Distributions of certain market observables in an on-line betting exchange

Stephen J. Hardiman[1], Steven T. Tobin[1], Peter Richmond[2] and Stefan Hutzler[1]

[1]School of Physics, Trinity College Dublin, Ireland
[2] Complex Adaptive Systems Laboratory,
University College Dublin, Belfield, Dublin 4, Ireland
hardimas@tcd.ie

January 24, 2011

Abstract

We analyze high frequency time series of odds and trading volume as extracted from markets on an on-line betting exchange, Betfair.com. As in financial markets, the probability distributions for the change in the market price (odds) are seen to exhibit fat tails. We find significant differences in the statistics of odds changes which occur during the in-play activity of sporting events, when we argue that news is driving market dynamics, and the statistics of odds changes when game-play activity has ceased. Furthermore, we investigate the distribution of the market trading volume as sampled at high frequency and find it can be fit very well to a log-normal distribution function.[1]

1 Introduction

Traditionally, if one wanted to make a bet on a certain outcome (such as the outcome of a football game), the bet would be placed with a bookmaker. The bookmaker sets the odds deemed appropriate for the expected probability of the outcome and the customer can place bets against those odds. Bookmakers can expect to make profits in the long run by providing odds to their customers that slightly overestimate the true probabilities of that given outcome.

A betting exchange operates in a manner more akin to a stock exchange. Here, there is no single agent acting as a bookmaker, but agents can play both

[1]Dynamics of Socio-Economic Systems, Vol.2 Number 1: 121-137. (2010). ISSN 1852-379x. Received: 16/12/09, Accepted: 30/08/10. http://www.dyses.org.ar/IJ-DySES

the roles of the customer and the bookmaker by choosing to either place a bet on an outcome ("*back*") or to offer a bet ("*lay*") that can be backed by another agent. An agent who bets $1 on an outcome (a "back' bet') at 3:1 odds can expect to get $4 back if they win ($3 profit), but lose their $1 stake in the event of a loss. The converse of backing is to lay a bet. In this case, the agent plays the role of the bookmaker by taking up the bet of another agent. In the same example, the layer must pay out $4 dollars to the successful bettor in the event of a winning outcome, but retains the $1 stake in the event of a loss. Laying is thus equivalent to betting against an outcome. The match-making is managed with a double auction order book as with a regular financial market, except the *bid* and *ask* columns are replaced with *back* and *lay*. It is important to note, however, that unlike a stock market, a betting market has a definite conclusion. The odds will inevitably move towards infinity or zero as the outcome becomes a certainty.

The data used in our study is extracted from the on-line betting website, Betfair.com[2], based in London, England. Since Betfair was launched in June 2000, it has become the largest on-line betting company in the UK and the largest betting exchange in the world. The website claims over 2 million clients from all over the world.

Betfair acts as a kind of prediction market [1, 2]. Prediction markets are speculative markets in which assets traded have a value that is tied to a particular event (e.g., will the next US president be a Republican). Previous studies into prediction markets have shown that the market price of a contract can provide a good reflection of the true odds of that event taking place, or at least the mean market belief [1, 3, 4, 5, 6, 7]. One of the oldest and most famous prediction markets, the Iowa Political Stock Market, has been shown to beat opinion polls at forecasting the outcomes of presidential elections [7], however differing viewpoints [8] do exist in the literature. One recent paper has studied the returns of the Iowa Presidential Stock Market with an aim to compare and contrast the dynamics of prediction market price changes with financial returns [9]. Here, we perform a similar analysis in the case of betting market returns.

Our recent study of the dynamics of Betfair betting markets [10] for Champions League football matches showed that the volatility of odds fluctuations exhibited long-range dependence, a phenomenon frequently observed for the volatility of financial price returns. Furthermore, by studying the market during half-time, we were able to identify statistical differences in market dynamics between times when the game-play of the football match was underway such that news of match events was reaching the market and times when news was suspended. In this paper, we complement these findings with an analysis of the probability distributions of betting market odds returns and we extend our analysis to markets for Wimbledon tennis matches.

[2] http://www.betfair.com/

2 Data analysed

The data set used in this analysis comprises time resolved snapshots of the order-book for the activity of Betfair markets for matches that took place during two different sporting competitions. We have analysed 146 soccer matches from the 2008 Champions League football tournament and 30 tennis matches from Wimbledon 2008. Also included with each order-book snapshot, is a value for the "odds" at which the last bet was matched during that second and the value of bets matched to date for that outcome. We have chosen to look at these particular data sets, as they are very popular, televised events which attract a high volume of trading activity on the Betfair website

It is important to note that the odds with which Betfair users choose to back or lay outcomes can only assume values from a specified set of numbers between 1.01 and 1000 imposed by the Betfair user interface. If a bet is matched at a Betfair 'odds' of 3.0, the backer will triple his/her money in the event of a payout, such a payout would be more traditionally represented in the gambling parlance as "2-to-1".

The market remains open for Betfair users to trade bets up until the very end of a match, when the conclusion is known and the market odds naturally move towards either 1.01 or 1000. A typical football data set is composed of about 6500 1-second records representing the state of the betting market at each second during the lifetime of the match. This typically includes ninety minutes of play time, approximately three minutes of injury time and fifteen minutes of half-time.

Since, unlike in football, the length of a tennis match is not fixed, a typical tennis data set varies from 5000-15000 1-second records. We rescale the time in our football and tennis matches to a dimensionless quantity between zero and one, by dividing by the total match time.

The data was collected using front-end software which interfaces with Betfair provided by Fracsoft[3].

3 Implied Probability Returns

We define the *implied probability* as the reciprocal of the current Betfair market "odds". Much like the price of a contract in a binary option market [9], the implied probability is bounded between zero and one and reflects the probability of that particular outcome. Following [9], we consider the change in the logarithm of this value with respect to time to be a stochastic variable analogous to financial price returns and call it the *implied probability log-return*, r. In Fig.(1), we have plotted the evolution of the implied probability for an example football match which took place on the 9th of April during the 2008 Champion's League football tournament. Manchester United, the home team, won the match with a goal by Carlos Tevez in the 70th minute. This is reflected by a

[3] http://www.fracsoft.com/

large change in the implied probability at roughly t = 1400 seconds remaining. The implied probability for the football teams Man United and Roma winning (and the draw) are seen to move towards approximately 1.0 and 0.0 respectively as the match ends and the result becomes known.

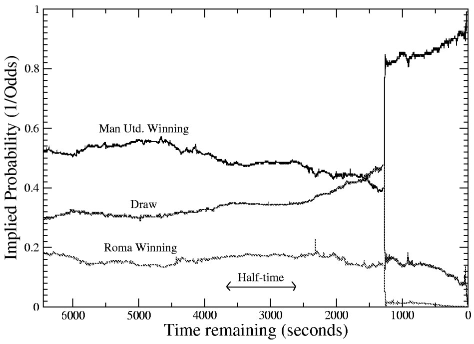

Figure 1: The evolution of the implied probability (reciprocal of match odds) for the three possible outcomes of a Champions League football match. As the match draws towards its conclusion, the implied probabilities for the outcomes tend towards 0 and 1 as the market odds tend towards 1000 and 1.01. We see a large probability change for each outcome at the $t \approx 1400s$ mark corresponding to a goal scored by Man Utd.

To compare the returns of markets with different volatilities, we normalise the returns, $r(t)$, in each market, M, by the standard deviation of the returns for that market.

$$r^*(t) = \frac{r(t)}{\sigma_M} \qquad \sigma_M^2 = \langle r(t)^2 \rangle_{r(t)\in M} - \langle r(t) \rangle^2_{r(t)\in M} \tag{1}$$

Such a normalisation procedure is frequently applied in financial market literature [11, 12]. The scatter plot of normalised market returns as a function of time as drawn from all football and tennis matches studied is shown in Fig.(2). As was seen in [9], the returns are non-stationary with a volatility that increases towards the conclusion of the market. This is visible in Fig.(3), in which we display the average volatility of normalised returns, r^*, as a function of match time remaining for the two ensembles. Also clearly visible in Fig.(2) and Fig.(3) is the half-time break in the case of Champions League data.

As in [9] we deal with this non-stationarity by employing a further normalising procedure. We divide the normalised return at time t, $r^*(t)$, by the local volatility average, $\sigma^*(t)$. We call this the *detrended return*, $\hat{r}$.

$$\hat{r}(t) = \frac{r^*(t)}{\sigma^*(t)} \tag{2}$$

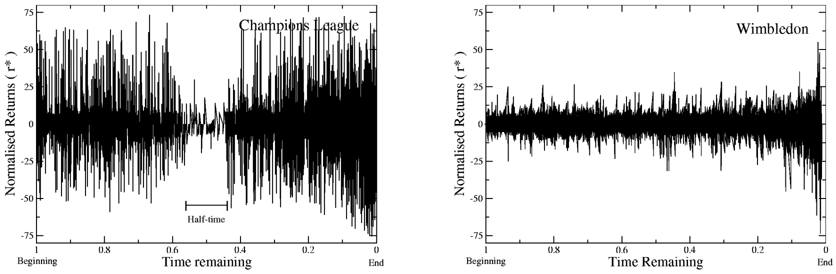

Figure 2: The ensemble of normalised implied probability log-returns, r^*, for the Champions League (left) and Wimbledon (right) markets. We see that the returns typically grow larger as a match reaches its conclusion. We see a very different behaviour in the football returns for the times between 0.45 and 0.55. This corresponds approximately to half-time.

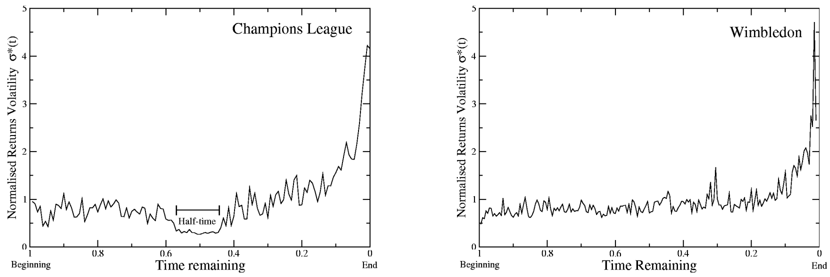

Figure 3: The volatility of the normalised implied probability returns as a function of time remaining till settlement for Champions League (left) and Wimbledon (right) markets. In these graphs, the volatility is measured as the standard deviation of normalised returns, $\sigma^*(t) = \sqrt{< r^{*2}(t) > - < r^*(t) >^2}$, as calculated over 500-second non-overlapping time windows. In both tennis and football, the volatility of returns is seen to steadily increase as the match reaches its conclusion.

$\sigma^*(t)$ is calculated by partitioning the time remaining till the end of the match into time windows, T.

$$\sigma^{*2}(t) = \langle r^*(t)^2 \rangle_{t \in T} - \langle r^*(t) \rangle_{t \in T}^2 \quad (3)$$

We use 65 non-overlapping time windows between rescaled time $t = 0$ and $t = 1$ for Champions League returns, and 50 non-overlapping time windows between rescaled time $t = 0$ and $t = 1$ for Wimbledon returns.

We now investigate the stationarity of the detrended returns for different time windows during the course of the matches using the Kolmogorov-Smirnov test on the tails of the distributions of returns drawn from each time window. The Kolmogorov-Smirnov statistic is a measure which quantifies the difference between two empirical probability distributions (see Appendix). If two samples are assumed to be drawn from the same distribution, the P-value is the probability with which we would expect to see a larger value of the KS-statistic than that which was observed. Low P-Values indicate a poor agreement.

The P-values for the KS-tests for Wimbledon and Champions League data are displayed graphically in Fig.(4). The length of the time windows were chosen as a compromise between maximising the temporal resolution of Fig.(4) and maximising the number of samples, ν, used in estimating the tails of the probability distributions of detrended returns drawn from each time window. ν is typically 300 for the Champions League data, and 100 for the Wimbledon data.

For the Champions League returns, we can identify two types of behaviour corresponding to in-play trading activity and trading activity that occurs during half-time. We also find that the distribution of detrended returns drawn from the end of the football matches ($t < 0.06$), differs significantly from that of returns drawn from other times in the matches, so we choose to neglect these returns in the analysis which follows. In Fig.(5), we have plotted the distribution of half-time returns along side in-play returns. We see a markedly different functional form for the distributions. In particular the distribution of returns which take place while the match is in-play exhibit fatter tails.

The difference in character between market returns that occur while the play of a football match is underway, and those that occur during half-time when no news is hitting the market has also been identified in [10]. We suggest that the distribution of returns drawn from half-time is representative of endogenous market behaviour, whereas the distribution of returns as sampled during the in-play activity is instead reflective of the changing probability of the outcome of the match which itself is dependent on the game-play and scoring of football. This difference cannot be ascribed to a lack of trading volume, since as seen in Section 4 (Fig. 8), trading volume remains significant during half-time.

If our hypothesis is true that in-play returns are indeed representative of the game-play and scoring of the sport on which the market is based, then we may expect two different sports to have a different probability distribution for in-play returns. The scoring in football is poorly resolved, matches will often end nil-nil or with a low number of goals scored. In tennis however, a match may be 5-sets long, each of which potentially contain 13 games, each of which

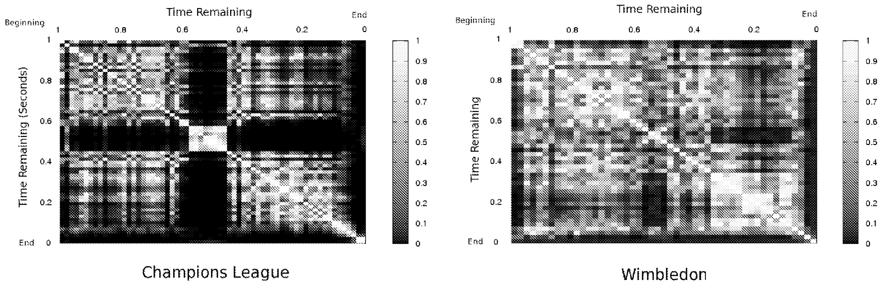

Figure 4: The P-Values for the Kolmogorov test on the tails of the distributions of detrended returns drawn from different time windows. The time windows for the Champions League data represent lengths of 100 seconds. The length of the time window for a Wimbledon tennis match is typically approximately 200 seconds. A black square ($P = 0$) in the matrix indicates a poor match between the distribution of returns drawn from the corresponding time windows. A white square indicates a very good agreement ($P = 1$). The black bands which criscross the matrix indicate that the market returns during half-time in a Champions League football match have a different distribution to those drawn from other times during the match.

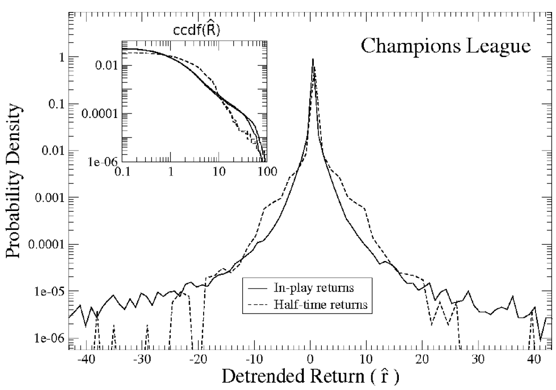

Figure 5: The distribution of detrended returns, $P(\hat{r})$, as drawn from the in-play activity of all champions league market compared with the distribution of returns which occur during half-time. The complementary cumulative distribution function, $ccdf(\hat{R}) = \int_{\hat{R}}^{\infty} P(\pm\hat{r})d\hat{r}$, for the positive and negative tails of both data sets are shown in the inset. We find the in-play activity exhibit fatter tails. We suggest that the returns distribution which results from returns during half-time are representative of endogenous market behaviour, whereas the returns we see while the match is underway reflect the game-play and mechanism of scoring in football.

may involve any number of points scored. Thus, we expect a goal scored in football to make a much more dramatic impact on the probability of a given side winning than the impact that a single point scored may have in tennis. We should then conclude that the distribution of changes in probability of a side winning in a football match should exhibit fatter tails than that for tennis. This is what is observed in Fig.(6).

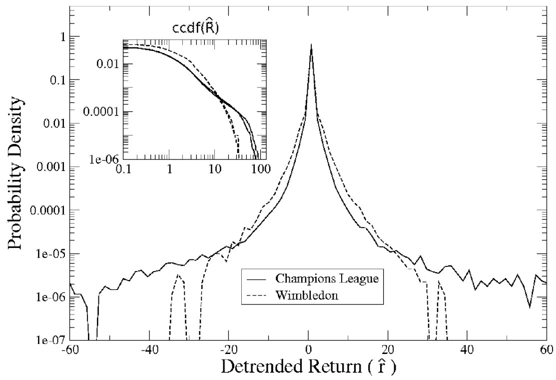

Figure 6: The distributions of the (in-play) detrended implied probability return for Wimbledon and Champions League markets. The complementary cumulative distribution functions, $ccdf(R) = \int_R^\infty P(\pm\hat{r})d\hat{r}$, are shown in the insets. We pool together the normalised returns from in-play (excluding half-time) to generate the histograms for Champions League returns. We see a different functional form for the returns distributions for the two sports. The Champions League distribution appears to have fatter tails, indicating that large changes in the probability of outcomes may be a more common occurrence in football than in tennis.

This reasoning is also supported by the observed difference between the distributions of returns taken from the end of a football match and those taken from the rest of the match. Due to the relatively low frequency with which goals are scored in football, it is likely that by the final minutes of the match (when a score-line can be often 2-0) the outcome becomes an almost certainty, and the market odds will be typically already at 1000.0 or 1.01 with little probability of changing. Furthermore, in the case of teams drawing, in the final minutes of the match, the implied probabilities can take dramatic changes as depicted in Fig.(1).

We now turn our attention to the returns, r_H, drawn from the region identified as half-time. We do not normalise and detrend these returns as defined by Eq.(4) and Eq.(2) but only normalise by dividing each return, r_H, by the standard deviation of half-time returns from the market, M, that it is drawn from.

$$r_H{}^*(t) = \frac{r_H(t)}{\sigma_H(M)} \qquad \sigma_H(M) = \langle r_H(t)^2 \rangle_{r_H(t)\in M} - \langle r_H(t) \rangle^2{}_{r_H(t)\in M} \quad (4)$$

The complementary cumulative distribution function, $ccdf(R_H^*) = \int_{R_H^*}^{\infty} P(|r_H^*|)d|r_H^*|$ for the magnitude of normalised half-time returns, $|r_H^*|$, from all Champions League football markets is shown in Fig.(7).

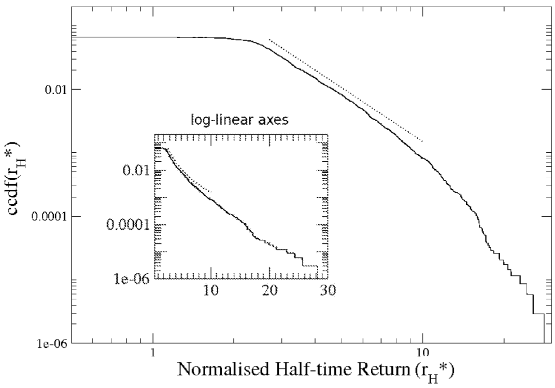

Figure 7: The tail of the complementary cumulative distribution function for the absolute value of normalised half-time returns r_H^*. The shape of the tail on the Log-Linear inset suggest the tail decays slower than exponential. The dotted line is the slope of the power law fit, $P(|r_H^*| > x) \sim x^{-\alpha}$, for $\alpha = 2.82$. We conclude that the power-law is not an adequate fit to the tail.

To investigate power-law scaling in Betfair returns, as observed in the case of returns from the binary option prediction market of [9], we attempt to fit a Pareto law, $P(|r_H^*| > x) \sim x^{-\alpha}$ to the tail of the distribution function for half-time returns.

Using maximum likelihood estimation, we estimate an exponent $\alpha = 2.82 \pm 0.022$ for normalised returns with $|r_H^*| > 2.70$. Despite a reasonable visual fit to the data, for $2.7 < r_H^* < 10$, the tail of the distribution decays more exponentially for $r_H^* > 10$ and using the more rigorous statistical analysis provided by the methods of [13] we must reject the null hypothesis of a power-law tail. As suggested in [9], this departure from power-law form may be a consequence of partisan trading, as follows.

It is unlikely that all traders on Betfair play the market objectively by backing and laying outcomes equally, based only on whether they believe the market odds of the outcome are undervalued or overvalued. There are likely to be biased traders who truly believe their favourite team is going to win. Since partisan traders are unable to or simply choose not to accomodate new information as it hits the market, they have a constant willingness to bet on their favourite team (or conversly lay the opposing team) which may prevent returns with a large magnitude from occurring. In [9], the authors propose the hypothesis that partisan trading leads to more exponentially decaying returns. A similar phenomenon may be occurring in the case of the Betfair implied probability returns.

4 Trading Volumes

In financial markets it is often found that the tails of the probability distribution function for share trading volume (number of shares traded per unit time on a particular asset) exhibits a power-law form [14]. In this section we aim to test that hypothesis on time series of trading volumes extracted from Betfair markets.

Since certain matches or outcomes are more popular and draw larger trading volume than others, following [11], we normalise the trading volumes, V, in each market, M, by their medians, $\tilde{V}_M$.

$$V^* = \frac{V}{\tilde{V}_M} \tag{5}$$

where the median is calculated over all 1-second trading volumes, V, recorded for a given market M. As in the previous section we investigate the stationarity of the trading volume observable in Fig.(8) by plotting the average trading volume as a function of time remaining till the conclusion of the market. We factor out the trend evident in Fig.(8), by dividing the normalised trading volume V^* by the local median, $\tilde{V^*}(t)$

$$\hat{V}(t) = \frac{V^*(t)}{\tilde{V^*}(t)} \tag{6}$$

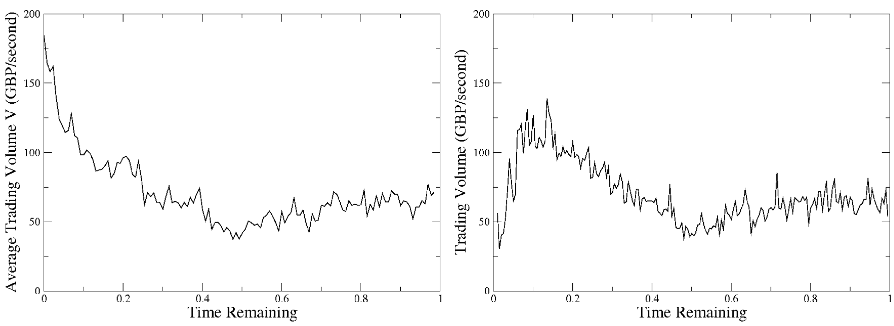

Figure 8: The average 1-second trading volume, V, as a function of time remaining till settlement for Champions League (left) and Wimbledon (right) markets. The average is calculated over all markets for 500-second non-overlapping time windows.

where $\tilde{V^*}(t)$ is calculated for 65 non-overlapping time windows between rescaled time $t = 0$ and $t = 1$ for Champions League data and 50 non-overlapping time windows between rescaled time $t = 0$ and $t = 1$ for Wimbledon data. We then pool together all the detrended trading volumes $\hat{V}$ from each market to create one large sample from which to calculate a normalised trading volume probability distribution, $P(\hat{V})$ for each sport studied.

We find the bulk of the resulting trading volume distributions, $P(\hat{V})$, for both sports can be described very well by a log-normal distribution.

$$P(\hat{V}) = \frac{1}{\hat{V}\sigma\sqrt{2\pi}} e^{-\frac{(ln\hat{V}-\mu)^2}{2\sigma^2}} \tag{7}$$

To perform the fits, we calculate the mean,

$$\mu = \langle U \rangle \tag{8}$$

and standard deviation,

$$\sigma = \langle U^2 \rangle - \langle U \rangle^2 \tag{9}$$

of the logarithms of the detrended 1-second trading volumes, $U = ln\hat{V}$. Since we expect the distribution of logarithms of log-normally distributed variables to be normal, we overlay the functional form of a Gaussian with mean, μ, and variance, σ

$$G(U, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(U-\mu)^2}{2\sigma^2}} \tag{10}$$

on the empirical distribution for the logarithms of the 1-second trading volumes, U. As seen in Fig.(9), we find very good agreement. The Kolmogorov-Smirnov statistic for the comparison between the empirical trading volumes and the fitted log-normal distributions is 0.022 in the case of Champions League trading volumes and 0.017 in the case of Wimbledon trading volumes. These values are too large to accept the null hypothesis of the Kolmogorov-Smirnov test, but it is important to note, that with a very large number of samples, as in the case of our Betfair trading volumes (The Champions League normalised trading volumes distribution was drawn from 587,993 points, and the Wimbledon distribution as drawn from 168,598 points), the Kolmogorov-Smirnov test will reject the null hypothesis with only minor deviations from the model distribution. A log-normal functional form is consistent with some experimental distributions of daily stock market trading volumes observed in [15, 16]

However, in Fig.(9), we see that the log-normal function which fits the bulk of the distribution misses the tails. We investigate the presence of a power-law tail, $P(\hat{V} > X) \sim X^{-\alpha}$, in the Champions League and Wimbledon trading volumes, using the methods of [13]. For Wimbledon trading volumes we must reject the hypothesis of a power-law tail. But for Champions League trading volumes, we estimate a tail exponent of $\alpha = -2.60 \pm 0.12$ for values of $\hat{V} > 226$ using maximum likelihood estimation (and bootstrapping) and cannot reject the hypothesis of a power-law tail for $\hat{V} > 226$ with a P-Value of 0.89.

However, since the Betfair market functions via the medium of the Betfair website, http://www.betfair.com/, we cannot rule out the effect of network traffic dynamics in shaping the statistics of trading volumes. In [17], the authors find a log-normal fit to the distribution for network traffic. Furthermore, they conclude that it is the dynamical process related to the inter-arrival time of network packets that plays the key role in the formation of the observed log-normal traffic distributions. It may be that the trading volume distributions we observe in Fig.(9) are influenced by network traffic dynamics.

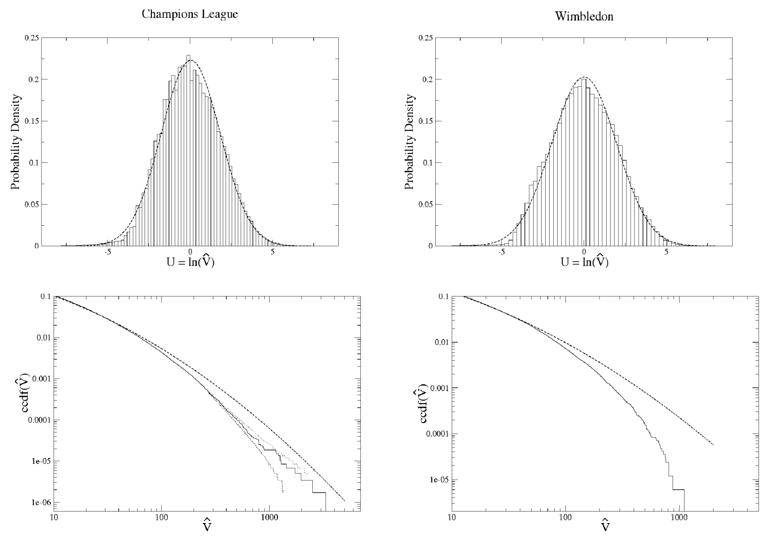

Figure 9: The probability distribution, $P(U)$, for the logarithm of the normalised 1-second trading volumes for Champions League (left) and Wimbledon (right) markets. The dashed lines represent log-normal fits with mean $\mu = 0.046$ and variance $\sigma = 1.79$ for Champions League data and $\mu = 0.016$ and $\sigma = 1.96$ for Wimbledon. To investigate power-law scaling in the tails of the distribution, we have plotted the complementary cumulative distribution functions, $ccdf(\hat{V}) = \int_{\hat{V}}^{\infty} P(\hat{V}) d\hat{V}$, below. The tail falls away exponentially in the case of Wimbledon trading volumes, however, for Champions League volumes, we cannot reject the null hypothesis of power-law behaviour for $V > 226$. The confidence intervals for the power law fit are depicted with dotted lines.

5 Summary and Conclusions

We have investigated some basic statistical properties of market observables in Betfair betting markets with an aim to compare and contrast with known results for financial market data. Unlike financial markets, in which traders buy and sell financial assets, on Betfair it is bets that are instead 'backed' or 'laid'. We draw parallels between the Betfair betting market and binary option prediction markets and employ a similar method presented in [9] to process the returns.

For football market returns, we identify a difference in tail behaviour between the distribution of returns drawn for times when game-play is underway and from half-time when match play has ceased. We argue that the returns we observe during half-time are representative of endogenous market behaviour, whereas when the match is underway and news is hitting the market, the returns are instead driven by the evolution of changing probabilities in a competitive event. We further compare the in-play returns of football with those drawn from tennis match markets. We find much fatter tails in the distribution of football match returns than in the distribution of tennis match returns, and this agrees qualitatively with expectations considering the differences in game-play and scoring of these two sports.

Following [9], we have investigated power-law scaling in the tails of the distribution of returns that occur at half-time, which we have associated with endogenous market behaviour. We find that although much of the distribution appears to obey a Pareto law with exponent $\alpha = 2.8$, the extreme tails are characterised by a more exponential decay. We hypothesise that partisan trading may be preventing large magnitude returns from occurring, as suggested by the authors of [9].

We have also analysed the distribution of the 1-second trading volumes for Betfair market trading. We find that this distribution can be fit well to a log-normal function in the case of markets for both sports studied. Furthermore, we observe power-law scaling in the tails of the trading volume distribution for football returns. However, since Betfair market trading occurs over the medium of the World Wide Web, via the Betfair website, http://www.betfair.com/, we must concede that the trading volume statistics we observe may be influenced by network traffic dynamics.

6 Appendix : Kolmogorov-Smirnov Test

The Kolmogorov-Smirnov test is a non-parametric test used to evaluate the likelihood that a given sample of empirical data is drawn from a given model distribution (or to evaluate the likelihood that two empirical distributions are drawn from the same distribution). The Kolmogorov-Smirnov statistic, D_n, is defined as the maximum difference between the empirical distribution function of a given sample of n observations X_i,

$$F_n(x) = \frac{1}{n} \sum_{i=1}^{n} I_{X_i \leq x} \tag{11}$$

and the cumulative distribution function $F(x)$ of the model distribution (or empirical distribution function of the second sample). Here, $I_{X_i \leq x}$ is the indicator function, equal to 1 if $X_i \leq x$, and equal to 0 otherwise.

$$D_n = \sup_x |F_n(x) - F(x)| \tag{12}$$

Under the null hypothesis that the sample data are drawn from the model distribution F(x), then

$$\lim_{n\to\infty} \sqrt{n} D_n = \sup_t |B(F(t))| \tag{13}$$

where $B(t)$ is a Brownian bridge, a stochastic process whose probability distribution is the conditional probability distribution of a Wiener process, given the condition that B(0) = B(1) = 0.

The P-value for the test is,

$$P = 1 - V(\sqrt{n} D_n) \tag{14}$$

Where $V(x)$ is the cumulative distribution function for the maximum deviation from zero of a Brownian bridge. We reject the null hypothesis at the 5% level, if $P < 0.05$, corresponding to 95% of the maximum deviations from zero, of a Brownian bridge being greater than our empirical value of D_n. If comparing two empirical samples of size n_1 and n_2, then

$$P = 1 - V(\sqrt{n_{eff}} D_n) \tag{15}$$

where

$$n_{eff} = \frac{n_1 n_2}{n_1 + n_2} \tag{16}$$

In this paper, we use the Kolmogorov-Smirnov test to evaluate how similar the distributions of normalised implied probability returns taken from different time windows in our data sets are. However, because of the limited resolution of the recorded Betfair market returns (e.g. the Betfair user interface only allows you to back and lay at specified odds: 1.01, 1.02, etc.), small changes in odds are forced to take on the nearest allowable values. This leads to a poorly resolved, 'spiky' implied probability returns at its centre. To avoid this issue at the centre of the distributions, we only compare the distributions for magnitudes of detrended returns greater than 1.5.

7 Acknowledgments

Stephen Hardiman is supported by the Irish Research Council for Science, Engineering and Technology: funded by the National Development Plan and Peter Richmond acknowledges support from the TCD Visiting Professorships and Fellowships Benefaction Fund. We would like to thank Betfair and Fracsoft for the provision of their data.

References

[1] R. Forsythe, F. Nelson, G. R. Neumann, J. Wright, *Anatomy of an experimental political stock market*, The American Economic Review **82** (5), 1142-1161 (1992).

[2] J. Wolfers, E. Zitzewitz, *Prediction Markets*, Journal of Economic Perspectives **18** (2), 107 (2004).

[3] J. Wolfers, E. Zitzewitz, *Interpreting Prediction Market Prices as Probabilities*, Institute for the Study of Labor, Discussion Papers 2092 (2006).

[4] S. Gjerstad, *Risk Aversion, Beliefs and Prediction Market Equilibrium*, Microeconomics (0411002) (2004).

[5] J. E. Berg, T. A. Rietz, *Information Markets: A New Way of Making Decisions*, AEI Press Washington DC (2006).

[6] J. E. Berg, R. Forsythe, F. D. Nelson, T. A. Rietz, *The Handbook of Experimental Economic Results Volume 1*, *North Holland* (2008).

[7] J. E. Berg, T. A. Rietz, *Prediction market accuracy in the long run*, International Journal of Forecasting **24**, 283-298 (2008).

[8] R. S. Erikson, C. Wlezien, *Are political markets really superior to polls as election predictors?* Public Opinion Quarterly **72**, 190-215 (2008).

[9] S. R. Majumder, D. Diermeier, T. A. Rietz and L. A. N. Amaral, *Price dynamics in political prediction markets*, Proceedings of the National Academy of Science, **106**(3), 679-684 (2009).

[10] S. J. Hardiman, P. Richmond and S. Hutzler, *Long-range correlations in an online betting exchange for a football tournament*, New Journal of Physics, **12** 105001 (2010).

[11] V. Plerou, P. Gopikrishnan, L. A. N. Amaral, M. Meyer and H. E. Stanley, *Scaling of the distribution of price fluctuations of individual companies*, Physical Review E **60** (6), 6519-6529 (1999).

[12] K. Matia, L. A. N. Amaral, S. P. Goodwin and H. E. Stanley, *Different scaling behaviors of commodity spot and future prices*, Physical Review E **66**, 045103 (2002).

[13] A. Clauset, C. R. Shaliz and M. E. J. Newman, *Power-law distributions in empirical data*, ArXiv e-print 0706.1062v1 (2007).

[14] P. Gopikrishnan, V.i Plerou, X. Gabaix and H. E. Stanley, *Statistical properties of share volume traded in financial markets*, Physical Review E **62**, 4493-4496 (2000).

[15] B. B. Ajinkya and P. C. Jain, *The behavior of daily stock market trading volume*, Journal of Accounting and Economics **11** (4), 331-359 (1989).

[16] I. Antoniou, Vi. V. Ivanov, Va. V. Ivanov, and P. V. Zrelov, *On the log-normal distribution of stock market data*, Physica A **331** (3-4), 617-638 (2004).

[17] I. Antoniou, V. V. Ivanov, V. V. Ivanov and P V Zrelov, *On the log-normal distribution of network traffic*, Physica D **167**, 72-85 (2002).