

PREDICTION MARKETS: Theory, Evidence and Applications.

By Leighton Vaughan Williams, Ming-Chien Sung and Johnnie Johnson

A convenient definition of a prediction market is a judgment-based market that is created or employed for the purpose of making predictions. The theoretical underpinning derives from the efficient market hypothesis, and from the view that relevant information concerning the likelihood of future events that is dispersed among many people can be effectively aggregated via market mechanisms to help forecast event outcomes. Prediction markets of various types have been employed to forecast outcomes ranging from open and closed-door election outcomes (Rhode and Strumpf, 2013; Vaughan Williams and Paton, 2015; Vaughan Williams and Reade, 2015) to the box office success of movies (Court and McKenzie, 2018; Escoffier and McKelvey, 2015), the forecasting of infectious disease activity (Farrow et al., 2017; Polgreen et al., 2007), and the probability of meeting project deadlines (Cowgill and Zitzewitz, 2015; Leigh and Wolfers, 2007).

The potential of these markets to forecast event outcomes has generated significant interest among social scientists and the business community, while corporations use prediction markets as a decision support tool (O’Leary, D.E., 2015, Spears et al., 2009). Prediction markets also have many potentially valuable applications for public policy (Paton et al., 2010; Tetlock et al., 2017) and for economic forecasting (Snowberg et al., 2013). A separate focus of attention has been the design of the market mechanisms used in prediction markets (Chen, 2011; Klingert and Meyer, 2018; Spann and Skiera, 2003). There has also been a developing interest in recent years in the idea of combining forecasts from a range of different forecasting methodologies (Graefe et al., 2013; Pathak et al., 2015; Rothschild, 2015), and some have questioned aspects of the performance of markets relative to other forecasting methodologies (Atanasov et al., 2017; Erikson and Wlezien, 2008; Fry and Brint, 2017). Others have addressed how these markets might be influenced by attempts at market manipulation (Buckley and O’Brien, 2017; Hanson and Oprea, 2009). Forecasts provided by prediction markets have also been examined for systematic biases (Cowgill et al., 2009; Ma et al., 2016; Restocchi et al., 2018) as well as the way in which new information is incorporated into market prices (Sung et. al, 2019; Vaughan Williams and Reade, 2016). This work complements related research into the efficiency of wagering markets and other forecasting methodologies (Croxson and Reade, 2014; Vaughan Williams et al., 2016; Vaughan Williams and Stekler, 2010).

The papers in this special issue address such topics and beyond. In the remainder of this introductory essay, we provide focused summaries of the papers chosen for inclusion and draw some conclusions about the overall state of research in this area and challenges for the future.

PAPERS IN THE SPECIAL ISSUE

1. *Joyce E. Berg and Thomas A. Rietz – Longshots, Overconfidence and Efficiency on the Iowa Electronic Market*

Berg and Rietz analyse the forecasting efficiency of binary prediction markets. Such markets, which predict probabilities of discrete outcomes, have become quite popular. However, researchers generally cite measures of forecasting efficiency based on linear prediction markets. Linear markets forecast the level of an outcome. To measure efficiency, forecasts can be compared to actual outcomes directly. This differs from binary markets where one must compare the forecast outcome probabilities to actual outcome frequencies under repeated, essentially identical, conditions. The authors use a unique set of repeated, binary prediction markets to study efficiency using frequency and logistic analysis. While they document a pricing bias at intermediate horizons, the bias disappears before the forecasted events occur. The observed bias conflicts with static models of bias based on the longshot bias, overweighting of low probability events, prospect theory and other models that predict over-pricing of low-probability events. The bias observed is transitory, affects only very high and low prices, and follows a pattern predicted by the information-based, over-reaction story of Daniel, Hirshleifer and Subrahmanyam (1998). Early in a market, when there is little information, prices appear unbiased. As traders observe information, prices respond, leading to transitory overpricing of high probability events (and underpricing of low probability events). As the event approaches and more information is reflected in prices, this bias fades and, ultimately, disappears. The paper makes two observations that are important in interpreting binary market prices: (1) Prices are forecast probabilities and cannot be compared directly to outcomes. (2) Prediction markets may appear overconfident (i.e., over-price very high probability events and under-price very low probability events) at intermediate horizons.

2. *Alasdair Brown and Fuyu Yang – The Wisdom of Large and Small Crowds: Evidence from Repeated Natural Experiments in Sports Betting*

Brown and Yang analyse the role that crowd size plays in the accuracy of prediction/betting market prices. They ask whether large markets with high participation rates provide more accurate forecasts than smaller markets. There are conflicting theoretical predictions in this area. On the one hand a large market populated with noise traders (De Long et al, 1990) may produce biased prices. On the other hand, large markets create the incentives for information acquisition as returns from informed trading are higher. Indeed, information may actually be dispersed amongst the crowd (Galton, 1907) and therefore as long as forecast errors are not correlated a larger crowd will produce more accurate predictions. The authors examine repeated natural experiments in tennis betting, where every other year the Queen's Club tennis tournament clashed with a major soccer tournament. Importantly, these clashing tournaments exogenously reduced the participation rate in tennis betting. They found that larger markets (without the clashing soccer tournament) produced more accurate predictions. This a result of practical importance, as many prediction market designers can set rules or incentives which determine whether the prediction market will be large or small. In this case, bigger is better.

3. *Chung-Ching Tai, Hung-Wen Lin, Bin-Tzong Chie and Chen-Yuan Tung – Predicting the Failures of Prediction Markets: A Procedure of Decision Making Using Classification Models*

Tai, Lin, Chie, and Tung provide a decision-support framework to help policy or decision makers who count on the forecasts of prediction markets. Regardless of the accuracy of prediction markets in general, sometimes the failures in their predictions will put decision-makers' stake at risk. Instead of blindly relying on markets' predictions, therefore, this paper offers a systematic way of assessing the credibility of the predictions. On the basis of previous studies about factors influencing the efficiency of prediction markets, the authors incorporated a list of variables into statistical models as well as machine learning algorithms to explore the underlying (and possibly nonlinear) relationships between these factors and the prediction results. More specifically, four classification models were used to classify the credibility of prediction market forecasts. They then used a combined forecasting technique to integrate the classification results from these four models. By doing so, they are able to provide a judgment about a specific prediction market's forecast prior to the predicted event. A large dataset involving 650 markets was used to train and test their combined forecasting framework. The results indicate that the method was able to discover models capable of predicting the failures of prediction markets according to different criteria. Most interesting is the flexibility of the proposed framework which allows decision-makers to build their own models with different sets of variables, classification models, and decision objectives.

4. *Andrew Grant, David Johnstone and Oh Kang Kwon – The Cost of Capital in a Prediction Market*

Grant, Johnstone and Kwon treat the contracts in prediction markets as financial assets. Contracts are "priced" from the traders' subjective perspective using utility theory or the finance expression of utility theory known as the capital asset pricing model (CAPM). The contract's price implies its discount rate or required rate of return. By exploring the price-implied discount rates of binary contracts, the authors clarify how a trader's rational required expected return reacts to the trader's subjective probability of winning. The general finding is that a trader who is more confident of the contract expiring in the money requires a lower expected return on that contract. The expected return required of a typical prediction market contract that pays zero or one is found to increase linearly in its ratio of ex ante perceived subjective payoff variance to payoff mean. A surprising but immediate conclusion is that the natural risk-aversion of well-informed traders induces a favourite-longshot bias in prediction markets, where longshot contracts cost more than they should relative to their low chance of winning and favourites are relatively too cheap. This bias is explained not by errors in traders' probability assessments, but by the posted bids and asks of traders whose probability assessments are accurate but who are risk-averse.

5. *Luis Felipe Costa Sperb, Ming-Chien Sung, Tiejun Ma and Johnnie Johnson – Keeping a Weather Eye on Prediction Markets: Improving Forecasts by Accounting for Environmental Conditions*

Costa Sperb, Sung, Johnson and Ma investigate the influence of environmental factors on the calibration of probabilistic estimates derived from prediction market prices. The literature suggests that environmental factors, such as weather and atmospheric conditions, can affect the information processing and cognitive ability of decision-makers, leading to sub-optimal estimations about uncertain events. This, in turn, has the potential to impair the effectiveness of prediction markets as a means of appropriately aggregating and weighting relevant information dispersed in the public. In the context of horserace betting markets, the authors found that even after the effects of these conditions on the performance of contestants (horses and jockeys) have been discounted, the accuracy of the probabilities derived from market prices is affected systematically by the prevailing weather and atmospheric conditions. By correcting for this phenomenon, they showed that significantly better forecasts can be derived from prediction market prices, and that these have substantial economic value. Importantly, the results of this paper suggest that when the purpose of a prediction market is to derive accurate probabilistic estimates from final contract prices, forecast accuracy can be improved greatly by identifying and correcting for conditions where prediction markets systematically under-perform.

6. *J. James Reade and Leighton Vaughan Williams – Polls to Probabilities: Comparing Prediction Markets and Opinion Polls*

Reade and Vaughan Williams consider the forecast performance of prediction markets alongside that of opinion polls in the context of US election outcomes. The first contribution of this paper is to evaluate a much wider range of prediction markets than previous studies have been able to, most notably considering the large commercial prediction markets, Intrade and Betfair. Prediction market forecasts are probabilistic in terms of which candidate will win, whereas opinion polling output tends to be reported in terms of vote shares. This makes it difficult to compare how close each type of forecast was to providing the most efficient forecast, where efficiency is measured in terms of information incorporated at the time the forecast was made. Although these two variables are related, their correspondence requires a set of assumptions, be they theoretical or statistical. Page (2008) provides theoretical assumptions, and the second contribution of this paper is to use the related statistical assumptions to make an empirical conversion between probabilities and vote shares. As both opinion polls and prediction market forecasts can be corrected for bias, it is more informative to think about the precision of forecasts. On this metric, prediction market forecasts of election outcomes perform better than opinion polls.

7. *Chen Di, Stanko Dimitrov and Qi-Ming He – Incentive Compatibility in Prediction Markets: Costly Actions and External Incentives*

Di, Dimitrov, and He consider the role of external incentives on the ability of prediction markets to aggregate information. Most papers on external incentives show that when external incentives exist, it is likely that prediction markets will not always accurately capture all agents' private information. However, papers that consider external incentives tend to assume agents' costs external of the prediction market are symmetric. The authors show that

even when agents may take actions external to the prediction market that influence the outcome of the traded event, if the actions external of the prediction market are present, prediction markets may indeed capture all agent's private information. In particular, so long as the desired external action is rewarded more than the undesired action, prediction markets do not incentivise undesirable actions. This insight hopefully addresses concerns some managers may have in using prediction markets in the workplace due to potentially incentivising undesirable actions.

8. *Thomas Auld and Oliver B. Linton – The Behaviour of Betting and Currency Markets on the night of the EU Referendum*

Auld and Linton identify the night of the UK EU Referendum of 2016 as providing a unique natural experiment to study the degree to which information was discounted in two parallel forms of prediction market: financial and betting markets. In particular, they explore the behaviour of the sterling dollar exchange rate as well as binary contracts listed on Betfair that pay out according to the outcome of the Referendum. The authors argue that for those few hours overnight the *sole* determinant of prices in these assets was the flow of information provided by the results of the vote. Using public information they construct an ex-ante joint prior distribution for the vote-share of every constituency that announced that night. Using a Bayesian methodology they update this prior as results arrive and compute a dynamic probability that the UK leaves the EU. They find that although both markets are slow in pricing the information contained in the results, by around 2 to 3 hours, the betting market is less inefficient than the financial market. Further, by constructing a theoretical model that links the prices of these contracts, independent of the outcome of the vote, they show that there were violations of weak market efficiency. There were apparently very profitable arbitrage opportunities involving selling the pound and placing hedging bets on Betfair that would pay out whether or not the UK voted to leave the EU.

9. *Christoph Lohrmann and Pasi Luukka – Classification of Intraday S&P500 Returns with a Random Forest*

Lohrmann and Luukka link stock markets to prediction markets and analyse the prediction of S&P500 intraday returns. In contrast to much of the existing literature, this is set up as a four-class problem and not as a regression or binary classification problem. Premised on these four return classes, four trading strategies are tested against a simple buy-and-hold strategy. All the suggested trading strategies conduct buy- and sell-decisions only based on certain predicted classes. The research indicates that the four classes differ in their contribution to the return on the strategies. In particular, the two extreme classes with 'Strong Positive' and 'Strong Negative' predicted returns lead overall to higher mean returns than the 'Slightly Positive' and 'Slightly Negative' return classes. This result holds true even when misclassified returns are included in the averages. Therefore, using strategies that act only on a subset of the predicted classes (such as the extreme predictions) can work to generate higher profits than following all the predictions.

10. Oliver Strijbis and Sveinung Arnesen – Explaining Variance in the Accuracy of Prediction Markets

Strijbis and Arnesen go beyond the current focus of the prediction market literature by combining observational and experimental analyses of prediction market errors. They investigate the prediction error of a real money prediction market using a logarithmic market scoring rule for 65 direct democratic votes in Switzerland. The authors distinguish between prediction market error due to the set-up of the market, features of the event to be predicted, and the participants involved. They find that the prediction market accuracy varies primarily according to the set-up of the market, while the features of the event and especially the composition of the participant sample hardly matter. Hence, those applying prediction markets should consider carefully the specific configuration of their market while they can remain more relaxed about the composition of their sample of traders.

11. Alasdair Brown, J. James Reade and Leighton Vaughan Williams – When Are Prediction Market Prices Most Informative?

Brown, Reade and Vaughan Williams consider the informativeness of prediction market prices. In the aftermath of information events, how do market prices react? This naturally matters, as prediction markets are extensively used to elicit forecasts about uncertain future events. The paper uses the Intrade exchange and the release of opinion polling data. They investigate, in particular, what happens in prediction markets in the immediate aftermath of opinion polling releases. Such releases are significant news events, and in the case of Gallup polls, occur on a regular schedule, at 1pm Eastern Time during most days during the 2012 presidential campaigns in the US. The authors find that in the immediate aftermath of a poll release, there is an increase in trading activity. However, much of this activity involves relatively inexperienced traders, and as a result market efficiency declines in those moments. Once more experienced traders return in the following hours, price efficiency recovers. These findings are of practical importance, as they give a sense of the extent to which prediction market prices might be relied upon for accurate forecasts in the aftermath of significant news events.

CONCLUSION

In summary, the papers presented in this issue demonstrate that there are several factors, such as the natural risk aversion of well-informed traders and the trading behaviour observed immediately after the release of new information that can induce biases in prediction market forecasts. Equally, that certain predicted classes (such as extreme predictions) can be under-priced or over-priced. In addition, the papers demonstrate that the accuracy and precision of forecasts derived from prediction markets can be affected by a range of factors, including environmental factors, the size of the market, the set-up of the market, the external incentives provided and when in the market cycle predictions are observed.

However, evidence is presented that some aspects of prediction markets, despite *a priori* expectations, do not influence their accuracy, such as features of the event and the

composition of the participants. In addition, provided markets are given sufficient time, accuracy of prediction market forecasts increases and forecasts derived from these markets outperform, on key metrics, other forecasting mechanisms such as opinion polls and financial markets.

Consequently, the evidence presented here should provide comfort to the potential users of prediction market forecasts. Taken together, the papers suggest that by thinking carefully about the design of prediction markets, by setting up procedures to adjust forecasts to account for known biases and by planning carefully when forecasts from prediction markets are used, they can be a powerful source of forecasting ability.

These findings reflect a notable increase in interest in the applications to which prediction markets can be put, as well as the broader body of evidence suggesting that these markets might produce better forecasts than alternative forecasting mechanisms. The insights gained have also been shown to have potentially valuable applications for policy, not least when accurate forecasts are required in relation to quantifiable targets. The findings also reflect a wider interest in the best way to design and implement prediction markets.

Even so, some have questioned how far prediction markets can outperform other methods of forecasting, or whether they are at best a supplement to more traditional forecasting methodologies. These doubts have attracted added focus in very recent years following some high-profile forecasting failures in the context of major political event outcomes. In particular, the 2016 EU referendum in the UK and the 2016 US presidential election produced results that were a surprise not only to the great majority of pollsters but also to followers of the prediction markets. There are various theories to explain why the markets failed in these big votes. One theory looks to the basic laws of probability. An 80 per cent favourite can be expected to lose one time in five, if the odds are correct. In the long run, according to this explanation, things should even out. A second theory to explain the surprise results is that something fundamental has changed in the way that information contained in political prediction markets is perceived and processed. One interpretation is that the hitherto widespread success of the markets in forecasting election outcomes, and the publicity that was given to this, turned them into something of an accepted measure of the state of a race, creating a perception which was difficult to shift in response to new information. This is a form of 'anchoring'. Linked to this is a herding hypothesis, which is that because the prediction markets had by 2016 become so firmly entrenched in conventional wisdom as an accurate forecasting tool, people herded around the forecasts, propelling the implied probabilities of favoured outcomes upwards. A third theory is that conventional patterns of voting broke down in 2016, primarily due to unprecedented differential voter turnout patterns across key demographics, which were not correctly modelled in most of the polling and which were missed by political pundits, political scientists, politicians and those trading the prediction markets. Finally, there has been widespread discussion of the impact of manipulation, not only of the markets themselves but also of the distribution of information and misinformation.

Future research into the uses and applications of prediction markets might focus on what can potentially be learned from recent forecasting failures to improve the efficiency of forecasts derived from prediction markets, as well as to identify when markets are likely to be most informative. The ways in which market forecasts can be adjusted to allow for systematic

biases is also likely to attract the continued interest of researchers. Besides comparing the efficiency of prediction markets with other forecasting methodologies, there is in addition a growing interest in the idea of combining forecasts derived from the different methodologies to generate improved forecasts. Increasing attention is also being paid to the design and incentive structure of markets and the contribution they can make as a corporate and policy decision support tool.

The papers selected for this special issue of the International Journal of Forecasting reflect this diverse research agenda, and offer an exceptionally strong contribution to the existing literature on prediction markets. Importantly, they provide a valuable framework upon which future work can build. It is with no little anticipation that those interested in theory, evidence and applications relating to prediction markets look forward with great anticipation to the outputs of further research in this area.

References

Atanasov, P., Rescober, P., Stone, E., Swift, S.A., Servan-Schreiber, E., Tetlock, P., Ungar, L. and Mellers, B. (2017). Distilling the Wisdom of Crowds: Prediction Markets vs. Prediction Polls. *Management Science*. 63 (3), March, 587-900.

Buckley, P. and O'Brien, F. (2017), The Effect of Malicious Manipulations on Prediction Market Accuracy. *Information Systems Frontiers*. 19 (3), June, 611-623.

Chen, Y. (2011), Mechanisms for Prediction Markets. In: Vaughan Williams, L. (ed.) *Prediction Markets: Theory and Applications*. New York: Routledge. Chapter 4. 35-51.

Court, D., Gillen, B., McKenzie, J. and Plott, C.R. (2018). Two Information Aggregation Mechanisms for Predicting the Opening Weekend Box Office Revenues of Films: Boxoffice Prophecy and Guess of Guesses. *Economic Theory*, 65 (1), January, 25-54.

Cowgill, B. and Zitzewitz, E. (2015). Corporate Prediction Markets: Evidence from Google, Ford and Firm X. *Review of Economic Studies*. 82, October, 1309-1341.

Cowgill, B., Wolfers, J. and Zitzewitz, E. (2009). Using Prediction Markets to Track Information Flows: Evidence from Google. Available at:
www.columbia.edu/~bc2656/GooglePredictionMarketPaper.pdf

Croxson, K. and Reade, J.J. (2014). Information and Efficiency: Goal Arrival in Soccer Betting. *The Economic Journal*. 124, 575, March. 62-91.

Daniel, K., Hirshleifer, D. and Subrahmanyam, A. (1998). Investor Psychology and Security Market Under- and Overreactions. *Journal of Finance*. 53 (6), 1839-1885.

De Long, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J. (1990). Noise Trader Risk in Financial Markets. *Journal of Political Economy*. 98 (4), August, 703-738.

Escoffier, N. and McKelvey, B. (2015). The Wisdom of Crowds in the Movie Industry: Towards New Solutions to Reduce Uncertainties. *International Journal of Art Management*, 17 (2), December, 52-63.

Farrow, D.C., Brooks, L.C., Hyun, S., Tibshirani, R.J., Burke, D.S. and Rosenfeld, R. (2017). PLOS. *Computational Biology*, March 10.

Fry, J. and Brint, A. (2017). Bubbles, Blind-Spots and Brexit. *Risks*. 5 (3), 37. 1-15.

Galton, F. (1907). *Vox Populi (The Wisdom of Crowds)*. *Nature*. 1949 (75), 450-451.

Graefe, A., Armstrong, J.S., Jones, R.J. and Cuzan, A.G. (2014). Combining Forecasts: An Application to Elections. *International Journal of Forecasting*, 30, 1, January-March, 43-54.

Hanson, R. and Oprea, R. (2009). A Manipulator Can Aid Prediction Market Accuracy. *Economica*. 76, 302, April, 304-311.

Klingert, F.M.A. and Meyer, M. (2018). Comparing Prediction Market Mechanisms: An Experiment-Based and Micro Validated Multi-Agent Simulation. *Journal of Artificial Societies and Social Stimulation*, 21 (1), 7.
<http://jasss.soc.surrey.ac.uk/21/1/7.html>

Leigh, A. and Wolfers, J. (2007). Prediction Markets for Business and Public Policy, *The Melbourne Review*, 3 (1), May, 7-15.

Ma, T., Tang, L., McGroarty, F., Sung, M., & Johnson, J. E. V. (2016). Time is Money: Costing the Impact of Duration Misperception in Market Prices. *European Journal of Operational Research*, 255 (2), 397-410.

O'Leary, D.E. (2015). User Participation in a Corporate Prediction Market. *Decision Support Systems*. Vol. 78, October, 28-38.

Page, L. (2008). Comparing Prediction Market Prices and Opinion Polls in Political Elections. *The Journal of Prediction Markets*. 2 (1), May, 91-97.

Pathak, D., Rothschild, D. and Dudik, M. (2015). A Comparison of Forecasting Methods: Fundamentals, Polling, Prediction Markets and Experts. *The Journal of Prediction Markets*, 9, 2, 1-31.

Paton, D., Siegel, D. and Vaughan Williams, L. (2010). Gambling, Prediction Markets and Public Policy. *Southern Economic Journal*, 76 (4), April, 878-883.

Polgreen, P.M., Nelson, F.D. and Newmann, G.R. (2007). Use of Prediction Markets to Forecast Infectious Diseases. *Clinical Infectious Diseases*, 44, January, 272-279.

Restocchi, V., McGroarty, F. and Gerding, E. (2018). The Temporal Evolution of Mispricing in Prediction Markets. *Finance Research Letters*. Available online 11 August. <https://www.sciencedirect.com/science/article/abs/pii/S1544612318303349>
In Press. Available online, 11 August.

Rhode, P.W. and Strumpf, K. (2013). The Long History of Political Betting Markets: An International Perspective. Section VI: Prediction Markets and Political Betting. In: Vaughan Williams, L. and Siegel, D.S. (eds.) *The Oxford Handbook of the Economics of Gambling*. New York: Oxford University Press. 560-588.

Rothschild, D. (2015). Combining Forecasting for Elections: Accurate, Relevant and Timely. *International Journal of Forecasting*. 31 (3), July-September, 952-964.

Snowberg, E., Wolfers, J. and Zitzewitz, E. (2013). In: Prediction Markets for Economic Forecasting. In: *Handbook of Economic Forecasting* (eds.), Elliott, G. and Timmermann, A., Vol. 2, Part A. 657-687.

Spears, B., LaComb, C., Interrante, J., Barnet, J. and Senturk-Dogonaskoy, D. (2009). Examining Trader Behavior in Idea Markets: An Implementation of GE's Imagination Markets. *The Journal of Prediction Markets*, 3, 1, 17-39.

Sung, M., McDonald, D. C. J., Johnson, J.E. V., Tai, C. C., & Cheah, E-T. (2019). Improving Prediction Market Forecasts by Detecting and Correcting Possible Overreaction to Price Movements. *European Journal of Operational Research*, 272(1), 389-405.

Tetlock, P.E., Mellers, B.A. and Scoblic, J.P. (2017). Bringing Probability Judgments into Policy Debates via Forecasting Tournaments. *Science*, 333, 6234, 481-483.

Vaughan Williams, L. and Paton, D. (2015). Forecasting the Outcome of Closed-Door Decisions: Evidence from 500 Years of Betting on Papal Conclaves. *Journal of Forecasting*, 34 (5), August, 391-404.

Vaughan Williams, L. and Reade, J.J. (2015). Forecasting Elections. *Journal of Forecasting*, 35 (4), July, 308-328.

Vaughan Williams, L. and Reade, J.J. (2016). Prediction Markets, Social Media and Information Efficiency. *Kyklos*, 69 (3), August, 518-556.

Vaughan Williams, L. and Stekler, H. (2010). Sports Forecasting. *International Journal of Forecasting*. 26 (3), July-September, 445-447.

Vaughan Williams, L., Sung, M., Fraser-Mackenzie, P.A.F., Peirson, J. and Johnson, J.E.V. (2016). Towards an Understanding of the Origins of the Favourite-Longshot Bias: Evidence from 500 Years of Betting on Papal Conclaves. *Economica*, 85, 338, April. 360-382.