The Role of Speculative Trade in Market Efficiency: Evidence from a Betting Exchange*

Alasdair Brown†
University of East Anglia

Fuyu Yang‡
University of East Anglia.

December 11, 2014

*We would like to thank Raphael Markellos, Peter Moffatt, Arnold Polanski and an audience at UEA for comments. Alasdair Brown would like to acknowledge financial support for this project 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.

†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

Does speculative trade reduce mispricing - and help create efficient markets - or drive prices further from fundamentals? We analyse betting exchange trading, on 9,562 U.K. horse races in 2013 and 2014, to find out. Crucially, as each race is run, the fundamental value of bets is unambiguously revealed. We find that the direction and volume of market order trade is predictive of fundamentals, suggesting that speculative trade is, on average, conducive to market efficiency. However, much of this effect is concentrated in the in-running period (during races) - when, even without trade, asset fundamentals would be revealed seconds later.

JEL Classification: G02, G12, G14

Keywords: market efficiency, trading volume, asset fundamentals, betting markets

1 Introduction

What role does speculative trade play in market efficiency? Do speculators reduce mispricing, by buying up cheap assets and selling those that are expensive? Or do speculators - with their propensity for overconfidence, erroneous beliefs and herd behaviour - buy and sell the wrong assets, drive prices further from fundamentals, and make markets less efficient? These have always been important questions in finance, but have assumed greater importance in light of the heightened trading activity of recent years. Chordia *et al.* (2011) document that share turnover on the New York Stock Exchange increased more than five-fold from 1993 to 2008. The Bank of International Settlements estimates that daily turnover in foreign exchange markets was $5.3 trillion in 2013, up from $4 trillion in 2010 (Triennial Central Bank Survey (2013)). If trade has a distortionary effect on asset prices, the case for increasing trading costs - perhaps through a transactions tax - becomes stronger.

The problem with establishing the role that speculative trade has on market efficiency is that asset fundamentals are seldom observed. Stocks, to take the most prominent financial asset, are infinitely-lived.[1] One possibility is to look at the immediate effect of trading volume (in a particular direction) on returns and, if these returns do not subsequently reverse,

[1] Even for financial assets that are finitely-lived, we cannot cleanly observe fundamentals. In option

conclude that trading volume reduces mispricing and increases market efficiency. However, there is a mechanical element to such a test - as trading creates price pressure unrelated to information - and any analysis is extremely sensitive to the time-horizon chosen for reversal and is therefore less than satisfactory.

Moreover, theory is ambiguous regarding the role of speculative trade in market efficiency. In Milgrom and Stokey (1982), no trade takes place in a speculative market as the willingness to trade signals private information (and adverse selection). On the other hand, add liquidity traders - who trade for non-informational reasons - as in Glosten and Milgrom (1985), and trade does take place, and in fact drives prices towards fundamental values. The presence of liquidity traders allows the informed to hide.

An alternative view is that noise traders may drive prices further from fundamental values. In De Long *et al.* (1990), noise traders - governed by sentiment - bid up the price of already over-priced assets. The arbitrageurs in their model, rather than correcting this mispricing, sometimes join in as short-term returns are higher for over-valued assets. This 'positive-feedback trading' drives prices further from fundamentals and creates larger bubbles. Another channel through which trading activity may be distortionary is when traders exhibit herd behaviour, as in Avery and Zemsky's (1998) model. In such a model, traders ignore their own private information signal, trade based on inferences from the trades of others, and therefore amplify any mispricings.

In short, there is no single clear prediction of whether, on average, speculative trade moves prices towards fundamentals (i.e. removes mispricing), or not.

In the spirit of Angrist and Pischke (2009), imagine the perfect empirical test. Suppose firms lived for only one day. Firms raised capital in equity markets at the beginning of each day, and these stocks were traded from morning until late afternoon. At the end of the afternoon, each firm was wound up and a liquidating dividend was paid to all of the final stock holders. If such a situation existed, we could sort all stocks by the amount of buying

markets, for example, suppose there is heavy buying of an out-of-the-money call option, which subsequently finishes in the money. The relationship between trading volume and fundamentals, however, is endogenous, as the counterparty to the option trade may have hedged their exposure in the stock market, thereby driving up its price, and increasing the likelihood that the option finishes in-the-money. This could be totally unconnected to firm fundamentals.

and selling volume that took place throughout the day. (For every trade there is a buyer and seller, and therefore it would be sensible to consider buying and selling only by market order traders, who trade at quoted prices in the book). We could then compare the returns for those stocks that received pronounced buying volume, with those that did not. If returns for the former were higher, trade would appear to drives prices toward fundamentals (measured by the size of the liquidating dividend) and increase market efficiency. Or, if returns for those with large buying volume were lower, then, on balance, speculative trade would appear to have a detrimental impact on market efficiency.

While such a scenario is far-fetched in equity markets, it closely resembles trading in U.K. horse race betting markets. Bets are traded throughout the day on Betfair, a public limit order book similar in structure to the major financial exchanges. Traders submit market and limit orders, and either bet on (i.e. buy) or against (sell) a given horse. The race is then run, and the fundamental value of each bet is unambiguously revealed.[2]

In this paper we analyse Betfair trading on 9,562 U.K. races run in 2013 and 2014. We use approximately 225 million observations of the limit order book to extract data on more than 60 million separate trades. Using a variation of the Lee and Ready (1991) algorithm, we classify back orders (betting on a horse) and lay orders (betting against) and calculate the net buying volume (back minus lay volume) by market order traders for each horse. We then divide trades up by comparing the winning probabilities of those bets that received greater net volume within their race (above the median), with those that did not. (Importantly, we are able to control for the well-known favourite longshot bias, by using nearest neighbour matching methods to pair up horses with very similar implied win probabilities (using the trade-weighted average betting prices)). In effect, we are asking whether a relatively large amount of buying volume is predictive of a horse's fundamental probability of winning? Put another way, if two horses are trading at the same price, but one receives a great deal more

[2] We are assuming that horse race fundamentals are unaffected by the trading volume prior to and during the race. In other words, we assume that the jockey/horse does not adapt his/her effort after observing a number of favourable/unfavourable punts. As our strongest results relate to the in-running period - when jockeys cannot check betting activity - we do not believe that this is a strong assumption. Note, that this does not preclude the fixing of races; it only precludes the decision to fix a race after having seen the betting (i.e. in the last 20 minutes prior to the race).

bets, does this indicate that the high volume horse is more likely to win than his/her odds suggest? If so, this would imply that speculative trade is, on average, conducive to market efficiency, as mispriced quotes attract a disproportionate amount of trade.

We also exploit two dimensions of variation in trading on the exchange. Trading takes place in the win market, where the identity of the first-placed horse is speculated upon, and also in the place market, where the identities of those that finish just behind (typically in places 2 to 4, but this depends on the number of runners) is the focus of betting. In some ways the place market can be likened to the stock market and the win market to a corresponding option market. Trade in the place market involves a broad bet that a horse will do well, in much the same way a buy trade in the stock market speculates that a stock is undervalued. Betting in the win market, however, is a much more precise bet that the horse will do *very* well, in much the same way that a trader can buy a heavily out-of-the-money call option to speculate that the underlying stock is grossly undervalued.

Our second source of variation is produced by a clean separation in trading period. Trading takes place in the hours before races, and also in the 1-5 minutes that a race is run. We liken the period prior to racing to trading prior to earnings announcements, for example, when all information that is publicly available is reasonably stale. The in-running period therefore corresponds to trading during and immediately after earnings announcements, when traders chew over the implications of the new information and incorporate them into asset prices.

These two sources of variation lead to a mixture of results. We find that the average trade prior to races is only slightly predictive of fundamentals. For example, in our preferred matching specification, a horse that receives a relatively large amount of net trading volume is 0.4% more likely to place than a horse with (almost) identical pre-race odds that did not receive such volume. The average trade is, however, more informative prior to races in the win market, where more precise predictions are made. A horse that receives a relatively large amount of net trading volume is 1% more likely to win than its low volume, but same price, equivalent.

The picture is very different when we consider trading during races (i.e. as information arrives). A horse that receives a disproportionate amount of bets is 3.6% more likely to win, and 12.2% more likely to place, than horses trading at similar prices but without the same level of net trading volume. In other words, while the average trade is only slightly informative regarding fundamentals when information is stale, it is highly predictive of fundamentals when information is fresh. Putting the two time periods together, there is evidence to suggest that speculative trade is at least conducive, and certainly not unfavourable, towards market efficiency. Perhaps the results suggest that Glosten and Milgrom's (1985) model of trading - albeit with different proportions of informed traders prior to and during information events - dominates the models of Milgrom and Stokey (1982), De Long *et al.* (1990) and Avery and Zemsky (1998) in our market.

This brings us to the external validity of our results. On the debit side, the most attractive property of the betting market is also its greatest weakness. As each market is short-lived, and fundamentals revealed before the day is out, positive-feedback traders of the type modelled in De Long *et al.* (1990) are unlikely to be found. These traders know that a stock is overvalued but, having observed that sentiment has taken over a section of noise traders, also buy stock, in the hope of capitalising on short-run price movements and exiting prior to any ultimate correction. Due to the brevity of trading on the horse-racing we study, it is unlikely that traders buy (short) assets that they know are over (under) valued.

On the credit side, we have a pure speculative market without many of the other trading motivations (e.g. hedging, risk-sharing) found in financial markets. While betting is a leisure activity for some, there are a sufficient number of informed, experienced professionals present in this market. Moreover, the incentives for efficiently predicting winners are clear. More than 6 billion GBP is staked on the horse races that we study. Finally, and perhaps most importantly, the market structure mirrors the dominant market structure of financial markets (Parlour and Seppi (2008)). The limit order book operated by Betfair allows for the backing (buying) and laying (shorting) of horses, the submission (and cancellation) of limit and market orders, and the use of algorithmic trading strategies.

It is this market structure that distinguishes this paper from earlier work on the efficiency of betting markets. Researchers have long noted that fundamentals are revealed in horse racing betting. Griffith (1949), Ali (1977) and Snowberg and Wolfers (2010), among many others, analysed betting returns in U.S. pari-mutuel markets. In these markets, the payout on each horse (in the case of a win) is inversely proportional to the amount wagered on that horse. In other words, prices cannot reflect information *without* trade. The existence of a positive favourite longshot bias - where returns to betting on favourites exceed those of betting on longshots - suggests that trade goes some way, but not quite far enough, to incorporating all information into asset prices. On an alternative path, Shin (1993) and Vaughan Williams and Paton (1997), among others, have looked at U.K. bookmaker markets. Here, prices can reflect fundamentals without any trade - as bookmakers can update their quotes irrespective of trading volumes - but the patterns of bets taken by bookmakers are not typically revealed. The betting exchange - by replicating the dominant financial market structure, and providing the transparency necessary to construct detailed data on the direction and volume of orders - allows for a clean test of the role of speculative trade in market efficiency.[3]

The rest of the paper is structured as follows. In Section 2 we describe the data, and the construction of trading volumes for each horse in each market. In Section 3 we present the analysis, which includes regression and nearest neighbour matching methods. In Section 4 we discuss the results in relation to recent empirical results in financial markets, and Section 5 concludes.

2 Data

The focus of our study is the Betfair betting exchange. The exchange is the largest of its kind in the world, with the majority of its customers based in the U.K.. Prior to the development of this exchange, and others of its kind, the U.K. betting market was comprised of bookmakers and the Tote, a pari-mutuel pool. On the exchange, bettors can bet on (i.e. back) or bet against (lay) a given horse, and can also submit both market orders and limit orders. Market orders execute with limit orders already in the book, and limit orders sit in the book until an offsetting market order arrives. Prior to the inception of the exchanges it was not possible for bettors to take short positions (lay) or submit limit orders.

Prices on the exchange are quoted in the form of odds, including the stake. For example, if a bettor backs a horse at 3, then 2 pounds will be returned for every pound staked if

[3] Smith *et al.* (2006) also studied the favourite longshot bias on Betfair - finding a slight positive bias - but did not consider the role of trade in market efficiency.

the horse wins. If a bettor lays a horse at 4, he/she is liable to pay 3 pounds to his/her counterparty for every pound accepted if the horse wins. The pricing grid ranges from 1.01 (i.e. 1 pence for each pound staked), to 1000 (999 pounds for each pound staked). There is no margin trading, at least not for the ordinary punter, as all liabilities must be lodged with the exchange prior to any order being submitted.

To illustrate the format of trading on the exchange, we present two screenshots of the Betfair limit order book in Figure 1. The screenshots are taken from the win market for the 14.00 race at Chepstow on the 30th September 2014. This race, while outside of our data-set, is used for illustrative purposes. Market order traders can back a horse on the left-hand-side of the book, or lay on the right-hand-side. The best three back and lay quotes are displayed, with the volume available at each price indicated just below. In the left panel we display the screen prior to the race, with the in-running screen in the right panel. As the race progresses, certain horses become more favoured and others drift out of contention. The favourite, Parish Business, narrowed from 3.8 (i.e. 2.8 to 1) prior to the race, to 2.1 (1.1 to 1) in our in-running example. On the other hand, Significant Move, priced at 30 (29-1) beforehand, drifted out of contention as no limit order traders were offering a price for a market order trader to lay this horse during the race.

We obtained data on 9,562 U.K. horse races from Fracsoft, a third-party provider of Betfair limit order book data. This is the full database of U.K. races between the 24th March 2013, and the 19th March 2014 inclusive. This data comprises second-by-second quotes of the best back and lay quotes, the last matched price, and the total matched volume on that bet (at that second). This data includes both the win and place markets, and both pre-race (from the morning of the race) and 'in-running' trading. We concentrate on the last 20 minutes of trading prior to each race, estimating that 80.4% of pre-race trading (by volume) takes place in this window. Certain markets, trading periods, and horses are missing from the database, particularly for the place market as it goes in-running. In total, however, we have approximately 225 million observations of the Betfair limit order book spread over a year.

This data-set is then merged with data provided by Betfair on the winners (and those that place) in each race.[4] Incidentally, Betfair also provide aggregate trading volume data, but this is not categorised into back/lay -initiated market orders, and therefore we rely on our own constructed data-set of orders. We find that there is a correlation of 0.97 between Betfair's measure of total pre-race volume - for each horse in each market - and our measures based on the Fracsoft data, even though our measures are based only on the last 20 minutes of pre-race trading. This suggests to us that our forthcoming descriptions of order flow are representative of larger trading windows. Data from Fracsoft and Betfair was merged according to the selection ids of each bet. Approximately 10% of the Fracsoft data could not be matched with Betfair win data, and was therefore dropped from the analysis.

The first step is to classify the direction of orders. All trade has a buyer and a seller, and therefore aggregate levels of trade will be uninformative about fundamentals. For example, if a buy trade is predictive of fundamentals, the offsetting sell trade of the counterparty is equally unpredictive of fundamentals, and therefore the informativeness of the buy trade is cancelled out. As a result, it seems reasonable to focus on market orders, which are executed at quotes (or limit orders) that currently reside in the book. Market order traders are presented with a range of limit orders - to back or lay each horse, in each market, in each race - and can trade against those quotes which they believe are mispriced. Their accuracy in identifying mispriced quotes will determine how we, in this study, evaluate the role of speculative trade in market efficiency.

To identify the direction and volume of market orders, we use a variant of the Lee and Ready (1991) algorithm commonly used in the market microstructure literature. We observe the total matched for a given bet each second. We can therefore identify that an order occurred if the total matched at time t is greater than the total matched at time $t-1$. We also observe the last matched price each second. We therefore classify the latest order as a back order if the last matched price at time t is closer to the best quoted back odds at time $t-1$ than it is to the best quoted lay odds at time $t-1$. Conversely, if the last matched price was closer to the best lay quote then we classify this as a lay order. If the last matched price at time t is precisely in the middle of the back-lay spread, then we cannot classify this order and it is ignored.

[4] This data can be found at the following link: http://www.betfairpromo.com/betfairsp/prices/index.php.

This method is successful at classifying the majority of trades. We estimate that there are 50,548,617 orders prior to races across the year, and 10,978,350 orders during races. Prior to racing, 59.5% of orders are back orders, betting on a given horse, and 37.7% are lay orders, betting against a given horse. The remaining 2.8% fall directly in the middle of the spread, and are therefore unclassified. During racing, the proportions lean back towards lay betting. 51.9% are classified as back bets, 46.7% as lay bets, with the remaining 1.4% unclassified. The average individual back bet is 104 GBP prior to racing, and 183 GBP during racing. Lay bets are of more consistent size across trading periods, averaging 117 GBP prior to racing, and 109 GBP during racing.

Our next step is to sum up all of these orders to give a picture of the net volume on each horse. We consider each horse, in each market (win or place), and in each time period (pre-running or in-running) separately. We then sum the volume of back orders and take away the sum of lay orders. Net volume (measured in 000s of GBP) is then summarised in the top panel of Table 1. We can see that net volume is, on average, positive. For example, prior to races in the win market, the average market order volume is 8,967 GBP more for back bets than for lay bets. There is a great deal of variation, however, with one bet receiving 710,000 more in back bet volume, and one receiving approximately 637,000 more in lay bet volume. In the bottom panel, we also consider total volume (the sum of back and lay bet volume). In general, there is more trade in the win market and prior to racing. Of course, the latter can be explained by the shorter duration of in-running betting.

We are primarily interested in whether or not a given horse received a disproportionate amount of back volume. But what is a disproportionate amount? This has to be measured relative to other horses in the same race, as certain races may be subject to more back/lay bets than others. With this in mind, we calculated the median net volume within each race. This was done separately for win/place markets and pre-race/in-running trading periods. We then created an indicator variable for whether or not a given bet received more net volume than the median net volume within its race. This cut-off varies substantially across races. The average cut-off is 5,371 GBP in the win market prior to races, and 1,298 GBP during races. The lowest median net volume is -92,266 GBP prior to races, and the highest is 81,301 GBP. In other words, it is possible to be classified as a horse that received a disproportionate amount of back bets even with a strongly negative net volume, and it is also possible to be classified as a horse that received a disproportionate amount of lay bets even with a strongly positive net volume. It all depends on the net volume received by the horse's competitors in that race, in that market, and in that trading period.[5]

It is worth noting at this stage that our empirical strategy jointly tests the informativeness of back (long) and lay (short) bets. Consider horses that received relatively large net volumes in their race (i.e. more/larger back bets). We will compare their actual win probabilities to horses with similar odds that received relatively less net volume (i.e. more/larger lay bets). If there is a difference in the win probabilities of these two sets of horses - and the high net volume horses win more frequently than the equivalent low net volume horses - then this prediction of fundamentals could have occurred one of two ways. Either bettors knew the high net volume horses were better than their odds suggest, and bet heavily on them, or bettors knew that the low net volume horses were worse than their odds suggest, and bet heavily against them. Either way, it will be captured by our analysis in Section 3.

Returning to the data, it is important to note the differences in the types of horses/bets that receive more net volume. In Table 2 we summarise the back odds (including the stake) for horses in our data. The back odds displayed are the weighted averages obtained by market order traders (weighting is by back bet volume). Bets are broken down by market (win and place) and trading period (pre-race and in-running). Bets are also broken down into those that received greater than the median net volume within their race, and those that did not. We can observe that there is a clear pattern of higher net volume for favourites. To give an example, the average odds in the win market prior to races were 10.84 for high net volume horses, and 62.18 for low net volume horses. Similar results emerge for the place market and in-running trading.

This complicates any analysis because of the favourite longshot bias, observed by Smith *et al.* (2006) for this betting exchange. In this bias, favourites (longshots) win more (less) often than their odds suggest. Therefore, if we simply compared the returns of high and

[5]We also designated the cut-off as the mean net volume within each race. This tends to shift the cut-off slightly to the right, due to the positive skew of net volume. However, the results were qualitatively similar to those that we present in the paper, and therefore we did not include them here.

low net volume bets, we could confound the predictive power of trading volume with the favourite-longshot bias.

To illustrate this point, consider the plots in Figure 2. We plot the actual win probabilities of horses/bets against the implied win probabilities inferred from the back odds described in Table 2. (If the Betfair odds are X (including the stake), then the implied win probability is 1/X). Implied win probabilities are rounded to two decimal places (just for this graph), in order that we have large enough buckets to give smooth average win probabilities. Win and place bets are also pooled, to ensure a breadth of implied win probabilities, and pre-race (in-running) bets are displayed in the left (right) panel.

For pre-race trading, implied win probabilities closely correspond to actual win probabilities. A 45 degree line is added for comparison. When it comes to in-running trading, however, there is a pronounced favourite longshot bias, with favourites (longshots) winning much more (less) than their odds imply. This raises two points. Firstly it necessitates the comparison of high/low net volume bets with the same (or very similar) implied probabilities, particularly when we look at in-running betting. Secondly, it hints at the results to come in Section 3. It would appear that limit order traders are slower than market order traders to receive and process the unfolding of each race. Our reasoning is as follows. At the beginning of races, actual win probabilities begin somewhere in the centre of the distribution, then drift toward 0 for losers and 1 for winners as races unfold. The pronouncement of the in-running favourite longshot bias therefore suggests that limit order traders - who provide the quotes - underestimate the extent to which in-running favourites (longshots) are already on their way to winning (losing). If limit orders traders are slow to update quotes, back orders will be placed at undervalued bets, and lay orders placed at overvalued bets. This is the pattern we observe. The fact that the favourite longshot bias is negligible prior to races, and yet pronounced during races, suggests that insiders are more common, and market order trade more informative about fundamentals, when information is arriving.

3 Analysis

The aim of our analysis is to compare the win probabilities of horses that received greater net volume (i.e. more backers) with those that received less. If the win probabilities are higher for high net volume bets - after controlling for implied win probability - this would suggest that speculative trade disproportionately targets mispriced quotes, and is therefore conducive to market efficiency. With this in mind, we begin by estimating a linear probability model. We use this type of model, rather than a logit/probit model, in order to facilitate coefficient comparisons with our subsequent nearest neighbour matching approach.

We regress an indicator variable equalling 1 if the bet won (i.e. the horse won or placed, depending on the market) on the implied win probability of the bet (inferred from the weighted average back odds) and an indicator variable equalling 1 if the bet received greater than the median net volume within its race. In Table 1 the results of this regression are displayed for win and place markets, and pre-race and in-running periods, separately. We cluster heteroskedasticity-consistent standard errors at the race level.[6]

The first point to note is that the linear probability model fits the data much better prior to races (the differences in R^2s notwithstanding), as testified by the coefficient close to 1 for implied win probability. This is to be expected given the relationship between actual and implied win probabilities displayed in Figure 2. Prior to races, high net volume bets are not significantly more likely to win, after controlling for implied win probabilities, than low net volume bets. This applies to both the win and the place markets. During races, high net volume bets are actually less likely to win than their low net volume equivalents in the win market, but more likely in the place market. These results have to be interpreted with some caution, however, as the coefficients associated with the implied win probability variable are comfortably above 1 for the in-running samples. This would generate some implausible predictions for win probabilities for extreme favourites (i.e, with implied win probabilities close to 1).

[6] An alternative to the linear probability model with clustering at the race level is the conditional logit model. This model assigns observations to groups (in our case, races), and the likelihood is calculated relative to the group. We choose the linear probability model, however, to aid comparison with the nearest neighbour estimates to follow in this Section.

We therefore require a model that does not specify a functional relationship between implied and actual win probabilities. Using nearest neighbour techniques, we can match a high net volume bet with implied win probability X, with a low net volume bet with the same implied win probability. In other words, this helps to rule out the favourite longshot bias as a confound. Nearest neighbour matching is predominantly used in the policy evaluation literature, as a technique to non-parametrically match treated individuals to non-treated control individuals. It is a form of selection on observables. Nearest neighbours are determined by measuring the Mahalanobis distance between individuals, based on observable matching variables. In our case, we wish to match 'treated' bets (i.e. those with greater than the median net volume in its race), with 'control' bets (i.e. those with net volume less than or equal to the median in its race), and match them based on implied win probability and some race variables. We can then assess whether the high net volume bets win more often. If they do, this would suggest that trading volume is disproportionately concentrated in picking off mispriced quotes, which aids the efficiency of markets.

We begin by matching only on implied win probability. For this, and all of the following matching procedures, we allow for treatment bets to be matched with more than one control bet, and we match with replacement (i.e. a control bet can be matched with more than one treatment bet). As shown in Table 2, high net volume bets have lower odds, and higher implied win probabilities, than low net volume bets. The task therefore is to organise treatment and control groups such that any differences in the matched pairs are insignificant. The nearest neighbour matching procedure is successful at this. We conduct t-tests on the matched pairs of treatment and control bets and find that, after matching, differences in implied win probability are insignificant (even at the 10% level). This applies to both the win and place market, and the pre-race and in-running trading periods. In other words, there is a large region of common support where both high and low net volume bets, with very similar or identical implied win probabilities, can be found.

Once this matching has taken place, we can then compare the actual win probabilities of the matched bets. Analysis for the pre-race and in-running periods can be found in Tables 4a and 4b respectively. Looking at the pre-race data first, we find that high net volume bets are 0.56% more likely to win than their low net volume equivalents in the win market. This is significant only at the 10% level, suggesting that the average speculative trade plays a small role in market efficiency when information is stale. While some trades may be highly informative, and others highly distortionary, the average trade is only slightly informative. In the place market, high net volume bets are actually less likely to win, but this difference is insignificant. The picture changes dramatically when we look at the in-running trading period. High net volume bets are 3.9% more likely to win than their low net volume equivalents in the win market (with significance at the 0.1% level), and 11.7% more likely to win in the place market (with significance also at the 0.1% level). Put simply, trade is highly predictive of fundamentals when information is arriving.

In the remainder of Tables 4a and 4b, we conduct a series of robustness exercises. Firstly, we add two additional matching variables: the median net volume within the race, and the median total volume within the race. The reasoning for this is as follows. An ideal test would be to compare the win probabilities of one high and one low net volume bet - with identical implied win probabilities - from exactly the same race. Unfortunately, there are few bets with the same implied win probabilities within the same race. Our next best option is to find treatment and control bets from races with similar betting patterns. We can judge race similarity by the location of the cut-off (the median net volume within the race), i.e. how much net volume did it take to be classified as a high net volume bet? Secondly, we can identify similar races by the overall betting interest: the median total volume across horses in the race. This helps to ensure that treatment and control bets from similarly liquid markets are matched. We present the results of these exercises in columns 3-6 of Tables 4a and 4b. Our final robustness exercise is to adjust for a bias in the matching estimators when there is more than one continuous matching variable (the addition of the last two matching variables took the number of continuous variables to 3). Following Abadie and Imbens (2011), the results of this estimation are in columns 7 and 8 of Tables 4a and 4b.

We find a similar pattern across all specifications. The average trade is slightly informative prior to races - and largely only in the win market - but is very informative during races, particularly in the place market. Using our preferred matching procedure in columns 7 and 8 of Tables 4a and 4b, we find that high net volume horses are 1% more likely to win and 0.4% more likely to place than their low net volume equivalents, in betting prior to races.

In betting during races, the roles reverse with high net volume horses 3.6% more likely to win and 12.2% more likely to place than their low net volume equivalents. In other words, the average speculative trade has only a little to contribute when information is stale (and markets are perhaps already efficient), but plays a key role in the incorporation of new information into asset prices.

To illustrate and summarise our results, we present Figure 3. This is a replication of Figure 2 but instead bets are broken down into those that received high net volume in that race (i.e. above the median) and those that did not. As in Figure 2, pre-race bets are presented in the left panel, and in-running bets in the right panel. We observe little difference in the win probabilities of high and low net volume bets - for each implied win probability - prior to races, but there are pronounced differences during races. High net volume bets pay out much more frequently, than their low net volume equivalents, during races. In other words, speculative trade is much more predictive of fundamentals as information arrives. This is most visible for implied win probabilities greater than 0.5, which more often than not are found in the place market.

Why is volume during races so predictive of fundamentals? It is quite possible that sections of the betting public are at the racecourses, witness the unfolding of races earlier, process this information quickly, and place large volume bets on the likely winners and placers. Indeed such in-running course side betting behaviour has been rumoured for a while (see the article in the *Guardian* on the 29th June 2011, 'Betfair's premium charge increase welcomes winners with open palms'). Does this diminish the external validity of our results? In our opinion, no. While many public economic announcements are deliberately released to all market participants at the same time, many are not. See, for example, the storm surrounding the 2 second advantage that traders could purchase from Reuters for access to the University of Michigan Consumer Sentiment Index (*New York Times*, July 12th 2013, 'Fair Play Measured in Slivers of a Second'). Bettors and traders alike look for slight edges either in the speed with which they obtain information, or the speed with which this information is processed and traded upon. Perhaps the only puzzle is that the willingness to trade during races does not signal the possession of private information and lead to the breakdown of markets (as predicted by Milgrom and Stokey (1982)).

In the next section we will discuss our results more broadly in relation to the theoretical literature, and also draw comparisons with recent empirical work in financial markets.

4 Discussion

On average, we find that trade is only slightly predictive of fundamentals when information is stale, but closely foreshadows terminal payoffs when information is arriving. Therefore, if you take the day's trading as a whole, speculative trade is, on balance, concentrated in mispriced assets and therefore, at least in part, drives asset prices towards their fundamental values. This would seem to support the Glosten and Milgrom (1985) picture of trade, where limit order traders (replacing the market-maker in their model) trade with a population of uninformed traders - who trade randomly (thereby creating noise) - and informed insiders, who trade in the direction of fundamentals. The difference here is that the proportion of informed differs substantially across trading periods.

We should be clear that this does not rule out the presence of other types of traders. It is quite possible that there are bettors that herd, and bettors that hold erroneous beliefs about the prospect of a horse in a given race (thereby buying (shorting) over (under) valued bets). For example, suppose there is an undervalued bet prior to a race. Perhaps informed bettors purchase a large quantity of this bet, knowing that the horse is more likely to win than quoted prices suggest. Perhaps, at the same time, there is another group of market order traders - making the opposite bet and betting against this horse - that believe the horse is over-rated. The informational content of the former's trades is therefore diluted by the erroneous beliefs reflected in the latter's trades. No matter. In essence, amongst all of the types of traders and trades that exist in this market, we are simply trying to get a picture of the fundamental information content of the average trade. This average trade is of slight information prior to races (perhaps unsurprisingly, as traders have had plenty of time to digest the public information available to them at that point), but is highly informative during races.

This brings us to a discussion of recent empirical results in financial markets. Kim and Stoll (2014) wanted to establish whether trading imbalances - of the type we consider here

- are predictive of the informational content of earnings announcements. Notwithstanding that earnings announcements are a somewhat short-term and noisy measure of fundamentals, they found that there was almost no relation between trading imbalances and the content (or surprise) in announcements. This seems to tally with our pre-race results. When information is stale, we cannot hope to garner much information on fundamentals by looking at the direction of order flow. On the other hand, there is a great deal of information contained in order flow that takes place during races. Moreover, trade seems to be the process by which some new information is incorporated into prices. This corresponds with work by Love and Payne (2008), who examine foreign exchange rates around macroeconomic announcements. They estimate that up to a third of the informational content of these announcements is driven into prices via trade (rather than liquidity providers simply updating their quotes).

The second source of variation in our data is provided by the separation of win and place markets. More precise predictions are required in the former, much as an option market allows for the precise predictions of future stock price movements. For example, if an investor knows that a stock is significantly undervalued, it would make sense to leverage this information using an out-of-the-money call option, rather than simply buy the stock. While it has been noted since Black (1975) that informed investors may favour the option market due to this built-in leverage, recent work has examined relative volumes in option and equity markets to test this proposition. Roll *et al.* (2010) examine the option/equity volume ratio (O/S ratio) around earnings announcements. They find that this ratio is higher prior to announcements - suggesting relatively more informed option market activity - and link this ratio to post-announcement returns. Similar work has found that the O/S ratio is predictive of returns prior to mergers and acquisitions announcements (Augustin *et al.* (2014)), and in advance of bankruptcy filings (Ge *et al.* (2014)).

This corresponds with our results from the win and place markets. Prior to races, when information is stale, speculative trade is more predictive of fundamentals in the win market. During races, when information arrives rapidly, speculative trade is informative in both markets, but particularly in the place market. In other words, it would appear that traders need to leverage inside information, and make more precise predictions, when all public information is stale, but trade in a much more coarse fashion when information is new and yet to be fully digested.

5 Conclusion

Trading volumes are at historical peaks in financial markets. While much of this trade may be for hedging and risk-sharing purposes, much is undoubtedly speculative in nature. The level of financial market activity has brought renewed calls for a financial transactions tax to curtail speculative activity. The main defence against such an imposition is two fold: transaction taxes will 1) reduce market liquidity and 2) hinder market efficiency.

The problem with the second argument is that we have no way of observing the true efficiency of most financial markets, as fundamentals are not observed. In this paper, we therefore use horse race betting on an exchange - which shares many of the same features as financial limit order books - as a laboratory to assess the role of speculative trade in market efficiency. Crucially, the fundamental values of horse race bets are unambiguously revealed as each race ends.

Aggregate speculative trade has, of course, zero informational content. For every buyer there is a seller, and if one's trade is informed, the other's is not. We therefore divide trade up in two ways. We consider only market order trades, and not the limit order trades/quotes of their counterparties. We also divide up horses into those that received greater net volume (i.e. more positive bets) in their race, and those that did not. We can then compare the actual win probabilities (or fundamentals) for high and low net volume bets with exactly the same implied win probabilities (inferred from the trade-averaged betting prices). If the high net volume bets win more often, trade is, on average, conducive to market efficiency.

We find that the predictive capacity of speculative trade varies substantially with the trading period. Prior to races, when information is stale and markets almost already efficient (see the left panel of Figure 2), the average speculative trade contains little information. The only significant trade-based information is in the win market where - as in option markets - precise predictions of fundamentals may allow for leveraging of small nuggets of information. When the race begins, however, speculative trade is an excellent predictor of fundamentals, particularly in the place market where more coarse predictions will suffice.

In short, speculative trade, while elevated even when it contains little or no information, plays an important role in the incorporation of new information into asset prices. In our market, market order traders seem to sprint ahead of limit order traders when information is arriving, and trade very quickly on the implications of this new information. That being said, the type of fundamental information that this trade reveals would be in the hands of all traders (and not require much in the way of processing), just a few seconds later.

References

- Abadie, A., Imbens, G., W., (2011). Bias-Corrected Matching Estimators for Average Treatment Effects. *Journal of Business and Economic Statistics*, 29, 1-11.

- Ali, M., M., (1977). Probability and Utility Estimates for Racetrack Bettors. *Journal of Political Economy*, 85, 803-815.

- Angrist, J., D., Pischke, J., S., (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

- Augustin, P., Brenner, M., Subrahmanyam, M., G., (2014). Informed Options Trading Prior to M&A Announcements: Insider Trading?
  Working paper, http://ssrn.com/abstract=2441606.

- Avery, C., Zemsky, P., (1998). Multidimensional Uncertainty and Herd Behavior in Financial Markets. *American Economic Review*, 88, 724-748.

- Black, F., (1975). Fact and Fantasy in the Use of Options. *Financial Analysts Journal*, 31, 36-41, 61-72.

- Chordia, T., Roll, R., Subrahmanyam, A., (2011). Recent Trends in Trading Activity and Market Quality. *Journal of Financial Economics*, 101, 243-263.

- De Long, J., B., Shleifer, A., Summers, L., A., Waldmann, R., J., (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy*, 98, 703-738.

- Ge, L., Humphery-Jenner, M., Lin, T., C., (2014). Informed Options Trading Prior to Bankruptcy Filings. Working paper, http://ssrn.com/abstract=2302888.
- 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.
- Griffith, R., M., (1949). Odds Adjustments by American Horse-Race Bettors. *The American Journal of Psychology*, 62, 290-294.
- Kim, S., K., Stoll, H., R., (2014). Are Trading Imbalances Indicative of Private Information? *Journal of Financial Markets*, forthcoming.
- Lee, C., M., C., Ready, M., J., (1991). Inferring Trade Direction from Intraday Data.*Journal of Finance*, 46, 733-746.
- Love, R., Payne, R., (2008). Macroeconomic News, Order Flows, and Exchange Rates. *Journal of Financial and Quantitative Analysis*, 43, 467-488.
- Milgrom, P., Stokey, N., (1982). Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26, 17-27.
- 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.
- Roll, R., Schwartz, E., Subrahmanyam, A., (2010). O/S: The Relative Trading in Options and Stock. *Journal of Financial Economics*, 96, 1-17.
- Shin, H., S., (1993). Measuring the Incidence of Insider Trading in a Market for State-Contingent Claims. *The Economic Journal*, 103, 1141-1153.
- Smith, M., A., Paton, D., Vaughan Williams, L., (2006). Market Efficiency in Person-to-Person Betting. *Economica*, 73, 673-689.
- Snowberg, E., Wolfers, J., (2010). Explaining the Favorite-Long Shot Bias: Is it Risk Love or Misperceptions? *Journal of Political Economy*, 118, 723-746.

- Vaughan Williams, L., Paton, D., (1997). Why is there a Favourite-Longshot Bias in British Racetrack Betting Markets? *The Economic Journal*, 107, 150-158.

Figures and Tables

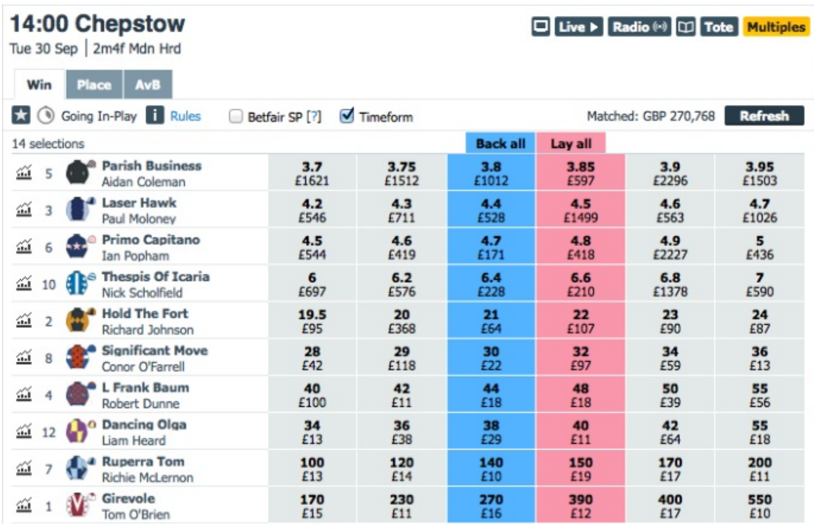

No.	Selection			Back all	Lay all		
5	Parish Business, Aidan Coleman	3.7 £1621	3.75 £1512	3.8 £1012	3.85 £597	3.9 £2296	3.95 £1503
3	Laser Hawk, Paul Moloney	4.2 £546	4.3 £711	4.4 £528	4.5 £1499	4.6 £563	4.7 £1026
6	Primo Capitano, Ian Popham	4.5 £544	4.6 £419	4.7 £171	4.8 £418	4.9 £2227	5 £436
10	Thespis Of Icaria, Nick Scholfield	6 £697	6.2 £576	6.4 £228	6.6 £210	6.8 £1378	7 £590
2	Hold The Fort, Richard Johnson	19.5 £95	20 £368	21 £64	22 £107	23 £90	24 £87
8	Significant Move, Conor O'Farrell	28 £42	29 £118	30 £22	32 £97	34 £59	36 £13
4	L Frank Baum, Robert Dunne	40 £100	42 £11	44 £18	48 £18	50 £39	55 £56
12	Dancing Olga, Liam Heard	34 £13	36 £38	38 £29	40 £11	42 £64	55 £18
7	Ruperra Tom, Richie McLernon	100 £13	120 £14	140 £10	150 £19	170 £17	200 £11
1	Girevole, Tom O'Brien	170 £15	230 £11	270 £16	390 £12	400 £17	550 £10

(a) Pre-race

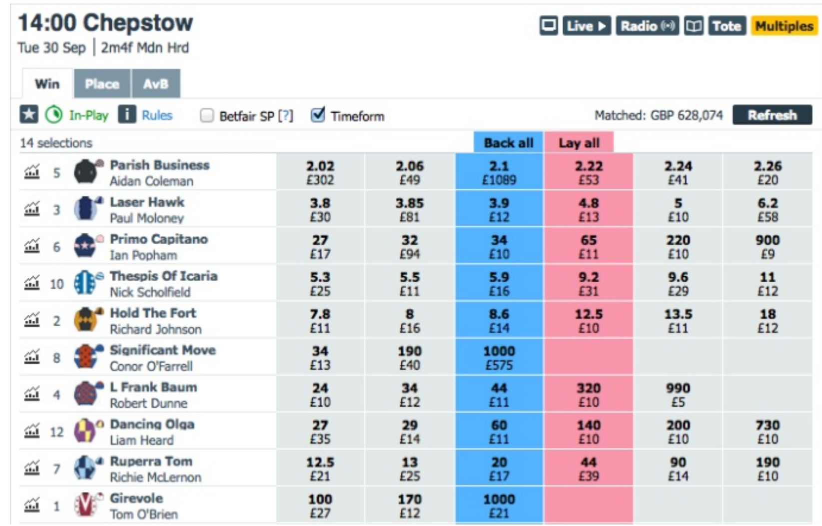

No.	Selection			Back all	Lay all		
5	Parish Business, Aidan Coleman	2.02 £302	2.06 £49	2.1 £1089	2.22 £53	2.24 £41	2.26 £20
3	Laser Hawk, Paul Moloney	3.8 £30	3.85 £81	3.9 £12	4.8 £13	5 £10	6.2 £58
6	Primo Capitano, Ian Popham	27 £17	32 £94	34 £10	65 £11	220 £10	900 £9
10	Thespis Of Icaria, Nick Scholfield	5.3 £25	5.5 £11	5.9 £16	9.2 £31	9.6 £29	11 £12
2	Hold The Fort, Richard Johnson	7.8 £11	8 £16	8.6 £14	12.5 £10	13.5 £11	18 £12
8	Significant Move, Conor O'Farrell	34 £13	190 £40	1000 £575			
4	L Frank Baum, Robert Dunne	24 £10	34 £12	44 £11	320 £10	990 £5	
12	Dancing Olga, Liam Heard	27 £35	29 £14	60 £11	140 £10	200 £10	730 £10
7	Ruperra Tom, Richie McLernon	12.5 £21	13 £25	20 £17	44 £39	90 £14	190 £10
1	Girevole, Tom O'Brien	100 £27	170 £12	1000 £21			

(b) In-Running

Figure 1: Examples of the Betfair limit order book, both prior to a race (left panel) and during a race (right panel). These screenshots are taken from the win market on the 14.00 race at Chepstow on the 30th September 2014. Punters can back (lay) horses by clicking on quotes on the left-hand-side (right-hand-side) of the book. The best three quotes for back and lay bets, with associated volumes underneath, are displayed. Bettors can also access the book algorithmically through the Betfair Application Programming Interface.

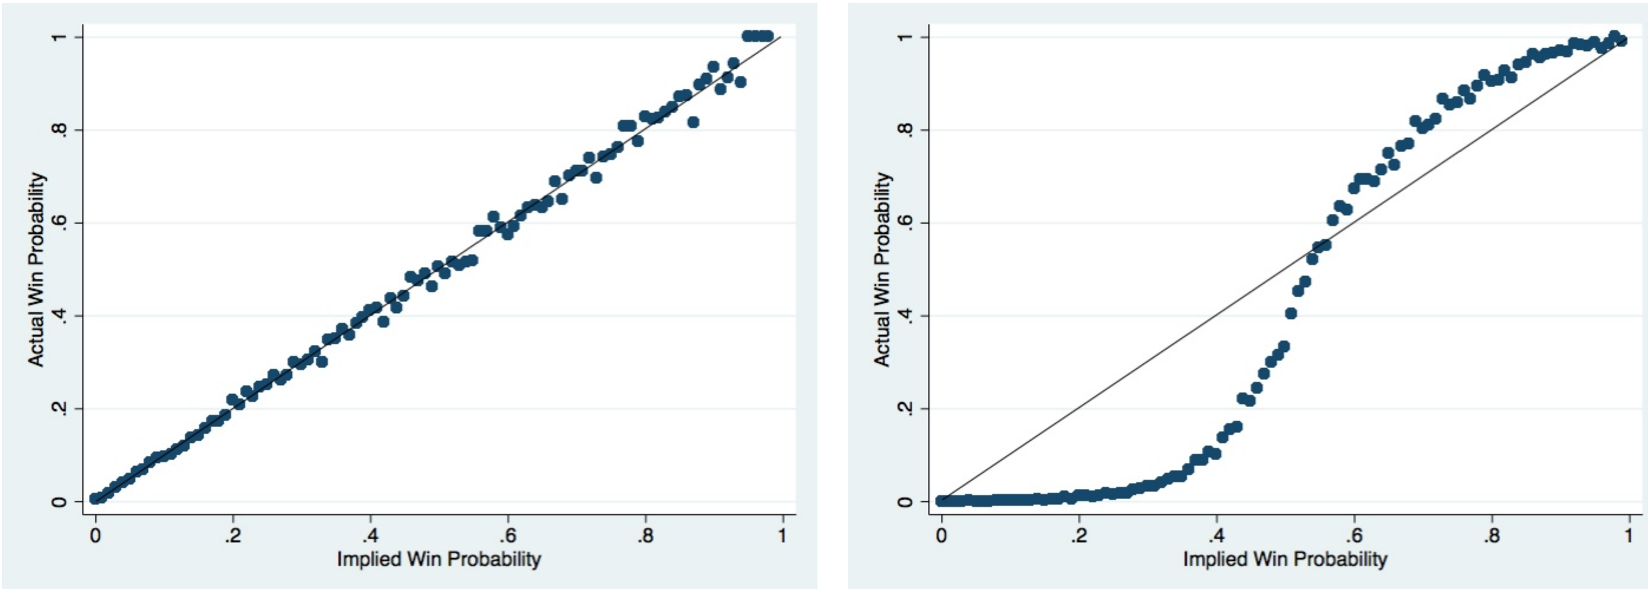

(a) Pre-race

(b) In-Running

Figure 2: Plots of implied win probability - inferred from the trade-weighted average back odds for each horse - against actual win probability. Win and place market bets are pooled for this illustration. The implied win probability is rounded to 2 decimal places to create large enough buckets for comparison with actual win probabilities. Pre-race betting is in the left panel, and in-running betting is in the right panel. A pronounced favourite-longshot bias - where favourites (longshots) win more (less) often than their odds suggest - can be observed in the latter.

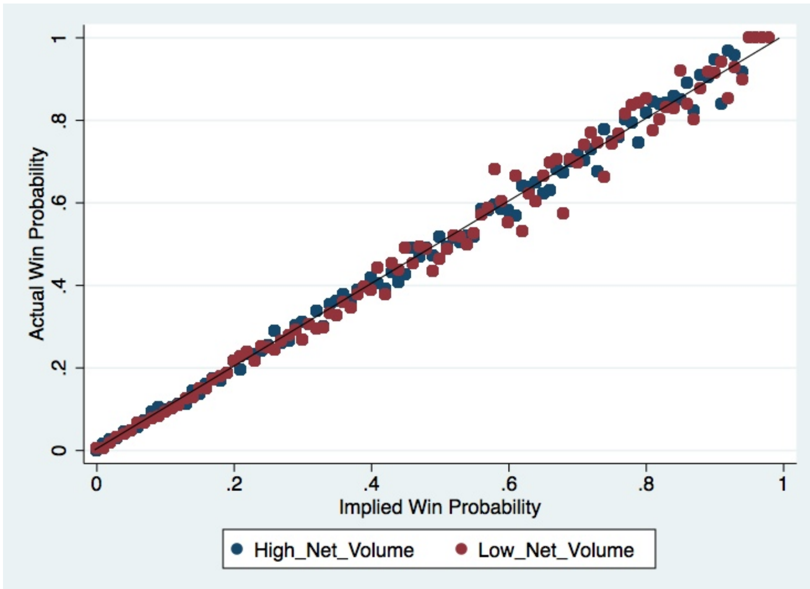

(a) Pre-race

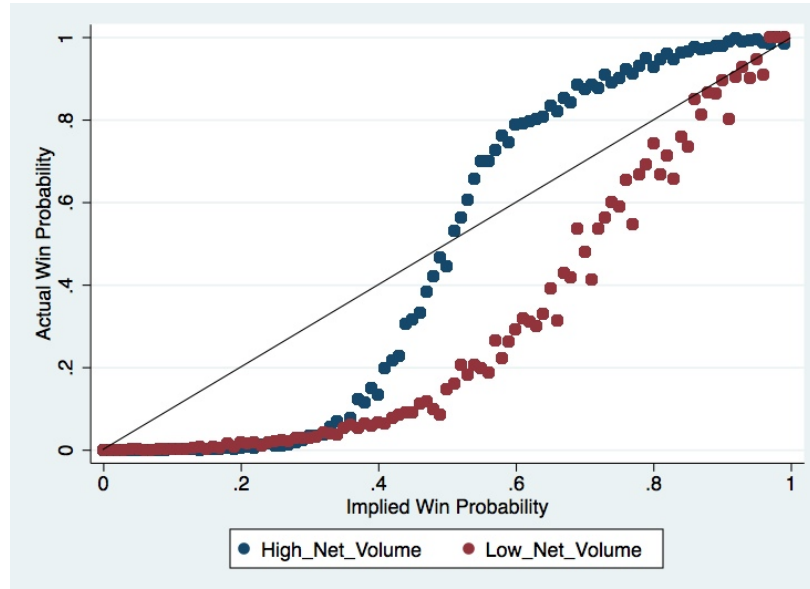

(b) In-Running

Figure 3: Plots of implied win probability - inferred from the trade-weighted average back odds for each horse - against actual win probability. For this figure we break down bets into those that received high net volume (defined as above the median net volume in that race) and those with low net volume. Win and place market bets are pooled for this illustration. The implied win probability is rounded to 2 decimal places to create large enough buckets for comparison with actual win probabilities. Pre-race betting is in the left panel, and in-running betting is in the right panel. While there is little difference between the win probabilities of high and low net volume bets prior to races, high net volume bets pay out much more frequently (i.e. are much more predictive of fundamentals) during races.

Table 1: Trading Volume					
Net Volume					
Pre-Race	Obs	Mean	Std. Dev.	Min	Max
Win Market	85670	8.966689	16.62787	-636.6945	710.0387
Place Market	84143	1.379245	2.547442	-46.26987	71.67953
In-Running	Obs	Mean	Std. Dev.	Min	Max
Win Market	84342	5.320234	13.39392	-116.2127	296.5063
Place Market	71267	.5224766	1.181332	-14.59328	41.55827
Total Volume					
Pre-Race	Obs	Mean	Std. Dev.	Min	Max
Win Market	85670	53.6023	94.91045	.0034	4154.702
Place Market	84143	9.052303	11.27837	.00246	242.636
In-Running	Obs	Mean	Std. Dev.	Min	Max
Win Market	84342	17.77687	34.99981	.002	1342.291
Place Market	71267	1.421691	1.897861	.002	52.37412

Summary statistics on trading volume (measured in 000s of GBP) for each horse. In the top panel, net volume - defined as the total back volume on each horse minus the total lay volume - is displayed. In the bottom panel, total betting volume (back volume plus lay volume) is displayed. For both net and total volume, bets are broken down by win/place market and pre-race/in-running trading periods. Back and lay bets were classified by a variation of the Lee and Ready (1991) algorithm described in Section 2.

Table 2: Back Odds					
Pre-Race	Obs	Mean	Std. Dev.	Min	Max
Win Market					
$\leq$ Median Net Volume within Race	44390	62.1786	112.3596	1.044498	990
$>$ Median Net Volume within Race	41280	10.84439	22.76577	1.025357	729.5147
Place Market					
$\leq$ Median Net Volume within Race	44411	10.3988	19.0894	1.017669	617.8196
$>$ Median Net Volume within Race	39732	3.655617	5.997309	1.029557	394.7709
In-Running	Obs	Mean	Std. Dev.	Min	Max
Win Market					
$\leq$ Median Net Volume within Race	43862	48.5026	85.3905	1.026501	990
$>$ Median Net Volume within Race	40480	12.34034	29.72359	1.01	826.9736
Place Market					
$\leq$ Median Net Volume within Race	33252	4.821749	9.08082	1.01	430
$>$ Median Net Volume within Race	38015	2.275229	4.276322	1.01	500

Summary statistics on horse odds, calculated as the weighted average price executed by back market orders (odds include the stake on Betfair, hence all are above 1). Horses are divided into those that received greater than the median amount of net volume in their race, and those that did not. The statistics are also divided up by win/place market and pre-race/in-running trading periods. Across all of these sub-samples, more heavily favoured horses - with lower odds - tend to attract greater net betting volume.

Table 3: Regression Analysis				
Dependent Variable: Win Indicator	1	2	3	4
Intercept	-0.0017349	-0.004045	-0.101814***	-0.3531056***
	(.0012144)	(.002067)	(.0006086)	(.0019605)
Implied Win Probability (from Odds)	0.9914628***	0.9995899***	1.477206***	1.307194***
	(.0132729)	(.0071592)	(.0040484)	(.0047325)
> Median Net Volume within Race (Indicator)	0.0030796	0.0031563	-0.0524268***	0.1287579***
	(.0025142)	(.0033667)	(.0013049)	(.0030434)
Market	Win	Place	Win	Place
Trading Period	Pre-Race	Pre-Race	In-Running	In-Running
R^2	0.1292	0.187	0.6688	0.5824
No. of Clusters (Race-Market Combinations)	9560	9534	9562	9537
No. of Observations	83942	84110	83170	71235

Regression analysis. An indicator, equalling 1 if the bet won (i.e. the horse won or placed, depending on the market), was regressed on the implied win probability of the bet (inferred from the weighted average back odds), and an indicator variable equalling 1 if the horse/bet received greater than the median amount of net volume within the race. The win (place) market is considered in regressions 1 and 3 (2 and 4), and pre-race (in-running) trading periods are considered in regressions 1 and 2 (3 and 4). Standard errors - heteroskedasticity-consistent and clustered at the race-market level - are in parentheses, and * indicates significance at the 0.1% level.**

Table 4a: Nearest Neighbour (Pre-Race)				
Dependent Variable: Win Indicator	1	2	3	4
> Median Net Volume within Race (Indicator)	0.0055872.	-0.0007431	0.009578*	0.0013316
	(.0032232)	(.0037954)	(.0041251)	(.0037543)
Market	Win	Place	Win	Place
Matching on Implied Win Probability (from Odds)	Yes	Yes	Yes	Yes
Matching on Median Net Volume within Race	No	No	Yes	Yes
Matching on Median Total Volume within Race	No	No	No	No
Bias Adjustment for Continuous Matching Variables	No	No	No	No
No. of Observations	83942	84110	83942	84110
Dependent Variable: Win Indicator	5	6	7	8
> Median Net Volume within Race (Indicator)	0.0139501*	0.0066223.	0.0099924.	0.0041198
	(.0056231)	(.003999)	(.0056229)	(.0039986)
Market	Win	Place	Win	Place
Matching on Implied Win Probability (from Odds)	Yes	Yes	Yes	Yes
Matching on Median Net Volume within Race	Yes	Yes	Yes	Yes
Matching on Median Total Volume within Race	Yes	Yes	Yes	Yes
Bias Adjustment for Continuous Matching Variables	No	No	Yes	Yes
No. of Observations	83942	84110	83942	84110

Nearest neighbour analysis, for the pre-race trading period. The average win indicator is compared for bets where the horse/bet received greater than the median amount of net volume within the race (i.e. the 'treated'), with those that did not (the 'control'). Treatment and control bets are matched, at various stages, on implied win probability, median net volume within their race, and median total volume within their race. In rows 7 and 8 we also adjust for biases that arise in estimates when there is more than 1 continuous matching variable (see Abadie and Imbens (2011)). Standard errors - heteroskedasticity-consistent - are in parentheses, and * and . indicates significance at the 5% and 10% level.

Table 4b: Nearest Neighbour (In-Running)				
Dependent Variable: Win Indicator	1	2	3	4
> Median Net Volume within Race (Indicator)	0.0386197***	0.1170756***	0.035181***	0.1206921***
	(.0035487)	(.0035718)	(.0036063)	(.0035926)
Market	Win	Place	Win	Place
Matching on Implied Win Probability (from Odds)	Yes	Yes	Yes	Yes
Matching on Median Net Volume within Race	No	No	Yes	Yes
Matching on Median Total Volume within Race	No	No	No	No
Bias Adjustment for Continuous Matching Variables	No	No	No	No
No. of Observations	83170	71235	83170	71235
Dependent Variable: Win Indicator	5	6	7	8
> Median Net Volume within Race (Indicator)	0.0392329***	0.1264547***	0.0356352***	0.1223487***
	(.0040659)	(.0038672)	(.0040635)	(.0038666)
Market	Win	Place	Win	Place
Matching on Implied Win Probability (from Odds)	Yes	Yes	Yes	Yes
Matching on Median Net Volume within Race	Yes	Yes	Yes	Yes
Matching on Median Total Volume within Race	Yes	Yes	Yes	Yes
Bias Adjustment for Continuous Matching Variables	No	No	Yes	Yes
No. of Observations	83170	71235	83170	71235

Nearest neighbour analysis, for the in-running trading period. The average win indicator is compared for bets where the horse/bet received greater than the median amount of net volume within the race (i.e. the 'treated'), with those that did not (the 'control'). Treatment and control bets are matched, at various stages, on implied win probability, median net volume within their race, and median total volume within their race. In rows 7 and 8 we also adjust for biases that arise in estimates when there is more than 1 continuous matching variable (see Abadie and Imbens (2011)). Standard errors - heteroskedasticity-consistent - are in parentheses, and * indicates significance at the 0.1% level.**