perform an initial test on winners and losers, and then introduce the financial market and simulated market groups, controlling for bet price. I do this for both a survivor-only and an attrition-adjusted sample.

Rather than incorporate all dropouts from the first bet, assigning bettors to nodes according to their inferred probabilities, I simply include customers who dropped out at the previous round of betting. Those that continued and transacted a bet in each round are included in the survivor sample. Those that dropped out at the previous node are included in the sample with a stake size of zero. The estimation results for losers are presented in Table 5.6, while the results for winners are presented in Table 5.7. For ease of exposition, I also plot the regression line of the two main estimation results in the previous two tables, for both winners and losers, in Figures 5.4 and 5.5.

Panel A of Table 5.6 shows the results for losing bettors for survivors only, while Panel B includes the attrition-adjusted sample. In Panel A, the initial model (1), shows the estimation for all losing bettors, irrespective of market type[1]. The main models of interest in this table are Models (4) and (7), those for losing financial market bettors and losing simulated market bettors. The results for winners are show in Table 5.7. Again, Panel A shows the results for winning bettors for survivors only, while Panel B includes the attrition-adjusted sample. In Panel A, the initial model (1), shows the estimation for all winning bettors, irrespective of market type[2]. The main models of interest in this table are Models (4) and (7), those for winning financial market bettors and winning simulated market bettors.

For losses, the slope and constant are no different in the financial and simulated market in models (2) and (5), however the coefficients on the bet price dummies are significant. When we add bet prices as controls in models (3) and (6), this relationship remains the same. The interaction terms in models (4) and (7) are overall not significant. For wins, both the coefficients on bet price, bet sequence and the interaction terms are significant in Financial markets, neither the interaction terms nor the constant are significant for simulated markets. While the constant is non-trivial to interpret on its own, as it indicates the stake size in each subsequent bet had the first bet stake been zero, it is instructive to examine the prediction for the full amount bet in subsequent bets, given the first bet size.

[1]The number of observations here includes losing bettors at each node after the first bet i.e. (2: 10,321), (3: 3,824), (4: 1,611) and (5: 772). The counts for the fifth bet are not visible in the decision trees presented earlier but are tabulated in Table 5.1. The total for the attrition-adjusted sample in the losing domain is 27,395 i.e. 18,805 + 5,885 + 2,268 + 437.

[2]The number of observations here refers to the number of bettors in the winning domain subsequent to the first bet, as follows: (2, 13,475), (3, 6,906), (4, 4,182) and (5, 2,832).

Table 5.6

Regression Estimation: Losses

The dependent variable is this model is the log of stake size, $ln(S_n)$, in USD. The sample includes only bettors in the losing domain who have had consecutive strings of losses. The model is estimated using the second and subsequent bets. Model 1 is estimated on the entire sample. Models 2 and 5 include financial market bettors and simulated market bettors, respectively and are estimated without controls. Models 3 and 6 include dummy variables for bet price with 5/1 bets as the base category. Models 4 and 7 include the bet price dummies, dummies for the second and subsequent bets and also interact the bet number dummies with the breakeven point. Model 8 is estimated on the pooled sample of losers with a dummy variable equal to zero for financial market bettors and 1 for simulated market bettors. Panel A is a survivor-only analysis, with Panel B including the attrition-adjusted sample. The number of observations (N=18,317) in Model (1) refers to the number of bettors in the losing domain subsequent to the first bet i.e. (L: 11,692), (LL: 4,132), (LLL: 1,688) and (LLLL: 805).

Panel A: Survivors only

	All	Financial markets			Simulated market			Pooled
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\ln(S_1)$	0.860***	0.856***	0.845***	0.850***	0.872***	0.855***	0.861***	0.852***
	(189.31)	(158.60)	(151.64)	(127.34)	(87.05)	(85.15)	(70.82)	(147.48)
Evens			-0.312***	-0.315***		-0.345***	-0.348***	-0.324***
			(-11.32)	(-11.44)		(-11.26)	(-11.33)	(-15.83)
5/6			0.110***	0.112***		-0.0278	-0.0257	0.0724***
			(4.86)	(4.95)		(-0.84)	(-0.78)	(3.91)
1/20			0.371***	0.382***		0.0863	0.0932	0.303***
			(6.51)	(6.71)		(1.15)	(1.24)	(6.50)
Bet=3				0.0757*			0.0497	0.0666**
				(2.39)			(1.40)	(2.75)
Bet=4				0.0644			0.0948	0.0731*
				(1.46)			(1.74)	(2.11)
Bet=5				0.173*			0.146	0.166**
				(2.48)			(1.78)	(3.06)
Bet=3 × $\ln(S_1)$				-0.0219			-0.0141	-0.0192
				(-1.65)			(-0.55)	(-1.69)
Bet=4 × $\ln(S_1)$				0.0171			0.0218	0.0158
				(0.92)			(0.67)	(1.00)
Bet=5 × $\ln(S_1)$				-0.0440			-0.0791	-0.0471
				(-1.44)			(-1.73)	(-1.82)
Fin/Vir								-0.0893***
								(-5.52)
Constant	0.276***	0.339***	0.318***	0.286***	0.163***	0.297***	0.270***	0.313***
	(28.99)	(26.29)	(15.72)	(12.96)	(11.16)	(11.23)	(9.68)	(16.68)
N	18,317	12,381	12,381	12,381	4,406	4,406	4,406	16,787
R2	0.701	0.701	0.710	0.710	0.704	0.714	0.715	0.722

Table 5.6: (continued)

Panel B: Attrition-adjusted sample								
	All	Financial markets			Simulated market			Pooled
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\ln(1+S_1)$	0.447***	0.462***	0.437***	0.555***	0.330***	0.310***	0.454***	0.535***
	(65.55)	(57.08)	(52.96)	(48.75)	(24.47)	(22.86)	(22.03)	(53.80)
5/1			-0.537***	-0.493***		-0.456***	-0.389***	-0.448***
			(-13.73)	(-12.92)		(-8.64)	(-7.90)	(-14.41)
5/6			-0.208***	-0.182***		-0.234***	-0.184***	-0.171***
			(-5.24)	(-4.70)		(-4.26)	(-3.58)	(-5.38)
Evens			-0.335***	-0.302***		-0.317***	-0.262***	-0.295***
			(-8.19)	(-7.56)		(-5.79)	(-5.13)	(-9.12)
Bet=3				-0.0144			0.0368	0.00507
				(-0.59)			(1.22)	(0.27)
Bet=4				0.0223			0.0804	0.0244
				(0.53)			(1.55)	(0.74)
Bet=5				0.125			0.0969	0.0919
				(1.86)			(1.18)	(1.76)
Bet=3 × $\ln(1+S_1)$				-0.275***			-0.300***	-0.284***
				(-16.41)			(-11.00)	(-19.73)
Bet=4 × $\ln(1+S_1)$				-0.162***			-0.249***	-0.174***
				(-5.80)			(-5.48)	(-7.24)
Bet=5 × $\ln(1+S_1)$				-0.105*			-0.169*	-0.103**
				(-2.49)			(-2.31)	(-2.82)
Fin/Vir								-0.0938***
								(-9.78)
Constant	-0.0335***	-0.0409***	0.308***	0.284***	0.00616	0.376***	0.297***	0.305***
	(-3.67)	(-3.43)	(7.87)	(7.17)	(0.42)	(7.04)	(5.65)	(9.32)
N	48,069	31,841	31,841	31,841	14,273	14,273	14,273	46,114
R2	0.195	0.196	0.207	0.253	0.131	0.144	0.205	0.256

Table 5.7

Regression Estimation: Wins

The dependent variable is this model is the log of stake size, $ln(S_n)$, in USD. The sample includes only bettors in the winning domain who have had consecutive strings of wins. The model is estimated using the second and subsequent bets. Model 1 is estimated on the entire sample. Models 2 and 5 include financial market bettors and simulated market bettors, respectively and are estimated without controls. Models 3 and 6 include dummy variables for bet price with 5/1 bets as the base category. Models 4 and 7 include the bet price dummies, dummies for the second and subsequent bets and also interact the bet number dummies with the first stake. Model 8 is estimated on the pooled sample of winners with a dummy variable equal to zero for financial market bettors and 1 for simulated market bettors. Panel A is a survivor-only analysis, with Panel B including the attrition-adjusted sample. The number of observations (N=32,165) in Model (1) refers to the number of bettors in the winning domain subsequent to the first bet i.e. (W, 15,907), (WW, 8,058), (WWW, 4,902) and (WWWW, 3,298).

Panel A: Survivors only

	All	Financial markets			Simulated market			Pooled
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\ln(S_1)$	0.849***	0.851***	0.847***	0.867***	0.847***	0.841***	0.866***	0.852***
	(264.67)	(227.76)	(228.90)	(178.62)	(110.52)	(110.63)	(93.80)	(147.48)
Evens			-0.262***	-0.251***		-0.192***	-0.169***	-0.324***
			(-7.26)	(-6.93)		(-5.22)	(-4.62)	(-15.83)
5/6			0.196***	0.184***		0.0706**	0.0630*	0.0724***
			(9.75)	(9.16)		(2.81)	(2.52)	(3.91)
1/20			0.469***	0.420***		0.322***	0.263***	0.303***
			(22.10)	(19.56)		(10.32)	(8.30)	(6.50)
Bet=3				0.102***			0.169***	0.0666**
				(4.53)			(5.14)	(2.75)
Bet=4				0.239***			0.260***	0.0731*
				(7.71)			(5.49)	(2.11)
Bet=5				0.327***			0.427***	0.166**
				(8.72)			(6.41)	(3.06)
Bet=3 × $\ln(S_1)$				-0.0238**			-0.0427*	-0.0192
				(-2.69)			(-2.35)	(-1.69)
Bet=4 × $\ln(S_1)$				-0.0553***			-0.0645**	0.0158
				(-4.88)			(-2.58)	(1.00)
Bet=5 × $\ln(S_1)$				-0.0585***			-0.0781*	-0.0471
				(-4.44)			(-2.51)	(-1.82)
Fin/Vir								-0.0893***
								(-5.52)
Constant	0.377***	0.396***	0.140***	0.0678***	0.332***	0.242***	0.149***	0.313***
	(48.61)	(40.73)	(7.49)	(3.42)	(23.58)	(11.25)	(6.49)	(16.68)
N	32,165	23,879	23,879	23,879	6,217	6,217	6,217	16,787
R2	0.724	0.725	0.737	0.739	0.716	0.726	0.730	0.722

Table 5.7: (continued)

Panel B: Attrition-adjusted sample								
	All	Financial markets			Simulated market			Pooled
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ln(1+S_1)	0.709***	0.711***	0.713***	0.727***	0.661***	0.661***	0.708***	0.535***
	(120.34)	(105.35)	(106.79)	(74.16)	(44.16)	(44.26)	(36.32)	(53.80)
5/1			-0.254***	-0.196***		-0.0328	0.0222	-0.448***
			(-9.89)	(-7.59)		(-1.02)	(0.69)	(-14.41)
5/6			-0.294***	-0.225***		-0.0587*	-0.0120	-0.171***
			(-19.42)	(-14.75)		(-2.06)	(-0.43)	(-5.38)
Evens			-0.395***	-0.321***		-0.224***	-0.165***	-0.295***
			(-18.49)	(-14.97)		(-7.30)	(-5.43)	(-9.12)
Bet=3				-0.0873**			-0.0119	0.00507
				(-3.12)			(-0.27)	(0.27)
Bet=4				0.103**			0.257***	0.0244
				(3.10)			(3.87)	(0.74)
Bet=5				0.309***			0.428***	0.0919
				(7.77)			(5.13)	(1.76)
Bet=3 × ln(1+S_1)				-0.111***			-0.169***	-0.284***
				(-6.73)			(-4.83)	(-19.73)
Bet=4 × ln(1+S_1)				0.0495**			-0.0277	-0.174***
				(2.89)			(-0.59)	(-7.24)
Bet=5 × ln(1+S_1)				0.0742***			0.0227	-0.103**
				(3.90)			(0.42)	(-2.82)
Fin/Vir								-0.0938***
								(-9.78)
Constant	0.145***	0.153***	0.341***	0.296***	0.116***	0.197***	0.120***	0.305***
	(14.90)	(12.80)	(22.13)	(14.91)	(5.90)	(7.27)	(3.79)	(9.32)
N	44,767	33,148	33,148	33,148	9,386	9,386	9,386	46,114
R2	0.402	0.396	0.405	0.430	0.350	0.354	0.384	0.256

I therefore plot the estimation results for the main models in each of the two main groups (Financial and simulated bets) for each bet price in Figure 5.4.

The dashed line in each plot indicates no subsequent change in stake. Across all bet prices and in both main groups of financial and simulated market bettors, there is a slight increase in stake sizes with each subsequent bet. In the case of 5/1 bets, stake sizes remain below the size of the first stake, but increase in the third, fourth and fifth bets relative to the base case of the second bet. For 1/20 bets, subsequent bets are higher than the first stake sizes and remain so with each subsequent round of betting.

In a general sense, the difference between the winning and losing domains is that there is a widening gap between the first stake and each subsequent stake. There is no significant difference between behaviour in the financial market and

Figure 5.4

Plots: Financial Markets

The following figures plot the coefficient estimations for Models 4 & 7 for financial bettors in the winning and losing domains presented in Tables 5.6 and 5.7. The x- and y-axes represents the first stake size and the next bet, respectively. The models estimates are plotted for Financial bettors, for each bet price and for the third and subsequent bets.

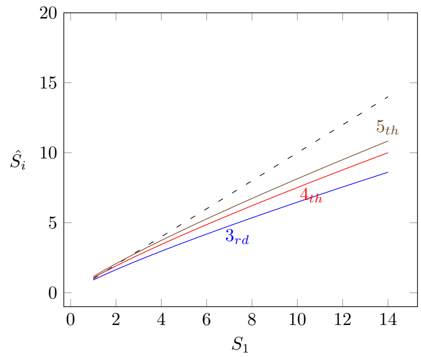

(a) Wins, Fins, 5/1 Bets

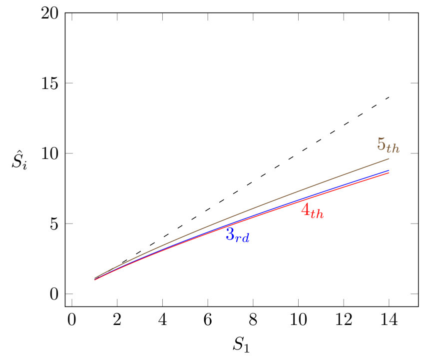

(b) Losses, Fins, 5/1 Bets

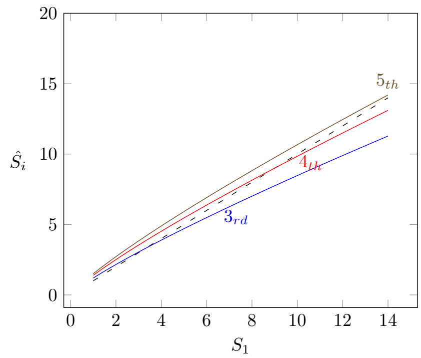

(c) Wins, Fins, Evens Bets

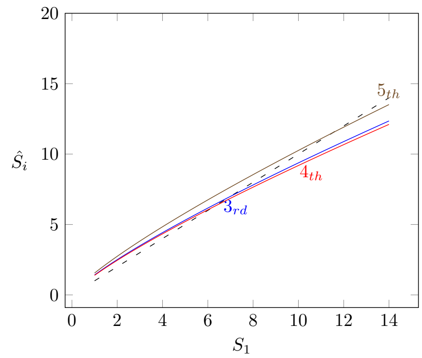

(d) Losses, Fins, Evens bets

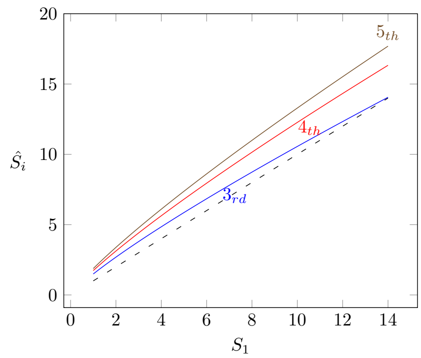

(e) Wins, Fins, 5/6 Bets

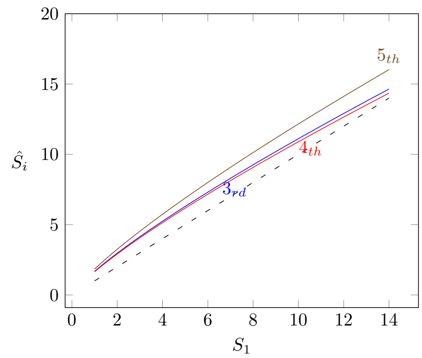

(f) Losses, Fins, 5/6 Bets

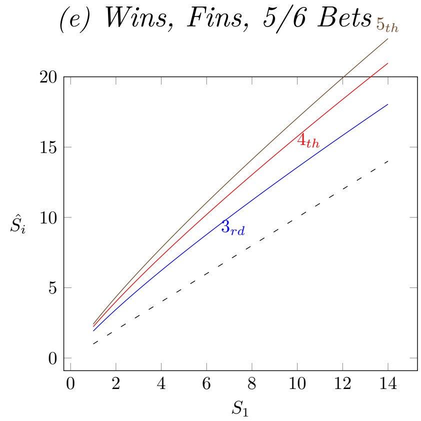

(g) Wins, Fins, 1/20 Bets

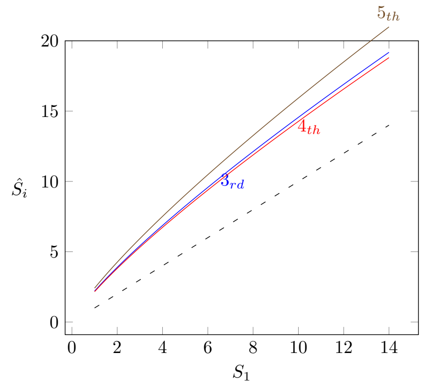

(h) Losses, Fins, 1/20 Bets

Figure 5.5

Plots: Simulated Market

The following figures plot the coefficient estimations for Models 4 & 7 for simulated market bettors in the winning and losing domains presented in Tables 5.6 and 5.7. The x- and y-axes represents the first stake size and the next bet, respectively. The models estimates are plotted for simulated market bettors, for each bet price and for the third and subsequent bets.

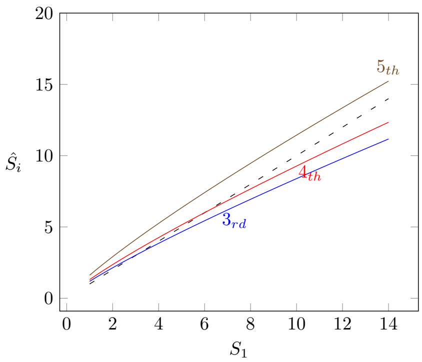

(a) Wins, Sim, 5/1 Bets

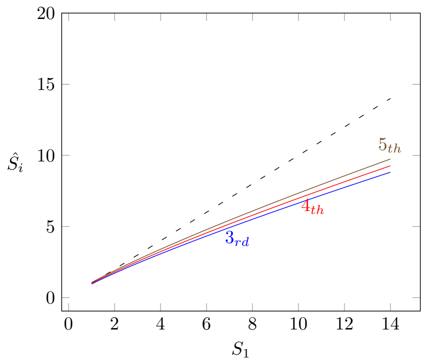

(b) Losses, Sim, 5/1 Bets

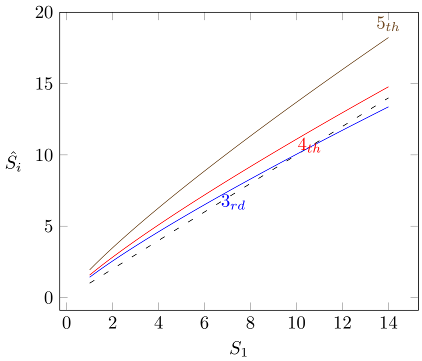

(c) Wins, Sim, Evens Bets

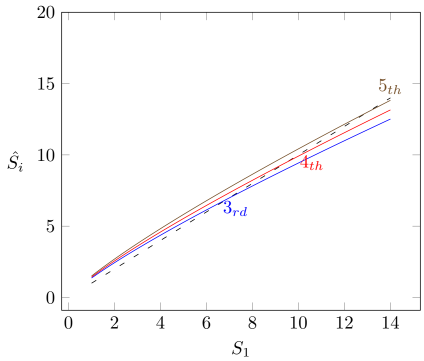

(d) Losses, Sim, Evens Bets

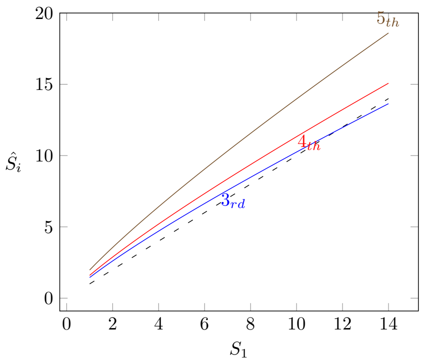

(e) Wins, Sim, 5/6 Bets

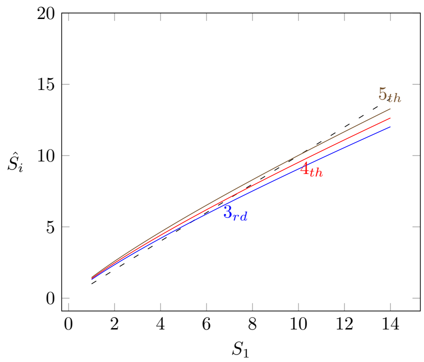

(f) Losses, Sim, 5/6 Bets

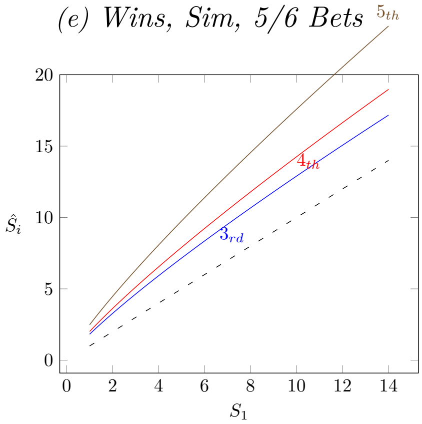

(g) Wins, Sim, 1/20 Bets

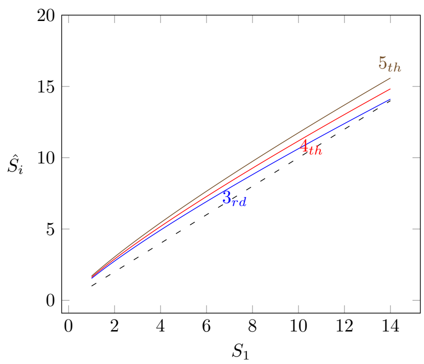

(h) Losses, Sim, 1/20 Bets

the simulated market. In Figure 5.4, it is clear that winners in both setting increase their stake sizes, however in the losing domains, the results are not as clear.

It may be that the decision of what magnitude to place on each subsequent bet in the losing domain is not dictated by the initial bet, but rather by the bettors current profit or loss. Although the concepts of 'learning about ability' and 'learning by doing' may be at play in the winning domain, the disposition to avoid losses may be what's driving the participation by losing bettors in the fourth bet[1]. In effect, losing bettors who learn in a rational Bayesian sense should be reducing the weight they apply to their own ability in the financial markets and decreasing their stakes accordingly or dropping out, but seem to increase their stake size regardless.

5.4.4 Robustness

We have seen that after strings of successive winnings bets, bettors increase their stake sizes. In the case of losing bets, there is also an increase but not to the same extent. In the plots presented earlier, winners increased their stakes at a higher rate with each subsequent bet, as evidenced by the increasing gap between the level of stake size with each subsequent bet. I now examine whether there is a tendency to increase bet size with each subsequent round of betting regardless, independent of having winning or losing strings of bets.

I present the results of further robustness tests in Table 5.8. Rather than focusing only on bettors who had successive strings or wins and losses, I expand the analysis again included bettors who have had mixed experiences of wins and losses. I restrict the analysis to the first five bets in order to include interaction terms on the third and subsequent bets, using the second bet as the base. Gervais and Odean (2001) also state that the greatest overconfidence in a trader's career is exhibited in the initial stages. In order to address endogenous attrition, I adjust for attrition from the previous bet only[2].

Instead of splitting the sample into winners and losers, I include *chain*, which is a score equal to zero if a bettor's initial bet is a losing bet and one otherwise. Thereafter, it is incremented with each winning bet and de-incremented with each losing bet. If a bettor has an equal amount of wins and losses, this value will be zero. It is therefore positive for those who have had more winning bets that losing bets and vice versa. A bettor who has had four winning bets in a row will

[1] I examine this premise in detail in the next chapter.

[2] The total number of observations, N=110,926 comprises 44,041, 27,599, 21,288 and 17,998 from the second and subsequent bets, respectively.

Table 5.8

Cross-sectional Regression

This sample contains the first five bets for all bettors. The number of observations in the first two models (N=110,926) includes the number of bettors at end round subsequent to the first bet and dropouts from the previous bet. Models 3 and 4 include Financial market bettors and simulated market bettors, respectively. Bettors who bet on both markets have been excluded. The model is estimated using the second and subsequent bets. *Chain* has a value of zero prior to the first bet and is incremented for winning bets and de-incremented for losing bets. The dependent variable is this model is the log of stake size, $ln(S_1)$, in USD.

	All	All	Fins	Sim
$\ln(S_1)$	0.640***	0.639***	0.646***	0.572***
	(99.96)	(100.28)	(86.92)	(38.84)
Bet sequence				
Bet=3	0.114***	0.108***	0.112***	0.094***
	(7.83)	(7.44)	(6.12)	(3.30)
Bet=4	0.233***	0.212***	0.234***	0.184***
	(14.69)	(13.44)	(11.67)	(5.91)
Bet=5	0.308***	0.272***	0.291***	0.264***
	(18.60)	(16.36)	(14.08)	(7.25)
Bet=3 × $\ln(S_1)$	0.080***	0.077***	0.083***	0.101***
	(8.71)	(8.35)	(7.82)	(4.42)
Bet=4 × $\ln(S_1)$	0.099***	0.093***	0.095***	0.133***
	(10.30)	(9.74)	(8.57)	(5.58)
Bet=5 × $\ln(S_1)$	0.105***	0.098***	0.099***	0.136***
	(10.82)	(10.09)	(9.04)	(5.04)
Bet price				
1/20	0.655***	0.455***	0.497***	0.291***
	(59.55)	(37.27)	(31.51)	(12.49)
5/6	0.372***	0.298***	0.326***	0.214***
	(43.52)	(34.19)	(27.11)	(15.04)
Evens	0.223***	0.172***	0.203***	0.127***
	(22.01)	(16.97)	(14.12)	(8.33)
Win/Loss Chain				
Chain		0.092***	0.090***	0.109***
		(38.73)	(30.42)	(25.66)
Constant	-0.258***	-0.175***	-0.221***	-0.088***
	(-24.85)	(-16.61)	(-15.48)	(-4.78)
Observations	110,926	110,926	78,100	25,761
R2	0.457	0.465	0.460	0.428

have a value of 4 in the fifth bet, for example.

The first model is very similar to the regression estimates presented earlier, but the sample is pooled across winners and losers and by product type. The second model includes the variable chain to allow for bettors with winning and losing bets and the coefficient on this variable is significant. Even after controlling for the variable chain, the coefficients on each subsequent bet are still significant. Thus, there is an increase in stake with each bet regardless of the feedback received as time passes. This is similar to the result found in Song et al. (2014), who found that overconfidence attenuates the effect of reinforcement learning. However, in this case it seems that overconfidence entices bettors to continue to bet in a simulated market in the face of losses. While a truly rational investor would never accept such a bet, we could expect even irrational bettors to exit the market after experiencing successive losses.

In models three and four, after pooling winners and losers and accounting for path dependency in result with the variable chain, the coefficient on bets three, four and five are still positive and significant, indicating that regardless of wins and losses, bettors put more stake at risk with each bet irrespective of product type. Indeed, the coefficient estimate for the variable chain on the virtual market is slightly larger than in the financial market, indicating that the result has a slightly higher effect in the simulated market.

To verify this, I present the results of a panel regression, and further expand the sample to include the first 100 bets for each bettor. Table 5.9 presents the estimation results for the main models in a panel regression setting with fixed effects. In this analysis, I do not restrict that sample to the first five bets, but rather include all bets for all bettors, including the initial bet.

The dataset contains features of both short panel and long panel types, in that there are many time periods and many individuals in the sample. Stake and cumulative loss are varying regressors, while the choice of Financial or Simulated market is time-invariant across agents. In the panel regression, I model the choice of next stake for all bettors, forecasting the next stake size using the lag of stake. If a bettor exits the market, the next bet is coded with a stake size of zero with the same characteristics as the previous bet (choice of financial or simulated market, bet price etc).

In the first pooled model in Panel A, the coefficient on chain is positive and significant. When run separately for financial market bettors and virtual market bettors, the coefficient is very similar. However when interacting chain with the market preference flag in the fourth model, the coefficient on the variable chain

Table 5.9

Panel Regression

This sample contains all bettors and all bets. The dependent variable is this model is the log of stake size, $ln(S_1)$, in USD. The sample includes only bettors in the winning domain who have had consecutive strings of wins. The model is estimated using the second and subsequent bets. Model 1 is estimated on the entire sample. Models 2 and 5 include financial market bettors and simulated market bettors, respectively are are estimated without controls. Models 3 and 6 include dummy variables for bet price with 5/1 bets as the base category. Models 4 and 7 include the bet price dummies, dummies for the second and subsequent bets and also interact the bet number dummies with the first stake. Model 8 is estimated on the pooled sample of winners with a dummy variable equal to zero for financial market bettors and 1 for simulated market bettors. Panel A is a survivor-only analysis, with Panel B including the attrition-adjusted sample.

Panel A: Survivors only

	All	Fins	Sim	All (Int)
L.Stake	-0.0110	-0.0138	0.000727	-0.0124
	(-1.62)	(-1.75)	(0.04)	(-1.72)
Win/Loss Chain				
Chain	0.0245***	0.0251***	0.0228***	0.0230***
	(9.90)	(8.53)	(4.66)	(7.83)
Bet price				
1/20	0.845***	0.977***	0.529***	0.864***
	(37.86)	(34.60)	(13.61)	(36.77)
5/6	0.303***	0.368***	0.206***	0.311***
	(24.13)	(21.65)	(10.51)	(23.38)
Evens	0.224***	0.265***	0.188***	0.235***
	(18.75)	(16.19)	(10.20)	(18.66)
1. Virts				0
				(.)
1. Virts × Chain				0.00961
				(1.69)
Constant	1.756***	1.841***	1.459***	1.779***
	(96.59)	(76.40)	(45.55)	(91.41)
N	82,703	58,920	17,512	76,432
R2	0.876	0.877	0.874	0.880

Table 5.9: (continued)

Panel B: Attrition-adjusted sample				
	All	Fins	Sim	All (Int)
L.Stake	0.470***	0.467***	0.464***	0.467***
	(82.85)	(72.50)	(34.45)	(80.07)
Win/Loss Chain				
Chain	0.00113***	0.000999***	0.00336***	0.000984***
	(4.20)	(3.34)	(4.95)	(3.30)
Bet price				
1/20	1.188***	1.317***	0.835***	1.211***
	(57.71)	(52.80)	(21.11)	(56.03)
5/6	0.385***	0.450***	0.256***	0.394***
	(35.85)	(33.22)	(13.51)	(34.92)
Evens	0.287***	0.327***	0.220***	0.295***
	(28.39)	(25.15)	(12.88)	(27.92)
1. Virts				0
				(.)
1. Virts × Chain				0.00228**
				(3.00)
Constant	0.616***	0.615***	0.523***	0.618***
	(40.43)	(32.61)	(19.18)	(38.95)
N	533,953	406,434	97,596	504,030
R2	0.785	0.782	0.762	0.786

is still positive[1]. This indicates that the virtual market is more sensitive to the chain of result history, something that is unexpected.

The results show increasing overconfidence with time, as evidenced by the interaction terms on the third and subsequent bets in the regression estimations. This is consistent with Gervais and Odean (2001), who describe a model in which overconfidence may 'wax and wane' during a trader's lifetime. Independent of the path of results, stake sizes are increasing in both the virtual and financial markets over time.

5.5 Conclusion

I find significant differences in bet stakes for those who have experienced strong positive or negative reinforcement in the form of successive wins and losses. In the winning domain, confidence manifests itself primarily in the form of increased stake sizes. In the losing domain, however, underconfidence is present in the form of attrition. Bettors who experience strings of losses exit the market, but those

[1] Across bettors, the product preference flag is a time-invariant regressor and thus is omitted.

who remain do not reduce their stake sizes. It seems that most bettors in the domain of losses learn about their ability and exit.

I found evidence of overconfidence in a market where such confidence was warranted but also in a setting where such behaviour was futile. Bettors in the winning domain in Financial markets increased their bet sizes, but simulated market bettors did so also, albeit to a lesser extent. This result was shown in a simple analysis of mean and median stake sizes across the evolution of betting in the first five bets, but also in a multivariate setting where bet price (a proxy for risk preference) and the stake range in the first five bets were used as controls.

When the analysis was expanded to include more bets, the result with regard to attrition was similar. As regards median stake sizes, those subject to strong positive reinforcement became optimistic and increased their stake sizes monotonically with each subsequent bet, however there was no such clear pattern in the losing domain. It is clear that strong reinforcement has a salient effect if it comes in the early stages of a bettors career. However, even negative reinforcement does not succeed in tempering the stake sizes of those in the losing domain with each subsequent bet. Negative reinforcement only causes bettors to exit the market if it happens in the initial rounds of betting.

Clear ex-post evidence of overconfidence was shown with an analysis of median profit amounts for bettors subject to strong positive reinforcement learning. Having increased their bet sizes, they experience lower median profit at each node, albeit better performance than those in the losing domain.

In financial market transactional data, it is not trivial to disentangle the relative magnitude and effect of rational and reinforcement learning. In a simulated market, however, there is no such confound. Empirical tests of learning theories in a laboratory setting infer how investors learn and behave in financial markets. I show that the effect of reinforcement learning on overconfidence may not be any stronger in financial markets than in a simulated market. This suggests that empirical tests may overstate the true level of biased learning in financial markets with ramifications for the existence, prevalence and magnitude of pricing anomalies and biases such as overconfidence. I compare learning in a baseline market where there can be no private information, no herding or information cascades, no tax-loss selling and no observation or imitation, (only private signals and randomness) with a market that has all of the preceding characteristics. This may be because agents are boundedly rational or because they 'lack the probabilistic information about the structure of payoffs necessary to successfully apply Bayes' rule (Wiseman, 2009).

I find that agents have the same propensity for overconfidence engendered by reinforcement learning in a simulated market treatment where outcomes are uncorrelated and 'learning by doing' is impossible as they do in a treatment where the literature suggests they should rationally learn from their past experiences. In effect, they place higher relevance on their previous success or losses and disregard public information and other signals in a type of 'reverse informational cascade'.

The results have broader implications for learning theories in financial markets. Bettors on the financial markets either have a preference for skewness (Golec and Tamarkin (1998)) and 'chase' their losses, or are subject to reinforcement learning. Bettors who show a preference for the random-number generated market, however, do not exhibit the same behaviour in the losing domain of the decision trees. We can see evidence of Bayesian learning in a setting where both 'learning by trading' and 'learning about ability' are possible (Financial markets) but also evidence for naïve reinforcement learning in the simulated market where they are not. This is evidenced by the increasing bet sizes for winners on Financial markets (rational Bayesian learning) and both increasing bet sizes for losers on Financial markets and bet size changes on the simulated market (naïve reinforcement learning).

Alternative explanations for stake size changes include changes in risk preferences after gains and losses, rather than overconfidence per se. Using stake size changes that resulted in negative returns was essentially an ex-post measure. Further extensions to the analysis in future research could include an alternative proxy for overconfidence. A variable such as *percentmaxbet*, bettors' current stake as a percentage of their maximum allowed bet size, could offer a convincing measure of ex-ante overconfidence. Nevertheless, such a variable would still necessitate the analysis of subsequent bettor P&L to determine whether an increase in the percentage of max stake wagered resulted in positive returns, indicative of skill or ability, or negative returns, indicative of overconfidence.

In addition, the existence of discrete bet prices may provide some insight into the difference between overconfidence and optimism. Overoptimism describes beliefs in the first moment of returns and overconfidence beliefs with regard to the volatility of such returns (Shefrin, 2008). An analysis of bets on the 5/6 outcomes (where bettors forecast the direction of the underlying relative to the current price) and 1/20 outcomes where investors forecast the magnitude of price move relative to the current price, may offer further opportunity to differentiate between these two biases.

Furthermore, overconfidence indicates an overestimation of one's abilities, whereas optimism suggests biased expectations about exogenous events. In that sense, simulated market bettors may be optimistic in the sense that they know that their own abilities cannot bring about positive outcomes, but they continue as they overestimate the likelihood of positive outcomes from their participation in betting on the future path of the random-number generated series.

Chapter 6

Learning theories and the disposition effect

6.1 Introduction

This chapter examines whether there is behaviour consistent with the prospect theoretical explanation for the disposition effect in a setting with terminal assets which essentially force liquidation at each decision node. In particular, I explore the reflection effect predicted by Kahneman and Tversky (1979). I define behaviour consistent with prospect theory preferences as risk-seeking in the domain of losses and risk-aversion in the domain of gains. While this is consistent with the disposition effect in stock market trading, further clarification is needed in this setting. I clarify how characteristics of a betting market, where positions are closed out automatically at expiry, differs from a traditional financial market context.

Prospect theory is only one of the many posited explanations for the disposition effect, some of which include mean reversion (Jiao, 2012), self-deception (Hirshleifer, 2001), self-justification (Kaustia, 2010) and cognitive dissonance (Chang, Solomon, and Westerfield, 2013). The difficulty in disentangling these effects has made for a lively debate in the literature, with empirical and experimental work offering competing explanations. This setting includes transactions on a simulated market, however, which should preclude a belief in mean reversion by rational agents. In addition, there are no tax considerations, as the dataset consists of financial bets rather than taxable financial instruments. I present evidence of risk-aversion in the domain of losses and risk-seeking in the domain of gains in this setting. I show evidence of an inverse disposition effect: those who are in-the-money in terms of P&L do not exit the market, while those who are out-of-the-money do. I also show that strong positive or negative reinforcement learning induces an inverse disposition effect.

A number of reasons have been suggested for winning investments being sold more readily than losers: differences in utility functions in the domain of gains and losses, belief in mean reversion, portfolio rebalancing and avoidance of higher transaction costs on losers, for example. The disposition effect, coined by Shefrin and Statman (1985), is posited as a behavioural bias which concerns the decision of whether to book, or hold on, gains and losses. When faced with new information such as a significant change in the price of an asset, the choice is between realising gains and losses or maintaining paper gains and losses. There is exhaustive evidence of its existence. Nevertheless, it is not a unified model of investor behaviour and does not incorporate 'the entirety of cognitive biases and rational limitations' (Machina, 1987). As I show, there are other push and pull factors driving behaviour, one of which is reinforcement learning.

As one of the explanations for the disposition effect, the shape of the prospect theory utility curve predicts risk seeking behaviour in the losing domain and risk averse behaviour in the positive domain. In either the positive or negative domains, this prediction is manifested firstly by attrition (i.e., agents quitting while ahead in the winning domain and staying in the market in the losing domain) and secondly by changes in risk (i.e., agents reducing risk in the positive domain and increasing risk in the losing domain). I suggest that there are two opposing factors affecting behaviour here: the decision to quit or stay driven by one effect; the decision to raise or reduce risk motivated by a second. In effect, reinforcement learning may account for the high (low) levels of attrition in the domain of losses (gains), while the reflection effect implied by prospect theory may be driving the risk-related decision.

I argue that a group of agents who have experienced consistently strong positive or strong negative feedback (winning streaks or series of consecutive losses) are more likely to exhibit behaviour consisted with reinforcement learning than the disposition effect. The dataset facilitates the isolation of the effect of these agents. Moreover, as reinforcement learning is a bias in itself, I disentangle both types of learning (Reinforcement Learning and Bayesian learning) with a focused test in two market settings: one in which Bayesian learning can take place, and a second where outcomes are independent and only naïve reinforcement learning is possible. The setting also has the advantage of having uncorrelated returns and hence, no possibility of mean reversion. I suggest that by not isolating the effect of reinforcement learning, empirical studies of the disposition effect may be understating the extent of this bias. In addition, a number of characteristics in this setting lend itself to a clean test of the disposition effect. Firstly, there are no tax considerations in this setting as these are not taxable investments. In addition, since agents know that returns are uncorrelated in the Virtual Market setting, there can be no belief in mean reversion by rational agents.

The structure of the rest of this chapter is as follows. Section 6.2 outlines the theoretical framework within which I propose the hypotheses. Section 6.3 presents the research questions. The dataset has been explored in detail in Chapter 4, however clarifications of any additions to or omissions from the sample are outlined in detail, where relevant. Section 6.4 presents evidence of the existence or absence of the reflection effect in the entire sample and provides support from the first set of hypotheses, arguing that reinforcement learning may mitigate this bias. I place the results in context with prior literature in this section. Section 6.5 concludes.

6.2 Theoretical framework

Utility curves proposed by Bernouli (1738), Friedman and Savage (1948), Markowitz (1952) and Kahneman and Tversky (1979) are presented in Figure 6.1. Bernouli (1738) expanded on the expected value rule to suggest a concave utility function and proposed a decreasing marginal utility (presented in Figure 6.1a), addressing some of the behavioural issues raised by the St. Petersburg Paradox.

Friedman and Savage (1948) posited a hypothesis which explained both the purchase of insurance and the participation in lotteries, but ever since, empirical evidence of behaviour inconsistent with this theory has been outlined in the literature. The Friedman and Savage (1948) utility curve is presented in Figure 6.1b. The point W_0 dominates W_1 to W_2 and is an example of the purchase of insurance. The point $W_0{}'$ dominates $W_1{}'$ to $W_2{}'$ and is an example of a lottery play. The ex-ante and ex-post reference points in this model are final states (of wealth, for example) rather than actual wins and losses, and the curve has both convex and concave regions. Friedman and Savage (1948) state that while gambles are often purchased in a pure form, when betting, an agent buys both a gamble in the theoretical sense (i.e., the participation in the mechanics of the game of chance) but entertainment also. While the two components can be separated: buying entertainment only by paying to participate in a game with 'stage money' or the gamble alone by issuing instructions to someone else to participate, it is critical to attempt to disentangle both components.

Markowitz (1952), however, argued that in some regions this utility curve implies 'behaviour which is not only not observed but would be generally considered

Figure 6.1

Utility curves

The following figures display utility curves from Bernouli (1738), Friedman and Savage (1948), Markowitz (1952) and Kahneman and Tversky (1979).

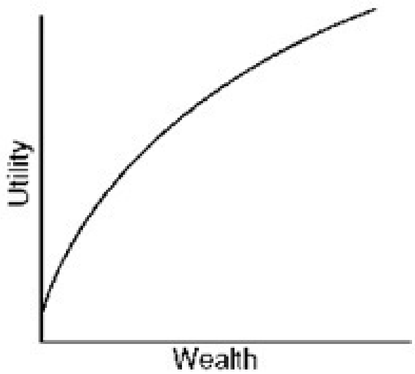

(a) Bernouli (1738)

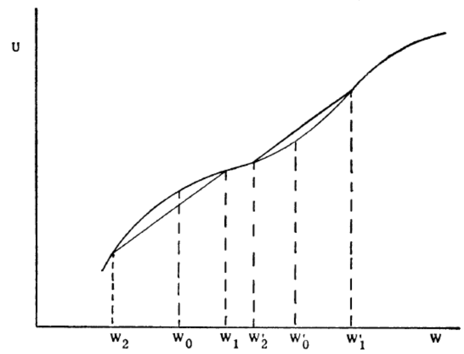

(b) Friedman and Savage (1948)

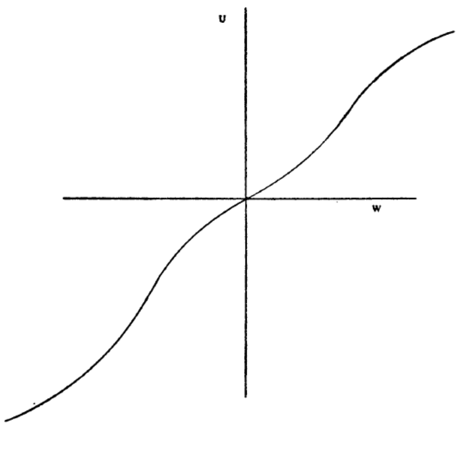

(c) Markowitz (1952)

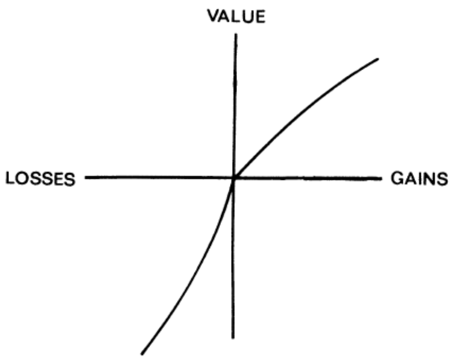

(d) Kahneman and Tversky (1979)

peculiar if it were' (Markowitz (1952), p. 152). He suggested a utility curve which was convex immediately below present wealth and concave about it, but concave and convex deep in the winning and losing domains, respectively. This indicated moderate risk seeking for moderate winners, however risk aversion for heavy winners, while at the same time predicating moderate risk aversion for moderate losers, and risk seeking by heavy losers. Rather than defining a utility function in only the positive domain, he proposed a curve (shown in Figure 6.1c) with three inflection points: the middle one at the origin representing current wealth, a first convex, then concave curve in the positive domain, and a concave, then convex curve in the negative domain. The distance between the two points where the utility curve changes from convex to concave in each of the positive and negative domains was a function of wealth. The introduction of this hypothesis accounted for some of the observed anomalies in agent behaviour.

Markowitz (1952) stressed that when analysing a series of consecutive bets, behaviour in one particular round of betting cannot be explained without reference to both the gains and losses which have already occurred and the probabilities of future bets. He stressed the importance of the outcome of the 'evening' of betting as whole which is consistent with the hypothesis of Golec and Tamarkin (1998) that bettors change their behaviour in consecutive rounds with a view to seeking skewness towards the end of the session. In a fair game, the distribution of outcomes is symmetric if bettors stake the same size each time, however, by reference to their previous performance, agents can vary their stakes in an attempt to change the shape of the distribution of outcomes for the entire betting session.

In short, Markowitz (1952) implied that the shape of the Friedman and Savage (1948) curve implies that agents would bet conservatively when winning and liberally when losing, and would be counter-intuitive. He states that stakes should be higher for moderate winners, but that larger winners would play lower stakes or stop playing. In his own words: 'In the vernacular, the heavy winner would have made his 'killing' and would wish to quit while winning' (Markowitz (1952), p. 156). Conversely, a moderate winner would be expected to play lower stakes or not play at all, while a person deep in the losing domain could be expected to continue with the game. The inference is that as stake size gets smaller, the expected utility increases, which Markowitz (1952) states is contrary to observed behaviour and a criticism of this hypothesis. In fact, in an experimental setting, Mosteller and Nogee (1951) showed that when asked to issue instructions as to how their own money should be bet by others, players exhibited conservative behaviour when losing and liberal behaviour when winning. This stylised fact is consistent with the observed behaviour in the sample presented in Section 6.4.

Kahneman and Tversky (1979) list a series of situations in which observed behaviour is contrary to that predicted by expected utility theory and propose a framework called prospect pheory as a model to describe decision making under risk. They proposed a framework based on gains and losses, where a decision weight is applied to each probability and a value weight applied to each outcome in the expected utility function. They suggest that agents are risk-seeking in the loss domain and risk averse in the domain of gains relative to a reference point.

The most relevant for this analysis are the 'certainty effect' and 'reflection effect' (i.e., choices made by respondents in the negative domain were the opposite of those made in the positive domain). The reflection effect describes a situation in which risk seeking choices in the positive domain are associated with the opposite choices in the negative. This implies a function (shown in Figure 6.1d) which is concave above the reference point and convex below it.

The difference between the Kahneman and Tversky (1979) utility curve and that proposed by Markowitz (1952) is the relative steepness of the curve to the right of the origin (the reference point). The reference point in this setting was the initial position and subsequent losses and gains. With the shape of this curve, Kahneman and Tversky (1979) suggest that the displeasure associated with losing a given amount is greater than the pleasure associated with winning an equivalent amount. The loss aversion component evident in the shape of the prospect theory utility curve has motivated extensive empirical and experimental research on investor behaviour.

While motivating the hypotheses with predictions from prospect theory, it may be instructive to point out some critical differences between this setting and the Kahneman and Tversky (1979) experimental design. Prospect theory was designed to explain anomalies with respect to one-shot bets. There was no reference to past returns or outcomes. With regard to the evolution of trading behaviour as subsequent rounds of betting evolve, Kahneman and Tversky (1979) have the following to say: 'The well known observation that the tendency to bet on long shots increases in the course of the betting day provides some support for the hypothesis that a failure to adapt to losses or to attain an expected gain induces risk seeking.' (Kahneman and Tversky (1979), p. 287). Again, this is consistent with the Markowitz (1952) reference to agents treating consecutive bets as being part of a larger 'betting session' rather than independent plays. The examples of aversion to losses and the contention that losses loom larger than gains, led to Shefrin and Statman (1985) coining the phrase the 'disposition effect'.

Shefrin and Statman (1985) examine agents disposition towards holding on to losing stock positions and realising gains relatively early. Their research question centered around whether investors behaviour was at odds with that proposed by Constantinides (1983) who proposed a tax avoidance strategy based on the difference between short-term and long-term tax rates for gains and losses. They show that investors book gains too early and hold on to losses too long, as predicted by prospect theory. The concept of mental accounting is also introduced by Shefrin and Statman (1985). Odean (1998a) tested for the existence of the disposition effect, using a stock's purchase prices as the reference point, and conclude the investors exhibit behaviour consistent with Shefrin and Statman (1985), even when controlling for other behavioural explanations such as mean reversion, portfolio rebalancing and optimal tax strategies.

Prospect theory, therefore, has been suggested as a cause of the Disposition Effect by Shefrin and Statman (1985). However, prospect theory preferences are only one possible cause of this anomaly and many others have been posited in the literature. Summers and Duxbury (2012) show that prospect theory in isolation may not bring about the disposition effect and that emotions such as elation and regret are a necessary pre-requisite for it to be present. Shefrin and Statman (1985) also included some discussion of pride seeking and regret avoidance. Kaustia (2010) discussed information-based motives (portfolio balancing, mean-reversion and private information), belief in target prices and psychological explanations such as mental accounting, regret aversion and self-control. Jiao (2012) tests the mean-reversion hypothesis in an experimental setting. Odean (1998a) incorporates portfolio rebalancing, transaction costs and private information. Both Shefrin and Statman (1985) and Thaler and Johnson (1990) discuss investors' reluctance to close mental accounts at a loss. Hirshleifer (2001) proposed self-deception as a motivation for observed behaviour. Kaustia (2010) concludes that prospect theory is unlikely to explain the disposition effect and argues against mean reversion as a cause also. Barberis and Xiong (2009) state that agents derived utility from realized gains and losses only, not from paper gains and losses. The existence of competing arguments as to the cause of this anomaly has led to a lively debate in the literature, and numerous attempts to refute the mean-reversion hypothesis.

Weber and Camerer (1998) propose the reference point and reflection effects as explanations for the disposition effect at the aggregate level. They present an experimental setting in which subjects sold fewer shares when both the purchase price and previous prices were lower than the current price, suggesting two possible reference points. Although price changes in their experimental design were positively autocorrelated, precluding mean reversion, they found that agents bought losers and sold winners.

As regards prospect theory, the disposition effect can be framed by agents who face a paper loss (in the 'editing phase' suggested by Kahneman and Tversky (1979)) as a choice between the certain prospect of realising a loss and the risky prospect of the stock rebounding with a certain probability or continuing to decrease at another. Prospect theory dictates that as the reference point (the current paper loss) is in the negative domain, the risky alternative will dominate.

Conflicting empirical evidence has, however, also been presented. Thaler and Johnson (1990) analysed decision making in the presence of prior outcomes, concluding that there was risk seeking in the presence of a prior gain, risk aversion in the domain of losses, and also a particular attraction in the losing domain for break-even prospects. Later in the analysis, I suggest cumulative losses as a driving force behind subsequent betting decisions. Barberis (2012, p. 49), in proposing a model that explains betting on skewed lotteries states: 'Our model makes a number of novel predictions – predictions that, we hope, will eventually be tested'[1].

Kaustia (2010) concludes that it is not trivial to rationalize the disposition effect with any popular explanation and proposes an interesting direction for future research. He suggests that a story based on the psychological motive for self-justification (manifested as regret-avoidance in Shefrin and Statman (1985) or self-deception in Hirshleifer (2001)) may offer a simple explanation. Moreover, Kaustia (2010) suggest that agents may be more inclined to hold on to losing stocks if they see some chance of breaking even and makes reference to the concept of escalation of commitment discussed in Staw (1981). Staw (1981) showed that people become 'entrapped' in a losing course of action and a significant contributory factor to entering such a spiral is the mistaken belief that one is close to a goal Rubin and Brockner (1975).

In summary, the evidence presented by Kahneman and Tversky (1979) has been greeted with caution, and a number of open questions remain. In this study, two issues are crucial if the results can be credibly generalised to wider financial market decisions. Firstly, we must assess the relative merits of conflicting explanations for this bias, only one of which is prospect theory preferences. There may be a number of alternative explanations for the behaviour we observe. Nevertheless, I suggest that this setting goes some way towards eliminating some of these explanations[2]. Secondly, we must decide whether we should indeed expect

[1] The final empirical chapter of this thesis presents such evidence in an experimental setting.

[2] There are no tax implications, hence no tax-loss selling. In the simulated market, there can be no belief in mean-reversion by rational agents.

behaviour consistent with, or opposite to, that predicted by prospect theory.

Barron and Erev (2003) stress the difference between 'decisions from description' and 'decisions from experience'. The certainty equivalents outlined in Kahneman and Tversky (1979) are an example of decisions from description as agents have full knowledge of outcomes and their related probabilities in advance, but no possibility of learning as they concern one-shot decisions. Decisions from experience refers to situations in which agents do not have information about the probabilities of all prospects, but rather learn from experience as they sample over a series of rounds from different outcomes [1]. Barron and Erev (2003) found risk seeking in the domain of gains and when the initial example in Kahneman and Tversky (1979) was presented in a learning-type environment, 63% of respondents chose the risky-option (as opposed to 80% in Kahneman and Tversky (1979)). Ludvig and Spetch (2011) suggest that the entire s-shaped curve relating objective value to subjective utility may be reflected when agents learn from experience.

According to Shefrin and Statman (1985, p. 777): 'Kahneman and Tversky's finding was obtained in a controlled experimental situation. Economists tend to treat experimental evidence with some caution and are reluctant to conclude automatically that similar features will be exhibited in real-world market settings. Indeed, it is important to look at market behavior in order to ascertain whether such behavior patterns can be discerned in actual trading.' Indeed, according to Kahneman and Tversky (1979, p.265): 'these experimental studies typically involve contrived gambles for small stakes, and a large number of repetitions of very similar problems. These features of laboratory gambling complicate the interpretation of the results and restrict their generality.'

I present evidence of risk seeking in the domain of gains and risk aversion in the domain of losses. I also show that agents subject to strong positive or negative reinforcement exhibit behaviour inconsistent with prospect theory. While this chapter examines panel data on individual transactions, the next chapter offers evidence from an experimental setting to add support to the arguments. I test the premise that sign of the disposition effect measure differs between smaller and larger investors. Finally, I use two measures of sophistication to examine whether agents who engage in more complex strategies exhibit a reduced disposition effect measure.

[1] The IOWA Gambling Task (Bondarenko and Bossaerts, 2000) is another example of such a setting.

6.3 Hypothesis development

The orthodox disposition effect setting concerns stock purchases, where a long position is held each time a paper gain or loss is marked, or the position is closed. If a position in a stock is liquidated, the customer no longer has a position and in market parlance is 'out'. In the case of a paper loss, it is indisputable that if the position is marked-to-market, the agent has lost real money, however loss averse agents can decide to continue and not 'pay up' in the hope of avoiding a certain loss. There is a parallel in a casino setting, where after the result of each bet, agents have won or lost money and the result is also indisputable.

Nevertheless, agents essentially have a paper gain or loss until they decide to quit and cash in their chips (what remains of, or what is currently in addition to, their initial endowment) at the cashier's cage. The online betting equivalent is also similar, where a betting account is funded with an electronic payment method and the agent plays until the account no longer has a balance, or in the case of winning bets, until a withdrawal is made back to the bank account that funded the initial endowment. In that sense, it is crucial to determine at which point an agent has exited the market.

There are limitations to this approach, however, as they are fundamental differences between keeping cash on account with a bookmaker and holding a position in a stock. A long position in a stock is intrinsically risk, and a paper gain or loss can change the value of a position, with triggers such as margin calls and position close-outs as a consequence. The value of casino chips withheld from betting, or cash on account with a bookmaker does not change in value unless a decision is made to actively bet.

Hartzmark and Solomon (2012) provide evidence of the existence of the disposition effect in a negative expected return gambling market, and argue that a parsimonious model of the disposition effect, which is usually discussed in the context of positive expected return assets, must take into account situations in which individuals are are locally risk-seeking. They stress that their results are consistent with the Barberis (2012) model of time-inconsistent prospect theory preferences, where agents accept negative outcome gambles at a casino as they intend leaving when the start losing, but rather leave the casino when they are winning and continue to gamble when they face losses.

Hartzmark and Solomon (2012) suggest that this behaviour is similar to a disposition effect, and that such results question standard explanations for the disposition effect such as prospect theory and a belief in mean reversion. They also address the argument that bettors are not profit maximizing rational traders and gain utility from gambling. They argue that for such preferences to distort prices, they must affect not only the decision to enter the market, but how they behave subsequently also. If they act as rational profit maximizers after having entered the market, this should not affect prices in a systematic manner.

Borghesi (2012) uses betting data from an online betting exchange to test for the disposition effect. He argues that using bets mitigates the joint hypothesis problem in traditional financial markets (where one cannot convincingly test for deviations from true value unless one has a model for determining the correct value), as the true price of a bet is known with certainty upon expiry. Betting data has also been used to test for the disposition effect and to analyse risk preferences by Andrikogiannopoulou and Papakonstantinou (2011), Brown and Yang (2015), Andrikogiannopoulou (2010) and Barseghyan, Molinari, O'Donoghue, and Teitelbaum (2015). In the section that follows, I first outline measures used in prior literature and then introduce some modifications.

Odean (1998a) analyses individual stock portfolios, and each day a sale took place, places each portfolio into one of four categories: paper gains, paper losses, realized gains or realized losses. In contrast to this, each round of betting in this setting effectively constitutes a trading day when gains or losses are booked and thus there are no omissions from the sample. Similarly, agents who experienced a paper gain or loss at each round can either continue to play or can quit or 'cash in their chips', converting a paper gain or loss into a realized one. I first analyse behaviour in the losing and winning domains and examine whether the propensity to quit is path dependent. Thereafter, I examine changes in risk preferences (in the form of bet prices), to test for path dependency in risk.

Odean (1998a), while analysing the disposition effect at the aggregate level, proposed the following measure:

$$\begin{gathered}\frac{Realized\ Gains}{Realized\ Gains + Paper\ Gains} = PGR\ (Proportion\ of\ Gains\ Realized)\\ \frac{Realized\ Losses}{Realized\ Losses + Paper\ Losses} = PLR\ (Proportion\ of\ Losses\ Realized)\end{gathered} \tag{6.1}$$

As a test of the existence of the disposition effect in his sample, Odean (1998a) hypothesised that PGR was larger than PLR. I follow a similar approach to Odean (1998a) but modify the definition of PGR and PLR as follows[1]:

[1]Feng and Seasholes (2005) state that the 'PGR' and 'PLR' ratios of sales for gains and sales for losses have a number of drawbacks in that such an approach works well at the aggregate level but not at the individual account level. As a result, I propose a modification.

$$\frac{\#\ Agents\ with\ paper\ gains\ who\ cash\ in\ their\ chips}{\#Agents\ with\ paper\ gains} = PGR$$

$$\frac{\#\ Agents\ with\ paper\ losses\ who\ cash\ in\ their\ chips}{\#Agents\ with\ paper\ losses} = PLR \tag{6.2}$$

For example, if a number of investors held portfolios which consisted in aggregate of four winning stocks and four losing stocks, and liquidated two of the winning stocks and only a single losing stock, the measures outlined in Odean (1998a) would result in $PGR = 1/2$ and $PLR = 1/4$. We would find identical measures if a round of betting consisted of four winning bettors and four losing bettors, and two of the winning bettors exited the market while all but one of the losing bettors continued to bet.

Such an approach is warranted in this setting as the underlying betting proposition constitutes a long or short position in a terminal asset[1]. The end of each round of betting is a decision node, where the decision to quit or continue is made. At each such node, positions are essentially liquidated rather than being merely marked-to-market. As such, in each domain, rather than comparing the proportion of liquidated positions to the total number of positions, I compare the proportion of agents who exited the market to the total number of agents in each domain at each round of betting.

Such definitions of paper gains and losses are certainly open to debate. Losing stocks which have not been sold may have been marked-to-market and while this constitutes a tangible loss of capital, there is some probability of the stock retracing its value if the position is held. In a betting market, losses are final at each mark-to-market point (upon settlement of bets). The only condition in which there is a probability to recoup losses is to participate in another round of betting. A condition in which both the stock market and betting market are analogous in this respect would be if agents' stocks were closed out at the end of each trading day (and agents were prohibited from investing in the same stock in the next trading session). In this case, the loss upon mark-to-market would be final and the only option left to investors would be to participate in another day of trading (with an identical or different stock).

In fact, this is the case for most day traders who may not wish to hold positions overnight (or may be prohibited from doing so by risk-managers at proprietary trading firms) or who may not have the capital to pledge sufficient margin to hold

[1]These are not 'wasting assets' in the sense that their value is constantly in decline with each subsequent round, however their expiry time is well-defined in advance.

positions overnight. In that sense, this setting is similar to that of day-traders in proprietary trading firms whose trading accounts must be 'flat' before the closing auction at the exchange.

The measure used to test for the existence or absence of a disposition effect in this setting relates purely to levels of attrition, with continuation defined as risk seeking behaviour and exit from the market classified as risk aversion. However, we can analyse both risk-seeking and risk-averse behaviour in both the domains of gains and losses by examining subtle changes in risk preferences with each subsequent bet. As there are a number of bet prices to choose from (from symmetric bets such as 5/6 and 1/1 to skewed bets such as 1/20 and 5/1), I now incorporate an analysis of the propensity to place safer or riskier bets in addition to examining attrition.

The first example in Shefrin and Statman (1985) describes the two alternatives open to an investor who is long a stock that has reduced in price by $10:

A. *Sell the stock now, thereby realizing what had been a $10 'paper loss'.*

B. *Hold the stock for one more period, given 50-50 odds between losing an additional $10 or 'breaking even.'*

In this setting, there are a number of bet prices (5/6, Even-money, 1/20 and 5/1) with associated levels of risk. If a bettor decides to continue to play, they have a choice whether to continue at the same level of risk or whether to increase or decrease risk with each subsequent bet. With reference to the proposal by Markowitz (1952), Kahneman and Tversky (1979) and Golec and Tamarkin (1998) that agents may base their behaviour on a full 'evening' of betting, agents who have been betting on 5/6 outcomes and have negative returns, for example, are faced with the following four propositions:

A. Quit: *Accept a certain loss to date with 100% probability.*

B. Same risk profile: *Play again with a 50% probability of winning 83% on a 5/6 bet and a 50% probability of losing 100% of the stake.*

C. Higher risk profile: *Play again with 11% probability of winning 500% on a 5/1 bet and 89% probability of losing 100% of the stake.*

D. Lower risk profile: *Play again with 42% probability of winning 100% on a 1/1 bet and 58% probability of losing 100% of the stake.*

In effect, the choice is between continuing and quitting (survival and hazard) and entails a number of distinct choices using the current state of winnings/losses as a reference point. In both the domain of gains and losses, the choice is between a certain and an uncertain payoff. For those in the positive domain: a sure gain (i.e., manifested by attrition and indicative of quitting while ahead) or continuing to play and accepting a risky gamble. For those in the negative domain: a sure loss (i.e., manifested by attrition and indicative of realizing a paper loss or sunk cost) or continuing and playing a further round, thereby accepting a risky gamble. It entails the following choices (using the current state of winnings/losses as a reference point):

- For those in the positive domain
  - A sure gain i.e. attrition (Quit while ahead)
  - A risky gamble i.e. continue to play (same risk profile)
  - A risky gamble i.e. continue to play (become risk seeking)
  - A risky gamble i.e. continue to play (become risk averse)
- For those in the negative domain
  - A sure loss i.e. attrition (Realizing a paper loss/sunk cost)
  - A risky gamble i.e. continue to play (same risk profile)
  - A risky gamble i.e. continue to play (become risk seeking)
  - A risky gamble i.e. continue to play (become risk averse)

As with any test of the disposition effect, the specification of the reference point is crucial. Similar to Weber and Camerer (1998), I argue that bettors 'frame' the decision as a choice between a certain loss, with a negative value, or keeping playing, accepting a gamble to either break even or lose again. In the initial test, I assume that the reference point is cumulative losses to date, however the reference point could be an arbitrary, subjective amount which differs from agent to agent or could even a dynamic, adaptive reference point. For robustness, I test a number of alternative specifications.

The following two hypotheses test for the existence disposition effect behaviour in the domains of gains and losses:

> **H1a.** *Losers are more likely to place riskier bets than winners in aggregate.*

> **H1b.** *Winners are more likely to place safer bets than losers in aggregate.*

Interestingly, Nolte (2016) presents evidence of a non-linear disposition effect measure, with inverse measures for small profits and losses and a positive disposition effect for larger gains and losses. In effect, larger investors exhibited a reduced disposition effect than smaller investors. Agents trading with more complex strategies also exhibited a lower disposition effect in Nolte (2016), which motivates a later robustness test in this chapter. I therefore split the sample according to the size of gain and loss, and examine whether larger loses and gains in fact lead to a positive disposition effect measure.

At each decision node, the disposition effect argues that winners are disposed to booking gains and losers tempted to wait for their positions to retrace. These are strong temptations with extensive empirical support. However, empirical studies (and the support for the hypotheses in the preceding chapter) have also shown that individuals are subject to reinforcement learning, such that they repeat actions which have resulted in positive feedback and shun actions which have resulted in negative feedback. In this case, what of agents who have experienced a number of days of positive returns in a row or who have had 'winning streaks'?

The following hypothesis tests the premise that agents subject to strong reinforcement will not exhibit a disposition effect:

H2a. *Losers in the reinforcement cohort are more likely to quit than winners.*

H2b. *Winners in the reinforcement cohort are more likely to continue than losers.*

I argue that reinforcement learning is a salient bias which interacts with the disposition effect. There is certainly a temptation to book gains, however there is also a temptation to continue in the domain of gains, repeating an action which has resulted in positive feedback. In the domain of losses, the thesis is similar: there is a temptation to 'wait for it to come back', however there is also the temptation to walk away and desist from repeating an action which has resulted in successive negative feedback.

Finally, I examine how sophistication of agents and the existence of 'near-miss' outcomes drive behaviour. Song et al. (2014) show that overconfident characteristics such as online trading (as opposed to offline or phone trading) have a moderating effect on overconfidence. Barber and Odean (2002) show that online investors trade more aggressively, with more risk and with reduced profit than offline investors. However in this case, those who had moved from offline to online trading had a preference for high-risk growth stocks. We therefore have some reason to believe that the choice of execution channel may also be an indicator of sophistication.

One final variable in the dataset could serve as a useful independent variable in such a model. Nolte (2016) found that agent using complex strategies exhibited a reduced disposition effect measure. Bettors in this dataset can choose to place a number of bet types, from single bets on Financial Markets or the simulated market, to more complex bets including multiple legs. Such bets could include a bet on the closing price of the DJIA, for example, combined with a bet on a sporting outcome. I incorporate this analysis with the following hypothesis:

> **H3.** *Sophisticated agents, as proxied by bet channel and bet type, are less likely to exhibit a disposition effect.*

Kaustia (2010) states that investors may be more committed to holding on to a losing stock if they perceive that they are likely to break even. Staw (1981) proposed that the closer an agent perceives a goal to be, the more committed he/she is likely to be to persist in losing strategies in order to try to achieve it. Thaler and Johnson (1990) also show that prospects which offer the opportunity to break even are particularly attractive. I test this premise by focusing on 'near-miss' outcomes.

I define a near-miss event as a losing bet where the settlement price was within a narrow range of the strike price of the bet upon expiry. For example, a 'bullish' (or 'higher-than') 5/6 bet on the FTSE 100 to expire at midday with a strike of 5,954.08, would be out-of-the-money unless the FTSE printed at least 5,954.09 at the exchange at 12.00 GMT. If the settlement price was any lower than 5,954.09, the bet would expire worthless and the bettor would lose their stake. Any higher, and the bettor would receive 1.8333 times their stake (i.e. the refunded stake and 0.8333 times the stake as winnings). I consider a near-miss outcome to be a bet which expired out-of-the-money, but with a settlement price within 0.001% of the strike price of the bet. In this example, no more than 6 ticks lower i.e. the FTSE printing between 5954.02 and 5954.08. The final hypothesis is therefore as follows:

> **H4.** *Agents who have experienced a near-miss outcome are less likely to quit and more likely to place a risky bet.*

The next section first presents an analysis of the orthodox disposition effect measure in aggregate. Thereafter, I focus on risk-seeking and risk-averse behaviour in the domains of gains and losses. I then analyse the effect reinforcement learning cues have on behaviour. Finally, I examine the effect proxies for sophistication and near-miss outcomes have on risk preferences.

6.4 Results

To embark on the analysis, I first examine to what extent a disposition to closing winning positions and hold on to losing positions exists in the sample in aggregate. In Table 6.1, I present the number of bettors exiting the market, the number of 5/6 and 1/20 bets (categorised as *safe* bets) and the number of 5/1 and even-money bets (categorised as *risky* bets) being placed. The sample includes all bettors out to the tenth round of betting.

Table 6.1

Aggregate disposition effect measure

This table presents the level of attrition, the number of *safe* (1/20, 5/6) and *risky* (1/1 and 5/1) bets being placed, the Proportion of Gains Realized (PGR), Proportion of Losses Realized (PLR) and the corresponding disposition effect measure for the first ten rounds of betting.

	Bet	1	2	3	4	5	6	7	8	9	10
Loss	Quit	8,499***	3,453	1,894	1,283	991***	726	633	510***	442	370
	Safe	4,096***	3,069	2,531	2,129	1,923	1,668	1,494	1,350***	1,303	1,177
	Risky	6,252***	5,290	4,620	4,256	3,789	3,586	3,265	3,012***	3,018	2,796
	N	18,847***	11,812	9,045	7,668	6,703	5,980	5,392	4,872***	4,763	4,343
Gain	Quit	3,586	1,435	766	476	370	281	248	202	168	166
	Safe	3,147	2,120	1,667	1,366	1,161	1,057	999	925	747	755
	Risky	10,335	8,463	7,464	6,772	6,289	5,844	5,516	5,275	4,884	4,688
	N	17,068	12,018	9,897	8,614	7,820	7,182	6,763	6,402	5,799	5,609
	PLR	0.45	0.29	0.21	0.17	0.15	0.12	0.12	0.11	0.09	0.09
	PGR	0.21	0.12	0.08	0.06	0.05	0.04	0.04	0.03	0.03	0.03
	DE	-0.24	-0.17	-0.13	-0.11	-0.10	-0.08	-0.08	-0.07	-0.06	-0.06

From the offset, there is an inverse disposition effect measure, which is in line with the attrition results presented in the previous chapter. A higher number of agents exit the market in the losing domain than in the winning domain, with the highest rates of attrition in the initial stages. In aggregate, the inverse disposition effect measure becomes less negative with each round of betting and stabilises to approximately 6% after the ninth bet[1]. We must therefore reject the premise that investors realize more winning positions than losing positions. There is, however, a precedent in the literature for an inverse disposition effect measure.

The sample in Odean (1998a) exhibited a negative disposition effect in December, motivated by the tax-loss selling outlined in Constantinides (1983). However, Dhar and Zhu (2002) showed that a significant number of investors in their sample displayed an inverse disposition effect measure during the rest of the year. Krause, Wei, and Yang (2009) also show evidence of the existence of a reverse disposition effect in the form of the length of loss-making and profitable invest-

[1]Further analysis not tabulated here showed a stable disposition effect measure from the tenth through to the twentieth bet.

ments. Barberis and Xiong (2009) presented evidence of a reverse disposition effect while testing an explanation based on prospect theory. They conclude that PGR is always lower than PLR in a two-period setting, and is lower approximately half the time when the number of period is greater than 10. Similarly, we see a reduction in the disposition effect measure over time, however it does remain negative throughout.

6.4.1 Univariate results

With Nolte (2016) in mind, I introduce a split according to the absolute size of the loss or gain in the next table in order to verify whether there is a difference in the disposition effect measure between smaller and larger bettors. At each round of betting, I split the outcomes into losses and gains with three categories of each according to the absolute size of the loss or gain. The buckets are not fixed, and customers can move between buckets as the rounds of betting progress and their P&L evolves. Panel A of Table 6.2 presents results for the entire sample.

In addition to analysing any difference in disposition effect between large and small bettors, we can also address the second set of hypotheses on the effect of strong positive or negative reinforcement. To disentangle the learning effect, I split the sample into those who have had consecutive runs of strong positive or negative feedback (three, four or five losses or gains in a row i.e., the agents most susceptible to reinforcement learning) and the remainder, who have only been subject to mixed feedback and weak reinforcement learning. Panel B shows the behaviour of the sample when those who have been subject to strong positive or negative reinforcement have been removed, whereas Panel C includes only those agents who have had two, three, four or five consecutive winning or losing bets.

In Table 6.2, we can analyse the percentages of bettors deciding to quit or changing the risk profile of their bets in a number of ways. Comparisons can be made between stake size categories in each panel, between rounds of betting, between categories of stake size, and between all of these statistics across panels.

I first focus on the behaviour of all bettors presented in Panel A of Table 6.2 which includes the entire sample of agents. Firstly, as regards attrition, it is clear that losers are more likely to quit than winners i.e., across all panels, ‘quit’ percentages are higher for losses than gains (irrespective of loss/gain size category and across consecutive bets). The probability of quitting also decreases with the number of bets irrespective of category, which may be evidence of a learning effect, with the most salient learning experience at the start. This initial result is contrary to the prediction of prospect theory as there is risk seeking in the domain of gains and risk aversion in the domain of losses.

Table 6.2

Disposition effect

The measure summarised in Panel A below addresses the question of whether the propensity to bet is the same in the face of a similar gain versus a loss. The dataset contains all bettors, irrespective of preference with regard to product type, bet price or bet duration. Each percentage represents the proportion of customers carrying a paper gain or loss at each round of betting who quit, continued playing a safe bet price (1/20, 5/6) or who continued playing a risky bet price (5/1, 1/1). As shown previously, bet size is highly skewed, thus three categories of gain and loss are introduced. Gain and loss categories were chosen to endeavour to have a similar number of bettors (in each of the first two categories at least).

		€0 - €5					**€5 - €50**					**€50+**				
		Bet 1	Bet 2	Bet 3	Bet 4	Bet 5	Bet 1	Bet 2	Bet 3	Bet 4	Bet 5	Bet 1	Bet 2	Bet 3	Bet 4	Bet 5
									Panel A: All bettors							
Loss	Quit	48.13***	31.69	22.95	18.44	16.16***	40.43	25.54	18.69***	14.60	13.51	26.59	19.06***	12.82	11.93	9.83
	Safe	24.21***	29.84	31.87	32.41	34.07	17.22	19.84	23.00***	21.96	22.43	10.90	12.04***	15.18	14.68	14.74
	Risky	27.65***	38.47	45.18	49.15	49.77	42.35	54.62	58.31***	63.44	64.05	62.51	68.90***	72.01	73.39	75.43
	N	12,914***	7,721	5,600	4,637	3,954	5,098	3,493	2,852***	2,486	2,220	835	598***	593	545	529
Gain	Quit	23.67	13.79	9.16	6.81	5.58	17.47	9.85	6.35	4.59	4.10	14.03	8.19	4.90	2.49	3.08
	Safe	21.99	22.08	21.22	20.52	19.43	13.99	13.08	12.86	12.12	11.32	7.69	6.81	6.84	6.12	6.27
	Risky	54.34	64.12	69.62	72.68	74.99	68.54	77.07	80.80	83.28	84.58	78.28	85.00	88.26	91.39	90.65
	N	10,368	6,793	5,349	4,465	3,927	5,581	4,212	3,671	3,266	3,048	1,119	1,013	877	883	845
	DE	-24.46	-17.90	-13.79	-11.63	-10.58	-22.96	-15.68	-12.34	-10.01	-9.41	-12.56	-10.87	-7.91	-9.44	-6.75
									Panel B: Learning Effect Removed							
Loss	Quit	***	25.61	19.77	17.12	15.25***		20.88	17.04***	14.10	12.90		13.53***	11.02	11.62	9.65
	Safe	***	29.15	29.72	31.12	33.11		19.67	21.96***	20.71	21.67		10.56***	14.41	14.52	13.78
	Risky	***	45.25	50.50	51.76	51.64		59.44	61.00***	65.19	65.44		75.91***	74.58	73.86	76.57
	N	***	3,702	4,081	3,978	3,627		1,901	2,213***	2,192	2,086		303***	472	482	508
Gain	Quit		19.31	11.24	8.19	6.47		14.52	8.91	6.07	4.91		13.70	6.50	3.64	3.89
	Safe		27.23	25.79	24.73	22.72		16.20	16.79	15.60	15.03		8.77	9.51	8.05	8.52
	Risky		53.46	62.97	67.08	70.81		69.29	74.30	78.33	80.06		77.53	83.99	88.31	87.59
	N		2,589	3,141	3,126	3,028		1,488	1,930	2,058	2,116		365	431	522	540
	DE		-6.30	-8.54	-8.93	-8.77		-6.37	-8.12	-8.02	-7.98		0.17	-4.52	-7.98	-5.76
									Panel C: Learning Only							
Loss	Quit	48.13***	37.30	31.47	26.40	26.30***	40.43	31.09	24.41***	18.37	23.13	26.59	24.75***	19.83	14.29	14.29
	Safe	24.21***	30.48	37.66	40.21	44.65	17.22	20.04	26.60***	31.29	34.33	10.90	13.56***	18.18	15.87	38.10
	Risky	27.65***	32.22	30.88	33.38	29.05	42.35	48.87	48.98***	50.34	42.54	62.51	61.69***	61.98	69.84	47.62
	N	12,914***	4,019	1,519	659	327	5,098	1,592	639***	294	134	835	295***	121	63	21
Gain	Quit	23.67	10.39	6.20	3.58	2.56	17.47	7.31	3.50	2.07	2.25	14.03	5.09	3.36	0.83	1.64
	Safe	21.99	18.91	14.72	10.68	8.34	13.99	11.38	8.50	6.21	2.90	7.69	5.71	4.26	3.32	2.30
	Risky	54.34	70.69	79.08	85.74	89.10	68.54	81.31	88.00	91.72	94.85	78.28	89.20	92.38	95.84	96.07
	N	10,368	4,204	2,208	1,339	899	5,581	2,724	1,741	1,208	932	1,119	648	446	361	305
	DE	-24.46	-26.90	-25.26	-22.82	-23.74	-22.96	-23.79	-20.91	-16.30	-20.88	-12.56	-19.65	-16.47	-13.45	-12.65

Agents are also less likely to quit with a large loss or gain than with a small loss or gain as evidenced by the lower attrition rates in the higher stake categories. The aggregate disposition effect measures are not positive for the higher categories of gains and losses as in Nolte (2016), however they are less negative than the smaller gain and loss categories. We can now address changes in risk profiles in addition to attrition. Conditional on staying, winners are more likely to place a risky bet than losers, contrary to the disposition effect prediction i.e. all 'risky' (safe) percentages are higher (lower) for wins than losses. Larger losses and gains make agents more likely to place a risky bet than with small losses and gains. In effect, the higher the gain, the riskier the agents bet, as evidenced by the higher percentage of winners moving into riskier bets as gains become more substantial. Again, this is not consistent with prospect theory. In the domain of losses, the probability of risky bets is lower than in the domain of gains. In this panel, therefore, the evidence in favour of a prospect theoretical explanation is contradictory.

In Panel B, having removed the effect of the learning cohort, losers are now only marginally less likely to place a risky bet than winners, which is in contrast to the behaviour of the full sample in Panel A, where there was more risk seeking in the domain of gains. In Panel B, there is more risk seeking than Panel A in the losing domain and more risk aversion than Panel A in the domain of gains. Indeed, the probability of booking gains and holding on to a loss is marginally lower in this 'weak reinforcement' group. This is as a result of the removal of the effect of the agents presented in Panel C, where losers (winners) are much less (more) likely to place a risky bet than in Panel A. It seems that when agents subject to strong reinforcement are removed, winners are more likely to book gains in Panel B than Panel A, however the effect is not significant and the disposition effect measure remains negative.

In Panel C, with strong reinforcement, the probability of quitting (staying) is higher in the losing (winning) domain compared to Panel A. In terms of the size categories, bigger wins cause agents to keep betting but with more risk seeking behaviour. On the loss side, there is more quitting and less risk seeking than Panel A. In the domain of gains, we see less quitting and more risk aversion than in Panel A. We can summarise the behaviour in all three panels as follows:

- Across all panels, agents in the losing domain are more likely to quit than winners (this effect is strongest with the learning group in Panel C). This is evidence supporting the rejection of Hypotheses H1 and H1b.
- Conditional on staying, losers are more likely to place riskier bets, with

winners more likely to place safer bets, as predicted by the disposition effect. This is evidence in favour of hypothesis H1a and H1b that losers place riskier bets than winners.

- In the weak reinforcement group (Panel B), there has been a shift in risk taking. Conditional on staying, losers place safer bets while winners place riskier bets.

- In the strong reinforcement group (Panel C), there has been an equivalent shift in risk taking. Losers place safer bets while winners place riskier bets.

There may be a number of explanations for this behaviour. Agents face the decision to quit, continue with more risk or continue with less risk after each bet. However, there may be an editing phase (similar to the coding, combination and cancellation operations proposed by Kahneman and Tversky (1979)), where this decision is framed as the decision to quit or stay, followed by the secondary decision to increase or reduce risk. If so, reinforcement learning may dominate the primary decision, while prospect theory preferences come into play when the decision concerns the choice of bet price.

Panel B has shown that those who have not been subject to reinforcement learning become more risk seeking in the face of losses, and slightly more risk averse in the domain of gains. Removing the reinforcement learning group has reduced the aggregate measure of the disposition effect in Panel B. Panel C has shown that those who have been subject to reinforcement learning quit in the face of losses, but do not quit while ahead. However, agents in the reinforcement group who have not learned to quit in the face of losses become more risk seeking, while those that have not learned to quit while ahead become more risk averse. The aggregate measure of the disposition effect is more negative for the survivors in this group.

There seems to be an interaction between the disposition effect and learning. Aggregate behaviour, which was examined in the Odean (1998a) and Barberis and Xiong (2009) papers, changes when pure reinforcement learners are removed. There may be a tug of war effect between the disposition effect and learning, but further analysis is necessary.

6.4.2 Multivariate results

In the previous section, we had to reject the hypotheses that losing bettors in the reinforcement cohort were more likely to change to riskier bets than winning bettors. This result was not consistent with a prospect theoretical explanation, however was consistent with competing empirical evidence from the literature on 'decisions from experience' outlined in Section 6.2.

In related literature, one possible reference point for a stock trade is the purchase price. While encoding gains and losses relative to a reference point, however, there are multiple possibilities for the reference point. In a stock purchase, it could quite easily refer to the previous day's close, to the high or low of the year, or indeed to any arbitrary price level established by way of technical analysis, for example. As a reference point, I use the cumulative loss or gain to date for bettors, *breakeven*.

In the analysis that follows, I focus on the saliency of the breakeven point for bettors. *Breakeven* is the equivalent amount an agent would have to wager (given the bet price being executed against) to neutralise losses incurred up to the current bet[1]. The intuition behind this choice is that bettor who have experienced losses may focus on their account balance to date, and base their future trading decisions around trying to recoup losses which have been incurred up to this point.

If behaviour in the losing domain is being driven by the disposition effect, cumulative losses to date should be a predictor of future betting behaviour. If we include cumulative losses as an independent variable in a regression where the dependent concerns the choice to quit or change risk, if the coefficient on cumulative losses is significant, the disposition effect argument may find support[2].

I introduce the natural split in the dataset at this point. Wright (2008) suggests an explanation based on responsibility, entrapment, escalation of commitment and locus of control. According to the author, those with an internal locus of control are more likely to feel responsibility, and hence a greater need to self-justify and consequently a larger disposition effect. Chui (2001) suggest that locus of control can partly explain the disposition effect in a modified Weber and Camerer (1998) experiment. As a result, we may some reason to believe that there may be a different effect in a market setting where agents have no reason to perceive a locus of control. As returns are uncorrelated in the simulated market, agents may perceive an external locus of control in this setting. As a result, we therefore do not expect to find evidence of such behaviour in the simulated market.

[1]For example, having lost a number of 5/1 bets and ending up with a negative account balance of -\$500, an agent would have to wager \$100 (given the return of 500% on a 5/1 bet) to recoup the cumulative loss to date at the next bet. The equivalent breakeven point for such losses accumulated with 5/6 bets would be \$602.41. There is no breakeven point for bettors before the first bet.

[2]As stated in Weber and Camerer (1998), this is actually a joint test which includes a test for the correct specification of the reference point.

Bettors who show a preference for financial markets may perceive an internal locus of control, and may have feelings of disappointment at losses and elation at wins. Bettors who bet on the simulated market, however, may not suffer from the same salient emotions, as they may perceive less responsibility for outcomes if they know they are uncorrelated. In that sense, the simulated market bettors may be as removed from the act of trading (and its associated biases) as are investors who delegate financial decisions to investment managers (and who show exhibit an inverse disposition effect in Chang et al. (2013)). In both cases, there is a shift in the locus of control: from investor to investment manager; from simulated market bettor to 'the lap of the gods'[1]. Further distinction can be made between the winning and losing domains. Summers and Duxbury (2012) show that regret (a product of responsibility) is necessary to induce investors to hold losing positions, elation alone (rather than rejoicing as a result of responsibility for a positive outcome) can provoke them to close winning positions.

As such, the results presented earlier may be being confounded by aggregating the sample. I therefore re-introduce the split according to product I presented in the previous chapter, and investigate whether learning or the disposition effect is driving behaviour in each group in a multivariate setting in order to control for a number of variables. For winners, the reflection effect of prospect theory predicts attrition. While any evidence of a disposition effect will manifest itself in attrition, we still expect to see aggregate reductions in stake size across that group, as those who leave the sample will have a stake size of zero. Again, while the literature predicts this to be the case for financial bettors, I hypothesise that this will not be evident for bettors on the simulated market.

In order to frame the next section in the literature, Figure 6.2 presents some assumptions governing the decisions faced by bettors. $S_i(1)$ and S_in refer to a bettors initial stake size and subsequent stake sizes, respectively. Again, the reason I focus on the first five bets is that the mean and median number of bets placed is the sample are 4.9 and 1. Gain and Loss refer to a series of winning and losing bets prior to execution of the fourth bet. $S^*{}_{breakeven}$ is the equivalent amount an agent would have to wager in the next bet to neutralise losses incurred up to that point.

At each round of betting, agents who have experienced gains are faced with the decision whether to bet again or whether to exit the market. I argue that those who decide to quit are realising gains, and such agents should be present in the subsequent bet with a stake size of zero. Those who bet again with a higher stake size than the first bet have been subject either to naive reinforcement learning by

[1] Langer and Roth (1975) named this mentality: 'Heads I win, tails it was chance'.

Figure 6.2

Decision paths for learning and the disposition effect

The following figure presents the decision faced by agents at each round of betting. S_{t-1}, *Gain* and *Loss* refer to a bettors previous stake sizes, and the gain or loss associated with the bet, respectively. Given either a gain or loss, each agent is faces with the decision to bet again (with an risky payoff) or quit (with a certain payoff). $S_t * (Bet\ price)$ refers to the net payoff resulting from the stake size and chosen bet price.

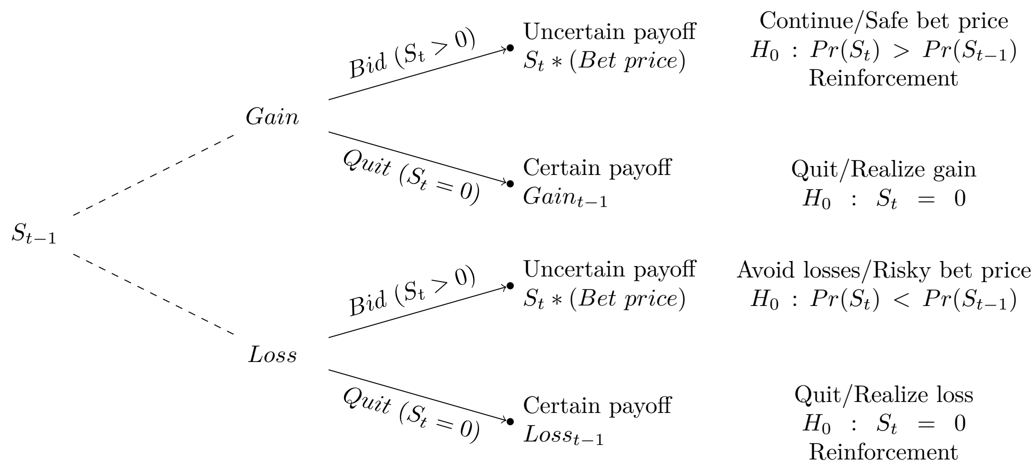

increasing their stake sizes as a result of positive reinforcement, or indeed rational Bayesian learning and have updated their stake sizes accordingly.

In the losing domain, those who bet again, and where their decision is driven by the cumulative loss to date, are showing behaviour consistent with prospect theory. If we find that the breakeven point for such agents is salient and significant, we may argue that they are remaining in the market and increasing risk in order to neutralise losses to date. For those agents in the losing domain that quit, I argue have either been subject to naive reinforcement in that they are avoiding behaviour which has been associated with negative feedback in the past, or are simply learning about their ability in a rational Bayesian manner and are quitting the market.

The unique characteristics of the simulated market facilitate disentangling some of the contradictory explanations above. In such a setting, there can be no rational Bayesian updating. Bettors in the winning domain who continue to bet, and who bet more than their initial stake size, can only have been subject to naive reinforcement learning. Those who quit are realising gains. In the losing domain, those who quit have been subject to negative reinforcement, while those that continue, and bet more than the cumulative loss to date are loss averse.

There is a caveat, however. In doing so, I assume that agents have cumulative losses to date as the reference point during the 'editing' phase of their betting decision. This may not be the case, and the reference point could be an arbitrary number exogenous to prior betting behaviour such as the amount an individual bettor would like to win or lose before ceasing to bet[1].

To analyse how the cumulative loss to date affects the choice of the next move for bettors, rather than using stake as the dependent, I model each subsequent decision directly as an ordered outcome dependent variable. In Table 6.2, the equivalent of this multinomial dependent variable had three categories: *quit*, *safe* and *risky*. In this section, I have coded the dependent variable as follows: *Quit=1*, *Safe=2*, *Risky=3* and estimated an ordered logit model. The variable of interest is the account balance to date, breakeven, which is alternative invariant, as the breakeven point will not change according to the choice made by bettors.

Table 6.3 presents a logistic regression estimation with odds ratios instead of coefficients (i.e., exponentiated coefficients $exp(\beta_1)$). The interpretation of the results concerns the effect of relatively subtle changes of the cumulative loss variable of dollar units, however we can also interpret the effect of x-unit dollar changes in the cumulative loss by taking $exp(x\beta_1)$, or simply or^x. Model (1) includes the entire sample, Model (2) is estimated without agents who have experienced consecutive strings of wins or loss, while Model (3) concerns only agents who have been subject to strong positive or negative reinforcement.

The odds ratio in Model 1, regressing the bettors next move with the cumulative loss or gain to date, is significant, but has a small economic effect. Even a ten-unit increase in the cumulative loss only increase the odds of in the odds of moving to higher categories of the *next move* variable (i.e., from *quit* to *safe*, or from *safe* to *risky*) by a factor of 1.07. In the case of agents with mixed experiences of wins and losses, it is not significant.

A one unit increase in the round of betting indicates that the odds of choosing a higher category of risk seeking behaviour are 1.12 times greater, suggesting that each additional round of betting in which agents participate increases their propensity to choose more risk seeking options (i.e., moving from exiting the market to continuing and placing safe bets, or changing from placing safe bets to placing riskier bets). It seems that the proportional odds ratio on the account balance to date does not affect risk and that this variable is not salient to bettors.

However a crude variable (not tabulated) indicating the sign of the bettors account balance[2], $Win/Loss$, indicates that a positive account balance to date increases the odds of risk taking by a factor of 1.36 for the entire sample. This

[1] We return to this issue in Section 6.4.

[2] Account balance in this case is the cumulative P&L on the bettor's account to date, rather than any net cash balance they hold with the bookmaker.

Table 6.3

Ordered Logit Regression

The dependent variable here is ordered in terms of risk seeking, with exit from the market coded as 1, placing of a *safe* bet coded as 2 and placing of a *risky* bet coded as 3. Odds ratios rather than coefficient estimations are output. *Cumulativeloss* refers to a bettor account balance. *Bet number* is the sequence in rounds of betting. *Pure financial* and *Pure simulated* refers to bettors who bet only on the financial market outcomes and on the simulated market, respectively. The base category in this case is bettors who bet on both types of market. Model (1) includes all bettors. Model (2) consists of bettors who had mixed experiences of wins and losses. Model (3) includes only bettors who had strings of wins and losses.

	(1)	(2)	(3)
Cumulative Loss	1.0007***	1.0001	1.0012***
	(3.98)	(0.39)	(3.59)
Bet number	1.1152***	1.0661***	1.0940***
	(56.34)	(21.94)	(19.47)
Fins/Vir=1	0.5342***	0.6490***	0.4042***
	(-47.36)	(-24.54)	(-42.80)
Fins/Vir=2	0.5094***	0.8106***	0.3167***
	(-37.68)	(-8.57)	(-42.37)
1.83	1.0693***	1.2302***	1.0050
	(5.19)	(9.58)	(0.26)
2	4.7307***	8.3570***	2.3355***
	(86.20)	(79.27)	(30.37)
6	3.4626***	9.1162***	1.5383***
	(60.67)	(70.71)	(15.45)
9	13.2996***	36.8445***	5.9728***
	(13.51)	(11.25)	(7.00)
Observations	166,597	97,291	69,306
R2	0.063	0.102	0.031

variable showed no significant change for agents with mixed wins and losses in Model (2) but is positive and significant for the learning cohort in Model (3). For those who have been subject to reinforcement, having a positive (negative) account balance increases (decreases) the odds of being in higher categories of risk by a factor of 3.22. This suggests that it is not the magnitude of the gain or loss to date that is salient, but simply whether it is positive or negative.

The *Pure*$-$ variable is equal to one if a bettor transacted only on Financial markets, equal to 2 if a bettor transacted only on the Virtual Market, with those that bet on both as the base category. There is no significant difference between the odds ratios on either of the Financial Market or Simulated Market flags, however compared to the base category of those bettors who transacted in both markets, the odds are approximately 0.64 to 0.72 times lower of becoming more risk seeking.

With regard to the proportional odds ratios on bet prices, compared to a bet on a 1/20 outcome (the base category), bets on 5/6 outcomes are 1.16 to 1.44 times more likely to lead to more risk seeking, whereas the skewed, asymmetric bet prices such as Even Money and 5/1 lead to odds of up to ten times higher than 1/20 bets in the learning only cohort. Given that the initial choice of such outcomes are risky in nature, it is natural that such choices would lead to further risk seeking behaviour in subsequent bets. In the next section, I re-classify the dependent variable to check whether this is the case.

6.4.3 Robustness: Sophistication and 'near misses'

Prior literature indicates that more sophisticated investors have less propensity to suffer from biases such as overconfidence or the disposition effect. I incorporate this into the analysis by using two proxies for sophistication. Firstly, bettors in this sample can execute bets online, via the company's mobile interface or over the phone. I now make use of this variable in order to test whether there is a difference in measure for online bettors. Table 6.4 presents a breakdown of stake size by bet channel in the dataset.

Table 6.4

Bet channels

This table presents stake sizes according to the channel through which bets were struck i.e. online, through the mobile interface or via a phone operator. All stakes are in USD.

	# Bets	**Stake**	**Mean**	**Std**	**P1**	**P50**	**P99**
		Panel D: Stake by Bet Channel					
Internet	1,658,571	73,007,125	44.02	195.90	0.08	8.80	600.00
Mobile	26,266	1,690,496	64.36	150.47	0.20	20.00	720.00
Phone	7,415	1,292,588	174.32	373.89	1.20	48.00	1,857.78

Clearly, online investors are in the overwhelming majority in the dataset, however there are some interesting details in the distribution of stake size. Bettors who transact bets over the phone place larger stake sizes. It is not clear whether this is due to specifics of phone bettors as opposed to online bettors, or whether there is a simpler explanation. Phone bettors may be allowed to transact higher stakes, for example, however there is no way of verifying this. Nevertheless, I add dummy variables for phone betting in the next model estimation.

Bet complexity may also proxy for sophistication. Firstly, an understanding of the probabilities associated with complex bet types may not be within the grasp of relatively unsophisticated agents. Secondly, sophisticated agents may wish to combine large bets which contain private information with vanilla bets on sporting outcomes with low stake sizes in order to obfuscate such behaviour. Table 6.5 presents summary stats on single bets, doubles and more complex accumulators. There were no double, treble or other accumulators in the dataset with more than one Financial Market or simulated market bet in any of the legs, therefore all non-singles bets contain bets on sporting outcomes[1]. As can be seen in Table 6.5, the overwhelming majority of bets were singles. As doubles, trebles and other accumulators represented such a small part of the dataset, this variable could not be used as a reliable proxy and as a result, the bet channel alone was used as a single proxy for sophistication.

Table 6.5

Bet complexity

This table presents stake sizes by bet type. All stakes are in USD. All non-singles bets contain one leg with a financial market or virtual market bet and one or more legs with sporting outcomes.

	# Bets	**Stake**	**Mean**	**Std**	**P1**	**P50**	**P99**
Panel D: Stake by Bet Channel							
Single	1,686,283	75,898,146	45.01	196.94	0.08	9.60	600.00
Double	3,451	62,936	18.24	54.21	0.02	3.00	300.00
Treble	1,185	19,870	16.77	46.12	0.01	3.05	240.00
Accumulator	1,151	8,121	7.06	20.74	0.01	1.60	100.00
Trixie	63	590	9.36	28.47	0.01	0.03	160.00
Lucky 15/31/63	39	209	5.37	6.08	0.03	3.00	30.00
Yankee	35	162	4.63	3.93	0.02	3.30	16.50
Patent	21	98	4.66	4.88	0.28	1.87	19.60
Canadian	12	41	3.46	2.41	0.06	4.53	6.24
Heinz/Super Heinz	12	36	3.00	2.96	0.76	2.28	11.40

I also address the hypothesis on near-miss outcomes in this section. Prior literature suggests that agents are more likely to escalate commitment if they perceive their particular goal to be closely within reach. I test this premise by examining *near-miss* bets. I define a near-miss event as a bet where the settlement price was within a certain range of the strike price of the bet and present histograms on near misses in Figure 6.3.

A possible weakness in the previous robustness tests was the strict ordering of the dependent variable into the *quit*, *safe* and *risky* categories. The intuition behind this order was to rank the dependent variable from the most risk averse to the most risk seeking alternative. In this section, however, I estimate a multinomial logit model in which the dependent variable is considered free of any rank or

[1]It is possible that because of the high correlation between such bets, bettors may not combine financial bets but must execute multiple 'singles' bets simultaneously.

Figure 6.3

Histograms of Strike - Settlement

The following histograms present the difference between bettors' strike prices and subsequent settlement prices.

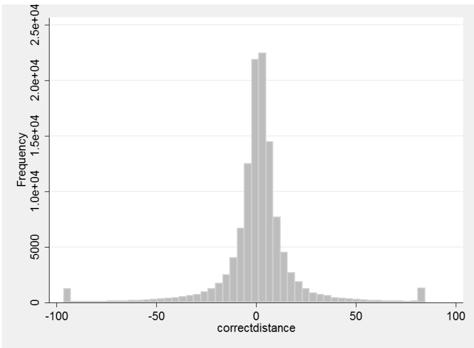

(a) DJIA

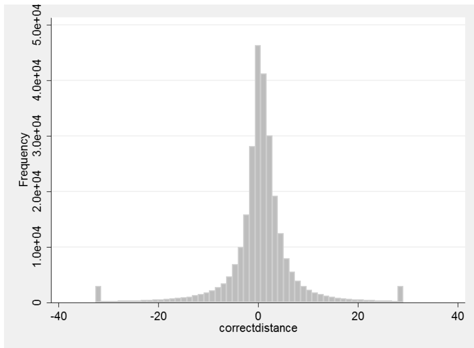

(b) FTSE

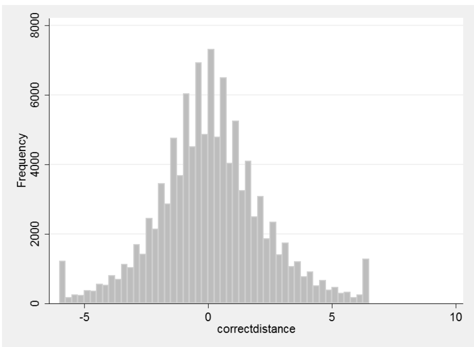

(c) Virtual Market

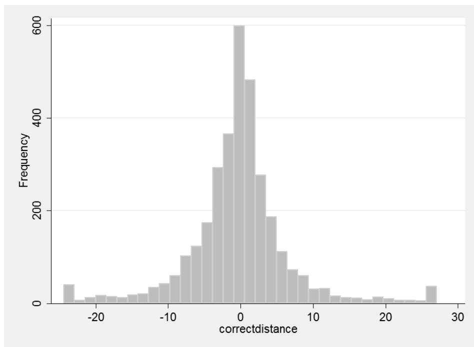

(d) S&P 500

order. I also re-classify the dependent variable in order to add a fourth category. The dependent variable encompassing bettors' subsequent decisions is now coded as follows: *1=Quit*, *2=Safer*, *3=Riskier*, *4=Same.* In contrast with the previous section, the dependent variable is coded as 2 if the implied probability of the bet is higher than the previous, and as 3 if the probability is lower. If bettors do not have a natural propensity to change bet prices and the choice of the next bet price is highly correlated with the previous, this re-classification should pick this up.

I report relative risk ratios (RRR) rather than coefficient estimations here, in order to be able to interpret the ratio of choosing each dependent variable category over choosing the base category of 'same'. I present the results in Table 6.6.

Table 6.6 contains a panel for each level of the dependent variable in comparison with the base category of *same*. The first model is estimated for all bettors, the second for bettors that experienced mixed results and the third for bettors

Table 6.6

Multinomial Logit Regression

The dependent variable here takes the value 1 if a bettor quit after the current bet, 2 if the next bet is safer than the previous, 3 if the next bet is riskier. The base category is continuation with no change of bet price. *Dial-a-bet* refers to bet transacted over the phone, whereas *Mobile* indicates bets that were placed using mobile devices. The base category in this case is online bets.

	(1) Next move		(2) Next move		(3) Next move	
Quit						
Fins/Vir=1	0.124***	(4.87)	−0.258***	(−8.68)	0.992***	(19.01)
Fins/Vir=2	0.563***	(20.25)	−0.0259	(−0.76)	1.560***	(28.61)
Dial-a-Bet	−0.667***	(−4.94)	−0.536***	(−3.34)	−0.927***	(−3.70)
Mobile	−0.0902*	(−2.04)	−0.183**	(−3.13)	0.0737	(1.04)
Near miss=1	0.963***	(75.37)	0.488***	(28.65)	1.683***	(84.22)
Riskier						
Fins/Vir=1	−0.695***	(−24.96)	−0.605***	(−16.73)	−0.817***	(−18.63)
Fins/Vir=2	−0.172***	(−5.44)	−0.0696	(−1.70)	−0.318***	(−6.38)
Dial-a-Bet	−1.648***	(−4.88)	−1.752***	(−4.24)	−1.444*	(−2.46)
Mobile	−0.461***	(−6.28)	−0.545***	(−5.83)	−0.313**	(−2.65)
Near miss=1	0.0492**	(2.76)	−0.0363	(−1.60)	0.195***	(6.48)
Safer						
Fins/Vir=1	−0.620***	(−21.76)	−0.536***	(−14.55)	−0.669***	(−14.69)
Fins/Vir=2	−0.131***	(−4.09)	−0.0645	(−1.55)	−0.172***	(−3.39)
Dial-a-Bet	−2.189***	(−5.31)	−2.002***	(−4.43)	−2.855**	(−2.83)
Mobile	−0.492***	(−6.78)	−0.605***	(−6.46)	−0.285*	(−2.46)
Near miss=1	1.308***	(68.99)	1.034***	(42.36)	1.765***	(58.09)
Observations	166,153		96,992		69,161	
R2	0.035		0.018		0.082	

who experienced strong reinforcement. For near miss outcomes relative to bets that settled far from the strike price, the relative risk of quitting relative to placing safe bets increases by a factor of 1.683, while the equivalent risk of placing a risky bet has no significant difference with the base category. The relative risk remains positive and significant irrespective of market preference, however is no different from zero for the *risky* category relative to safe bets.

6.5 Conclusion

What can be done to attenuate this bias? Given the evidence shown above, it seems that introducing further reinforcement learning cues to a market setting will (a) attenuate the level of the disposition effect by making losers quit and winners continue, however (b) will amplify the disposition effect by encouraging more risk seeking in the losing domain for those who do not learn to quit in the face of significant losses. Conversely, reducing the amount of reinforcement learning cues will (a) encourage less attrition by both losers and winners (possibly enhancing liquidity in markets in aggregate) while (b) encouraging more risk aversion in the face of losses and risk seeking in the face of gains.

As stated earlier, the definition of paper gains and loses are not trivial when dealing with terminal assets such as financial bets. In addition, given a strict interpretation of Shefrin and Statman (1985), the disposition effect refers to actions taken in response to an active position in financial stocks. The arguments in the hypothesis development section of this chapter define risk seeking in the domain of losses and risk aversion in the domain of gains as being similar to the disposition effect seen in stocks. However this is open to debate. To remain consistent with the original definition of the disposition effect, it may be instructive to refer to changes in risk preferences as a result of differing paths of wins and losses rather than the existence or absence of the disposition effect in future analysis.

In Chapter 5, we saw that there was evidence of rational Bayesian learning in the financial markets, however reinforcement learning in the simulated market. I used attrition rates and stake size changes to test the hypotheses. In this chapter, I introduced the disposition effect and predictions from prospect theory into the analysis, and examine changes in risk profiles in addition to stake sizes. I suggested that there is an interaction between learning and the disposition effect. When simulated market bettors and financial markets bettors were isolated, we saw some support for the hypotheses, and though the evidence was mixed, those subject to strong reinforcement showed no sign of a disposition effect.

I have compared and contrasted reinforcement learning and Bayesian learning in a number of settings, but the result that the disposition effect may be attenuated by reinforcement cues is a more interesting result. Given the evidence contradicting the disposition effect we have seen in the those subject to reinforcement, it begs the question whether this finding is particular to this setting or whether it could be generalised? To disentangle reinforcement learning from the disposition effect, I proceed with a focused experiment in the final chapter, it an attempt to discern which effect is the stronger. If we subject participants in an experiment to a setting where than can perceive no locus of control and in which they have successive wins or losses, will winners quit or continue? Will losers 'make peace with their losses' or exhibit a disposition to hold on to them? This forms the basis of the final empirical chapter.

Chapter 7

Learning and loss aversion: Experimental evidence

7.1 Introduction

This chapter describes an experiment designed to examine whether reinforcement learning forces individuals in the losing domain to quit, or whether loss aversion entices them to continue. Equally, in the winning domain, I investigate whether successive winning bets cause winners to continue to bet, or whether their book their gains and exit the market. In contrast with the previous two chapters, rather than using a panel dataset, I employ methods from experimental economics to address the hypotheses. Thus, the dataset for the current chapter was collected in the lab, rather than in the field.

Homo oeconomicus is expected to evaluate economic decision using the expected value rule and to rank alternative options in strict order. Bernouli (1738) challenged this to suggest a concave utility function and a decreasing marginal utility and addressed some of the behavioural issues raised by the St. Petersburg Paradox. Friedman and Savage (1948) explained the purchase of insurance and lotteries by way of a utility curve with both convex and concave regions, however Markowitz (1952, p. 152) argued that in some regions it implies 'behaviour which is not only not observed but would be generally considered peculiar if it were'. He suggested a utility curve which was convex immediately below present wealth and concave about it, but concave and convex deep in the winning and losing domains, respectively. This indicated moderate risk seeking for moderate winners, however risk aversion for heavy winners, while at the same time predicating moderate risk aversion for moderate losers, but risk seeking by heavy losers.

Empirical research in the area of behavioural economics and finance has brought normative predictions from Utility Theory into question. In particular, the observed behaviour in the winning and losing domains of economic decisions has received continued attention in the literature. This anomaly was coined the reflection effect by Kahneman and Tversky (1979) and was outlined as part of their Prospect Theory framework. In empirical research, a number of reasons have been suggested for winning investments being sold more readily than losers, for example: differences in utility functions in the domain of gains and losses, belief in mean reversion, portfolio rebalancing and avoidance of higher transaction costs on losing stocks.

Kahneman and Tversky (1979) proposed a framework based on gains and losses, where a decision weight is applied to each probability and a value weight applied to each outcome in the expected utility function. They suggest that agents are risk-seeking in the loss domain and risk averse in the domain of gains relative to a reference point. However, Thaler and Johnson (1990) analysed decision making in the presence of prior outcomes, concluding that there was risk seeking in the presence of a prior gain, risk aversion in the domain of losses, while noting a particular attraction in the losing domain for break-even prospects.

The loss aversion component evident in the shape of the Prospect Theory utility curve has motivated empirical and experimental research on investor behaviour. Shefrin and Statman (1985) show that investors book gains too early and hold on to losses too long, coining the phrase the disposition effect. Odean (1998a) tested for the existence of the disposition effect, using a stock's purchase prices as the reference point, and concluded that investors exhibit behaviour consistent with Shefrin and Statman (1985), even when controlling for other behavioural explanations such as mean reversion, portfolio rebalancing and optimal tax strategies.

Central to the explanation for the existence of this bias in Shefrin and Statman (1985) are the reference-point effects described in Kahneman and Tversky (1979) and the concept of mental accounting outlined in Thaler (1985).Weber and Camerer (1998) propose the reference point and reflection effects as explanations for the disposition effect at the aggregate level. They present an experimental setting in which subjects sold fewer shares when both the purchase price and previous prices were lower than the current price, suggesting two possible reference points. Although price changes in their experimental design were positively autocorrelated, precluding mean reversion, they found that agents bought losers and sold winners. Related literature has examined numerous alternative explanations for this bias, including Odean (1998a), who rejects portfolio rebalancing and avoidance of trading costs as reasons, however does not disentangle the reference-point effects of prospect theory from a belief in mean reversion.

The range of empirical evidence for behaviour consistent with loss aversion is wide. Loss aversion manifests itself in the endowment effect. Agents are willing to pay less for items than they are willing to accept to part with them, a phenomenon known as the endowment effect. Camerer, Babcock, Loewenstein, and Thaler (1997) show that New York taxi drivers are loss averse and exhibit wealth-reducing target-seeking behaviour, which they learn to avoid with experience. Also young children have been shown to exhibit the endowment effect (Harbaugh, Krauseb, and Vesterlund, 2001) and loss aversion has even been extended to non-monetary outcomes (Blavatskyy, 2011). Indeed, in a novel experiment, (Chen, 2006) introduced fiat currency to a group of Capuchin monkeys and found that they exhibited biases such as reference dependence and loss aversion.

In short, the range of evidence for this bias is compelling. However, Apicella, Azevedo, Christakis, and Fowler (2014) present evidence that such biases may not be universal. They examine one common heuristic, the endowment effect, with a natural experiment on a hunter-gatherer population in Tanzania. They found that members of this indigenous people living in isolation were willing to trade items with a probability of 50% (no endowment effect), whereas those who had increasing contact with the outside world were willing to part with goods with only a 25% probability. They note that this is in contrast to other studies such as Harrison and List (2004) in that the least experienced agents show the most rationality, as opposed to those with experience who may have 'unlearned'' it. Their experimental design suggests that cultural factors can have an influence however questions posed by Dommer and Swaminathan (2013) related to social self-threat and identity reveal that we should not hasten to underestimate the prevalence of such biases.

While Prospect Theory has been successful addressing some of the empirical regularities that are inconsistent with the Expected Utility Theory, it still leaves many questions unanswered. In particular, how should investors apply Prospect Theory following prior gains/losses (Thaler & Johnson, 1990)? Do they adapt to prior gains/losses or see them as part of the next investment? Odean (2002) provides evidence that investors more readily realise paper gains by selling their 'winning' stocks, but hang on to their 'losing' stocks longer. This is consistent with loss aversion, which is a prediction of Prospect Theory. However, would an investor behave the same way if he or she were subjected to strong reinforcement? For example, how long would the investor hold on to a stock that is day-by-day losing value? Conversely, would an investor rush to sell a stock that has yielded positive returns in each month during the past year? In fact, Barberis and Thaler (2013) state that applying prospect theory to the laboratory is confounded by the fact that Kahneman and Tversky (1979) gave little guidance on how to determine the reference point.

There is ample evidence of loss aversion in the literature, and in particular, evidence supporting investors tendancies to exhibit asymmetric responses to losses and gains, or the 'predisposition toward get-evenitis' (Shefrin, 2000), formally labelled the disposition effect in Shefrin and Statman (1985). However, there is a paucity of evidence relating conditions of strong reinforcement with loss aversion.

In this chapter, I define loss aversion as the stronger psychological effect of losses than gains documented in Kahneman and Tversky (1979), manifested in a steeper value value function for losses than gains ($\lambda = 2.25$). As pointed out by Brink and Rankin (2013), however, risk preference in the losing domain and loss aversion are distinct characteristics, as one can be loss averse even in the absence of risk if one prefers a certain gain of $20 over a certain loss of $20. Some of the literature argues that risk preferences are not stable over time, with Baucells and Villasis (2010) presenting experimental evidence that 63% of subjects changed their risk preferences over time.

Reinforcement learning theory, or the 'law of effect', dictates that agents will repeat behaviour that has been associated with positive feedback and avoid behaviour that has resulted in negative feedback. It dictates that agents should stick to given choices as long as they generate rewards, otherwise they should switch (Roth and Erev, 1995). Rational learning incorporates both private signals and public information, updating beliefs about payoffs accordingly. For example, Bayesian learning refers to weighing both 'experienced' and 'observed' outcomes equally, whereas reinforcement learning over-weighs experienced outcomes. In contrast with a pure 'stay/switch' reinforcement model, Bayesian belief-learners rationally learn from experience (Camerer and Ho, 1999).

Kaustia and Knüpfer (2008) examine the relationship between returns on previous IPO subscriptions and the likelihood of subsequent participation in further IPOs. They conclude that personally experienced returns are an important determinant of future activity which is consistent with reinforcement learning theory. Seru et al. (2010) investigate how individual investors are affected by two distinct types of learning: learning about their own abilities and learning by trading. They conclude that investor performance improves with experience, however they highlight that attrition due to investors learning about their lack of ability may be an important factor. Chiang et al. (2011) examine whether investors improve their ability by rational learning or whether their performance deteriorates due to reinforcement learning. They contend that the Kaustia and Knüpfer (2008) results are also consistent with rational Bayesian learning, as those investors who experience positive returns will tend to participate more often than those who have experienced negative returns.

Chapter 5 concerns the interaction between reinforcement learning and the disposition effect. By using a natural experiment consisting of the betting activities of in excess of 15,000 agents, it suggests that reinforcement learning dominates predictions made by Prospect Theory in agents who have been subject to purely positive or negative reinforcement. A natural extension of this research is to conduct a focused experiment to disentangle the effects of reinforcement learning and loss aversion in the laboratory. That is the focus of this chapter.

This study contributes to our understanding of decision making under risk. In an experiment, I find some support for the hypothesis that subjects exposed to a reinforcement-learning treatment are more prone to exhibiting behaviour consistent with loss aversion than the control group. I endowed subjects with £25, presented them with a trading interface and asked them to make a series of bets. The trading interface contained a financial chart with an underlying time-series generated by a random-number generator. All bets were fair and there was no commission, margin or over-round. They were faced with a decision node after the second and subsequent bets at which point they could decide to stop betting and be paid their earnings to date, or continue betting with a chance of losing or winning more.

The structure of the rest of this chapter is as follows. The hypotheses are outlined in Section 7.2 where I link the tests to the literature. Thereafter, Section 7.3 outlines the experimental design and presents the experimental interface the subjects interacted with. A discussion of the results is contained in Section 7.4. Section 7.5 concludes.

7.2 Hypotheses

The disposition effect argues that winners book gains and losers hold on to losing positions. These are strong temptations with extensive empirical support. However, empirical studies have also shown that individuals are subject to reinforcement learning, such that they repeat actions which have resulted in positive feedback and shun actions which have resulted in negative feedback. In this case, what of agents who have experienced a number of days of positive returns in a row or who have had 'winning streaks'? I argue that reinforcement learning is a salient bias which interacts with the disposition effect.

I argue that a group of agents who have experienced consistently strong positive or strong negative feedback (winning streaks or series of consecutive losses) are more likely to exhibit behaviour consisted with reinforcement learning than the disposition effect. The treatment group is therefore agents who have had successive wins or losses, while the control group consists of those subjects who have had mixed experiences of wins and losses. I therefore proposed the following hypotheses:

> **H1.** *Probability of attrition for winners in the treatment group is lower than in the control group.*
>
> **H2.** *Probability of attrition for losers in the treatment group is higher than in the control group.*

To reiterate the hypotheses, in the positive domain, there is a temptation to book gains, however there is also a temptation to continue repeating an action which has resulted in positive feedback. In the domain of losses, the thesis is similar: there is a temptation to hold on to a losing position, however there is also the temptation to walk away and desist from repeating an action which has resulted in successive negative feedback.

7.3 Experimental Design

This section describes the design of an experiment conducted in order to address the hypotheses. Experimental economics attempts to introduce subjects in a laboratory to conditions which mimic real-world decisions and endeavour to control for as many variables as possible. The function of an experimental design is to craft the instructions and environment subjects are presented with in the lab, in order to properly motivate the economic decisions they make and so that the results may be generalisable outside the laboratory. The motivation is not to replicate real-world decisions as closely as possible, but to structure the session so that conclusions drawn for subjects' decisions and actions in the lab are reasonable and credible.

A single experimental session in an economics laboratory with the same subjects will usually consist of a number of trials, each with a separate set of instructions. Within each trial, subjects are invited to participate in a number of replications of the same economic decisions. Some may involve individual decision making, while others may involve subjects collaborating in groups with the other group members known to subjects' or anonymous, where they interact with each other via an experimental interface. Full details of the experimental interface, including screenshots of the are presented in the appendix to this chapter in Section A.4.

In this setting, there was a single experimental session, effectively a pilot, consisting of a single trial with five periods. The institution in the experiment was a bookmaker, proxied by an experimental interface written in z-Tree (Fischbacher, 2007). The experiment encompasses the decisions made by 28 subjects interacting with the bookmaker. Additional treatments would have facilitated a within-subjects design, however with each agent designated to a control or treatment group only once, this was a between-subjects design.

In this pilot experiment, I only had the opportunity to administer a single treatment, therefore I controlled with randomization and by holding certain variables constant. Variables included stake size (held constant), the number of bets (a focus variable), random assignment to one of two treatment groups (or the control group) and the time-series underlying the individual charts presented to each subject, for each round (all 'nuisance' variables). Ideally, all variables affecting subject decisions would be controlled for by both directly and indirectly, however, I did not identify any confounded variables and variables such as stake size that were controllable were set as a constant. This mitigated any interaction between stake size and earnings on decisions to stop betting or continue betting. For example, an A-B-A crossover design with stake held constant in both A settings, while being allowed to vary in the B treatment would have been appropriate in this case.

There was no bankruptcy consideration in this setting. Although subjects could end up with a zero payment, by design, this could only happen at the end of the fifth bet, at which point all relevant decisions had been made. Moreover, there was no possibility of any subject losing more than their initial endowment. As subjects could not reach bankruptcy before the end of the experiment, there could be no extreme risk-seeking behaviour which may have taken place in the face of bankruptcy without the requirement of a net payment to the institution.

Risk is an intrinsic characteristic which, while it varies across subjects, is essentially fixed for each agent (Friedman and Sunder, 1994, p.44). I had envisaged conducting a Holt and Laury (2002) risk elicitation test during the questionnaire section of the experiment, but decided against this as subjects would be both fatigued after the experiment and also not subject to reward incentives directly related to the test. No incentives were offered for participation in or completion of the questionnaire. As differences in behaviour exhibited by agents with differing risk preferences was not the focus of the hypotheses, I controlled for risk preferences in the same way as other unobserved characteristics, by randomizing assignment to the treatment groups.

In this respect, I use a novel way to assign subjects to the control and treatment groups. Since the probabilities of winning or losing were equal at each round of betting (given that observations were drawn from a random number generator), each subject had the same probability of being assigned to each group irrespective of his/her relative risk preferences. In general, all unobserved heterogeneity between subjects was unobserved and uncontrolled. The unobserved characteristics may seem to present a problem, however, induced-value theory (Smith, 1976) posits that appropriately designed reward mechanisms can induce characteristics in subjects, rendering their unobserved differences irrelevant. A brief note on the conditions outlined by (Smith, 1976) as essential for the successful control of heterogeneous preferences, may be instructive at this point:

- **Saliency**: The text in the brochure recruiting subjects to register for the experiment and the poster (presented in Appendix A) highlighted the maximum payout to be earned by subjects. It was made clear that payouts would be based on subjects decision during the experiment. There was no indication that outcomes had been decided in advance, and the saliency of earnings were highlighted on the interface during the experiment (see Figures A.1 and A.1. The continuation screen in particular highlighted both the amount that subjects would be paid if they stopped betting, and also the amount the they could win or lose if they continued.

- **Dominance**: I did not offer a show-up fee as students were being offered a workshop on behavioural finance as part of the experimental session. I endowed subjects with £25. Given the probabilities associated with betting outcomes and the opportunity to stop betting after the second bet, this constituted expected earnings of £25 more than twice the opportunity cost even for students . In effect, as the payoffs were considerable, reward was considering to be the main motivating factor in subjects' decisions. In addition, I did not reveal any details related to the hypotheses, which may have caused certain subjects to attempt to pro-actively disrupt or aid in addressing the research question.

- **Monotonicity**: As subjects were offered monetary rewards, and that they did not (presumably) fit the profile of ultra-wealthy individuals, this property can be considered dealt with.

- **Parallelism**: The external validity of conclusions made via an experiment with so little subjects is bound to raise issues of external validity. In that respect, parallelism is crucial. Both the interface, probabilities of winning or losing, the betting proposition and even the random-number generated process creating the price series underlying the charts presented to subjects were designed to mirror the current betting offering of the bookmaker who provided the data for Chapter 5 and 6 of this thesis. A screenshot of the bookmaker's offering was shown to participants during the workshop following the experiment. Nevertheless, I did doggedly pursue realism with an attempt to fully replicate the bookmaker's offering: the specifics of the experiment were designed in order to address the hypotheses. Moreover, it would not be feasible to introduce the full range of characteristics of the market as it would have presented too-steep a learning curve for subjects.

- **Privacy**: Subjects were instructed not to talk during the experiment and this was made clear both in the written and verbal instructions. Results were not announced publicly. Subjects were not told the result of each bet verbally, nor was there any feature such as a score- or leader-board or tournament. Results were delivered directly to each participants screen. While it would have been possible for subjects at adjoining workstations to move and look at each other's screens to view their results, I did not observe this happening as a matter of course.

Friedman and Sunder (1994, p.13) state that if these conditions are satisfied, subjects' unobserved heterogeneous characteristics have been controlled for. A pilot experiment with eight PhD students was performed a week in advance of the experiment. The results of the dry-run were not analysed as they were doctoral students, and as such, potentially too familiar with the topic and could have induced the hypotheses.

7.3.1 Recruitment

Candidates were recruited during an event called Innovative Learning Week (ILW) at the University of Edinburgh. The experiment was intended to be run prior to an hour-long workshop on behavioural finance and experimental economics, where subjects would be told of the motivation behind the experiment and presented with some initial results based on the decisions they had made. Details of the experiment were published in the online and printed brochure distributed to university students and staff. A maximum of 40 places were made available and all candidates signed up online with their student or staff login details. A waiting list of approximately 40 more participants was compiled by staff in the ILW office. An email was circulated prior to the event requesting that subjects who could not participate on the day inform the ILW office by email and as a result, five subjects from the waiting list were contacted in advance of the event. A list of participant names and student numbers were circulated to the experimenters prior to the event.

Students as subjects

The subject pool involves only students. There are advantages and disadvantages to this. Firstly, in defence, I did not recruit any doctoral students who may have inferred the hypotheses and acted to support or negate the hypotheses. Secondly, there was a steep learning curve as regards the instructions, betting interface and institutional rules, and while undergraduate students generally are flexible enough to rapidly assimilate information in a new setting, it is arguable whether the general population has the same skills. Friedman and Sunder (1994, p.40) state that the average consumer or investor may not have the abilities routinely expected of laboratory subjects and also present evidence on why professionals may be unsuitable candidates for experiments, in that their prior experience may force them to behave under the professional circumstance they are familiar with, as opposed to operating under the rules of the experiment.

7.3.2 Funding

I applied for funding of £1400 directly to the Institute for Academic Development, the department responsible for accepting bids for events during Innovative Learning Week. The result of this application was unsuccessful. The funding committee were of the opinion that the payout per subject was too large and suggested offering another incentive for participation other than money. However, adequate reward was necessary in order for subjects to consider that they were making meaningful economic decisions during the experiment. A subsequent application to the Undergraduate Office and Accounting and Finance Group in the Business School was successful and both entities part-funded the payout pool for subjects.

7.3.3 Remuneration

Endowment in the experiment was symmetric, with subjects given an initial lump sum payment of £25 with which to trade and subjects all played the same role during the session. They were not paid a show-up fee, however the experiment was held as part of an event which included a workshop and presentation on behavioural finance, therefore I considered the educational benefits of participation in the workshop as a proxy for a payment for showing up. In that sense, the latter part of the session constituted a pedagogical exercise. Subjects had the opportunity of winning a maximum of £50 or finishing the experiment with £0, depending on the decisions they made during the experiment. The maximum possible loss was a subject's endowment. (Friedman and Sunder, 1994, p.50) state that an 80-90% affirmative response (having being paid) to the question of future participation by subjects' should be an indication that the level of reward offered during the experiment was sufficient relative to subjects' opportunity costs. The response rate at the end of the survey indicated a 100% intention to participate in subsequent experiments.

7.3.4 Deception & Ethics

Firstly, ethical issues were addressed by submitting a standard Ethical Checklist to the Institute for Academic Development and also the Business School at the university. No significant ethical issues were identified. (Friedman and Sunder, 1994, p.3) state that dominance and salience are lost if subjects doubt the announced relation between decisions and rewards or if they actively hedge against them. (Friedman and Sunder, 1994, p.17) While I could have designed the experiment in order to generate subsequent observations in the time-series underlying the charts presented to subjects which had an opposite sign to their choices in each of the treatment groups, this would not have been ethically sound. In addition, although it would have significantly increased the number of observations in each of the treatment groups (in particular after the third and forth bets), it would have necessitated a 'debriefing' statement after the experiment which would have affected salience and dominance in any future session with the same subjects. Such a practice would have been unknowable to subjects, however, the design would not have been consistent with the declarations made in the Ethical Checklist.

7.3.5 Lab log

Subjects were requested to log into their workstations and to open z-Leaf, after which the treatment was started with z-Tree from the experimenter's workstation. The projector was turned off so that subjects could not see the tables displayed on the z-Tree instance. There were no issues of note during the experiment, and while there were some rudimentary questions about the interface, there were no questions related to the research hypothesis or motivation behind the experiment at this point. None of these questions placed replicability at risk and were all dealt with publicly.

However it became clear that a number of subjects did not understand that observations in the time-series being displayed in charts on the graphical interface were independent. The experimental design was structure so that subjects would not finish out the experiment in their own time, but rather that results for each round would be displayed to subjects at the same time after all of the subjects had made their betting choices. A number of participants were examining the charts in great detail, taking notes and doing calculations, and some participants took a significant amount of time to make decisions. This was one drawback of the experimental design, as participants who had made their decisions first could not see the result of their bet until the last participants had also made their decision. An alternate structure for the rounds of betting in future experiments will address this issue.

On completion of the experiment, a brief questionnaire was activated from the experimenter workstation and participants were asked to fill in payment details, demographics and complete a CRT (Frederick, 2005) test. When the experiment and questionnaire were completed, a payment file was created with z-Tree and individual receipts were printed and distributed to participants during the presentation. Subjects were given the choice of electronic payment by the Business School or an Amazon Gift Voucher. A £1 premium was applied to the choice of an Amazon Gift voucher. Of the 28 subjects in total, only eight opted for the Amazon voucher with the £1 premium, with the remaining 20 subjects opting for payment by electronic transfer.

7.4 Results

If the experiment has indeed been well designed, drawing inferences from the data should be a formality (Friedman and Sunder, 1994, p.85). Firstly, a descriptive overview of the results is necessary, before embarking on the inferential statistics which will test the hypotheses. Summary statistics are therefore presented in Table 7.1.

7.4.1 Descriptives

A total of 122 bets were placed during the experiment, with a mean and median of 2.8 and 3.0 bets out of five, respectively. The mean payout was £23.21 per subject, with a minimum of £10 and a maximum of £40 and a total payout to subjects of £650. Some students did not answer any of the CRT questions correctly, however the mean and median were two questions correct out of three. The subject pool was homogeneous: all students, and all between the ages of 18

Table 7.1

Summary Statistics

I conducted an experiment consisting of a single treatment with 28 subjects. Participants were shown a trading screen with a chart (presented in Figure 7.2) and asked to place 'higher' or 'lower' bets. All bets were symmetric and there was no commission or margin. Subjects earnings were displayed along with the result of each bet. After the second and subsequent bets, subjects faced a decision node where they could stop betting and realize their earnings to date, or continue betting generating subsequent losses or gains. Summary statistics on betting and demographics are presented below.

Variable	N	Mean	SD	25th	Median	75th	Min	Max
Bets	122	2.80	1.39	2.00	3.00	4.00	1.00	5.00
Earnings	28	£23.21	£12.41	£10.00	£20.00	£37.50	£10.00	£40.00
Age	28	21.89	3.28	20	21	22	18	31
Female	28	0.25	-	-	-	-	-	-
CRT Score	28	2.00	1.12	1.00	2.00	3.00	0.00	3.00

Figure 7.1

Betting variables

The following histograms summarise aggregate betting variables. The distribution of earnings at the end of the experiment, the number of winning and losing bets, and the distribution of bet direction ('up' and 'down' bets) are presented in the panels below.

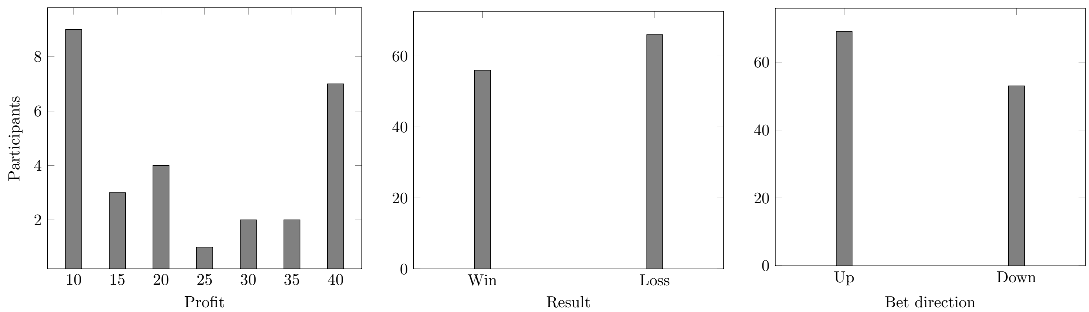

and 30.

Figure 7.1 presents the distribution of profit, the number of wins and losses, and the distribution of bullish or bearish bets placed by subjects. On balance, there were slight more losing bets than winnings bets, and more bullish bets than bearish bets, however the differences between categories were not statistically significant.

The average payout amount is quite high relative to similar experiments. For comparison, Table 7.2 presents the sample size and average payout in a number of related papers. A payout of £25[1] compares quite well to other studies and is at the higher end of average payouts. It is also likely higher than the opportunity cost for students and more consistent with the level of payout made in experiments involving market professionals or traders.

[1]Approximately €33.75 as of 05 October 2015.

The sample size is quite small in comparison, however, giving low statistical power to any inferences we make from the results. Nevertheless, I would argue that while the sample size may be of some concern, the average payout is not and hence we can be confident that subjects were adequately motivated to follow the experimental instructions. As a result, it is unlikely that any other strategy dominated that of maximizing wealth by rationally choosing to continue or quit using current wealth as a reference point.

Table 7.2

Sample size & average payout

This table presents the topic, sample size and average payout in related experiments. Papers are listed in chronological order.

Paper	Topic	Subjects	Expected payout
Weber and Camerer (1998)	Disposition effect	103	15.04 DM
Kruse and Thompson (2001)	Money and class point rewards in experiments	197	\$5.50
Kruse and Thompson (2003)	Low probability risk valuation	93	\$5.50
Haigh and List (2005)	Myopic loss aversion in students and traders	59	\$30.00
Hanson, Oprea, and Porter (2006)	Info aggregation and manipulation	197	\$37.20
Alevy, Haigh, and List (2007)	Information cascades in students and traders	109	\$30.00
Ackert, Mazzotta, and Qi (2011)	Asset pricing in segmented markets	99	\$41.13
Kirchler, Huber, and Kleinlercher (2011)	Microstructure under a Tobin tax	384	€15.00
Koessler, Noussair, and Ziegelmeyer (2012)	Info aggregation and belief elicitation	176	€14.00
Lindner (2014)	Entry decisions in market experiments	160	€16.25
Fullbrunn, Rau, and Weitzel (2014)	Ambiguity aversion in experimental markets	232	€20.00
Fullbrunn et al. (2014)	Ambiguity aversion in experimental markets	192	€20.00
Huber, Kirchler, and Stefan (2014)	Uncertainty and skewness in double-auctions	270	€17.00
Dutcher, Balafoutas, Lindner, Ryvkin, and Sutter (2015)	Relative performance incentives	216	€9.36
Stockl, Huber, Kirchler, and Lindner (2015)	Hot hand & gambler's fallacy in teams	360	€14.00

Finally, Table 7.3 presents some detail on the subjects' profiles with a list of the university faculties represented among the subjects. The group were relatively diverse in terms of their courses of study, which ranged from Mathematics & Statistics to Science & Religion, however the majority were pursuing programmes in which at least a foundation level of probability is taught, including business, economics, engineering, physics and mathematics. As a result, we might conclude that the subjects were relatively sophisticated in terms of their understanding of the independence of outcomes and probabilities.

Surprisingly, this was not the case and I revisit this issue in Section 7.4.3 when I discuss subjects' responses to a survey conducted after the experiment.

7.4.2 Attrition

As the stake was held constant at £5, the only decisions subjects had to make were to bet higher or lower and to continue or quit at each node. The first main result in this study is presented in Table 7.4, along with a breakdown for the treatment and control groups. On the basis of each subsequent bet, subjects were randomly assigned to a strong positive reinforcement treatment,

Table 7.3

Demographics: Level of education

Course of Study	Participants
Business and Management	5
Economics (and Accounting/Politics)	4
Biological Sciences (with Management)	4
Accounting & Finance	2
Artificial Intelligence	1
Astrophysics	1
Computer Science	1
Engineering	1
Geography & Economics	1
Law, Business & Philosophy	1
Mathematics & Statistics	1
MSc in Sustainable Energy Systems	1
MSc International Development	1
Physics	1
Science and Religion	1
Sound Design	1
Missing	1
Total	28

a strong negative reinforcement treatment or a control group. The assignment was essentially random and independent of subjects' decisions as the process driven the underlying market was random. With each round of betting, the number of subjects in each of the groups declined as the cumulative probability of having an increasing number of winning or losing bets in a rows decreased. For example, there are 17 (11) subjects in the strong negative (positive) reinforcement treatment group after the first bet. After the second bet, this reduces to nine and six for the negative and positive treatments, respectively. Nevertheless, there is more attrition in the losing treatment group and the subject pool is loss averse.

Table 7.4

Attrition: Treatment & control groups

Attrition counts are presented in parentheses below. There are no observations in the control group at the end of the first bet as there are only two possible outcomes up to that point.

Period		**Control**			**Losers**			**Winners**	
	N	Att	%	N	Att	%	N	Att	%
1	-	-	-	17	(0)	0	11	(0)	0%
2	13	(1)	7.7%	9	(2)	22.2%	6	(0)	0%
3	16	(1)	6.2%	4	(1)	25.0%	5	(1)	20%
4	20	(3)	15%	1	(0)	0%	1	(0)	0%
5	19	(-)	-	0	(-)	-	0	(-)	-

The lack of sufficient observations in each treatment group (and associated standard errors) precludes standard statistical inference, however we do see some support for the hypotheses. The rates of attrition are higher in the losing domain than either the control group or the winning domain with each subsequent bet[1]. Subjects with consecutive strings of winning bets are exiting the market at a lower rate than either the losing or control groups. This is contrary to the prospect theoretical prediction. Those in the domain of wins should be dropping out at a higher rate than losers. However there is empirical support for such as result.

This is consistent with results presented in Thaler and Johnson (1990) in which students were offer a choice between the status quo and a fair gamble and in which risk seeking (aversion) was observed in the presence of prior wins (losses). However, there is a difference between Experiment 4 in Thaler and Johnson (1990) and this setting. Theirs was a two-stage decision, whereas this experiment was a multi-stage setting. Subjects had up to four previous results with which to base their next decision of whether to continue to bet or quit. Nevertheless, the results suggest that not only do prior results matter, but that predictions from prospect theory related to one-shot gambles may not be effective in the presence of feedback.

Battalio, Kagel, and Jiranyakul (1990) also present evidence of risk seeking in conditions where prospect theory predicts risk aversion. The result here is also consistent with this. In addition, they report more conservatism in terms of risk seeking in both the winning and losing domains when real, rather than hypothetical, payoffs are used, however I did not have any provision for testing this hypothesis in the experimental design.

Barron and Erev (2003) stress the difference between what they refer to as 'decisions from description' and 'decisions from experience'. Decisions from description are the standard certainty equivalents presented in Kahneman and Tversky (1979). They are one-shot decisions where full information is given to respondents, including outcomes and associated probabilities. Decisions from experience refer to settings in which no such information is given to respondents and they must learn by exploration and from feedback on their decisions which the optimal strategy is. Barron and Erev (2003) presented evidence of risk seeking in the domain of wins and risk aversion in the domain of losses. In a discussion of the 'description/experience gap', Ludvig and Spetch (2011) suggest that the entire s-shaped curve mapping from objective value to subjective utility may be reflected when subjects must learn from experience.

The experimental design shares characteristics of both settings. There is full

[1] In terms of the absolute number of subjects exiting the market, there is no difference in attrition between the winning and losing domains after the third bet

information given on outcomes and probabilities, however there is also feedback. In effect, subjects can choose to be influenced by their prior outcomes or treat the experiment as a series of one-shot decisions. I suggest that a strong reinforcement learning setting may have a very salient effect on subjects, focusing their attention on their prior results, which may account for the result here.

In order to test the hypothesis empirically, and to facilitate intuitive interpretation of the results given the small number of observations in each of the treatment groups, I modify an approach taken by Strahilevitz et al. (2011), and model the evolution of betting as a decision tree in Figure 7.2. As in Chapter 5, this is envisaged as a more natural representation of such as setting than a traditional tabular format. For further analysis, this result is summarized graphically in Figure 7.2.

As before, attrition counts are presented in parentheses. There are no observations in the control group at the end of the first bet as there are only two possible outcomes up to that point. Attrition counts and percentages in the losing and winnings control groups are based on observations in the topmost (W, WW, WWW, WWWW) and lowest (L, LL, LLL, LLLL) nodes of the decision tree. The control group constitutes subjects with mixed reinforcement i.e. WLW, LWWL etc. For the fourth and fifth bet, the control group encompasses two and three nodes, respectively.

The lack of observations in each group (especially in the latter rounds of betting) are of concern, however there is some support for the hypothesis that agents subject to strong negative (positive) reinforcement realize their losses (gains) at a higher (lower) rate than predicted by prospect theory. At the first decision node (after the second bet), there was attrition of 22.2% from the strong negative reinforcement group, 7.7% from the control group and no attrition from the strong positive reinforcement group. At the second decision node (after the fourth bet), there is attrition of 25% from the strong negative reinforcement group, 10% from the control group and 20% from the strong positive reinforcement group [1].

7.4.3 Survey

After the experiment, I elicited responses to a Frederick (2005) CRT test during the survey. As stated in Friedman and Sunder (1994, p.6), one caveat to bear in mind when analysing these results is that there was no performance-based payment directly related to the questionnaire and effectively no economic motivation for subjects to answer the questions correctly. The survey consisted of a number

[1]For completeness, a complete record of all subjects' decisions are given in Appendix A.

Figure 7.2

Evolution of betting: Attrition

The following figure summarises betting over the course of five rounds. It is a modification of the approach presented in Strahilevitz et al. (2011), and models the evolution of betting summarized in Table 7.4. Each node displays the number of subjects who continued betting at each round and the level of attrition at that round. The number of winning and losing bets, and the result history up to that point are displayed as join nodes between each bet. Attrition counts and percentages in the losing and winnings control groups are based on observations in the topmost (W, WW, WWW, WWWW) and lowest (L, LL, LLL, LLLL) nodes of the decision tree. The control group constitutes subjects with mixed reinforcement i.e. WLW, LWWL etc. For the fourth and fifth bet, the control group encompasses two and three nodes, respectively.

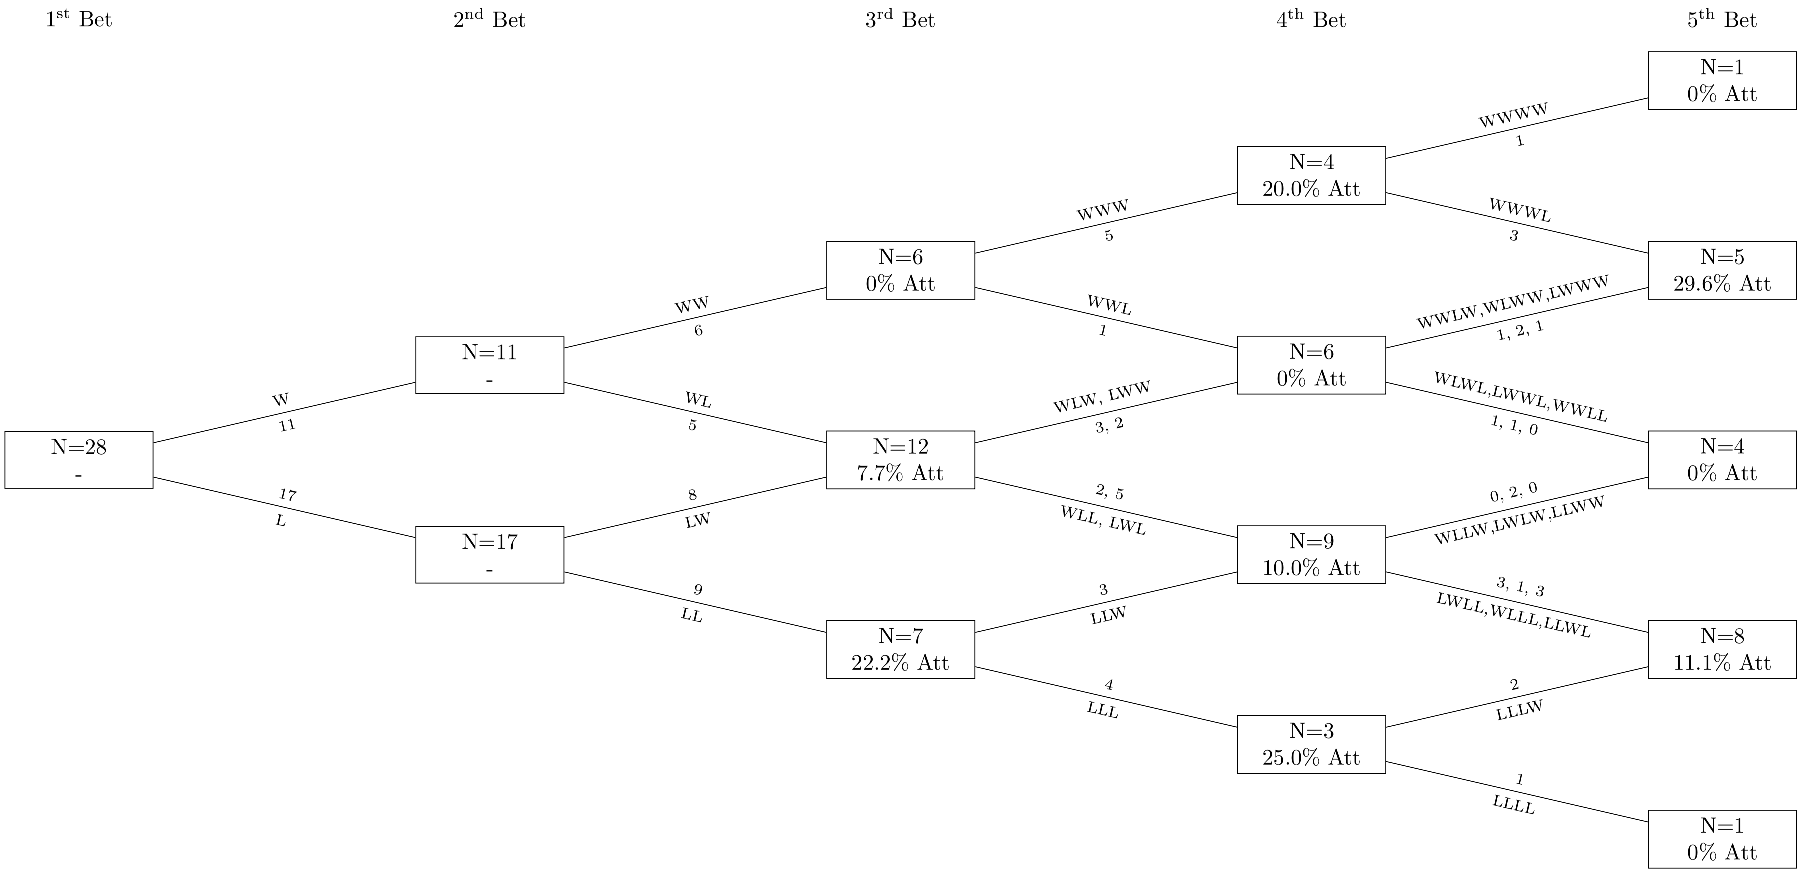

of questions to elicit demographics from subjects, the CRT test and a free-form comment section. Figure 7.3 presents the results of this test, in which the correct answer in each case happened to be the modal one.

Figure 7.3

Cognitive Reflection Test

The following histograms summarise the subjects responses to a Frederick (2005) CRT test.

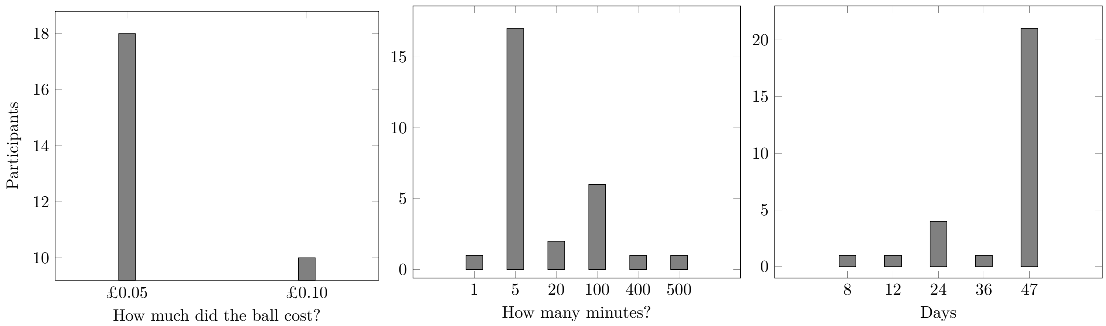

As shown in Table 7.5, there was no evidence of a correlation between a high score in the CRT test and high earnings in the experiment, as expected. However the free-form comment section of the survey did yield some interesting results.

Table 7.5

Earnings by CRT Score

Earnings are tabulate by CRT score in this table. There was no difference in earnings between the lowest and highest scoring groups. A two sample t-test for a difference in means between the lowest and highest scoring groups had a p-value of 0.888. There were only four observations in the lowest scoring group in any event. However, since results were independent, this result is intuitive.

Total Correct	Participants	Mean Earnings	Total Earnings
0	4	£21.25	£85.00
1	5	£28.00	£140.00
2	6	£22.50	£135.00
3	13	£22.30	£290.00

It was common knowledge that the process driving the underlying in this experiment was random. Moreover, subjects were made aware that results were uncorrelated. However, although it was common knowledge that the time-series underlying the charts on their trading interfaces were generated by a random-number generator, some subjects believed either that the next draw in the series could be predicted, or that the charts represented the stock prices of actual shares. This can be seen in some of the comments in Table 7.6. However, this is not a surprising result and is in fact consistent with the literature.

Table 7.6

Free-form comments

Earnings	Comment
£40.00	-
£40.00	I really enjoyed the experiment. Thank you very much!
£40.00	Great experiment! And I acted riskier than I wanted to.
£40.00	The experiment was very interesting, hope this is similar to real life.
£40.00	This experiment was very exciting! If there is another opportunity, i sincerely hope to participate in that again!!
£40.00	-
£40.00	Because I bet almost in the last, I guess whether it gives me more opportunity to win(actually I do win in the last 4 times due to this strategy). Hope to teach how to find the underlying tendency of graph.
£35.00	The experiment design, more rounds?
£35.00	Very interesting. Apart from the answer - Win or Lose, if you tell us why it reduced or increased, it would help more.
£30.00	-
£30.00	-
£25.00	-
£20.00	Very fun.
£20.00	-
£20.00	It's very interesting
£20.00	Interesting experiment
£15.00	-
£15.00	I take it these were real stock prices used - Taylor Wimpey showed up. Maybe the stock prices could be from other stock markets rather than LSE so that price charts we may have seen before don't show up during the experiment.
£15.00	-
£10.00	Enjoyable experiment.
£10.00	-
£10.00	-
£10.00	No
£10.00	-
£10.00	Thanks :D
£10.00	-
£10.00	-
£10.00	-

Barberis et al. (1998) examine how investors form beliefs and investigated the representativeness heuristic[1]. They presented subjects with earnings of an asset which followed a random walk but did not make them aware of this. They found that investors oscillated between belief in mean-reversion and trending

[1]The tendency of experimental subjects to view events as typical of some prior experience and to disregard probability as a result (Kahneman and Tversky, 1973).

when basing beliefs on a random series. The qualitative findings indicate similar patterns in investor beliefs.

De Bondt (1993) performs an experimental analysis examining intuitive assessment of time-series data in which subjects who were presented with random data based their predictions on past price trends. Subjects included both reversals in prices and momentum in their forecasts, leading De Bondt (1993) to conclude that subjects attempted to discover patterns in series in which they should have been aware were independent. Bloomfield and Hales (2002) also presented subjects with series of earnings results and elicit predictions for subsequent earnings announcements. Despite being instructed that the process generating the returns was a random walk, they found evidence that subjects exhibited behavior consistent with mean-reversal and momentum belief.

Durham et al. (2005) replicated the Bloomfield and Hales (2002) using betting market data. They suggest the gambler's fallacy as an explanation for the reason their subjects expected reversals after long periods of continuation. They posit that it is the representativeness heuristic that causes investors to detect patterns in time-series which are actually driven by a random process. Thus, there is a precedent in the empirical literature for the result here.

7.4.4 Limitations & alternative explanations

It is possible that some characteristic of the experimental design served to temper subjects' loss aversion. In particular, the standard setting for field experiments in this field is in stock transactions. In that case, a deliberate action is required to book gains on a winning position and stop losses on a losing investment. In this setting, bets expire at the end of each round and no intervention is required from subjects to stop losses or book gains. It may be that having losing positions effectively 'stopped-out' automatically is the cause for absence of loss aversion here. In fact, Weber and Camerer (1998) examined the disposition effect in an experimental setting, concluding that it is significantly reduced when positions are automatically closed out at the end of each trial.

Belief in mean reversion may cause similar behaviour to that observed here. Rather than booking gains and holding on to losses, investors may believe that losing stocks will rebound and winning stocks will revert to the mean. The experimental design attempted to remove a belief in mean reversion as a cause for observed behaviour by using a random process to generate a time-series and settle bets. In doing so, however, I may have introduced a confound.

Subjects were presented with a financial chart, however the time-series underlying the chart was random. I endeavoured to induce subjects to focus only on their past performance and current wealth as a reference point. Nevertheless, the qualitative results in the survey indicate that some subjects did base their decisions (irrationally) on patterns within the time-series. A possible remedy may be to elicit subjects' beliefs at each round of betting with regard to the chart i.e. mean reversion, continuation/momentum or independence[1].

It is possible that subjects regarded the initial endowment as 'house money'. In order to attempt to preclude house money as an explanation, the endowment could be distributed to subjects' a number of weeks in advance of the experiment, or alternatively, they could be given a real-effort task at the start of the experiment in order to accumulate some wealth with which to bet. Thaler and Johnson (1990) show that prior gains increase the attractiveness of gambles, while prior losses may reduce the willingness to take risks. There was no provision in this iteration of the experimental design for eliciting whether subjects considered their initial endowment a gain from the offset, which induced them to accept the symmetric bets they were offered.

7.4.5 Further research

In a section of the questionnaire, all of the participants answered that they would participate in further experiments run in the Business School in the future. After having done a preliminary literature review and a draft experimental design, but prior to the experiment, it was clear that the topic of this experiment would offer a rich ground for further research. As a result, an application for funding under an Early Career Venture Fund administered by the University of Edinburgh Business School to continue with research on this topic was successful. The funding will be used in order to facilitate the recruitment of a larger and more heterogeneous subject sample through the Behavioural Laboratory at the University of Edinburgh (BLUE).

The project will entail a series of behavioural economics experiments to be conducted at the Behavioural Laboratory at the University of Edinburgh (BLUE). A total of 162 subjects will be recruited by BLUE from their client base and it is envisaged that a total of seven half-days (18 subjects per session, three sessions per day, plus a pre-test half-day) will need to be booked at BLUE to facilitate conducting the experiments.

As in this study, the treatment group will be subjected to strong reinforcement (strings of wins or strings of losses). The aim is to investigate whether strong reinforcement mitigates or exacerbates loss aversion. In particular, it will be

[1]Introducing such a feature requires careful design so as not to serve as a 'nudge' to subjects.

examined whether strong reinforcement will lead to loss aversion by subjects in the treatment group compared to those in the control group.

7.5 Conclusion

The results support the hypotheses proposed and tested in earlier chapters. As a final empirical chapter, this study follows on naturally from the preceding two, and offers evidence from an alternate data source in order to add weight to the hypotheses tested in previous chapters.

In Chapter 5, we saw that attrition rates were highest for bettors who had experienced strings of losses, and lowest for those who had experienced successive wins in a row. In a focused experiment, I have presented results which exhibit the same pattern of attrition in the winning and losing domains. As I have held stake size constant, I have examined the 'quit/continue' decision independent of the decision to alter stake size. In Chapter 6, we saw a pattern of risk taking in a multi-round setting which was in contrast to that predicted by Kahneman and Tversky (1979) for one-shot decisions. In the experimental treatment outlined in the current chapter, however, I did have the opportunity to test this premise[1].

Given further treatments, I could test the interval validity of the results presented in this chapter, however the preceding chapters using field data from a natural experiment in financial betting, offer the opportunity to test the external validity of the results in a comparable setting. While the data in previous chapters was essentially field-happenstance data from a natural experiment, this is laboratory-experimental data. The experiment was designed in order to have sufficient internal validity to be able to attempt to draw causal inference, as well as to make generalisation about behaviour.

[1] I did not introduce differing bet prices in order to allow subjects in the winning or losing domains to change the variance of their bets.

Chapter 8

Conclusion

8.1 Thesis summary

The first section of this chapter provides a terse summary of result presented in the thesis, and is meant to serve as a synopsis of the results, rather than an evaluation of their contribution. The implications of the results are discussed in Section 8.3, where their contribution to the literature will be outlined.

Chapter 1 was prefaced by the thesis abstract, and include a brief executive summary, a terse overall motivational section, an outline of the structure of the thesis and a summary of results. Between the abstract and Chapter 1, all the concepts and questions under analysis in subsequent chapters were tentatively introduced.

Chapter 3 provided a broad literature review with scope across the aggregate topic of financial betting markets. The main aim of the chapter was to motivate this setting for the testing of behavioural finance theories. There was discussion of test of market efficiency, individual trader behaviour and asset pricing. Literature from diverse settings such as traditional financial markets, prediction markets, sports betting, financial betting, the binary option industry and the market microstructure literature (with a focus on the literature on market making) was outlined. This chapter served to provide a framework for the targeted literature reviews included in each of the three main empirical chapters.

Chapter 4 served as an introductory setting to the three main empirical chapters. The betting proposition was outlined in detail, relevant regulatory issues were presented, a detailed description of the mechanics of financial market betting was provided, and payoff and P&L diagrams of bets were presented. There also was an outline of the timeline of a financial bet: from the receipt of the price of the underlying by a bookmaker, the use of proprietary algorithms to offer a range of strike prices, the dissemination of those prices to agents via the bookmakers online betting interface, the placement of a bet, bet settlement and resulting. The dataset was also introduced in this chapter. A list of variables with descriptions was outlined and an initial data exploration in the form of summary statistics was presented. As an adjunct to this data exploration, in each of the three main empirical chapters, summary statistics with particular relevance to the hypotheses under analysis were also introduced.

The first main empirical chapter, Chapter 5, hypothesised that different types of learning would be evident in betting on the financial and simulated markets. The first hypothesis related to path-dependent behaviour in terms of stake sizes. There was support for this hypothesis in the form of differing levels of attrition in the winning and losing domains of both market settings, as well as path-dependent stake size changes which were evident from the decision trees presented in the results. The second hypothesis testing a proposal in the literature that agents are willing to continue in the face of successive losses in the hope of 'learning by doing'. There was support for the hypothesis in the form of lower attrition rates in the losing domain of the financial market setting. The notion that changes in stake size and levels of attrition could be instead driven by the disposition effect was introduced in the discussion section of the chapter.

Chapter 6 introduced predictions from prospect theory and motivated an analysis of whether behaviour in the losing and winning domains of the sample was motivated by the disposition effect, rather than purely being driven by Bayesian and reinforcement learning. I hypothesised that there would be evidence of a disposition effect, manifested in less attrition by losers, and more risk seeking by losers: a standard test in the literature. However, I split the sample into agents who had been subject to reinforcement learning and those with mixed experiences of wins and losses. I hypothesized that the reinforcement cohort would exhibit behaviour contrary to that predicted by prospect theory and expected higher attrition by losers and less attrition by winners. There was support for this hypothesis and evidence of an interaction between learning and the disposition effect, motivating the experimental setting in the final empirical chapter.

Chapter 7 presented laboratory evidence from a financial betting experiment which took place at the University of Edinburgh, consisting of five rounds of betting on a simulated market and a brief questionnaire. The experimental design was outlined in detail, along with concerns about the control of some of the unobserved heterogeneity of subjects. There also was a very brief laboratory log with a note on the administration of the experiment. The treatment group consisted of agents who had experienced a series of consecutive wins and losses, while the control group contained agents who had mixed betting results. I hypothesised that those subject to strong positive or negative reinforcement would exhibit behaviour inconsistent with the predictions of prospect theory and the disposition effect. Admittedly, the sample size was quite small, however I nevertheless found some support for this hypothesis. Table 8.1 lists the hypotheses tested in the thesis along with a note on whether there was clear support, mixed support or clear rejection for each.

Table 8.1

Hypothesis support/rejection

This table lists the hypotheses tested in the thesis along with a note on whether there is clear support, mixed support or clear rejection of each hypothesis.

Chapter	Hypothesis	Support
Chap. 5	H1: A string of wins/losses will lead to an increase/decrease in the stake size in the Financial Market.	Clear support in domain of gains. Mixed support in domain of losses. Winners increase their stake sizes. Losers who remain do not.
	H2: A string of wins/losses will not affect the stake size in the simulated Market.	Clear rejection. There is evidence in both attrition rates and median/mean stake size changes that this hypothesis must be rejected.
	H3: The magnitude of change in the stake size is positively related to the length of the string of wins/losses.	Clear support in terms of both mean and median stake size changes.
Chap. 6	H1a: Losers are more likely to place riskier bets than winners in aggregate.	Clear support in the form of higher probabilities of placing riskier bets than safer bets for losers.
	H1b: Winners are more likely to place safer bets than losers in aggregate.	Clear support. This result suggests that the House Money Effect may not be salient here.
	H2a: Losers in the reinforcement cohort are more likely to quit than winners.	Clear support in the form of higher attrition rates for losers when agents with mixed results are removed.
	H2b: Winners in the reinforcement cohort are more likely to continue than losers.	Clear support. Across the first five bets, winners are more likely to continue, having been subjected to strong positive reinforcement.
	H3: Sophisticated agents, as proxied by bet channel and bet type, are less likely to exhibit a disposition effect.	Weak support for this hypothesis as the coefficient on the the bet channel dummy was not significant in the 'reinforcement-only' model.
	H4: Agents who have experienced a near-miss outcome are less likely to quit and more likely to place a risky bet.	Clear rejection. The 'near miss' coefficients indicate a higher probability of placing safer bets and a lower probability of placing riskier bets, versus the base category of no change in bet price.
Chap. 7	H1: Probability of attrition for winners in the treatment group is lower than in the control group.	Clear support, albeit with a small sample size. The attrition rate for agents who have experienced successive losers is higher than the control group.
	H2: Probability of attrition for losers in the treatment group is higher than in the control group.	Clear support, albeit with a small sample size. As above, strong positive reinforcement learning induces agents to remain in the sample.

8.2 Limitations

The main limitation of Chapter 5 is the absence of additional control variables for observations in the sample. Demographics, gender, level of education, proxies for investor sophistication or market knowledge would have offered the opportunity to test the hypotheses holding these factors constant. Unfortunately, no such variables were available. The dataset was extensive and had the rare property that individual (anonymised) accounts could be tracked over time, but there were no additional agent-specific variables.

In Chapter 6, a crucial element in any empirical test of the disposition effect is the specification of the reference point. For the losing domain, the analysis focused on using cumulative losses to date as the reference point. In the literature, a number of alternate reference points have been used, including the previous closing price of the stock, the daily high and low for buy and sells, respectively, or the purchase price of the stock, for example. Also, a feature increasingly common in the financial betting industry which would have offered a more intuitive test of the hypotheses, bet buybacks[1], are not offered by the data provider for this research.

The evidence presented in Chapter 7 is based on a pilot study with a rather small number of observations. As expressed in the discussion of the experimental design, an alternate design could have increased the counts in each of the treatment and control groups, however this would have introduced an element of deception and necessitated a debriefing statement. Any such deception would have had an effect on the saliency of rewards for future subject pools. Moreover, such a design may not have gained ethical approval. The Innovative Learning Week event at the University of Edinburgh presented a unique opportunity to administer such an experiment, however, notwithstanding the relatively small number of participants. The funding received also facilitated relatively large monetary rewards for subjects at £25 per subject. Further funding would facilitate re-running the experiment with a large enough sample to enable robust statistical inference based on the results.

Table 8.2 lists the theories underlying the hypotheses tested in the thesis along with a note on whether the results were consistent with, their predictions. Some of the results in Chapter 6 may be due to the setting used to test for the disposition effect. Analysing the behaviour of agents who are trading terminal assets, where they are essentially forced to liquidate their portfolios at the end of each round, may see a reduced level of risk seeking in the domain of losses. As bettors' positions are automatically closed out at the end of each round of betting, they cannot hold on to losing positions in the same vein as losing stock

[1]'Buybacks' allow bettors to close out their bets prior to expiry. Once a bet is transacted at a certain odds price, the current P&L of the bet is displayed on-screen, updating in real-time and customers have the opportunity to either take profit on winning bets or stop losses on losing bets, instead of waiting for the bet to expire and be settled.

Table 8.2

Theoretical Framework

This table lists the theories underlying the analysis in each chapter with a note on whether the results are consistent with, or in contrast with each theory.

Chapter	Theory	Support/Rejection
Chap. 5	Reinforcement learning theory	Clear evidence that strong positive and negative reinforcement causes path-dependent behaviour.
	Overconfidence theory	Evidence that wins cause agents in both Financial and simulated markets to increase their bet sizes and that such behaviour reduces wealth.
Chap. 6	Disposition effect	Risk aversion by losers in contrast to predictions for behaviour in stock investments. Inverse disposition effect.
	Prospect theory	Risk aversion by losers could be caused by a house money effect.
	House money effect	Evidence that positive returns induced continuation.
Chap. 7	Reinforcement learning	Strong evidence that agents subject to strong negative (positive) reinforcement quit (continued).
	Loss aversion/Reflection effect	Losers cut their losses and quit, the opposite of that predicted by prospect theory.
	House money effect	Issues with the large endowment before the experiment could be causing winners to continue due to a house money effect.

investors can. Weber and Camerer (1998) present evidence of a disposition effect in an experimental setting, but show a great reduced measure when positions were automatically closed at the end of each period. The inverse disposition effect measure evident in Chapter 6 may be as a result of using bets rather than financial stocks which are not liquidated each time a position is marked-to-market.

Generalising the results from Chapter 7 may be confounded by a number of alternative explanations and some of the features of the experimental design. The endowment given at the start of the experiment may have been considered a 'windfall' (Thaler and Johnson, 1990), which resulted in increased risk seeking in the winning domain as a result of a House Money effect. In addition, there was no threat of bankruptcy for subject, which may be induced further risk seeking by subjects as a result of the decreased risk to their own wealth.

8.3 Contribution

One of the topics covered by the early literature on market microstructure covered the activities of market makers or 'specialists' at the NYSE. Considerations of informational asymmetry and adverse selection were common. The activities of sports bookmakers are a very close analogy to this, and is a topic that has received considerable attention in the literature. However, there was a gap where these two worlds collide, which this thesis has endeavoured to fill. The data for this study consisted of the same products offered by market makers in financial markets, but in this case the market was made by online sports bookmakers. In addition, and in strict contrast with a continuous double-auction or order book, trades do not commonly affect prices. A detailed examination of risk management by the bookmaker was prohibited by the data licence governing my use of the data, however, the analysis was nevertheless one of the first explorations of the activities of relatively recent entrants into the market-making sphere: traditional sports bookmakers.

Firstly, in broad terms, the thesis introduced a unique, proprietary dataset of transactional panel data to the literature where individual accounts could be tracked and analysed through time. This facilitated an individual-level, rather than an aggregate level testing of the hypotheses, and such data is not commonly available, nor routinely present in the empirical behavioural finance literature.

The introduction of bets on the simulated market offered two contributions. A first, incidental result, was the extent to which customers of the data provider were willing to place bets on this market, which was surprising. In fact, as outlined in Chapter 4, the simulated market was the third highest market in terms of turnover for the data provider. Secondly, the incorporation of bets on the simulated market enabled me to perform tests with a level of control not possible with other financial products. The characteristics of this market setting offered a stark contrast to the financial market setting, in that information and experience was useful in one setting, but irrelevant in the other.

The mean and median bet sizes in the sample were quite small, however the level of betting for some agents was considerable. The highest bet placed on a five-minute expiry was £28,800. Given a dataset replete with such bets would have made the task of making generalisations from this setting to financial markets easier, however agents in this dataset were indeed using their own funds. In each case, they were performing identical, consecutive financial decisions and the use of their own money departed from the artificiality of laboratory data. Moreover, when the thesis did make use of experimental data in Chapter 7, the levels of expected remuneration for agents was considerable, and certainly exceeded the opportunity cost for the students participating in the session.

Another characteristic of the dataset was that the time period between action and feedback was considerably short. This added to the strength of the arguments put forward in the first two empirical chapters. With regard to learning, since bets were expiring every one or two minutes, the argument that learning of various kinds was taking place had some credibility, which would not have been the case had agents been placing bets once a day, or once a week, for example. In the case of the chapter on the disposition effect, the fact that bet placement and bet settlement followed so closely, gave weight to the argument that cumulative wins or losses to date was a credible reference point with which to edit subsequent decision of whether to continue betting or to quit.

Chapter 5 tested for path-dependent behaviour in financial markets and an absence of path-dependent behaviour in a simulated market. In effect, the hypotheses tested for rational Bayesian learning, one of the assumptions of models of individual behaviour. The existence of evidence of naive reinforcement learning would serve to reject this hypothesis. I offered evidence in support of rational learning in financial markets, however had to reject in favour of reinforcement learning in the simulated market. The contribution of this result is two-fold. Firstly, the results support the rational learning hypothesis in financial markets, in that winners have a lower rate of attrition, either learning about ability or learning by doing. Secondly, the results from the simulated market indicate irrational behaviour. There should have been no path-dependent behaviour in terms of stake in the simulated market setting, nor any evidence of asymmetric attrition in the setting. We will return to this notion of irrationality when discussing the results of Chapter 7.

Chapter 6 presented evidence of an interaction between learning and the disposition effect. There were a number of possible contributions here. Firstly, by not isolating the effect of agents who have experienced successive wins and losses, the literature may be understating the prevalence and magnitude of this bias. By extension, controlling for the number of successive prior months/years with the same sign of returns may offer an additional control to empirical studies of the disposition effect with panel data. Secondly, I showed that investors did not exhibit behaviour consistent with the disposition effect in the simulated market. This is evidence in favour of the proposition that agents are aware that this setting is one in which they do not have an internal locus of control and adds weight to the suggest that delegating financial decisions to a third-party may be a way for individual investors to mitigate the effect of this bias.

In Chapter 7, I hypothesised that agents subject to strong positive or negative reinforcement would exhibit behaviour inconsistent with prospect theory. In fact, I found support for the hypothesis that the behaviour observed would be the opposite of that predicted by the disposition effect. Although the sample size was admittedly small, I found support for this hypothesis, suggesting that reinforcement learning may play a stronger part than anticipated in the behavioural finance literature. How does this generalise to a broader market setting? The disposition effect suggest that investors faced with positive returns have the propensity to book gains, but to hold on to losses. The results suggest that given a series of successive monthly gains, for example, investors would be more prone to holding on to their positions, while investors facing successive losses would likely liquidate their position because the influence of reinforcement learning in the sample is stronger than that of the disposition effect.

As an aside, there were also some qualitative insights offered by the experiment. In a free-form comment section, some subjects indicated that they did not understand that they were betting on a time-series created by a random-number generator. Some respondents were under the impression that there was an element of skill involved and attributed their positive returns to their own ability to 'forecast' subsequent draws from a random-number generator. One subject in particular remarked that they recognised one of the randomly generated stock charts as being that of a stock traded on the NYSE at a certain time in the past. Such results may be an indictment of the robustness of the experimental design, in that the characteristics and conditions of the experiment and the underlying process generating time-series charts was not made quite clear to subjects. However, there is evidence of similar irrationality in the literature, where subjects have been presented with charts of a random walk and prompted to predict subsequent observations. I therefore attribute these response to irrationality by subjects in the experiment.

Appendix A

Appendix to Chapter 7

- A.1 Poster for experiment
- A.2 Informal instructions & consent form
- A.3 Payment receipt
- A.4 Experimental interface
- A.5 Raw data

INNOVATIVE LEARNING WEEK 2015

Business School

Could You Beat the Stock Market and Win £50

Test your skills in predicting stock market prices and win up to £50 (real money) in this unique experiment. No experience or money required!

10:00 to 11:00 on Tue 17 Feb 2015

Room 1.12 (Computer Lab), Business School

Book now: www.ed.ac.uk/innovative-learning

DIVE IN!

THE UNIVERSITY *of* EDINBURGH

INNOVATIVE LEARNING WEEK 2015

Business School

Could You Beat the Stock Market and Win £50?

Instructions

You have been given an endowment of £25 with which to bet in this decision-making experiment. In each round of betting, examine the simulated prices and determine whether the market will end up above or below its current level when it reaches the red vertical line. Once you have made your choice, please select either the ABOVE or BELOW button. Your stake is fixed at £5 and the odds are even. This means that if you win/lose, your earnings will go up/down by £5 and the chances of winning and losing are equal.

After the second round of betting, you will be able to STOP BETTING or CONTINUE BETTING. If you stop betting, you will realise your earnings and you will not be able to bet in subsequent rounds. If you continue betting, you will proceed to the third round of betting and your earnings will go up/down by £5 depending on the outcome of your third bet.

The subsequent rounds of betting are identical to the second round, such that at the end of each round you can either stop betting or continue betting. The fifth round is the final round.

A possible scenario:

Suppose that after the first two rounds, you have a win and a loss, such that your earnings are £25 + £5 - £5 = £25. At this point, if you stop betting, you will be paid £25. If you continue betting and win your third bet, you can stop betting and get paid £25 + £5 = £30 or continue to the fourth bet. On the other hand, if you continue betting and lose your third bet, you can stop betting and get paid £25 - £5 = £20 or continue to the fourth bet. And, so on.

INNOVATIVE LEARNING WEEK 2015

Consent Form

Please complete and sign below to consent to the data collected during this experiment being used for research purposes.

Name ....

Student/Staff Number ....

Signed ....

UNIVERSITY OF EDINBURGH
Business School

Dr. Ufuk Güçbilmez
University of Edinburgh Business School
29 Buccleuch Place
Edinburgh
EH8 9JS, UK
Email: u.gucbilmez@ed.ac.uk

Payment Receipt

Workstation: UEBS-TLAB-001

Email, Student/Staff ID: s1348032@sms.ed.ac.uk, s1348032

Subject: #9

I took part in the ILW Economics Experiment. I consent to my responses being stored in a computer database and being used anonymously for research purposes. I have been assured that my responses will be revealed to other people only in anonymous form. I accept that the experiment was carried out fairly.

Please tick your choice of payment below

☐ **Option A: Amazon.com voucher (by email) in the amount of £40.00 + £1**

☐ **Option B: Electronic payment in the amount of £40.00**

If you opt for Option B, please enter your Bank Sort Code and Account Number on the CASUAL PAYMENT VOUCHER. The School Accounts team have requested that you also email accounts@business-school.ed.ac.uk from your University Email address confirming your name, UUN, postal address and bank details using the Subject Line 'ILW – Stockmarket Event'. If there is any delay or problem with your payment, please feel free to contact Dr. Güçbilmez.

Signed: . . . .

Date: . . . .

29 Buccleuch Place
Edinburgh
EH8 9JS
Scotland

t +44 (0)131 650 8074
f +44 (0)131 651 3197
e business.school@ed.ac.uk
www.business-school.ed.ac.uk

THE UNIVERSITY OF EDINBURGH
CASUAL PAYMENT VOUCHER

Please fill out your name, address, postcode, bank sort code and account number.

PAY TO: UEBS-TLAB-001

***NAME**	
***ADDRESS**	
***POSTCODE**	

BLOCK LETTERS PLEASE

Email address, UUN s1348032@sms.ed.ac.uk, s1348032

***Sort Code** __ __ - __ __ - __ __ ***Account Number** __ __ __ __ __ __ __ __

Description of payment: ILW Stockmarket Event

OFFICE USE ONLY

SEPARATE HANDLING		INVOICE DATE	
INVOICE NO.		DUE DATE	

If documents are to be forwarded with payment enter 01 in separate handling field and enclose documents with this form, otherwise leave blank.

Cost Centre	Account Code	Job Code	NET Amount £	p	VAT Amount £	p	VAT Code
		Totals					

GROSS TOTAL **40.00**

	START	END
PAYMENT PERIOD	................/............../.............	................/............../.............

Prepared by: .... **Verified:** ....

Date Prepared 17 Feb 2015 .... **Authorised:** ....

(Authorised Signatory/Head of Dept.)

February 2011

A.4 Experimental interface

The experiment was coded in z-Tree (Fischbacher, 2007). Figure A.1 presents screenshots of the chart and betting interface, and the bet confirmation screen. At the start of the experiment, subjects were shown a chart with 250 observations. A further ten observations at the right-hand side of the chart were empty and populated with a 'Bet Higher' and 'Bet Lower' button spanning the rest of the chart.

The current price was highlighted at the end of the chart and served as the separator between the red and green betting buttons. The starting price, maximum, minimum and current price in the series were highlighted and displayed in a legend above the chart. The x-axis contained 260 observations in total (with 10 blank) using the total number of trading days in a year as the full range for the x-axis. Two boxes were displayed in panels at the bottom of the chart. A box on the lower left hand side displayed subjects' prior betting decisions and was not populated for the first bet. In the lower right corner, a betting variable panel simple outlined the fixed stake of £5 and the symmetric returns of £5 for both winning and losing bets[1].

Once participants had made their choice of a higher or lower bet, they clicked the corresponding green or red area of the chart. A very simple bet confirmation screen was then shown in order to acknowledge that subjects' bets had been placed. Subjects who had placed bets first were shown this screen until all subjects had made a choice. Thereafter, a resulting screen was shown to all participants at the same time. At the end of the first bet, this displayed their initial endowment (in order to highlight this as a reference point), the result of the bet and subjects' total earnings to date. A table was populated showing the round number, stake, betting price, betting direction, subsequent settlement price and profit in that round. This table was reproduced on the graph/betslip screen from this point onwards.

Betting during the second round continued in the same way as previous, however at the result screen, subjects were asked to make a decision whether to continue to bet with the possibility of winning or losing in subsequent bets, or stopping betting and realising their earnings to date. This screen was displayed at the end of the second, third and fourth bets. Figure A.1 displays a screenshot of the graph and betslip with the betting history panel populated, the re-

[1]Thaler and Johnson (1990) state that skewed bets may serve to complicated matters and focus attention on expected value of a bet rather than on the simple decision of whether to quit or continue. A fair bet removed this extra cognitive load on subjects and diverted their attention from a pure expected-value maximising strategy.

sult/continuation screen presented at the end of the third and fourth bets, the screen presented to subjects who decided to stop betting, and the screen presented to those who continued betting throughout the entire experiment.

Figure A.1 displays images of the questionnaire presented to subjects. The first screen concerned details necessary for the creation of individual payment receipts where subjects were asked to enter their university ID number and email address. The second screen elicited some simple demographics from the subjects. The third consisted of a Frederick (2005) CRT test, with the following questions:

1. A bat and a ball cost 1.10 in total. The bat costs 1.00 more than the ball. How much does the ball cost?

2. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?

3. In a lake, there is a patch of lily pads. Every day, the patch doubles in size.If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?

A Holt and Laury (2002) risk elicitation test was also designed in z-Tree, however plans for administering the test were shelved as it was decided that participants would be fatigued by such a detailed section after having participated in the experiment prior to the questionnaire. As a result, the final questionnaire was kept as brief as possible and only contained two screens with a minimum of questions on each screen.

Verbal/written instructions

A total of 28 subjects were present on the day of the experiment. Dr. Ufuk Güçbilmez, Tomás Ó Briain and Youyan Fu were in place as experimenters and moderators. A brief introduction to the experiment and workshop was given by Dr. Ufuk Güçbilmez. Thereafter, the following verbal instructions were read aloud by Tomás Ó Briain.

> *This is a decision-making experiment that involves betting on stock price movements. There are in total five rounds of betting. The instructions are simple, and you may earn a considerable amount of money. In particular, you have been given an endowment of £25 with which to bet. Depending on how many rounds you bet and the outcome in each round, your final earnings can reach up to £50. However, there is also the risk that you may lose part or all of your initial*

Figure A.1

Betting interface

The following figures display screenshots from the z-Tree program the subjects used to place their bets. The initial graph and betslip, the bet confirmation screen, the screen presented at the end of the first bet and the continuation decision screen presented at the end of the second bet are shown in panels below.

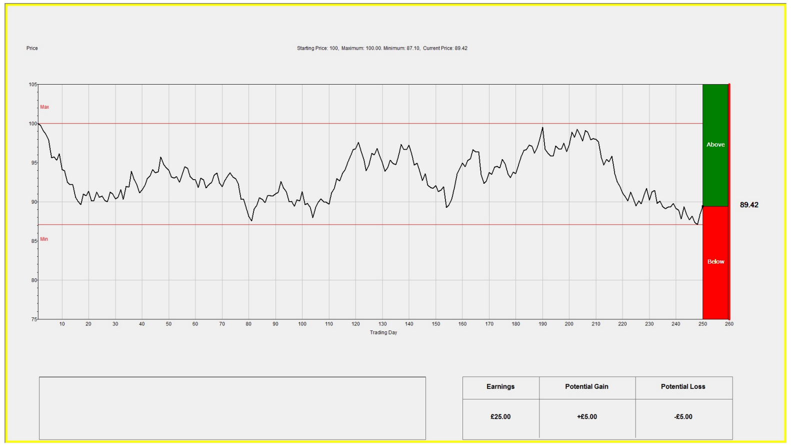

(a) *Graph and betslip*

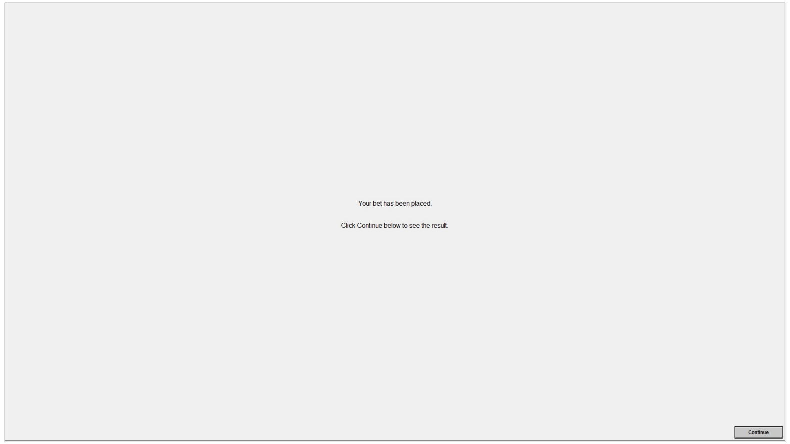

(b) *Bet confirmation*

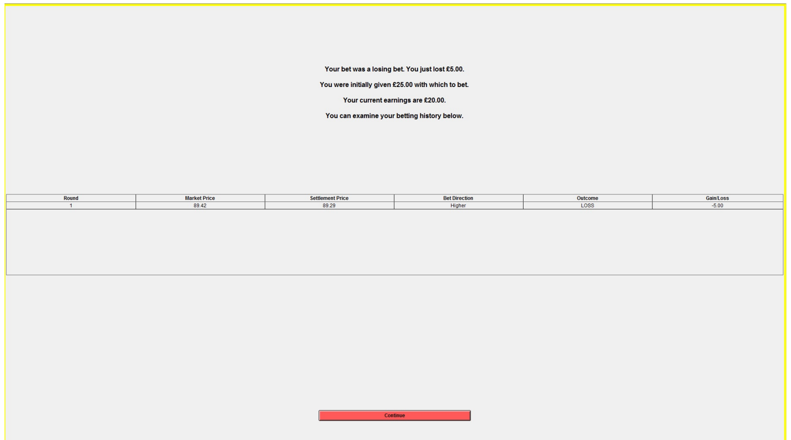

(c) *Result: Bet 1*

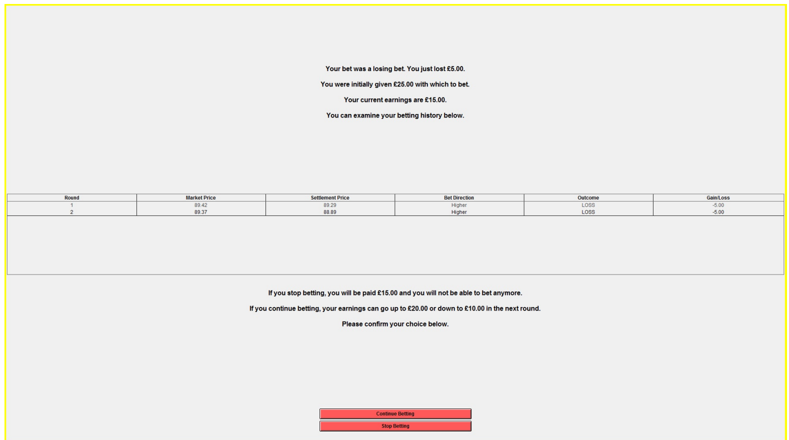

(d) *Result: Bet 2, Continuation decision*

Figure A.1 (Cont.)

Continuation screens

The following figures display screenshots from the z-Tree program the subjects used to place their bets. The graph and betslip including bet history, the continuation decision screen presented at the end of the third and fourth bets, the screen presented to those that decided to quit and the screen presented to those that continued betting until the end are presented in panels below.

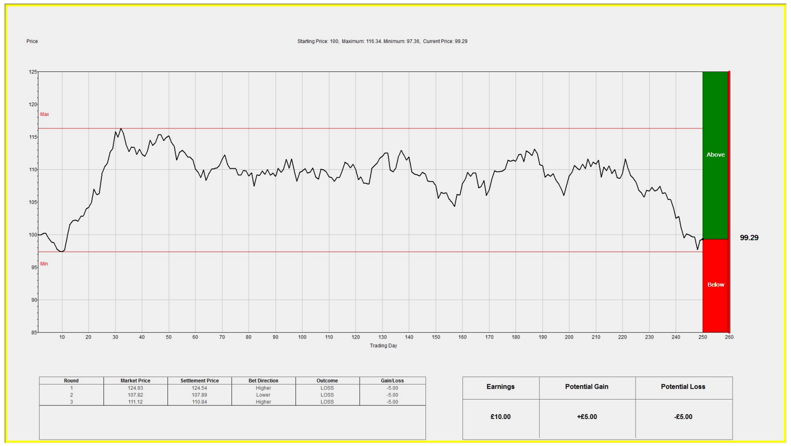

(e) Graph and betslip: Betting history

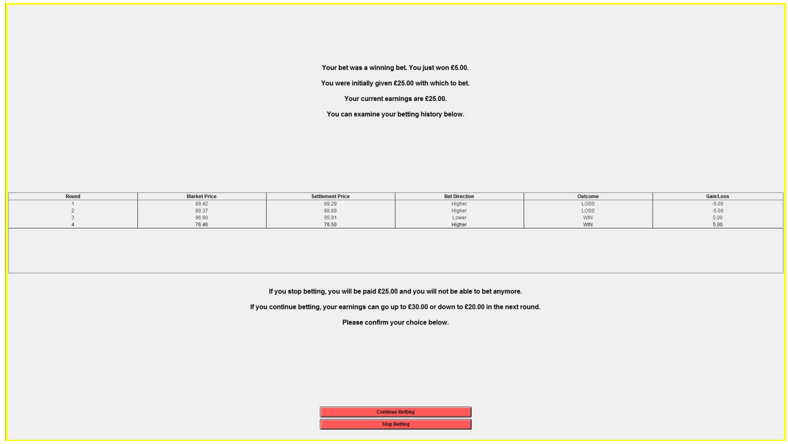

(f) Result: Bet 3-4, Continuation decision

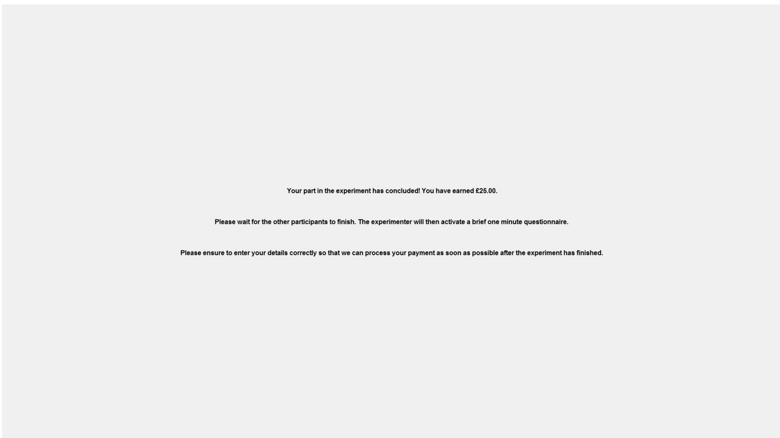

(g) Quit screen: Bets 2-5

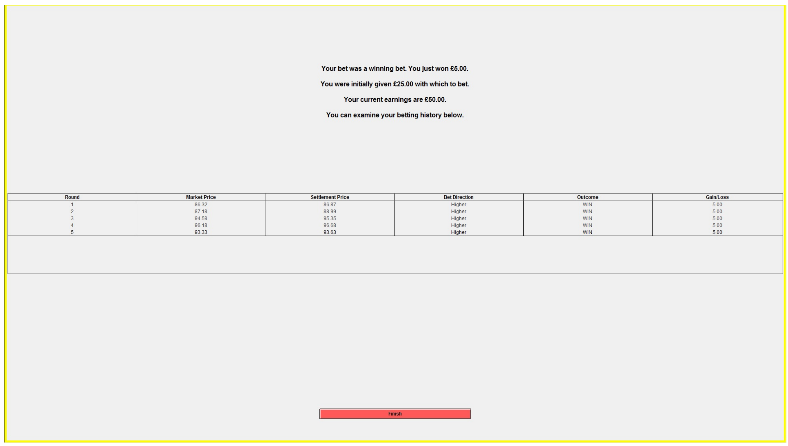

(h) Finishing screen: Bet 5

Figure A.1 (Cont.)

Questionnaire design

The following figures display screenshots from the z-Tree program the subjects used to place their bets. The graph and betslip including bet history, the continuation decision screen presented at the end of the third and fourth bets, the screen presented to those that decided to quit and the screen presented to those that continued betting until the end are presented in panels below.

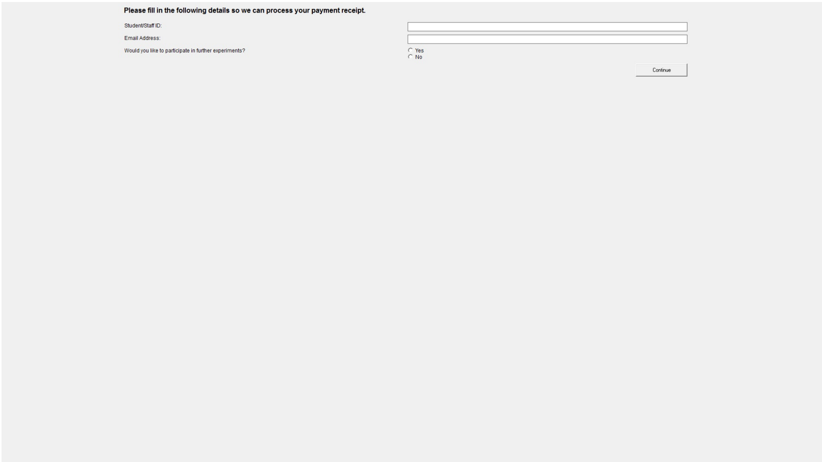

(i) Payfile details

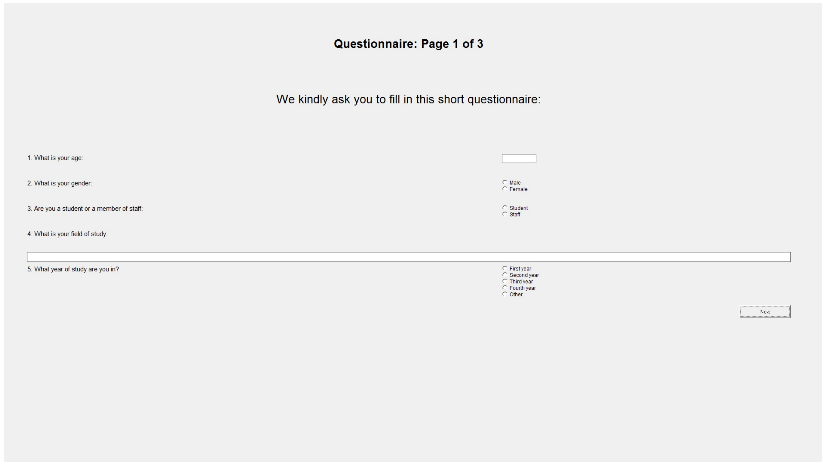

(j) Demographics

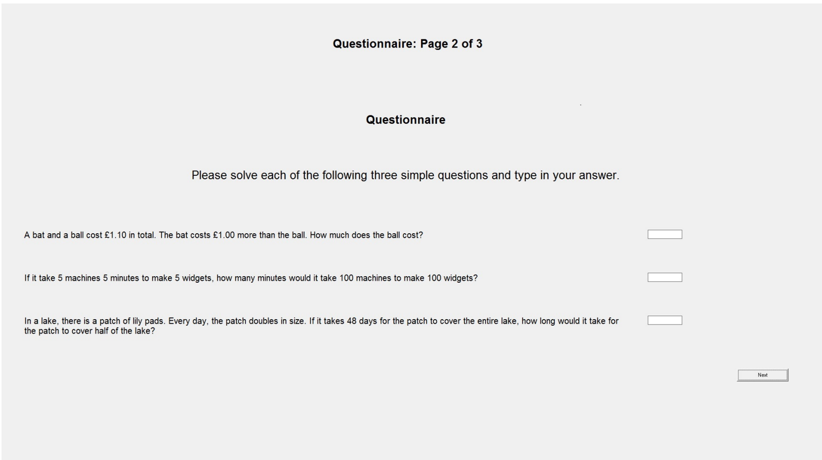

(k) CRT test

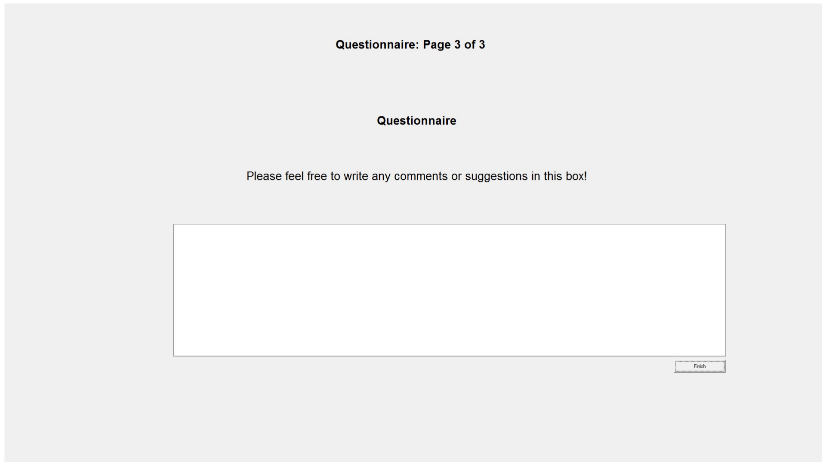

(l) Comments & Suggestions

> *endowment. Therefore, the amount of money you earn at the end of the experiment will be between £0 and £50.*
>
> *In each round of betting, you will be presented with a chart that contains simulated prices designed to have similar volatility to a stock index. The chart will be different for each participant and for each round. Therefore, your bets are completely independent from each other, so please do not talk or in any way communicate with other participants during this experiment. If you have a question or problem at any point in today's experiment, please raise your hand and one of us will come to you.*

Written instructions were also placed at each workstation along with a consent form, however subjects were instructed that they could participate without giving consent to their data being used for research purposes. There was no issue with leakage of information prior to the verbal instructions being read out: all subjects had the same role, endowment and rewards. In order to re-affirm the salience of rewards, an illustrative example of the expected payout was given in the written instructions. The example, however, contained a balanced pairs of examples, with contain no obvious behavioural suggestion. The written instructions were checked for loaded words or phrases. In Game Theory, a fact is considered common knowledge if each agent knows it, knows that all other agents know it, and knows that all other agents know that all agents know it, while it is acceptable in experiments to consider publicly announced information common knowledge (Friedman and Sunder, 1994, p.212). Both of the documents, along with a poster for the experiment and a sample payment receipt are included in Appendix A). The contents of the verbal and written instructions constituted 'common knowledge'.

Table A.1

Raw Data

Each subject was endowed with £25 at the beginning of the experiment. Subjects bet a fixed stake of £5 during each round, but were given the opportunity to quit at each point after the second bet. Their earnings to date were displayed on the continuation screen in other to highlight this as a reference point. This table presents each subjects total earnings and the point at which the subject quit.

Participant	Round1	Round2	Round3	Round4	Round5
1	£20	£15	£20	£15	-
2	£30	£35	£40	£35	-
3	£30	£35	£30	£35	-
4	£20	£25	-	-	-
5	£20	£15	£10	£15	£10
6	£20	£15	£20	£15	£10
7	£20	£25	£20	£25	£20
8	£30	£25	£30	£35	£40
9	£30	£35	£40	£35	£40
10	£20	£25	£30	£25	£30
11	£20	£15	£10	-	-
12	£20	£25	£20	£15	£10
13	£30	£25	£20	-	-
14	£20	£15	£10	£15	£10
15	£20	£25	£20	£15	£20
16	£30	£25	£30	£25	£20
17	£20	£25	£20	£15	£10
18	£30	£25	£20	£15	£10
19	£20	£15	-	-	-
20	£20	£25	£30	£35	£40
21	£20	£15	£20	£15	£10
22	£30	£35	£40	£45	£40
23	£20	£15	-	-	
24	£30	£25	£30	£35	£40
25	£30	£35	£40	£35	£40
26	£20	£25	£20	£25	£30
27	£20	£15	£10	£5	£10
28	£30	£35	£40	-	-

Appendix B

Appendix to Thesis

- B.1 Fixed-odds betting on campus
- B.2 University media coverage of the experiment in Chapter 7

Figure B.1

Fixed-odds betting on campus

I was passing the library at lunch some time in February of last year. Two students studying Actuarial Science in Heriot-Watt had a 'bucket shop' set up outside the library. The bet offering was as follows: *'Choose whether a stock or commodity will go up or down in the next five minutes. If you win, you get a profit of 60% on your investment. If you lose, you get 10% back.'* The minimum and maximum bets were £1 and £5, with a bet duration of five minutes. Trading prices and settlement prices were provided by TradePlus500. It was cleverly framed, but with the probability of an up or down move effectively p=0.5, and an expected value of -£1 + (£1.60 * 0.5) + (£0.1 * 0.5) = -0.15, the were changing a 15% margin/over-round for an even-money bet! I declined the bet, but wished them luck.

Figure B.2

ILW Experiment media coverage

The following image was taken by one of the participants during the experiment outline in Chapter 7 in order to cover the event in a blog during ILW at the University of Edinburgh on 17th February 2015 (Shrove Tuesday). Available at: `https://ilwuofe.wordpress.com/2015/02/18/quite-the-pancake-day/`

> Not quite the usual first task I undertake on a Tuesday morning, however *Could You Beat the Stock Market and Win £50?* proved a very exciting and engaging event. In short, all participants were allocated £25, and in the subsequent five rounds were shown a stock price chart and asked to decide whether they thought the price would rise or fall in the short term. With everyone given different charts there was no way of peeking at another screen to see what their choice was! Overall, the results showed that it was almost 50/50 between winning and losing each round, a result helping to show that in the short term the stock market can't be predicted!

(a) Photo courtesy of George Wood, University of Edinburgh

References

Ackert, L. F., Mazzotta, S., Qi, L., 2011. An experimental investigation of asset pricing in segmented markets. Southern Economic Journal 77 (3), pp. 585–598.

Alevy, J. E., Haigh, M. S., List, J. A., 2007. Information cascades: Evidence from a field experiment with financial market professionals. The Journal of Finance 62 (1), pp. 151–180.

Anderson, A., 2008. Is online trading gambling with peanuts? SIFR Research Report Series: Institute for Financial Research.

Andrikogiannopoulou, A., 2010. Estimating risk preferences from a large panel of real-world betting choices. Working Paper: Princeton University.

Andrikogiannopoulou, A., Papakonstantinou, F., 2011. State-dependent risk preferences: Evidence from online sports gambling. Working Paper: Imperial College London.

Apicella, C., Azevedo, E. M., Christakis, N. A., Fowler, J. H., 2014. Evolutionary origins of the endowment effect. American Economic Review 104 (6), 1793–1805.

Avery, C., Chevalier, J., 1999. Identifying investor sentiment from price paths. Journal of Business 72 (4), 493–521.

Barber, B. M., Odean, T., 1999. The courage of misguided convictions. Financial Analysts Journal 55 (6), 41–55.

Barber, B. M., Odean, T., 2000. Trading is hazardous for your wealth: The common stock investment performance of individual investors. Journal of Finance 116, 261–292.

Barber, B. M., Odean, T., 2001. Boys will be boys: Gender, overconfidence, and common stock investment. The Quarterly Journal of Economics 116 (1), pp. 261–292.

Barber, B. M., Odean, T., 2002. Online investors: Do the slow die first? Review of Financial Studies 15 (2), 455–488.

Barberis, N., 2012. A model of casino gambling. Management Science 58 (1), 35–51.

Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. Journal of Financial Economics 49, 307–343.

Barberis, N., Xiong, W., 2009. What drives the disposition effect? An analysis of a long-standing preference-based explanation. Journal of Finance 64, 751–784.

Barberis, N. C., Thaler, R., 2013. The Behaviour of Individual Investors. In: G. M. Constantinides, M. Harris and R. M. Stulz (eds.), Handbook of the Economics of Finance. Elsevier North Holland.

Barone-Adesi, G., Mancini, L., Shefrin, H., 2012. A tale of two investors: Estimating risk aversion,optimism, and overconfidence. Working paper, Swiss Finance Institute.

Barron, G., Erev, I., 2003. Small feedback-based decisions and their limited correspondence to description-based decisions. Journal of Behavioral Decision Making 16 (3), 215âĂŞ233.

Barseghyan, L., Molinari, F., O'Donoghue, T., Teitelbaum, J. C., 2015. Estimating risk preferences in the field. Journal of Economic Literature.

Battalio, R. C., Kagel, J. H., Jiranyakul, K., 1990. Testing between alternative models of choice under uncertainty: Some initial results. Journal of Risk and Uncertainty 3 (1), 25–50.

Baucells, M., Villasis, A., 2010. Stability of risk preferences and the reflection effect of prospect theory. Theory and Decision.

Bay, T., Sjödin, U., McGoun, E. G., 2011. Gaming the system. International Journal of Critical Accounting.

Bengtsson, C., Persson, M., Willenhag, P., 2005. Gender and overconfidence. Economics letters 86 (2), 199–203.

Benoît, J.-P., Dubra, J., 2011. Apparent overconfidence. Econometrica 79 (5), 1591–1625.

Benoit, J.-P., Dubra, J., Moore, D., 2014. Does the better-than-average effect show that people are overconfident? Two experiments. J. Europ. Econ. Assoc.

Benoit, J.-P., Dubra, J., Moore, D. A., 2008. A proper test of overconfidence. MPRA Working Paper.

Benos, A. V., 1998. Aggressiveness and survival of overconfident traders. Journal of Financial Markets 1 (3), 353–383.

Bernouli, D., 1738. Specimen theoriae novae de mensure sorti. Commentarii Academiae Scientiarun Imperiales Petropolitanae 5, 175-192. Translated in: Econometrica 22 (1954), 23–36.

Biais, B., Hilton, D., Mazurier, K., Pouget, S., 2005. Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. The Review of economic studies 72 (2), 287–312.

Blavatskyy, P. R., 2011. Loss aversion. Economic Theory 46 (1), 127–148.

Bloomfield, R., Hales, J., 2002. Predicting the next step of a random walk: Experimental evidence of regime-shifting beliefs,. Jornal of Financial Economics 65, 397–414.

Bondarenko, O., Bossaerts, P., 2000. Expectations and learning in iowa. Journal of Banking & Finance 24, 1535–1555.

Borghesi, R., 2012. A clean test of the impact of the disposition effect on assets prices. Working Paper: University of South Florida.

Brav, A., Heaton, J. B., 2002. Competing theories of financial anomalies. Review of Financial Studies 15:2, 575–606.

Brier, G. W., 1950. Verification of weather forecasts expressed in terms of probability. Monthly Weather Review 78, 1–3.

Brink, A. G., Rankin, F. W., 2013. The effects of risk preference and loss aversion on individual behavior under bonus, penalty, and combined contract frames. BEHAVIORAL RESEARCH IN ACCOUNTING.

Broihanne, M., Merli, M., Roger, P., 2014. Overconfidence, risk perception and the risk-taking behavior of finance professionals. Finance Research Letters 11 (2), 64–73.

Brown, A., Yang, F., 2015. Salience and the disposition effect: Evidence from the introduction of ‘cash-outs in betting markets. Working Paper: University of East Anglia.

Brown, R., Sarma, N., 2007. CEO overconfidence, CEO dominance and corporate acquisitions. Journal of Economics and business 59 (5), 358–379.

Camerer, C., Babcock, L., Loewenstein, G., Thaler, R., 1997. Labor supply of New York City cab drivers: One day at a time. Quarterly Journal of Economics 112 (2), 407–441.

Camerer, C., Ho, T. H., 1999. Experienced-weighted attraction learning in normal form games. Econometrica 67, 827–874.

Camerer, C., Lovallo, D., 1999. Overconfidence and excess entry: An experimental approach. American economic review, 306–318.

Cervellatia, E. M., Pattitoniac, P., Saviolib, M., 2013. Entrepreneurial underdiversification: Over optimism and overconfidence. Working Paper, The Rimini Centre for Economic Analysis, Italy.

Chang, T., Solomon, D. H., Westerfield, M. M., 2013. Looking for someone to blame: Delegation, cognitive dissonance, and the disposition effect. Journal of Finance (forthcoming).

Chen, M. K., 2006. How basic are behavioral biases? Evidence from capuchin monkey trading behavior. Journal of Political Economy 114 (3), 517–537.

Chiang, Y.-M., Hirshleifer, D., Qian, Y., 2011. Do investors learn from experience? Review of Financial Studies 24 (5), 1560–1589.

Chui, P. M. W., 2001. An experimental study of the disposition effect: Evidence from Macau. Journal of Psychology and Financial Markets 2 (4), 216–222.

Constantinides, G. M., 1983. Capital market equilibrium with personal tax. Econometric 51, 611–636.

Croxson, K., Reade, J. J., 2007. Information and efficiency: Goal arrival in soccer betting. Discussion Papers: Department of Economics, University of Birmingham.

Daniel, K. D., Hirshleifer, D., Subrahmanyam, A., 2004. A theory of overconfidence, self-attribution, and security market under-and over-reactions. Tech. rep., EconWPA.

De Bondt, W. F. M., 1993. Betting on trend: Intuitive forecasts of financial risk and return. International Jouranl of Forecasting 9, 355–371.

De Bondt, W. F. M., Thaler, R., 1985. Does the stock market overreact? The Journal of Finance 40:3, 793–805.

De Bondt, W. F. M., Thaler, R., 1990. Do security analysts overreact? American Economic Review 80:2, 52–57.

De Bondt, W. F. M., Thaler, R. H., 1987. Further evidence on investor overreaction and stock market seasonality. The Journal of Finance 42:3, 557–581.

Deaves, R., Lüders, E., Luo, G. Y., 2008. An experimental test of the impact of overconfidence and gender on trading activity. Review of finance, rfn023.

Dhar, R., Zhu, N., 2002. Up close and personal: An individual level analysis of the disposition effect. Yale ICF Working Paper 02 (20).

Di Guisto, R., Brown, A., Maughan, N., 2013. Rationality and risk intelligence in binary betting. Working Paper, Cayetano Gaming / AQR Capital Management.

Dommer, S. L., Swaminathan, V., 2013. Explaining the endowment effect through ownership: The role of identity, gender, and self-threat. Journal of Consumer Research 39 (5), 1034–1050.

Dorn, D., Sengmueller, P., 2009. Trading as entertainment. Management Science 55 (4), 591 – 603.

Durham, G. R., Hertzel, M. G., Martin, J. S., 2005. The market impace of trends and sequences in performance: New evidence. Journal of Finance 60 (5), 2251–2569.

Dutcher, E. G., Balafoutas, L., Lindner, F., Ryvkin, D., Sutter, M., 2015. Strive to be first or avoid being last: An experiment on relative performance incentives. Games and Economic Behavior 94, 39 – 56.

Egon Franck, E. V., Nuesch, S., 2010. Prediction accuracy of different market structures: Bookmakers versus a betting exchange. International Journal of Forecasting 26, 448–459.

Fabre, B., François-Heude, A., 2009. Optimism and overconfidence investors' biases: A methodological note. Finance 30 (1), 79–119.

Fama, E. F., 1970. Efficient capital markets. Journal of Finance, 383–417.

Feng, L., Seasholes, M., 2005. Do investor sophistication and trading experience eliminate behavioural biases in financial markets? Review of Finance 9, 305–351.

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

Fong, W. M., 2014. The lottery mindset: Investors, gambling and the stock market. Palgrave.

Frederick, S., 2005. Cognitive reflection and decision making. Journal of Economic Perspectives 19 (4), 25–42.

Friedman, D., Sunder, S., 1994. Experimental Methods: A Primer for Economists. Cambridge University Press.

Friedman, M., Savage, L. J., 1948. The utility analysis of choices involving risk. The Journal of Political Economy 56:4, 279–304.

Fullbrunn, S., Rau, H. A., Weitzel, U., 2014. Does ambiguity aversion survive in experimental asset markets? Journal of Economic Behavior & Organization 107, 810 – 826.

Gervais, S., Heaton, J., Odean, T., 2002. The positive role of overconfidence and optimism in investment policy. Rodney L. White Center for Financial Research Working Paper.

Gervais, S., Heaton, J., Odean, T., 2003. Overconfidence, investment policy, and executive stock options. Rodney L. White Center for Financial Research Working Paper 15 (02).

Gervais, S., Heaton, J. B., Odean, T., 2011. Overconfidence, compensation contracts, and capital budgeting. The Journal of Finance 66 (5), 1735–1777.

Gervais, S., Odean, T., 2001. Learning to be overconfident. Review of Financial studies 14 (1), 1–27.

Gilbert W. Bassett, J., 1981. Point spreads versus odds. Journal of Political Economy 89.

Gjerde, O., 1994. Market efficiency in scandinavian racetrack betting. Scandinavian Journal of Management 10, 409–420.

Glaser, M., Weber, M., 2007. Overconfidence and trading volume. Working Paper, Universit ÌĹat Mannheimt.

Glosten, L. R., Milgrom, P. R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100.

Goel, A. M., Thakor, A. V., 2008. Overconfidence, CEO selection, and corporate governance. The Journal of Finance 63 (6), 2737–2784.

Golec, J., Tamarkin, M., 1998. Bettors love skewness, not risk, at the horse track. Journal of Political Economy 106, 1, 205–225.

Goodman, J. S., Blum, T. C., 1996. Assessing the non-random sampling effects of subject attrition in longitudinal research. Journal of Management 22:4, 627–652.

Grinblatt, M., Keloharju, M., 2009. Sensation seeking, overconfidence, and trading activity. The Journal of Finance Vol. 64, No. 2, 549–578.

Grinblatt, M., Keloharju, M., Ikäheimo, S., 2008. Social influence and consumption: Evidence from the automobile purchases of neighbors. The Review of Economics and Statistics Vol. 90, No. 4, 735–753.

Grinblatt, M., Keloharju, M., Linnainmaa, J., 2011. IQ and stock market participation. The Journal of Finance Vol. 66, No. 6, 2121–2164.

Haigh, M. S., List, J. A., 2005. Do professional traders exhibit myopic loss aversion? An experimental analysis. The Journal of Finance 60 (1), pp. 523–534.

Hanson, R., Oprea, R., Porter, D., 2006. Information aggregation and manipulation in an experimental market. Journal of Economic Behavior & Organization 60 (4), 449 – 459.

Harbaugh, W. T., Krauseb, K., Vesterlund, L., 2001. Are adults better behaved than children? Age, experience, and the endowment effect. Economics Letters 70, 175–181.

Harrison, G. W., List, J. A., 2004. Field experiments. Journal of Economic Literature 42 (4), 1009–1005.

Hartzmark, S. M., Solomon, D. H., 2012. Efficiency and the disposition effect in nfl prediction markets. Working paper, University of Chicago.

Hausch, D. B., William, 1995. Efficiency of Sports and Lottery Betting Markets. North-Holland, Amsterdam, Ch. In: Robert A. Jarrow, V. M., Ziemba, W. T. (Eds.), Handbook of Operations Research and Management Science, pp. 545–580.

Hirshleifer, D., 2001. Investor psychology and asset pricing. Journal of finance 56 (4), 1533–1597.

Hirshleifer, D., Luo, G. Y., 2001. On the survival of overconfident traders in a competitive securities market. Journal of Financial Markets 4 (1), 73–84.

Hirshleifer, J., 1966. Decision under uncertainty: Applications of the state-preference approach. The Quarterly Journal of Economics 80 (2), 252–277.

Hodges, S. D., Tompkins, R., Ziemba, W. T., 2003. The favorite/long-shot bias in S&P 500 and FTSE 100 index futures options: The return to bets and the cost of insurance. EFA 2003 Annual Conference Paper No. 135 . Sauder School of Business Working Paper.

Hoffmann, A. O. I., Post, T., 2013. How does investor confidence lead to trading? Theory and evidence on the links between investor return experiences, confidence, and investment beliefs. Working paper, Maastricht University.

Holt, C. A., Laury, S. K., 2002. Risk aversion and incentive effects. American Economic 92 (5), 1644–1655.

Huber, J., Kirchler, M., Stefan, M., 2014. Experimental evidence on varying uncertainty and skewness in laboratory double-auction markets. Journal of Economic Behavior & Organization 107, Part B, 798 – 809, empirical Behavioral Finance.

Jiao, P., 2012. Belief in mean reversion and the disposition effect: An experimental test. Working paper: Claremont Graduate University.

Kahneman, D., Tversky, A., 1973. On the psychology of prediction. Psychological Review 80, 237–51.

Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica, 47:2, 263–292.

Kaustia, M., 2010. Prospect theory and the disposition effect. The Journal of Financial and Quantitative Analysis Vol. 45, No. 3, 791–812.

Kaustia, M., Knüpfer, S., 2008. Do investors overweight personal experience? Journal of Finance 63, 2679–702.

Keloharju, M., Knüpfer, S., Linnainmaa, J., 2012. Do investors buy what they know? Product market choices and investment decisions. The Review of Financial Studies Vol. 25, No. 10, 2921–2958.

Keloharju, M., Knüpfer, S., Torstila, S., 2008. Do retail incentives work in privatizations? The Review of Financial Studies Vol. 21, No. 5, 2061–2095.

Kelvey, R. M., Zovoina, W., 1975. A statistical model for the analysis of ordinal level dependent variables. Journal of Mathematical Society 4, 103–120.

Kirchler, M., Huber, J., Kleinlercher, D., 2011. Market microstructure matters when imposing a tobin tax: Evidence from the lab. Journal of Economic Behavior & Organization 80 (3), 586 – 602.

Ko, K. J., Huang, Z. J., 2007. Arrogance can be a virtue: Overconfidence, information acquisition, and market efficiency. Journal of Financial Economics 84 (2), 529–560.

Koessler, F., Noussair, C., Ziegelmeyer, A., 2012. Information aggregation and belief elicitation in experimental parimutuel betting markets. Journal of Economic Behavior & Organization 83 (2), 195 – 208.

Kotze, A., Joseph, A., 2009. Double barrier cash or nothing options: A short note. Financial Chaos Theory, Johannesburg, South Africa.

Krause, A., Wei, K.-C. J., Yang, Z., 2009. Determinants of disposition and reverse disposition effects. Working Paper: University of Bath.

Kruse, J. B., Thompson, M. A., 2001. A comparison of salient rewards in experiments: money and class points. Economics Letters 74 (1), 113 – 117.

Kruse, J. B., Thompson, M. A., 2003. Valuing low probability risk: Survey and experimental evidence. Journal of Economic Behavior & Organization 50 (4), 495 – 505.

Kyle, A. S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.

Lange, J., Economides, N., 2005. A parimutuel market microstructure for contingent claims. European Financial Management 11, 25–49.

Langer, E. J., Roth, J., 1975. Heads i win, tails it's chance: The illusion of control as a function of the sequence of outcomes in a purely chance task. Jouranl of Personality and Social Psychology 32.

Lefevre, E., 1923. Reminiscences of a Stock Operator. John Wiley & Sons.

Levitt, S. D., 2004. Why are gambling markets organised so differently from financial markets? The Economic Journal 114, 223–246.

Lindner, F., 2014. Decision time and steps of reasoning in a competitive market entry game. Economics Letters 122 (1), 7 – 11.

Linnainmaa, J., 2010. Do limit orders alter inferences about investor performance and behavior? The Journal of Finance Vol. 65, No. 4, 1473–1506.

Linnainmaa, J. T., 2011. Why do (some) households trade so much? Review of Financial Studies 24 (5), 1630–1666.

Linnainmaa, J. T., Saar, G., 2012. Lack of anonymity and the inference from order flow. The Review of Financial Studies Vol. 25, No. 5, 1414–1456.

Ludvig, E., Spetch, M., 2011. Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? PLoS One.

Machina, M. J., 1987. Choice under uncertainty: Problems solved and unsolved. Journal of Economic Perspectives 1 (1), 121–154.

Madhavan, A., 2002. Market microstructure: A practitioner's guide. Financial Analysts Journal 58.

Mahani, R., Bernhardt, D., 2007. Financial speculators' underperformance: Learning, self-selection, and endogenous liquidity. The Journal of Finance 62, 1313–1340.

Malmendier, U., Tate, G., 2005a. CEO overconfidence and corporate investment. The Journal of Finance 60 (6), pp. 2661–2700.

Malmendier, U., Tate, G., 2008. Who makes acquisitions? CEO overconfidence and the market's reaction. Journal of financial Economics 89 (1), 20–43.

Malmendier, U., Tate, G. A., 2005b. Does overconfidence affect corporate investment? CEO overconfidence measures revisited. European Financial Management 11 (5), 649–659.

Malmendier, U., Taylor, T., 2015. On the verges of overconfidence. Journal of Economic Perspectives 29 (4), 3–8.

Manski, C., 2004. Interpreting the predictions of prediction markets. NBER Working Paper 10359.

Markowitz, H., 1952. The utility of wealth. The Journal of Political Economy 60-2, 151–158.

Marshall, B. R., 2009. How quickly is temporary market inefficiency removed? The Quarterly Review of Economics and Finance 49, 917–930.

Menkhoff, L., Schmeling, M., 2010. Whose trades convey information: Evidence from a cross-section of traders. Journal of Financial Markets 13, 101–128.

Menkhoff, L., Schmidt, U., Brozynski, T., 2006. The impact of experience on risk taking, overconfidence, and herding of fund managers: Complementary survey evidence. European Economic Review 50 (7), 1753–1766.

Merli, M., 2014. Trading activity and overconfidence: First evidence from a large european database. Working Paper: Universite de Strasbourg.

Mosteller, F., Nogee, P., 1951. An experimental measurement of utility. Journal of Political Economy 59:5, 371–404.

Nolte, I., 2016. A detailed investigation of the disposition effect and individual trading behavior: A panel survival approach. Working Paper, Financial Econometric Research Centre, University of Warwick.

Odean, T., 1998a. Are investors reluctant to realize their losses? Journal of Finance Vol. LIII, No. 5, 1775–1798.

Odean, T., 1998b. Volume, volatility, price, and profit when all traders are above average. The Journal of Finance 53 (6), pp. 1887–1934.

Pastor, L., Veronesi, P., 2009. Learning in financial markets. Annual Review of Financial Economics 1, 361–381.

Qiu, J., 2007. Loss aversion and mental accounting: The favorite londshot bias in parimutuel betting. Jena Economic Research Papers.

RBC Capital Markets, 2001. Trade a short currency position with a european binary call.
URL `https://www.rbccm.com/global/file-410660.pdf`

Rob Bauer, M. C., Eichholtz, P., 2008. Option trading and individual investor performace. Working Paper: Maastricht University.

Roth, A. E., Erev, I., 1995. Learning in extensive form games: Experimental data and simple dynamic models in the intermediate run. Games and Economic Behaviour 8, 164–212.

Rubin, J. Z., Brockner, J., 1975. Factors affecting entrapment in waiting situations: The Rosencrantz and Guildenstern effect. Journal of Personality and Social Psychology 31 (6), 1054–1063.

Russell, T., Thaler, R., 1985. The relevance of quasi rationality in competitive markets. The American Economic Review 75:5, 1071–1082.

Sauer, R. D., 1998. The economics of wagering markets. Journal of Economic Literature.

Scheinkman, J. A., Xiong, W., 2003. Overconfidence and speculative bubbles. Journal of political Economy 111 (6), 1183–1220.

Sebehela, T., 2011. Binary options and real assets. Acuitas - Journal of Management Research 1 (2), 45–54.

Seru, A., Shumway, T., Stoffman, N., 2010. Learning by trading. Review of Financial Studies 23 (2), 705–739.

Shefrin, H., 2000. Beyond Greed and Fear: Understanding Behavioral Finance and the Psychology of Investing. Oxford University Press.

Shefrin, H., 2008. A Behavioral Approach to Asset Pricing. Elsevier.

Shefrin, H., Statman, M., 1985. The disposition to sell winners too early and ride losers too long: Theory and evidence. The Journal of Finance 40 (3), 777–790.

Shive, S., 2010. An epidemic model of investor behavior. The Journal of Financial and Quantitative Analysis Vol. 45, No. 1, 169–198.

Shleifer, A., Vishny, R. W., 1997. The limits of arbitrage. Journal of Finance 52 (1).

Skala, D., 2008. Overconfidence in psychology and finance: An interdisciplinary literature review. Bank i Kredyt 1 (4), 33–50.

Smith, M. A., Vaughan-Williams, L., 2010. Forecasting horse race outcomes: New evidence on odds bias in UK betting markets. International Journal of Forecasting 26, 543–550.

Smith, V. L., 1976. Experimental economics: Induced value theory. AMerican Economic Review 66 (2), 274–9.

Smith, V. L., Suchanek, G. L., Williams, A. W., 1988. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. Econometrica: Journal of the Econometric Society, 1119–1151.

Song, R., Jang, S., Hanssens, D. M., Suh, J., 2014. When overconfidence meets reinforcement learning. Working Paper, Kansas State University.

Statman, M., 2001. Lottery traders. Financial Analysts Journal: Forthcoming.

Statman, M., Thorley, S., Vorkink, K., 2006. Investor overconfidence and trading volume. Review of Financial Studies 19 (4), 1531–1565.

Staw, B. M., 1981. The escalation of commitment to a course of action. The Academy of Management Review 6 (4), 577–587.

Stockl, T., Huber, J., Kirchler, M., Lindner, F., 2015. Hot hand and gambler's fallacy in teams: Evidence from investment experiments. Journal of Economic Behavior & Organization 117, 327 – 339.

Strahilevitz, M. A., Odean, T., Barber, B. M., 2011. Once burned, twice shy. Journal of Marketing Research 48, 102–120.

Summers, B., Duxbury, D., 2012. Decision-dependent emotions and behavioural anomalies. Organizational Behaviour and Juand dDecision Processes 118, 226–238.

Tetlock, P. C., 2004. How efficient are information markets: Evidence from an online exchange. Yale University working paper.

Tetlock, P. C., 2008. Liquidity and prediction market efficiency. Yale University working paper.

Thaler, R., Johnson, E., 1990. Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. Management Science 36,6, 643–660.

Thaler, R. H., 1985. Mental accounting and consumer choice. Marketing Science 4 (3), 199–214.

Thaler, R. H., Ziemba, W. T., 1988. Anomalies: Parimutuel betting markets: Racetracks and lotteries. Journal of Economic Perspectives 2, 1 61– 174.

Trevelyan, R., 2008. Optimism, overconfidence and entrepreneurial activity. Management Decision 46 (7), 986–1001.

Vaughan-Williams, L., 1999. Information efficiency in betting markets: A survey. Bulleting of Economic Research 51, 1–39.

Vaughan Williams, L., 2005. Information Efficiency in Financial and Betting Markets. Cambridge University Press.

Weber, M., Camerer, C. F., 1998. The disposition effect in securities trading: An experimental analysis. Journal of Economic Behavior & Organization 33, 167–184.

Willimas, R. J., Wood, R. T., Parke, J., 2012. History, current worldwide situation, and concerns with internet gambling. In: Willimas, R. J., Wood, R. T., Parke, J. (Eds.), Routledge International Handbook of Internet Gambling. Routledge, pp. 3–27.

Wiseman, T., 2009. Sequential choice and non-bayesian observational learning. International Game Theory Review 11:03.

Wolfers, J., Zitzewitz, E., 2004. Prediction markets. Journal of Economic Perspectives 18, 107–126.

Wolfers, J., Zitzewitz, E., 2006a. Interpreting prediction market prices as probabilities. FRB of San Francisco Working Paper No. 2006-11.

Wolfers, J., Zitzewitz, E., 2006b. Prediction markets in theory and practice. NBER Working Paper No. 12083.

Woodland, L. M., Woodland, B. M., 1994. Market efficiency and the favourite-longshot bias: The baseball betting market. Journal of Finance 49, 269–79.

Wright, B. T., 2008. An experimental test of self-justification in stock sale decisions: Evidence of a reverse disposition effect in unsophisticated investors. Working Paper: Harvard College.

Yeoh, L.-Y., Wood, A., 2011. Overconfidence, competence and trading activity. Working Paper, Unviersity of Essex.

Zitzewitz, E., 2006. Price discovery among the punters: Using new financial betting markets to predict intraday volatility. Preliminary Paper: Dartmouth College.