Framing Effects and the Market Selection Hypothesis: Evidence from Real-World Isomorphic Bets*

Alasdair Brown[†]
University of East Anglia

Fuyu Yang[‡]
University of East Anglia

May 10, 2021

*We would like to thank the Editor, Georgia Kosmopoulou, and two referees for very helpful comments and suggestions. We are also grateful to Eberhard Feess, Glenn Harrison, Luis Muga, Amrish Patel, Rafael Santamaria, Tom Stafford, Bob Sugden, Ferdinand Vieider and audiences at Swansea University, Universidad Pública de Navarra Pamplona, the University of Reading, the University of Sheffield (Psychology), and the Center for the Economic Analysis of Risk (CEAR) Sports Betting Workshop in Cork for earlier feedback. Alasdair Brown acknowledges financial support from the British Academy and the Leverhulme Trust (Ref: SG140097). The research presented in this paper was carried out on the High Performance Computing Cluster supported by the Research and Specialist Computing Support service at the University of East Anglia. The usual disclaimer applies.

[†]Corresponding Author. School of Economics, University of East Anglia, Norwich NR4 7TJ, U.K.. Email: alasdair.brown@uea.ac.uk

[‡]School of Economics, University of East Anglia, Norwich NR4 7TJ, U.K.. Email: fuyu.yang@uea.ac.uk

Abstract

We collect data on 75 million GBP of tennis bets over a 6 year period to analyse whether participants in high-stakes environments recognise simple framing differences. The structure of this market means that we can place the same bet at the same time in two different ways. These two isomorphic bets are framed differently, and often priced differently. We find that bettors make frequent mistakes, choosing the worse of the two bets in 35% of cases. However, bettors who choose the inferior price earn higher returns from their bets, suggesting that their effort has been focused on fundamental information acquisition rather than bet execution. The net result is that market selection may, if anything, slightly favour those who are unable, or unwilling, to recognise framing differences.

JEL Classification: D01, D81, D91, G14

Keywords: framing, market selection, betting

1 Introduction

One of the classic results in psychology and behavioural economics is the effect of framing. In the famous Tversky and Kahneman (1981) disease problem, subjects chose to gamble on a riskier treatment programme when the outcomes were framed in terms of lives lost, but took the less risky treatment programme when the same outcomes were framed in terms of lives saved. Similar results were found when subjects were presented with financial losses and gains. These results are inconsistent with rational choice theory, where preferences are independent of the formulation of the decision problem.

As Fryer *et al.* (2012) noted, if people are susceptible to framing manipulations this can be used for one of two ends. It can either be used productively, by a policymaker for example, to increase savings for retirement or reduce gambling expenditure. Alternatively, it could be used manipulatively by firms or individuals seeking to sell poor financial products or low-yielding lotteries and gambles.

But are participants in real-world markets susceptible to framing manipulations? Framing effects may be diluted in real markets compared to laboratory settings because of the more high-powered incentives, because individuals have had an opportunity to gain experience of the environment and therefore recognise framing differences, and also because many markets will select out those with a weakness for framing. In addition, for a firm to profitably exploit framing effects, individuals must be willing to pay a premium when certain frames are applied.

In this paper we analyse whether participants in a high-stakes market recognise simple framing differences. We study 35 million prices and more than 600,000 bets placed on Betfair on Wimbledon tennis matches between 2008 and 2013. The structure of this market means that we can place the same bet at the same time in two different ways. We can place a 'back' bet that player 1 will win. Or we can place a 'lay' bet that player 2 will lose. Although framed differently, both bets will produce a positive payout if player 1 wins, and a negative payout if player 2 wins. However, due to non-zero bid-ask spreads, these two isomorphic bets often do not have the same price. Discrepancies in these prices mean that there is almost always a strictly dominant option, out of the two bets, for the bettor to choose.

We find that bettors make frequent mistakes when choosing which bet to take: 35% of the 660,261 bets in our sample are placed at the worse of the two prices on offer. These choices are violations of first-order stochastic dominance. It would appear that participants often fail to convert two simple gambles into a common frame in what Kahneman and Tversky (1979) labelled 'the editing phase'.

Our initial suspicion was that this may stem from bettors reluctance to use the unfamiliar lay bets. For more than two centuries, gamblers have for the most part placed back bets, leaving the lay side of betting to professional bookmakers. It may also simply be more intuitive for bettors to wager on the winner rather than place bets on the loser. Although such behaviour would still reflect a costly failure to recognise a simple framing difference, it may limit the external validity of our results for settings which do not have such a pronounced difference in the familiarity of different frames. In our results, we find that there is indeed an element of familiarity at play. The proportion of 'mistakes' reaches 48% when the optimal choice is the unfamiliar lay bet. However, there are still a significant proportion of mistakes (19%) when the optimal choice is the familiar back bet. These bettors are ignoring the more familiar and intuitive back bet and yet still choosing the inferior gamble.

We then ask whether market selection winkles out those with an inability to recognise framing differences. According to this line of thinking, those who fail to recognise framing differences will lose money and be slowly driven out of the market, leaving only those able to recognise framing differences. This does not seem to be the case. The proportion of mistakes stays constant, at roughly 35%, over the 6 years. We find that bettors who choose the inferior price actually earn higher returns from their bets. This suggests that their effort has been focused on fundamental information acquisition rather than picking the best of the two prices. The net result is that market selection may, if anything, slightly favour those who are unable, or unwilling, to recognise framing differences.

2 Related Literature

Our results contribute to a literature on framing in field settings. Due to the difficulty of producing two isomorphic choice sets – which differ only in terms of their frame – researchers have typically relied on experiments. Researchers have predominantly used these experiments to see if loss aversion (Kahneman and Tversky, 1979) can be harnessed to increase effort. This approach has been applied in a number of contexts: Hossain and List (2012) and Hong *et al.* (2015) studied factory workers, Fryer *et al.* (2012) studied teachers, De Quidt (2017) studied Mechanical Turk workers, List and Samek (2015) studied children's food choices, Romanowich and Lamb (2013) studied a group looking to quit smoking, and Gächter et al. (2009) studied academic economists. The evidence has been mixed, with only Hossain and List (2012) and Fryer *et al.* (2012) finding increased effort in the loss frame for the whole experimental sample. When effort is indistinguishable across loss and gain frames, this implies that real-world subjects are not particularly affected by frames.

Perhaps most closely related to our paper is Sonnemann *et al.* (2013), who study framing effects in a series of laboratory and field betting markets. They find that market participants assign higher probabilities to outcomes when the event space is randomly partitioned into a larger number of sub-intervals. When presented with a potential (and perhaps unlikely) outcome, subjects cannot help but assign some probability to it. Therefore, greater partition in the event space leads to inflated probability judgements. This is a clear framing effect.

Our paper is different from these earlier studies in two main ways. Bettors on the exchange are presented with both frames at the same time. We are therefore not examining whether subjects respond differently to different randomly-assigned frames (as the experiments do). Instead, we examine whether subjects are able to convert two simple gambles, observed simultaneously, into a common frame and then accurately rank them. A second difference is that in our setting isomorphic bets are produced naturally without the intervention of an experimenter. This allows us to examine a scale of activity – more than 75 million GBP worth of bets traded over 6 years – beyond even the most well-funded experiments.

Our work is also related to a literature on the market selection hypothesis. This hypothesis is inspired by Darwin's theory of natural selection (Darwin, 1859), and is often attributed first to Alchian (1950) and Friedman (1953). The idea is, as Cootner (1964, p. 80) puts it:

'If any group of investors was consistently better than average in forecasting stock prices, they would accumulate wealth and give their forecasts greater and greater weight. ... Conversely, investors who were worse than average in forecasting ability would carry less and less weight.'

The implication is that rational investors will survive and prosper at the expense of irrational or biased investors, and therefore in the long-term will come to dominate markets, making asset prices more efficient. Oberlechner and Osier (2012) use survey-based estimates to show that overconfident foreign currency dealers are not necessarily driven out by the market selection process, as inexperienced and experienced dealers show similar levels of overconfidence. They then speculate that overconfidence may confer advantages (e.g. persistence under adversity) which offset any bias in forecasting. Following the model of Blume and Easley (1992), much of the focus in the literature has been on whether unbiased or biased traders consume more. If biased traders consume less than unbiased traders, then they may actually survive at the expense of unbiased traders simply by picking up the risk premium generated by higher levels of investment. This is an empirical question, which Kendall and Oprea (2018) tackle in the laboratory.

We do not focus on the relative investment and consumption patterns of biased and unbiased traders. Instead, we look at the correlation between 1) the ability to recognise framing differences and 2) fundamental forecasting ability. We cannot observe the types of traders that enter the market, and the long-term efficiency of the market will undoubtedly be determined not simply by the exit of traders but also by the entry of new traders. But we can observe the direct disciplining effect of the market in the form of the returns of current market participants. The interesting aspect of our results is that bettors who are superior at recognising framing differences are less adept at forecasting match outcomes. It is the mildly irrational bettors – unable or unwilling to convert two gambles into a common frame – who produce the more accurate (long-term) forecasts. It may therefore be these mildly irrational traders who prosper due to market selection.

3 Data

3.1 Framing

The setting for our study is Betfair, the world's largest betting exchange. Betfair was founded in 2000 and allows bettors to take long and short positions on an outcome. Bettors can provide liquidity/quotes for others, or place bets at prices already quoted by others. This limit order book design is common in financial market settings (Parlour and Seppi, 2008), but represents a widening of choice in sports betting. Prior to the development of Betfair, bettors could take only the quotes provided by bookmakers, and could take only a long position on an outcome. Because Betfair allow for long and short positions, in two-outcome events such as tennis it is possible to place the same bet in one of two ways: either bet on player 1 to win, or bet on player 2 to lose.

To illustrate this market format, consider the screenshot in Figure 1. Odds are quoted on the outcome of the match between John Isner and Philipp Kohlschreiber. If I wish to bet on John Isner to win, I could place a back bet at the best odds of 1.63. This means that I will receive 1.63 back for each 1 GBP I stake, if Isner wins. If Isner loses, I lose my stake. Alternatively, I could place a lay bet on Philipp Kohlshreiber to lose at odds of 2.58. This means that for every 1 GBP I accept on this bet, I would be liable for 1.58 GBP (2.58 GBP - 1 GBP) if Kohlshreiber wins. If Isner wins, I get to keep the stake.

These bets are equivalent. Both bets accrue a profit if Isner wins but incur a loss if Kohlschreiber wins. However, the relative prices (odds) of these bets mean that one strictly dominates the other. Consider Figure 2, where we have tinkered with the stakes to illustrate the point. A back bet on Isner of 15.80 GBP at odds of 1.63 returns a profit of 9.95 GBP if Isner wins. If Kohlschreiber wins, the stake of 15.80 GBP is lost. An equivalent lay bet on Kohlschreiber at odds of 2.58 returns 10 GBP if Isner wins, and loses 15.80 GBP if Kohlschreiber wins. The lay bet strictly dominates the back bet as it incurs the same loss in the event of a Kohlschreiber win (15.80 GBP), but earns a higher profit (10 GBP versus 9.95 GBP) if Isner wins.

In Figure 3 we produce another screenshot for illustration. From a different match, we have two options for betting on Jeremy Chardy to beat his opponent Gilles Simon. We can either place a back bet on Chardy at odds of 2.52, or place a lay bet on Simon at odds of 1.66. The back bet is the dominant choice in this example, as it pays 10.03 GBP (as opposed to 10 GBP) if Chardy wins, and loses 6.60 GBP if Simon wins (same as the lay bet).

Why are these price differences not arbitraged away? The reason is that there are non-zero bid-ask spreads. Consider the Figure 3 example again. In a textbook arbitrage, the arbitrageur would take a long position in the underpriced asset and a short position in the overpriced asset. In this case, the back bet on Jeremy Chardy is underpriced and the lay bet on Gilles Simon is overpriced. Taking a long position on Chardy is straightforward at odds of 2.52. But, in order to take a short position on the Simon lay bet, we would need to cross the spread and place a back bet on Simon. Even if the spread was at the minimum tick size of 0.01 and you could back Simon at 1.65, this 'arbitrage' would involve back bets on Chardy and Simon at 2.52 and 1.65 respectively, implying a sum of win probabilities of $(1/2.52) + (1/1.65) = 1.003$. This is not a case of a 'dutch-book', and instead guarantees a (small) loss for the arbitrageur.

We collected data on 6 years (2008-2013) of Wimbledon Men's Tennis Championship betting from Fracsoft, a third-party provider of Betfair data. This data-set involves second-by-second observations of the best back and lay odds for each player in every match. We also observe the last transaction price and total betting volume, measured each second. The data-set includes pre-match betting, and inplay betting as the matches unfolded. We have 640 matches in our sample. This gives us more than 35 million observations.

Using this data we calculate two prices:

$$Price_1 = 1/BackOdds_1 \tag{1}$$

where $BackOdds_1$ are the best back odds quoted on player 1, and

$$Price_2 = 1 - (1/LayOdds_2) \tag{2}$$

where $LayOdds_2$ are the best lay odds quoted on player 2. Price 1 gives us the implied win probability of player 1 inferred from the back odds on that player, and Price 2 gives us the implied win probability of player 1 inferred from the lay odds on his opponent.

We then calculate the difference between these two prices (implied win probabilities):

$$PriceDifference = Price_1 - Price_2 \tag{3}$$

If the price difference is negative, then the dominant choice is to back player 1. If the price difference is positive, the dominant choice is to lay player 2. Consider Figures 1 and 2 again. For a bet on John Isner to win the price difference is $= 1/1.63 - (1 - 1/(2.58)) = 0.001$. This is positive and therefore, as we saw in the example, the dominant choice is therefore to lay Philipp Kohlschreiber.

We calculate the price difference for every player in every second (thereby switching the designations of players 1 and 2). In Table 1 we summarise this data. We also plot a histogram of the price difference in Figure 4 (truncated at an absolute price difference of 0.1). The absolute price difference is, on average, 0.01. This is quite a substantial price discrepancy, particularly when it is scaled up with a high volume bet. In one case, the price difference is as large as 0.898. There are non-zero price differences in 98.5% of observations. In almost all circumstances there is a dominant option for the bettor to choose.

Our next step is to reconstruct the orders placed by bettors. We are interested in market orders – placed at quotes already in the book – as bettors who traded with market orders made a choice between the dominant and the dominated option. We use the Lee and Ready (1991) algorithm to classify back and lay orders. This means that we classify that an order has taken place if the total betting volume for that player is higher at time t than at time

$t-1$. This order is classified as a back order if the transaction price at time t is closer to the best back price at $t-1$ than it is to the best lay price at $t-1$. If it is closer to the best lay price, we designate it a lay order. If the transaction price is precisely in the middle of the preceding back and lay quotes, then it cannot be classified.

Using this classification, we calculate that there were back orders in 3.61% of seconds, and lay orders in 3.25% of seconds (Table 1). 0.1% of seconds involved an order that we could not classify, as the transaction price was equidistant from the preceding back and lay quotes. We also estimate that the average back order was for a stake of 860 GBP, and the average lay order involved a stake of 929 GBP. This is likely to be an overestimate as it assumes that only one order took place in each second. Nevertheless it gives a rough idea of the sums wagered in this market.

In Figure 5 we plot a local polynomial regression of back and lay order frequencies against the price difference between the two isomorphic bets. When the price difference is negative, the dominant choice is to back the player. When the price difference is positive, the dominant choice is the lay the opponent. We find that bettors more often than not get this correct, but there are a significant number of mistakes. Bettors frequently lay when they should back (on the left-hand-side), and back when they should lay (on the right-hand-side).

We therefore create a trading score, which equals 1 if the dominant choice is taken, and 0 if the dominated choice is taken. For this trading score, we consider only bets where the full volume of the bet could have been accommodated at the best price on either side of the book. This condition is satisfied for 660,261 of the 2.48 million bets in our sample. These bets represent a small fraction of the total volume (75 million out of 2.1 billion GBP), but with these bets we can be certain that the prices we considered in Equations 1 and 2 were in fact available for the full volume that the bettor wished to bet. In the remaining cases the bettor would have had to execute some of the bet at the 2nd/3rd/4th best price, at least on one side of the book, and we do not observe all of these prices. For the 666,261 bets we consider, the average stake size is a more realistic 115 GBP.

The value of the trading score is summarised in Table 2. The average trading score is 0.653 from 660,261 bets, indicating that as much as 35% of bets are mistakes, where bettors opt for the inferior, dominated choice. In Figure 6 we plot a local polynomial regression of the trading score against the price difference. As might be expected, the trading score increases as the absolute price difference increases – and it becomes clearer that one option is superior to the other – but there is still a substantial proportion of mistakes. These mistakes are not simply a result of bettors having insufficient time to evaluate the two options before prices change. The probability that a price changes in any given second (and therefore the probability that the optimal choice *may* change) is just 0.46% pre-match, i.e. less than 1 in every 200 seconds, and yet the pre-match trading score is 0.54 (indicating 46% of pre-match bets are mistakes).

The average loss from choosing the dominated option in our sample is 0.35 GBP per bet. We calculate this loss by multiplying the absolute price difference by the stake size for each of our inferior bets (trading score=0). In one instance, a bettor incurred a loss of 211.40 GBP from choosing the wrong bet. A total loss of 80,528 GBP – spread over 229,293 bets – was incurred in our sample. This is an estimated lower bound as we have only included bets where the full stake could be wagered at the best price on either side of the book.

In the remainder of Table 2 we summarise the trading score for two different cases. We summarise the trading score for cases when the dominant choice involved a back bet, and for cases when the dominant choice involved a lay bet. One argument could be that bettors make mistakes only due to an unfamiliarity with the lay bet frame, as back bets have been the dominant frame for betting with bookmakers for over 200 years. As the asymmetry in Figure 6 implies, bettors do appear to make mistakes, in part, because of unfamiliarity with the lay bet frame. The average trading score is only 0.519 when the dominant choice involves a lay bet. However, it should be pointed out that there is still a substantial percentage of mistakes (19%) when the dominant choice involves the more familiar back bet.

We also break the sample down by early years (2008-2010) and later years (2011-2013). It could be argued that familiarity with the lay bet will increase over the years, reducing the aversion to this frame. We do find some evidence of this. In the early years, the trading score when the dominant choice involved the lay bet frame was 0.498, but this increased to 0.537 in the later years. However, this does not appear to be caused by the market selecting out those with a weakness for framing, as the trading score for cases when the dominant choice involved the familiar back bet frame decreased over the same period, from 0.837 to 0.786. We illustrate these two results in Figure 7. The overall level of mistakes did not decrease from 35%. Instead, over the years bettors displayed a greater willingness to accept the unfamiliar lay bet frame, both when it was optimal to do so, and when it was not.

3.2 Market Selection

So why did market selection not reduce the frequency of errors? It is natural to think that bettors who fail to recognise framing differences would lose money and be slowly driven out of the market, to be replaced by bettors able to identify when two isomorphic bets are merely framed differently. The extent of this market selection effect, however, depends on how an ability to recognise framing differences corresponds with an ability to forecast match outcomes (long-term fundamentals). It is possible that bettors who fail to recognise framing differences are actually superior at forecasting match outcomes. This would potentially explain why the proportion of 'mistake' bets – placed at the worse of the two prices – does not decline over time. We estimate an equation of the following form:

$$y_i = \beta_0 + \beta_1 x_i + \beta_2 z_i + \epsilon_i \tag{4}$$

y_i is an indicator variable, equalling 1 if player i won and 0 otherwise, x_i is the implied probability of outcome i as measured by the average of $Price1$ and $Price2$, z_i is the trading score, and ϵ_i is an error term. A variant of this equation, without the z_i term, is often labelled the Mincer-Zarnowitz regression (Mincer and Zarnowitz, 1969). When estimating this equation, we consider only cases when one of the 660,261 bets were executed.

The idea behind this equation is as follows. A bettor in this market has two considerations: 1) the position they wish to accumulate, and 2) whether to accumulate this position with a back bet on one player or a lay bet on the other player. Bettors can increase their returns by 1) selecting players who will go on to win, and 2) executing this bet at the better of the two prices. We want to see if those who can recognise framing differences (trading score=1) are better or worse at predicting match outcomes than bettors who cannot, or do not, recognise framing differences (trading score=0). If those who recognise framing differences are better at forecasting match outcomes, then $\beta_2 > 0$. If they are worse, then $\beta_2 < 0$.

The results of our regressions are presented in Table 3. After controlling for prices, bets placed by those who *can* recognise framing differences (trading score=1) win 0.8% *less* often than bets placed by those who cannot (trading score=0). The results are particularly pronounced for bets on favourites (-1.7%), a finding which we illustrate in Figure 8. Overall, there is a small negative correlation between an ability to predict fundamentals and an ability to recognise framing differences. Our results are robust to using a logit specification rather than OLS, and qualitatively similar when we add the magnitude of the price difference and the stake size of the bet as control variables.

In Table 4, we break our sample down by pre-match trading and inplay trading. During inplay trading, there is a regular presence of 'courtsiders': bettors who place quick bets in the stadium to adversely select bettors watching on a delayed signal at home (Brown and Yang, 2016). We question whether this negative correlation is concentrated in inplay betting. If a bettor has a fleeting informational advantage, as courtsiders do, they may not have time to care about picking the best price. This may mechanically explain why those who forecast outcomes better are inferior at identifying framing differences. However, we find that the negative correlation is actually concentrated in pre-match trading. This suggests that bettors who fail to recognise framing differences are instead concentrating on fundamental research. In the case of tennis betting this could involve researching the form of the two players on this surface, the results of historical match-ups, and other factors.

But which of these two effects dominate? Which bettors earn higher returns? Those able to recognise framing differences? Or those who neglect framing differences but who display a superior ability to forecast match outcomes? To answer this question, we compare the returns for bets executed at the worse of the two prices (trading score=0) to returns for bets executed at the better of the two prices (trading score=1). We find that returns are higher for bets with a trading score of 0 (-10.4%) than bets with a trading score of 1 (-12.02%). Although both sets of returns are negative – suggesting that, on average, both types of trader will be slowly pushed out of the market – the rate will be slower for traders who fail to convert the two bets into a common frame. In other words, in this case market selection may slowly lead to the survival of mildly irrational traders, who commit violations of first-order stochastic dominance.

4 Conclusion

In this paper we study whether individuals who participate in real markets are able to recognise simple framing differences. Despite all of the factors that might limit framing effects in real markets – incentives, opportunities to gain experience, and the market selecting out those with a weakness for framing – real market participants are indeed unable, or unwilling, to spot differences in frames. This result is particularly striking because of the size of the stakes involved, and because an inability to convert two gambles into a common frame (Tversky and Kahneman, 1986) is costly in this environment.

However, those unable to convert the two gambles into a common frame actually perform better than those who can. Bettors who choose the inferior price earn higher returns from their bets, suggesting that their effort has instead been focused on fundamental information acquisition rather than bet execution. The net result is that market selection may, if anything, favour those who are unable, or unwilling, to recognise framing differences.

Another way to look at this is that the real-world, satisficing, boundedly-rational punter (Simon, 1955) – who forecasts fundamentals well but does not recognise framing differences – is guilty of the type of violation of first-order stochastic dominance that rational choice and many behavioural models of decision under risk would find inconceivable. What is more, it is this punter whom the process of market selection may slightly favour.

References

- Alchian, A., A., (1950). Uncertainty, Evolution, and Economic Theory. *Journal of Political Economy*, 58, 211-221.

- Blume, L., Easley, D., (1992). Evolution and Market Behavior. *Journal of Economic Theory*, 58, 9-40.

- Brown, A., Yang, F., (2016). Slowing Down Fast Traders: Evidence from the Betfair Speed Bump. Working paper.

- Cootner, P., H., (1964). *The Random Character of Stock Market Prices.* MIT Press, Cambridge, Massachusetts.

- Darwin, C., (1859). *On the Origin of Species by Means of Natural Selection.* John Murray, London.
- De Quidt, J., (2017). Your Loss is My Gain: A Recruitment Experiment with Framed Incentives. *Journal of the European Economic Association*, forthcoming.
- Gächter, S., Orzen, H., Renner, E., Starmer, C., (2009). Are Experimental Economists Prone to Framing Effects? A Natural Field Experiment. *Journal of Economic Behavior and Organization*, 70, 443-446.
- Friedman, M., (1953). *Essays in Positive Economics.* University of Chicago Press.
- Fryer, R., G., Jr., Levitt, S., List, J., Sadoff, S., (2012). Enhancing the Efficacy of Teacher Incentives Through Loss Aversion: A Field Experiment. NBER working paper, 18237.
- Hong, F., Hossain, T., List, J., A., (2015). Framing Manipulations in Contests: A Natural Field Experiment. *Journal of Economic Behavior and Organization*, 118, 372-382.
- Hossain, T., List, J., A., (2012). The Behavioralist Visits the Factory: Increasing Productivity Using Simple Framing Manipulations. *Management Science*, 58, 2151-2167.
- Kahneman, D., Tversky, A., (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47, 263-292.
- Kendall, C., Oprea, R., (2018). Are Biased Beliefs Fit to Survive? An Experimental Test of the Market Selection Hypothesis. *Journal of Economic Theory*, 176, 342-371.
- Lee, C., M., C., Ready, M., J., (1991). Inferring Trade Direction from Intraday Data.*Journal of Finance*, 46, 733-746.

- List, J., A., Samek, A., S., (2015). The Behavioralist as Nutritionist: Leveraging Behavioral Economics to Improve Child Food Choice and Consumption. *Journal of Health Economics*, 39, 135-146.

- Mincer, J.,A., Zarnowitz, V., (1969). *The Evaluation of Economic Forecasts*, in "Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance", pp. 1–46,

- Oberlechner, T., Osier, C., (2012). Survival of Overconfidence in Currency Markets. *Journal of Financial and Quantitative Analysis*, 47, 91-113.

- Parlour, C., A., Seppi, D., J., (2008). Limit Order Markets: A Survey. *Handbook of Financial Intermediation and Banking*, edited by Boot, A., W., A., Thakor, A., V., Elsevier.

- Romanowich, P., Lamb, R., J., (2013). The Effect of Framing Incentives as Either Losses or Gains with Contingency Management for Smoking Cessation. *Addictive Behaviors*, 38, 2084-2088.

- Simon, H., A., (1955). A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 69, 99-118.

- Sonnemann, U., Camerer, C., F., Fox, C., R., Langer, T., (2013). How Psychological Framing Affects Economic Market Prices in the Lab and Field. *Proceedings of the National Academy of Sciences*, 110, 11779-11784.

- Tversky, A., Kahneman, D., (1981). The Framing of Decisions and the Psychology of Choice. *Science*, 211, 453-458.

- Tversky, A., Kahneman, D., (1986). Rational Choice and the Framing of Decisions. *Journal of Business*, 59, 251-278.

Figures and Tables

(a)

Figure 1: A screenshot of the Betfair exchange. A bettor can bet that John Isner will win the match in one of two ways: either i) place a 'back' bet on Isner (at odds of 1.63), or ii) place a 'lay' bet on his opponent Kohlschreiber (offering odds of 2.58).

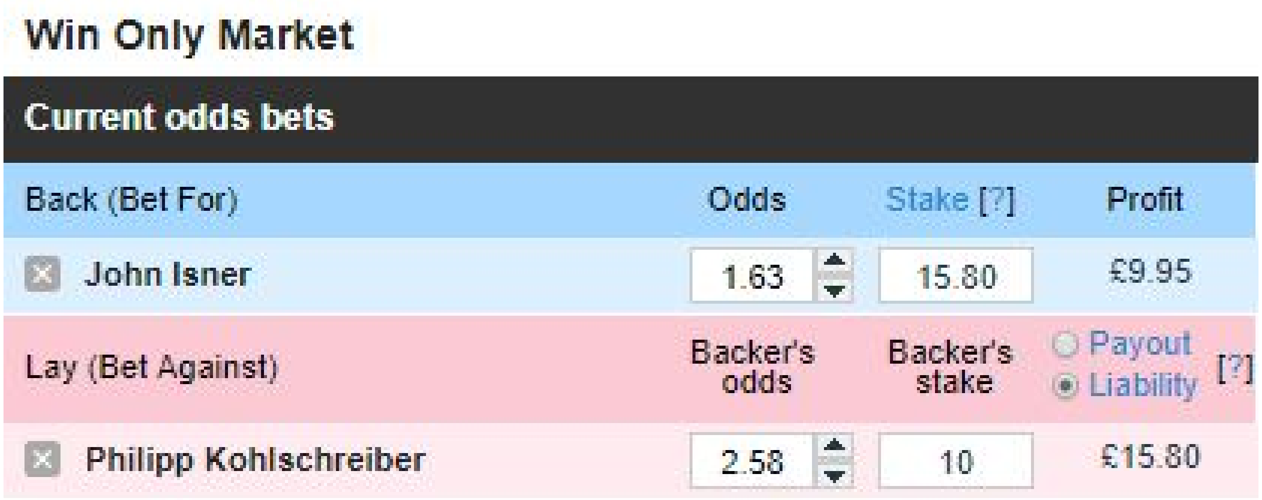

Figure 2: A comparison of the payoffs for a back bet on Isner or a lay bet on Kohlschreiber. The lay bet is the dominant choice in this example, as it pays 10 GBP in the event that Isner wins (as opposed to a 9.95 payoff for the back bet), and reaps the same loss as the back bet (15.80 GBP) if Kohlschreiber wins.

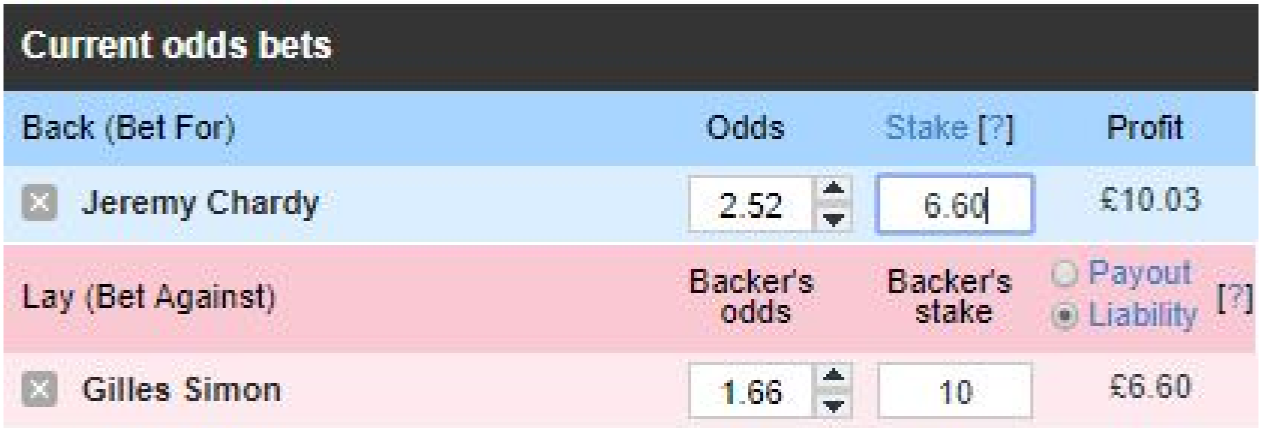

Figure 3: A comparison of the payoffs for a back bet on Jeremy Chardy or a lay bet on his opponent Gilles Simon. The back bet is the dominant choice in this example, as it pays 10.03 GBP in the event that Chardy wins (as opposed to a 10 GBP payoff for the lay bet), and reaps the same loss as the lay bet (6.60 GBP) if Simon wins.

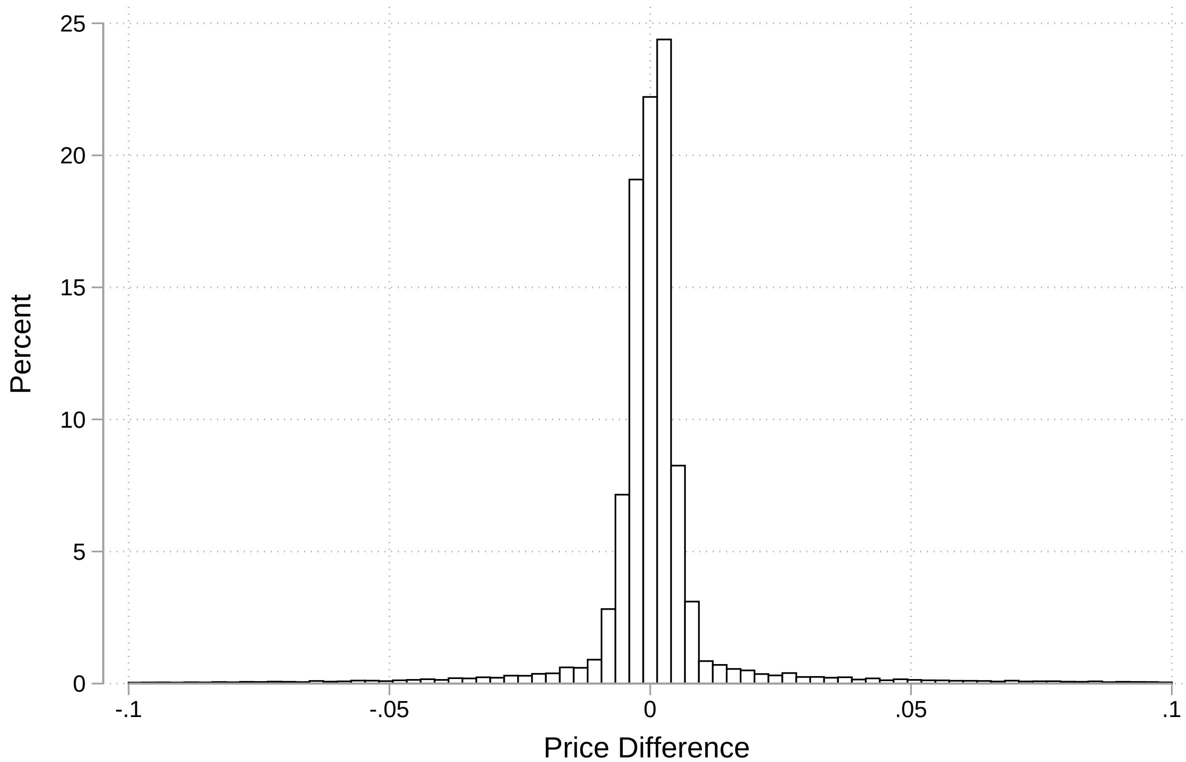

Figure 4: A histogram of the price difference between the two isomorphic bets.

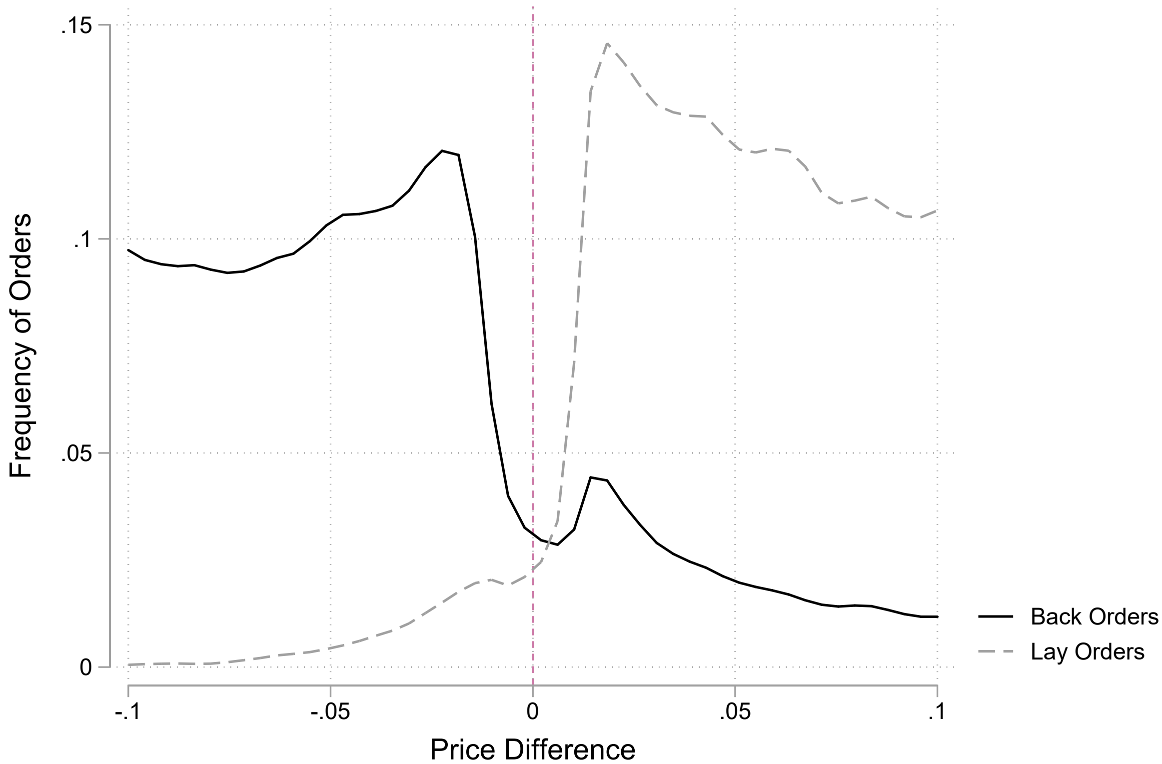

Figure 5: A local polynomial regression of the frequency of back/lay orders as a function of the price difference between the two isomorphic bets. When the price difference is negative, it is optimal to back the player. When the price difference is positive, it is optimal to lay his opponent.

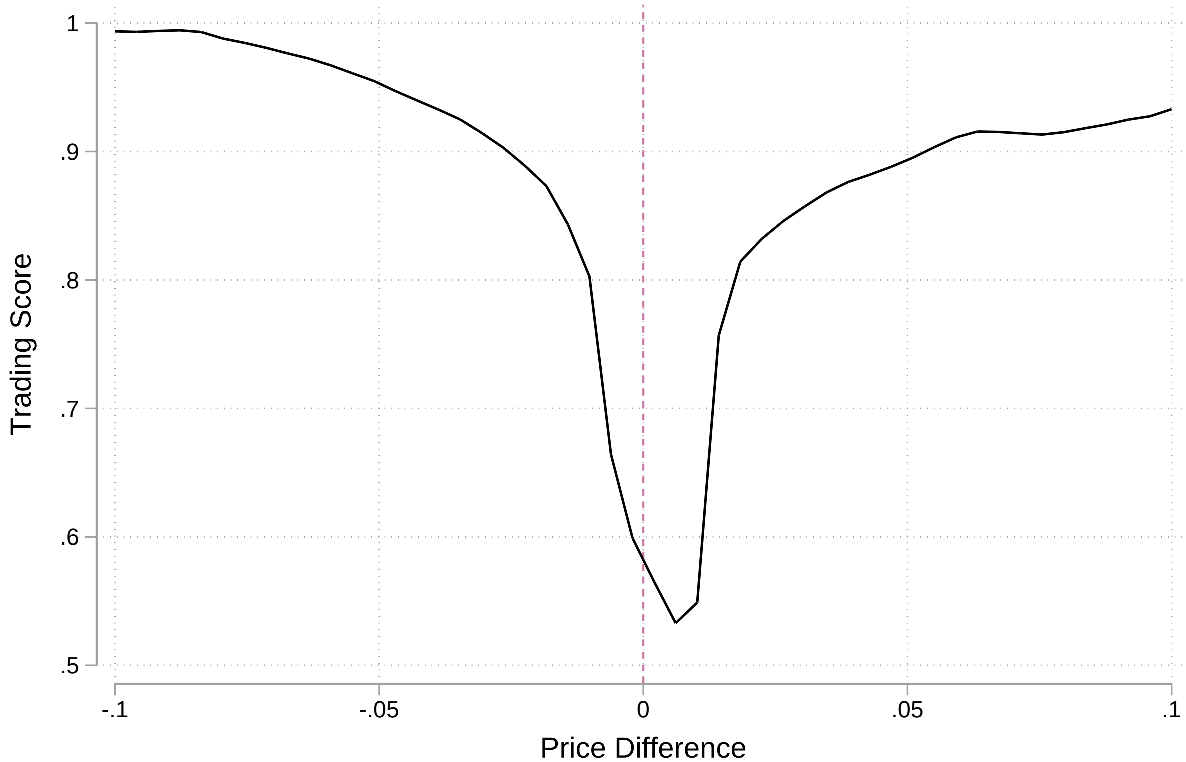

Figure 6: A local polynomial regression of the trading score as a function of the price difference between the two isomorphic bets. The trading score equals 1 if the dominant choice was taken, and 0 if the dominated choice was taken. On the left-hand-side, it is optimal to back the player. On the right-hand-side, it is optimal to lay the opponent.

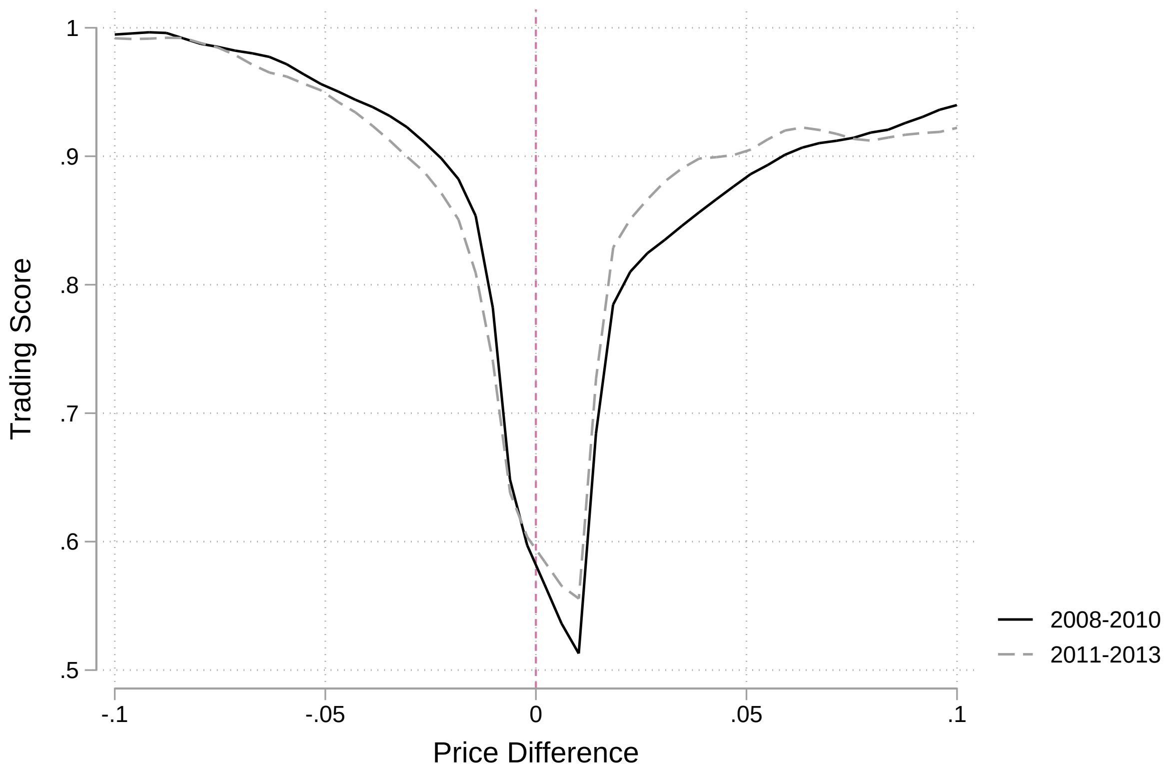

Figure 7: A local polynomial regression of the trading score as a function of the price difference between the two isomorphic bets. The trading score equals 1 if the dominant choice was taken, and 0 if the dominated choice was taken. On the left-hand-side, it is optimal to back the player. On the right-hand-side, it is optimal to lay the opponent. This time, separate estimations are carried out for the early years (2008-2010) and the later years (2011-2013).

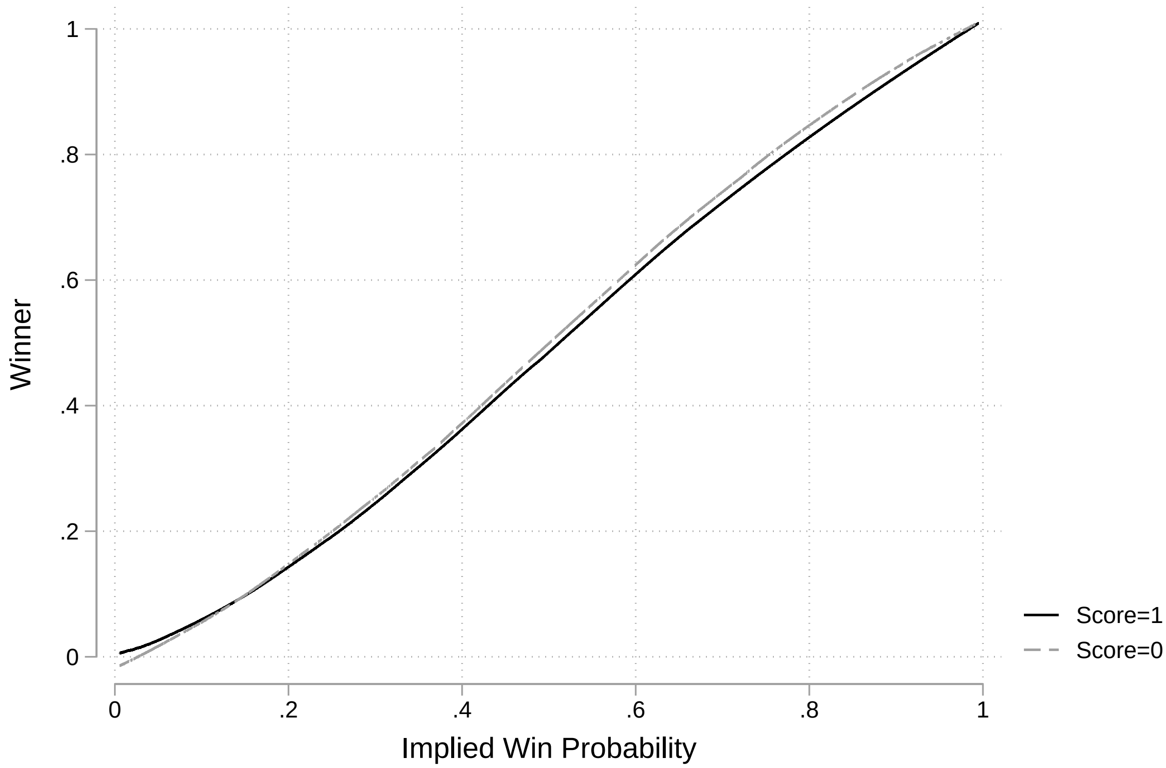

Figure 8: A lowess smoothed plot of actual win probability (average of the win indicator) against implied win probability. A comparison is made for bets with a trading score of 1 (executed at the best of the two prices), and bets with a trading score of 0 (executed at the worse of the two prices). Implied win probability is calculated as the average of Price 1 and Price 2 at the time of the bet.

Table 1: Summary Statistics

	(1)	(2)	(3)	(4)	(5)
Prices	N	mean	sd	min	max
Price 1	35,073,331	0.507	0.299	0.00100	0.990
Price 2	33,443,290	0.508	0.304	0.00990	0.999
Price Difference	33,194,428	0.00236	0.0357	-0.660	0.898
Absolute Price Difference	33,194,428	0.0113	0.0339	0	0.898
Order Indicators	N	mean	sd	min	max
Back Order	35,466,658	0.0361	0.186	0	1
Lay Order	35,466,658	0.0325	0.177	0	1
Unclassified	35,466,658	0.00127	0.0356	0	1
Order Size	N	mean	sd	min	max
Back Order Size	1,278,985	0.860	4.731	0.00199	744.7
Lay Order Size	1,153,339	0.929	4.614	0.00199	463.3

Summary statistics on prices, order frequencies, and order sizes. We display the prices for the two isomorphic bets, either betting on a player (Price 1) or betting against the opponent (Price 2). We also display the difference and the absolute difference between the two prices. Order sizes are measured in thousands of GBP.

Table 2: Trading Scores

	(1)	(2)
Trading Scores	N	mean
All Years (2008-2013)		
All	660,261	0.653
Back Bet Frame	303,781	0.810
Lay Bet Frame	356,480	0.519
Early Years (2008-2010)		
All	317,472	0.652
Back Bet Frame	144,088	0.837
Lay Bet Frame	173,384	0.498
Later Years (2011-2013)		
All	342,789	0.653
Back Bet Frame	159,693	0.786
Lay Bet Frame	183,096	0.537

Summary statistics on the trading score. The trading score equals 1 if the dominant choice was taken, and 0 if the dominated choice was taken. We compare the trading score for instances when the dominant choice involved a back bet, and for instances when the dominant choice involved a lay bet. In three separate panels we display statistics for all years (2008-2013), early years (2008-2010) and later years (2011-2013).

Table 3: Forecasting Match Outcomes

	(1)	(2)	(3)
VARIABLES	Winner	Winner	Winner
Average Price	1.125***	1.034***	1.012***
	(0.00139)	(0.00542)	(0.00403)
Trading Score	-0.00845***	-0.0174***	-0.00418***
	(0.00102)	(0.00164)	(0.00131)
Constant	-0.0655***	0.0162***	-0.0474***
	(0.00103)	(0.00458)	(0.00125)
Sample	All	Favourites	Longshots
Observations	660,261	270,869	389,392
R-squared	0.366	0.105	0.135

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Forecasting Match Outcomes. An indicator variable equalling 1 if the player won and 0 if he lost, is regressed on the average of the two prices at the time of the bet, and the trading score which equals 1 if the superior price of the two was taken and 0 if the inferior price was taken. Separate estimations are carried out for the full sample, bets on the favourite, and bets on the longshot. Estimation is by OLS.

Table 4: Forecasting Pre-Match and Inplay

	(1)	(2)
VARIABLES	Winner	Winner
Average Price	1.116***	1.124***
	(0.00281)	(0.00162)
Trading Score	-0.0250***	0.00121
	(0.00200)	(0.00121)
Constant	-0.0415***	-0.0757***
	(0.00214)	(0.00118)
Sample	Pre-Match	Inplay
Observations	183,155	477,106
R-squared	0.356	0.368

Robust standard errors in parentheses

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

Forecasting Pre-Match and Inplay. An indicator variable equalling 1 if the player won and 0 if he lost, is regressed on the average of the two prices at the time of the bet, and the trading score which equals 1 if the superior price of the two was taken and 0 if the inferior price was taken. Separate estimations are carried out for pre-match bets and inplay bets. Estimation is by OLS.