

Forecasting with Social Media: Evidence from Tweets on Soccer Matches *

Alasdair Brown[†] Dooruj Rambaccussing[‡]
University of East Anglia University of Dundee

J. James Reade[§]
University of Reading
Institute for New Economic Thinking, Oxford

Giambattista Rossi[¶]
Birkbeck, University of London

May 20, 2017

Abstract

Social media is now used as a forecasting tool by a variety of firms and agencies. But how useful are such data in forecasting outcomes? Can social media add any information to that produced by a prediction/betting market? We source 13.8m posts from Twitter, and combine them with contemporaneous Betfair betting prices, to forecast the outcomes of English Premier League soccer matches as they unfold. Using a micro-blogging dictionary to analyse the content of Tweets, we find that the aggregate tone of Tweets contains significant information not in betting prices, particularly in the immediate aftermath of goals and red cards.

JEL Classification: G14, G17

Keywords: forecasting, social media, prediction markets, wisdom of crowds, soccer

*This paper was previously circulated under the title ‘Using Social Media to Identify Market Inefficiencies: Evidence from Twitter and Betfair’. We would like to thank Rob Simmons, the Editor, and two anonymous referees for very helpful comments. We also also indebted to Ken Benoit, and participants at the LSE Applied Quantitative Analysis Workshop, the UEA Statistics and Data Science in the Digital Age programme, the European Economic Association 2016 Annual Congress in Geneva, and the Gijon 2016 conference on Sport and Media for feedback. Alasdair Brown acknowledges financial support from the British Academy and the Leverhulme Trust (Ref: SG140097), Dooruj Rambaccussing acknowledges financial support from an SIRE Early Career Grant, James Reade thanks the Open Society Foundations (OSF) and Oxford Martin School (OMS) for financial support, and Dooruj Rambaccussing, James Reade and Giambattista Rossi acknowledge support from Twitter through their #DataGrants scheme. The research presented in this paper was carried out on the High Performance Computing Cluster supported by the Research and Specialist Computing Support service at the University of East Anglia. The usual disclaimer applies.

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

[‡]Economic Studies, University of Dundee, Dundee DD1 4HN, UK. E-mail: d.rambaccussing@dundee.ac.uk

[§]School of Economics, University of Reading, Reading RG6 6AA, UK. E-mail: j.j.reade@reading.ac.uk

[¶]Department of Management, Birkbeck College, University of London, Malet Street, Bloomsbury London WC1E 7HX, U.K. Email: g.rossi@bbk.ac.uk

1 Introduction

Social media content – for example that produced on Twitter or Facebook – is increasingly used as a forecasting tool. For example, Hollywood studios use data from social media to forecast demand for new films.¹ Financial firms extract sentiment from Twitter to predict stock returns, and design funds to algorithmically trade based on this information². And social media is now even used for economic forecasting: in 2012, the Australian Treasury department launched a division to harness social media data to forecast workforce participation and retail sentiment, among other things.³

But how useful and accurate a forecasting tool is social media? On the one hand, social media content can harness the opinions and beliefs of a wide group of participants. Thus the ‘wisdom of the crowd’ (Galton, 1907, Surowiecki, 2005) may produce accurate forecasts. On the other hand, the incentives for an individual on social media to provide accurate information for forecasting may arguably be weak. Unlike in markets, accurate social media forecasts may enhance an individual’s reputation, but are not directly profitable. And worse, there are many incidences of misinformation on social media. For example, a hoax Tweet on the Associated Press Twitter feed in 2013, misreporting an explosion at the White House in Washington, briefly wiped \$136 billion off the S&P 500 Index.⁴

In this paper we evaluate the accuracy of social media forecasting in a fast-moving, high-profile environment: English Premier League soccer matches. We study 13.8 million Tweets, an average of 5.2 Tweets per second, during 372 matches that took place during the 2013/14 season. Our primary aim is assess whether information contained in these Tweets can predict match outcomes. Furthermore, we also aim to assess whether the forecasting capacity of social media is concentrated during large events (such as the scoring of a goal or the issuance of a red card) – which would indicate that social media helps to ‘break’ news – or whether any forecasting capacity is to be found in the aftermath of such events, in which case social media helps in the interpretation of information.

¹‘Hollywood Tracks Social Media Chatter to Target Hit Films’, Brook Barnes, *New York Times*, Dec 7th 2014.

²‘Hey Finance Twitter, You Are About to Become an ETF’, Eric Balchunas, *Bloomberg*, Oct 15th 2015.

³‘Treasury to Mine Twitter for Economic Forecasts, David Ramli, *Financial Review*, Oct 30th 2012.

⁴‘A Fake AP Tweet Sinks the Dow for an Instant’, Jared Keller, *Bloomberg*, April 23rd 2013

One problem is that social media content does not easily translate to probability forecasts. For example, we cannot state that X number of Tweets on a particular team, in a given interval, maps to a prediction that the team has a Y% chance of victory. Our solution, therefore, is to ask whether Twitter content can add information to probability forecasts produced by a prediction/betting market, Betfair. This is a high bar, as prediction markets have been found to outperform tipsters in the context of sports (Spann and Skiera, 2009), and outperform polls and experts in the context of political races (Vaughan Williams and Reade, 2015). Prediction markets have even performed well when illiquid, as was the case in the corporate prediction markets studied by Cowgill and Zitzewitz (2015), and have also performed well when attempts have been made to manipulate prices, as was the case in the presidential betting markets studied by Rhode and Strumpf (2004) and Rothschild and Sethi (2015). In addition, and of particular relevance to our setting, prediction/betting markets have been found to accurately digest information on events (goals) almost immediately (Croxson and Reade, 2014).

There are two broad theoretical reasons why Twitter activity might forecast outcomes even after controlling for betting market prices. Firstly, there may be limitations to the accuracy and amount of information that finds its way into betting prices. A large literature in finance has documented the limits of arbitrage, whereby markets may be inefficient due to, for example, risk-aversion or borrowing constraints on the part of informed arbitrageurs (see Gromb and Vayanos, 2010, for a survey). Social media, therefore, may provide an alternative repository for information that spectators do not have the risk-appetite or resources to impound into betting prices. Alternatively, information may not be held by an individual (who, of course, has the option to bet), but may instead be dispersed amongst the crowd. Since Galton's (1907) famous ox-weighing survey at the West of England Fat Stock and Poultry Exhibition, social psychologists have documented many instances when aggregated crowd forecasts have proved more accurate than individual, sometimes expert, forecasts. As put by Larrick *et al.* (2012), "combining judgments takes individual imperfection and smoothes the rough edges to isolate the collective's view of the truth. Or, to put it more mathematically and mundanely, averaging cancels error". Social media is a platform to harness the views of a large crowd. On the other hand, it would be remiss not to mention the possibility that

Twitter activity could be an erroneous predictor of outcomes, if Twitter is used to spread rumours as in the model of Van Bommel (2003).

In fact, we find that Twitter activity does contain accurate information not in betting prices. We measure the aggregate tone of all Tweets – using the micro-blogging dictionary of Nielsen (2011) – for each team, in each second, of each match during the 2013/14 season.⁵ We find that an overall positive tone in Tweets indicates that the team in question is 3.39% more likely to win than contemporaneous betting prices imply. Further tests suggest that much of this effect is indeed due to the ‘wisdom of crowds’, as the tone of a randomly selected Tweet does not have the same predictive power as the aggregated tone of all Tweets. In addition, we find that much of the informational content of Tweets can be found in the aftermath of large events. A positive tone in aggregate tweeting indicates that a team is 8.12% more likely to win than contemporaneous betting prices imply in the immediate aftermath of goals and red cards. In short, we find that Twitter content is most useful for harnessing the wisdom of crowds in the interpretation of new information.

We also examine the speed with which social media information is subsequently incorporated into betting prices. We find that stale Tweets still have information not in betting prices 60 seconds after the Tweet and later. Only when a Tweet follows a significant market event – such as a goal or a red card – is this social media information quickly digested by the market (approximately 10 seconds later). In other words, when social media output is not a response to a salient market event, its informational content often does not find its way into betting prices in a timely manner.

What are the returns to a strategy that exploits social media information? This will give us a rough idea of the magnitude of fundamental information that is embedded in social media content, but not in betting prices. To calculate returns, we first need an implementable strategy, which requires us to take two factors into account. Firstly, as evaluated in Brown and Yang (2015), Betfair operate a speed bump which creates a 5-9 second delay between the time at which an order is submitted and the time it is logged on the exchange. (This is to ensure that bettors in the stadia cannot adversely-select bettors at home watching the match

⁵We label our measure of Tweet content the ‘tone’ of the Tweet, as the label ‘sentiment’ is more often than not used to capture non-fundamental information, e.g. as in Edmans *et al.* (2007).

with a delay). Secondly, we need to allow time for a hypothetical arbitrageur to execute their trade; this means scraping the information from Twitter, calculating the tone of the Tweets (if applicable), and then placing an automated trade on Betfair. Once we account for these two factors, we find that a simple strategy of betting when Tweets on a team are positive yields average returns of 2.28% from 903,821 bets. These returns – for an investment duration of no more than 90 minutes – compare very favourably with average returns for all bets of -5.41%, and indicate that the marginal information contained in social media data is substantial.

Before extrapolating these results to other settings, there is an important caveat to bear in mind. Social media in our setting is accompanied by a betting/prediction market. There may be other settings where such a market does not exist; in which case the forecaster must confront the original problem of constructing probability forecasts from social media output. Nevertheless, our results suggest that social media is a useful forecasting tool in fast-moving environments – at least when combined with a prediction market – and is particularly helpful in the interpretation of new information.

The rest of the paper is structured as follows. In Section 2 we briefly summarize more of the related literature. In Section 3 we introduce the Twitter and Betfair data, and in Section 4 we conduct our analysis. Section 5 concludes.

2 Related Literature

Our paper is most closely related to the work of Vaughan Williams and Reade (2016). Vaughan Williams and Reade combined Betfair and Twitter data in order to examine whether breaking news reached betting market prices or Tweets first. Their focus was the ‘bigotgate’ incident in 2010, when the then U.K. Prime Minister Gordon Brown was caught off-camera insulting a prospective voter. News of that incident broke on Twitter first, and it was a number of hours before the news was incorporated into betting market prices. In the same vein, Asur and Huberman (2010) show that Twitter activity is a better predictor of box-office revenues, for 24 different movies, than market-based predictors. In our study we have many more significant events (goals/red cards) to analyse, and also many more (match) outcomes

to forecast than in these two papers.

Our work is also related to the finance literature on the use of social media to predict stock returns. Chen *et al.* (2014), Sprenger *et al.* (2014), and Avery *et al.* (2015) find that opinions on three different social media sites – Seeking Alpha, Twitter, and CAPS respectively – predict subsequent cross-sectional stock returns, and, in the case of Seeking Alpha, posts also predict earnings announcement surprises. Zhang *et al.* (2011) find that Twitter content can even predict returns at the aggregate index level. The problem with such analysis, however, is that returns are only a proxy for fundamental information, and any effect is inevitably sensitive to the time-horizon chosen. The ability of social media to predict returns may be solely due to its capture of short-run sentiment, and not due to fundamentals. Furthermore, if market participants believe that social media can be mined for asset-relevant information, then the belief becomes self-fulfilling, as their buying and selling pressure drives asset prices and returns. Even earnings announcements – which may be less vulnerable to such endogeneity issues – are still only a noisy, and, to a degree, manipulable, measure of asset fundamentals. In our paper we are only interested in forecasting fundamentals, and for this we have a very clean measure: the outcome of the match.

3 Data: Twitter and Betfair

The Tweets for our analysis were provided by Twitter as part of their #datagrants programme. We were granted access to the data via Gnip, a firm that is now a subsidiary of Twitter. Gnip provide both real-time access to Tweets, and also historical sweeps based on hashtags and tweeter IDs. We obtained Tweets for every English Premier League soccer match in the 2013/14 season. Many different hashtags are used for each team. For example, for Liverpool Football Club, the hashtags include #lfc and #ynwa (where ynwa stands for You'll Never Walk Alone). In addition to sourcing according to hashtags, we also obtained Tweets, and apportioned them to a team, based on the tag of well-known twitter accounts. In the case of Liverpool again, these include @thisisanfield, @liverpoolfc and others. Tweets on soccer teams can occur throughout the week, but we focus on Tweets that occur as the team in question is playing. There are some examples where more than one team is tagged in a

particular Tweet. For example, when Liverpool play Chelsea Football Club you may observe a hashtag #lfcvfc. These Tweets are dropped from our analysis as it is unclear which team the Tweet is focused upon. Our sweep of historical Tweets includes original Tweets, and the retweeting (sharing) of an original Tweet. Including retweets, we have 13.8 million Tweets in our sample.

Table 1 provides summary statistics on the number of Tweets per second. This is measured at the team-second level, so there are two observations, one for each team, each second. On average a team receives 2.6 Tweets per second. This distribution is highly positively skewed, with one instance of a team receiving 264 Tweets in a given second. Of these a slight majority are original, with an average number of original Tweets of 1.37 per team per second. (Retweets come in at an average of 1.23 per team per second). In the middle panel of Table 1 we also summarize an indicator variable, equalling 1 if there was at least one Tweet for that team in that second, and 0 otherwise. A look at this indicator variable tells us that teams receive at least one Tweet 44.1% of the time. A team receives an original Tweet 36.9% of the time, and a retweet 29.6% of the time. It is important to note that we only consider the inplay period (as matches are being played), so attention on these teams at that time is typically quite high. Having said that, the statistics we describe encompass all teams in the 372 matches in our sample, not simply the matches that are televised. Some matches, of course, receive much more attention on social media than others.

We are implicitly assuming that tweeting about a particular team is a positive signal about a team's prospects. In each match, the social media user has the choice of which team to tag in their Tweet. Tagging one team, but not the other, may suggest that the user rates the team's prospects. This assumption depends, of course, on the content of the Tweet. Tweeters may be using the Tweet to voice their dissatisfaction with the team's performance, and predict that they will lose the match. Therefore, we cannot argue that a Tweet – tagged to a team with their hashtag – is equivalent to a buy signal on a stock message board. With this in mind, we use the Nielsen (2011) dictionary for micro-blog content to establish the tone for each team each second. As mentioned in the Introduction, we label our resultant measure of Tweet content the 'tone' of a Tweet, as one alternative label, 'sentiment', often implies that the information is unrelated to fundamentals. It is precisely these fundamentals

(match outcomes) that we want to capture. Within the Nielsen dictionary, the majority of positive and negative words receive a score of 2 and -2 respectively. The dictionary also captures obscenities, allowing for scores of -4 or -5.

In Table 1 we present summary statistics on this tone measure. We calculate the aggregate tone for each team in each second; this is the sum across all Tweets (which is itself the sum of all words within the Tweet). As captured in the top panel of Table 1, aggregate tone is, on average, positive. There is, however, great variation across teams and across time. Aggregate tone ranges from -144 to +471. We also create an indicator variable, equalling 1 if aggregate tone is positive and 0 otherwise. From here we can see that 20.9% of team-second observations have positive tone. Note, however, that there are only Tweets in 44.1% of team-second observations, so tone is positive in approximately 46.5% of cases when there are Tweets. Conditional on there being at least one Tweet, aggregate tone is negative in 18% of cases and is neutral (i.e. aggregate tone=0) in 35.5% of cases. This variation justifies our reluctance to state that a Tweet tagged to a given team is an unambiguously positive signal about the team's prospects.

Perhaps the best way to describe this dictionary is to give a few examples. In April 2014, Liverpool played Chelsea in an important match for Liverpool's hopes of winning the Championship. At 0-0 in the first half, the Liverpool captain Steven Gerrard slipped, allowing the Chelsea forward Demba Ba to run through and score. The reaction on Twitter ranged from positive for Chelsea:

Hahahahahaha yes demba ! Chin up liverpool #CFC

This received a score of +1 in the Nielsen dictionary. Others focused more negatively on the outcome for Liverpool:

WTF #YNWA

This Tweet produces a tone score of -4, but is topped by the obscenities (censored in this paper but not on Twitter) in the next negative Tweet which gets a tone score of -9:

F*** u Steve u c*** #lfc

Like any dictionary, there will be instances when the meaning of the Tweet is not accurately captured. For example, the dictionary does not detect the sarcasm in the following Tweet:

Captain fantastic and all that... #LFC

This last Tweet got a score of +4.

The dictionary provided by Nielsen (2011) is appealing for our research setting, as it is primarily intended to classify micro-blog output. This means that it captures colloquialisms, such as ‘WTF’, and obscenities, which are often used in online soccer discussions. We do, however, recognise the problems that may arise when using a general dictionary in a specific context. As Loughran and McDonald (2011) illustrated in their study of financial text, words can have very different connotations in different settings. (One example of theirs was that the term ‘liability’ is negative in the majority of contexts, but less so in a financial context). An alternative is to devise our own dictionary specific for Twitter conversations on soccer. Our concern with this approach is that this inserts the researcher more closely into the data-generating process, and perhaps our classifications would subconsciously be biased in the direction of our prior hypotheses. Another point we can make is that, in this type of research, the success of a dictionary in capturing tone is, in part, revealed by the extent to which this tone tells us something that betting prices do not. As we will see in our later analysis, the dictionary of Nielsen is quite an informative predictor of match outcomes.

We marry our social media data with betting price data from Betfair, a U.K. betting exchange. The exchange operates as a standard limit order book of the type used by most financial exchanges. Bettors can place limit orders, which act as quotes for other bettors, or place market orders, which execute at prices currently quoted by others. Bettors can wager that a particular team will win (via a ‘back’ bet), or bet that a team will lose (via a ‘lay’ bet). We obtained Betfair limit order book data from Fracsoft, a third-party provider of historical data. This data include the best back and lay quotes (and associated volumes), measured each second, throughout 372 matches in the 2013/14 season. 8 matches are missing from the Fracsoft database, so we discard the Twitter data on these.

Our main measure from the betting data is the implied win probability. This is defined as

$(1/B_t + 1/L_t)/2$, where B_t are the best back quotes at time t , and L_t are the best lay quotes at time t . This is the midpoint of the back-lay (bid-ask) spread. We use this measure because we want to see whether Tweets can add information (or misinformation) to the probability forecasts produced by the betting market. To be specific, we are asking, does the presence of a Tweet for a team indicate that the team is more or less likely to win than betting market prices imply? (Later, when we see that a Tweet indicates that a team is more likely to win, we will use back bet prices to construct returns). The average implied win probability from the midpoint of the spread is 0.377, with a range from 0.001 to 0.985. This figure is higher than 0.33, as we exclude draws from the analysis (as Tweets are seldom tagged #draw), and draws occur less frequently than the other two outcomes (home and away wins). In Table 1 we also summarise an indicator variable equalling 1 if the team ultimately won, which averages 0.399. On the face of it, it may seem that there are positive returns to a blanket strategy of betting on all teams, as actual win probabilities are higher than implied win probabilities. However, this is partly because we have 802,445 missing observations for implied win probability – in instances when there are not quotes on both sides of the book – and partly because, at this stage, we are using the midpoint of the spread to calculate implied win probabilities. As we will show in our later Table 9 analysis, when we use back bet prices there are, on average, negative returns for a blanket strategy of betting on all teams.

To illustrate the Twitter and Betfair data together, we created Figure 1, which describes the aforementioned match between Liverpool and Chelsea in April 2014. In the left panel we have betting prices and the number of Tweets for Liverpool each minute of the match. In the right panel we have the same information for Chelsea. (We will examine prices and Tweets each second in our later analysis, but that level of granularity is too detailed for a plot). The Gerrard slip, and Demba Ba goal, occurred in the 48th minute (in stoppage time at the end of the first half). There is little indication from the number of Tweets that this goal was anticipated, but it did set off a spike in Tweets for both teams (albeit more for the scoring team Chelsea). Similarly, when we plot aggregate tone, instead of the number of Tweets, in Figure 2, we find a similar pattern. Aggregate tone did not appear to predict the goal, but tone certainly spiked for the scoring team Chelsea, who went on to win the match. There was a more modest uptick in tone, some of it perhaps defiant, for Liverpool, the team that

conceded. In the next section, we will exploit the full granularity of the data to establish whether this pattern applies across the full set of goals and matches.

4 Analysis

Throughout the analysis section we predominantly estimate an equation of the following form:

$$y_i = \beta_0 + \beta_1 x_{it} + \beta_2 z_{it} + \epsilon_{it} \quad (1)$$

y_i is an indicator variable, equalling 1 if team i won the match, x_{it} is the implied win probability of team i winning as measured from the odds at time t , z_{it} is an indicator variable capturing some element of Twitter behaviour for team i at time t , and ϵ_{it} is an error term. This equation, minus the z_{it} term, is commonly referred to as the Mincer-Zarnowitz regression (Mincer and Zarnowitz, 1969), and is often used in the estimation of the favourite-longshot bias (see Vaughan Williams and Reade, 2015, for example). We cluster our standard errors at the selection level (e.g. Liverpool to beat Chelsea on April 27th 2014). This means that there are two clusters per match, one for each team.⁶ We estimate the equation by OLS, but the results are qualitatively similar if we use probit or logit models. We are most interested in the β_2 coefficient, as significance here would indicate that social media content can improve on the forecasts produced by the betting market. Put another way, a significant β_2 coefficient would indicate that the betting market forecasts are inefficient.

In our first regression, z_{it} is an indicator variable equalling 1 if there is at least one Tweet for that team in that second. We can use this first analysis as a baseline with which to compare our later analysis of Tweet tone; comparing the two will tell us how much information is in the actual content of the Tweets (rather than just the tagging). We use an indicator variable, rather than count the number of Tweets in a second, as the latter would lead to a substantial

⁶As robustness exercises, for our main analysis we also clustered standard errors by match, by time (second), and used Newey-West autocorrelation-consistent standard errors up to lag 60. Our upcoming results are all robust to these alternative specifications. In fact, the standard errors are smaller for each of these alternatives than for the results we report with clustering at the selection level.

number of predicted values outside the unit interval (recall that there was one instance of 264 Tweets for a team in one second), and potentially bias our estimates. The results of our first regression can be found in Table 2. Firstly, there is little evidence of a favourite-longshot bias, as the coefficient associated with implied win probability is insignificantly different from 1.⁷ The coefficient associated with the Tweet indicator suggests that teams win more often than betting prices imply when there is a Tweet, but the difference is not statistically significant. In other words, for the full inplay period, there is little evidence that the mere presence of a Tweet conveys information not in market prices.

At present we have lumped together both original Tweets and retweets (the sharing of original Tweets). One way to break down our analysis is to separate these two types of Tweet out. It is possible that while original Tweets contain information on match outcomes, retweets do not. Alternatively, it may be that a Tweet is not relevant until validated by the another user who shares the information. In Table 2 we therefore repeat our first regression, but this time use an indicator variable for whether there was an original Tweet or a retweet, in regressions 2 and 3 respectively. We find little difference between the predictive power of original Tweets and retweets; neither have a significant effect across the full inplay data. There appears to be as much (or as little) information contained in the sharing of a Tweet as there is in the decision to tweet in the first place. This holds for all of our analysis, so, to save on space, for the remainder of the paper we only describe the results for all Tweets, not original Tweets and retweets separately.

Our final analysis in Table 2 is on the tone of Tweets. As outlined in Table 1, there is great variation in the tone of a Tweet, as measured by the Nielsen (2011) dictionary. We would expect that if a Tweet is to convey information not already in betting prices, then much of this information will be in the content of the Tweet and not simply in the existence of a Tweet. With this in mind we regress our win indicator variable on the implied win probability, and an indicator variable equalling 1 if aggregate tone – across all Tweets for that team in that second – is positive. Positive tone occurs in 20.9% of team-second observations and, conditional on there being at least one Tweet for that team, 46.5% of

⁷Betting exchange prices typically exhibit less of a favourite-longshot bias than bookmaker prices, as shown in Smith *et al.* (2006).

team-second observations. We find that positive tone does predict match outcomes in a way not fully captured by betting prices. Teams with positive tone tweeting in a given second are 3.39% more likely to win than a team without positive tone, after controlling for the implied win probability of contemporaneous betting prices. In other words, tone extracted from aggregate tweeting is a useful predictor of outcomes across the full inplay time period.

We summarise the main results from Table 2 in Figure 3. We estimate local polynomial regressions (of degree 3) of the relationship between the win indicator and the implied win probability (inferred from the midpoint of the bid-ask spread). The local polynomial estimation does not impose the linear structure assumed by our earlier regressions. In each panel we display a 45 degree line – which represents a perfectly efficient market – for comparison. (Note that we are using the midpoint of the bid-ask spread at this stage, so deviations from the 45 degree line do not demonstrate that there are profit opportunities; this will come later with our analysis of returns). In the two panels we can see that teams win slightly more often than prices imply when accompanied by a Tweet, and win even more often when accompanied by positive tone in aggregate tweeting. We display 95% confidence intervals which, due to the very high number of observations, are very narrow.

One issue is that favourites tend to receive more Tweets, and indeed more positive Tweets, and therefore we may be confounding the information in social media content with a simple favourite-longshot bias. Although the β_1 estimate in Table 2 does not indicate a bias, it does appear from Figure 3 that there are some non-linearities in the data. Our solution is to conduct nearest neighbour matching, as described in Todd (2008) and used in a betting market context by Brown and Yang (2017). We compare the average win indicator for observations where aggregate tone is positive to observations where the aggregate tone is negative. Importantly these ‘nearest neighbours’ are matched by minimising the difference in implied win probability. In other words, we compare the actual win probability of a team with an implied win probability X *and positive tone*, to the actual win probability of a team with implied win probability X *without positive tone*. We analyse a random 1/10 of the sample in this analysis, as it is too computationally intensive to use this methodology on the full sample. The results are found in Table 3. We find that teams with positive tone Tweets win 2.84% more often than teams with the same implied win probability but without positive

tone. This is slightly lower than our Table 2 estimates but still quite substantial. Indeed, with this specification we also find that a simple tagging of team with a Tweet indicates that a team is significantly more likely to win than prices imply. (For nearest neighbour matching we cannot cluster by selection, which partly explains this result). In short, social media content can predict match outcomes even after non-parametrically accounting for betting prices.

Is this evidence of the ‘wisdom of crowds’, or an indication that individuals, on their own, produce information on Twitter that is not in betting prices? To answer this question, in each second with more than one Tweet on a team we randomly selected one Tweet. We then created an indicator variable equalling 1 if the tone of that random Tweet was positive according to the Nielsen dictionary. Importantly, there is little chance of spillovers in information within each second, as it is unlikely to be technologically feasible to parse a Tweet from someone else and then post a new Tweet in the same second. In Table 4, we now regress our win indicator variable on the implied win probability, and this new indicator variable. We find that the tone of random individual Tweets does predict match outcomes. Specifically, a team is 1.88% more likely to win than betting prices imply if the tone of this random Tweet is positive. This suggests that there are indeed limits to the accuracy and amount of information that individuals impound into betting market prices. However, the effect is only a little over half the size of the aggregate positive tone effect, shown again in regression 1 of Table 4 for illustration. This suggests that the value of social media content for forecasting is, in part, due to the ability of such platforms to harness the ‘wisdom of crowds’.

When is the predictive power of social media at its most prominent? Is social media quicker to reveal information (goals, red cards) than betting markets? Or does social media help in the interpretation of information? We begin by analysing whether social media output reveals events (i.e. ‘breaks news’) quicker than betting markets. Note that Vaughan Williams and Reade (2016) found that Twitter broke news of the ‘bigotgate’ scandal before the betting markets. Firstly, we identify a series of positive match events in our data. Positive events can be scoring a goal, receiving the award of a penalty, or the opposition having a player sent off. We classify the occurrence of a positive event if the betting market is suspended – as it is after goals, penalty awards, and red cards – and subsequently reopens with a higher

price (implied win probability) for that team. If we find that Twitter activity predicts such events in the very near future (i.e. the next 5 seconds), it can be said that Twitter is faster to break news than the betting market. In the first regression of Table 5 we regress an indicator variable equalling 1 if, according to betting prices, a positive event occurs for team i in the next 5 seconds, on an indicator variable equalling 1 if there was at least one Tweet for that team (in the 5 seconds preceding the event). We find that the presence of at least one Tweet for a team does not predict positive events up to 5 seconds ahead. In fact, positive tone is actually a negative predictor of such positive events.⁸ In short, there is little to suggest that social media breaks news faster than betting markets.

In another series of tests, we then see whether social media activity predicts positive events in the next 60 seconds. With this horizon, we are not testing whether social media breaks news faster than betting markets, but whether social media users can predict imminent events. Again, we find that the presence of at least one Tweet for a team does not predict events up to 60 seconds ahead, and positive tone for a team is again a negative predictor of the occurrence of these events. Any ability of social media to predict match outcomes does not therefore stem from an ability to break news quicker, or predict the imminent occurrence of significant events such as goals and red cards.

Perhaps social media, and the tone that can be extracted from its content, is more useful in the interpretation of news, rather than in the breaking of news. There is evidence in the finance literature that markets sometimes over-react to information and sometimes under-react (Brooks *et al.*, 2003, Chan, 2003). Indeed, there is evidence of both over-reaction and under-reaction to news (goals) in precisely this market (Choi and Hui, 2014). Perhaps social media output can help to interpret new information and ameliorate these inaccurate market reactions. In Table 6 we repeat the regressions of Table 2 but this time focus on the time period either up to 5 seconds after a market suspension, or up to 60 seconds after a market suspension. The aim is to establish whether social media content can help in the interpretation of news in its immediate aftermath (up to 5 seconds afterwards) or soon after

⁸We also find that positive tone is a negative predictor of negative events – such as the concession of a goal – so there is nothing to suggest that positive tone indicates a bet is over-priced and negative events are about to occur.

(up to 60 seconds afterwards). We find that the existence of at least one Tweet for a team conveys information in the immediate aftermath (regression 1), but a Tweet up to 60 seconds after a market suspension does not (regression 2). Specifically, a team is 4.27% more likely to win, than betting market prices imply, if there is at least one Tweet tagged with that team in the immediate aftermath of the event (regression 1). Positive tone also suggests that a team is 8.12% more likely to win than betting prices suggest immediately after an event (regression 3). Positive tone conveys information even if the Tweet occurs up to 60 seconds after the event (see regression 4), though the magnitude of this effect is diminished relative to the immediate aftermath of the event (regression 3).⁹

We summarise the main Table 6 results in Figure 4. Figure 4 is a replication of Figure 3, except this time we only focus on the 5 seconds in the immediate aftermath of market suspensions. As detailed above, a team wins more often than betting prices imply, after a market suspension, when there is a Tweet for that team, and particularly when the aggregate tone of tweeting is positive.

In short, social media output aids in the interpretation of news. But is this because a Tweet signifies that a team is more likely to win, than prices suggest, after a positive event? This would imply that Tweets can point out when prices have under-reacted to information. Or does a Tweet signify that a team is more likely to win, than prices suggest, after a negative event? This would indicate that Tweets can point out when prices have over-reacted to information. Of course, Tweets could be useful for identifying mispricing in both positive and negative scenarios.

In Table 7 we break down our Table 6 analysis into instances when the news was positive and instances when the news was negative. (We only consider the 5 seconds after an event, as the effects in Table 6 were stronger in this time period). We find that Tweets predict a team is more likely to win than prices suggest after both positive and negative events.¹⁰ A

⁹In addition to this slice of the data – focusing on the periods immediately after goals and red cards – we also broke the whole sample down into home/away teams, and matches played on a Saturday at 3pm (not televised) and at other times (likely televised). We found no significant differences amongst these sub-samples so we do not report the results here.

¹⁰We have more observations following negative events (7,609) than following positive events (6,230). The main reason for this is that we classify positive/negative events by whether the price increased/decreased

team is 4.87% more likely to win than contemporaneous prices imply if there is at least one Tweet tagged with that team in that second, after a positive event (regression 1), and 5.31% more likely to win than prices imply if there is at least one Tweet tagged with that team in that second, after a negative event (regression 2). In other words, the mere presence of a tagged Tweet can inform us that the market has under-reacted to positive news, and/or over-reacted to negative news. Similar effects are found for positive tone tweeting, though such tone is more predictive of under-reaction to positive news.

One issue with our results so far is that we have ignored the Betfair speed bump. Betfair impose an artificial delay between the time at which an order is submitted and the time at which it is logged on the exchange. This delay, put in place to protect bettors watching at home from being adversely-selected by bettors in the stadia, varies from 5-9 seconds depending on the match. The implications of the delay, for our study, are that betting prices may not be able to fully reflect information immediately, because of the time it takes for an order to reach the exchange. While bettors can cancel orders immediately, new orders – which reflect new information – will take 5-9 seconds to reach the exchange.

With this in mind, perhaps it is more accurate to compare prices with the Tweets that occurred 5-9 seconds into the past, with the precise adjustment depending on the duration of the order processing delay. In Table 8 we do something along those lines. We regress an indicator variable, equalling 1 if the team won and 0 otherwise, on the implied win probability, and an indicator variable equalling 1 if there was an executable Tweet for that team $X+1$ seconds in the past. X is the duration of the order processing delay for that particular match, and an extra 1 second is added to simulate the time it might take to source data from Twitter and execute a trade algorithmically. In effect, we are checking whether there is a mispricing on Betfair, pointed out by activity on Twitter, that could generate abnormal profits. The analysis in Table 8 mirrors the results in Table 2. There is executable tone (revealed in Tweets $X+1$ seconds in the past) that could also indicate betting market mispricing. The similarity between the coefficients in Tables 2 and 7 create the view that the type of mispricing that Twitter reveals decays, but does not decay rapidly. For example, a team is 3.39% more

after a market suspension. If there are no quotes following a market suspension, we do not define the event as either negative or positive.

likely to win than contemporaneous betting market prices suggest if there is positive tone on Twitter (regression 4 Table 2), and is 3.38% more likely to win than prices at time $t + X + 1$ suggest if there is positive tone at time t (regression 2 Table 8).

A natural next step is to more generally assess the decay of social media information. We know that social media can still predict match outcomes – after accounting for betting prices – up to 10 seconds after the Tweet (or after the order processing delay has elapsed), but what about 60 seconds after the Tweet, for example? To answer this question, we regress an indicator variable, equalling 1 if the team won and 0 otherwise, on the implied win probability, and an indicator variable equalling 1 if there was a Tweet for that team precisely T seconds in the past. In Figure 5 we then plot these coefficients for $T = 0$ all the way to $T = 60$.¹¹ We display the results for all Tweets, and aggregate positive tone tweeting in the left and right panels respectively. In blue we use all Tweets, and in red we display only Tweets after a market suspension/event (such as a goal or a red card). We find that social media information after a goal/red card decays quite rapidly, with much of the content in prices after 10 seconds. However, there is little or no decay in social media information if the Tweet did not follow a major market event. What is more, the coefficients are still positive after 60 seconds, and with little sign of decline, suggesting that some social media information may not find its way into prices in a timely manner.

Our final step is to estimate the returns that might be available to an arbitrageur looking to capitalise on the mispricing we describe. This arbitrageur could extract information from Twitter, calculate the tone of the Tweets (if necessary), and then algorithmically execute a bet on Betfair. These bets will be executed at the *best back quote*, not at the midpoint of the back-lay spread as previously analysed. In Table 9 we present summary statistics on the returns to three strategies. (All strategies allow $X+1$ seconds for execution, where X is the duration of the order processing delay). The first strategy is to bet on all teams, in all seconds. This is the benchmark return and gives us an idea of the margins on the exchange. The second strategy is to bet when there has been at least one Tweet tagged with the team in question in that second. And the third and final strategy is to bet when aggregate tone for

¹¹Only cases where the match was still ongoing – and back and lay quotes were still available at $T = 60$ – are used, so the coefficients for $T = 0$ differ slightly from those displayed in Table 2.

that team in that second is positive. Betfair charge commission of 2-5% on profits within each market, with the rate depending on the historical activity of the bettor. We therefore present the returns to the four strategies for three different commission rates: 0% (a hypothetical rate), 2% (the minimum rate, available to the most active customers), and 5% (the maximum rate, available to a first-time user of the exchange).

There are substantial returns available for betting on the basis of positive aggregate tone. To take the highest commission rate of 5% as an example, the average returns to the positive tone strategy are 2.28% (from 903,821 bets). This compares with average returns of -5.41% across all 4.44 million bets. Given that the strategy returns are for, at most, a 90 minute investment, the magnitudes are quite striking. Of course, if we wanted to properly establish the economic significance of these returns, we would need to calculate the volume available for each of these bets, and use a model to predict the price impact of each of our trades. Nevertheless, even our back-of-the-envelope calculations illustrate the magnitude of information on social media that is not in betting prices.

5 Conclusion

The modern forecaster has a number of tools at their disposal. Two tools, in particular, have proved extremely popular: prediction markets and social media. We know that prediction markets generally lead to accurate forecasts, and outperform individual experts and polls in many settings. But does social media have anything to add? Can we combine probability forecasts from prediction markets with social media output to improve our predictions?

In this paper we analyse 13.8 million Twitter posts on English Premier League soccer matches, and compare them with contemporaneous betting prices available on Betfair, a popular U.K. betting exchange. We follow both throughout 372 matches in the 2013/14 season. We ask some simple questions: does social media activity predict match outcomes, after accounting for betting market prices? Does the aggregate tone of Tweets predict that a team is more likely to win than the market implies?

We find that Twitter activity predicts match outcomes, after controlling for betting market prices. Much of the predictive power of social media presents itself just after significant

market events, such as goals and red cards, where the tone of Tweets can help in the interpretation of information. To give an idea of the magnitude of social media information not in betting prices, a strategy predicated on social media activity can generate returns of 2.28%, for a series of less than 90 minute investments, on the betting exchange. In short, social media activity does not just represent sentiment or misinformation, but, if sensibly aggregated, can – when combined with a prediction market - help to improve forecast accuracy.

References

- Asur, S., Huberman, B., A., (2010). Predicting the Future With Social Media. Working paper.
- Avery, C., N., Chevalier, J., A., Zeckhauser, R., J., (2015). The ‘CAPS’ Prediction System and Stock Market Returns. *Review of Finance*, forthcoming.
- Brooks, R., M., Patel, A., Su, T., (2003). How the Equity Market Responds to Unanticipated Events. *Journal of Business*, 76, 109-133.
- Brown, A., Yang, F., (2015). Slowing Down Fast Traders: Evidence from the Betfair Speed Bump. Working paper.
- Brown, A., Yang, F., (2017). The Role of Speculative Trade in Market Efficiency: Evidence from a Betting Exchange. *Review of Finance*, 21, 583-603.
- Chan, W., S., (2003). Stock Price Reactions to News and No-News: Drift and Reversal after Headlines. *Journal of Financial Economics*, 70, 223-260.
- Chen, H., De, P., Hu, Y., Hwang, B., H., (2014). Wisdom of Crowds: The Value of Stock Opinions Transmitted Through Social Media. *Review of Financial Studies*, 27, 1367-1403.
- Choi, D., Hui, S., K., (2014). The Role of Surprise: Understanding Overreaction and Underreaction to Unanticipated Events Using In-Play Soccer Betting Market. *Journal of Economic Behavior and Organization*, 107, 614-629.

- Cowgill, B., Zitzewitz, E., (2015). Corporate Prediction Markets: Evidence from Google, Ford, and Firm X. *Review of Economic Studies*, forthcoming.
- Croxson, K., Reade, J., J., (2014). Information and Efficiency: Goal Arrival in Soccer Betting. *Economic Journal*, 124, 62-91.
- Edmans, A., Garcia, D., Norli, O., (2007). Sports Sentiment and Stock Returns. *Journal of Finance*, 62, 1967-1998.
- Galton, F., (1907). Vox Populi, *Nature*, 1949, 450-451.
- Gromb, D., Vayanos, D., (2010). Limits of Arbitrage. *Annual Review of Financial Economics*, 2, 251-275.
- Larrick, R. P., Mannes, A. E., Soll, J. B., (2012). The Social Psychology of the Wisdom of Crowds. In Krueger, J. I. (Ed.), *Frontiers in Social Psychology: Social Judgment and Decision Making* (pp. 227-242). New York: Psychology Press
- Loughran, T., McDonald, B., (2011). When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *Journal of Finance*, 66, 35-65.
- Mincer, J.,A., Zarnowitz, V., (1969). *The Evaluation of Economic Forecasts*, in “Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance”, pp. 1–46, *NBER*.
- Nielsen, F., A., (2011). A new ANEW: Evaluation of a Word List for Sentiment Analysis in Microblogs. *Proceedings of the ESWC2011 Workshop on 'Making Sense of Microposts*, 93-98.
- Rhode, P., W, Strumpf, K., S., (2004). Historical Presidential Betting Markets. *Journal of Economic Perspectives*, 18, 127-142.
- Rothschild, D., M., Sethi, R., (2015). Trading Strategies and Market Microstructure: Evidence from a Prediction Market. Working paper.
- Smith, M., A., Paton, D., Vaughan Williams, L., (2006). Market Efficiency in Person-to-Person Betting. *Economica*, 73, 673-689.

- Spann, M., Skiera, B., (2009). Sports Forecasting: A Comparison of the Forecast Accuracy of Prediction Markets, Betting Odds and Tipsters. *Journal of Forecasting*, 28, 55-72.
- Sprenger, T., O., Tumasjan, A., Sandner, P., G., Welpe, I., M., (2014). Tweets and Trades: the Information Content of Stock Microblogs. *European Financial Management*, 20, 926-957.
- Surowiecki, J., (2005). *The Wisdom of Crowds: Why the Many are Smarter Than the Few*. Abacus, 3rd edition.
- Todd, P., E., (2008). Matching Estimators. *The New Palgrave Dictionary of Economics*, Second Edition, Eds. Steven N. Durlauf and Lawrence E. Blume, Palgrave Macmillan.
- Van Bommel, J., (2003). Rumors. *Journal of Finance*, 58, 1499-1519.
- Vaughan Williams, L., Reade, J., J., (2015). Forecasting Elections. *Journal of Forecasting*, forthcoming.
- Vaughan Williams, L., Reade, J., J., (2016). Prediction Markets, Twitter and Bigotgate. *Kyklos*, forthcoming.
- Zhang, X., Fuehres, H., Gloor, P., A., (2011). Predicting Stock Market Indicators Through Twitter ‘I hope it is not as bad as I fear’. *Procedia - Social and Behavioral Sciences*, 26, 55-62.

Figures and Tables




Figure 1: **Number of Tweets.** Plots of implied win probability and the number of Tweets for each team, calculated each minute, in a match between Liverpool and Chelsea played at Anfield on the 27th April 2014.




Figure 2: **Tone**. Plots of implied win probability and aggregate tone for each team, calculated each minute, in a match between Liverpool and Chelsea played at Anfield on the 27th April 2014.

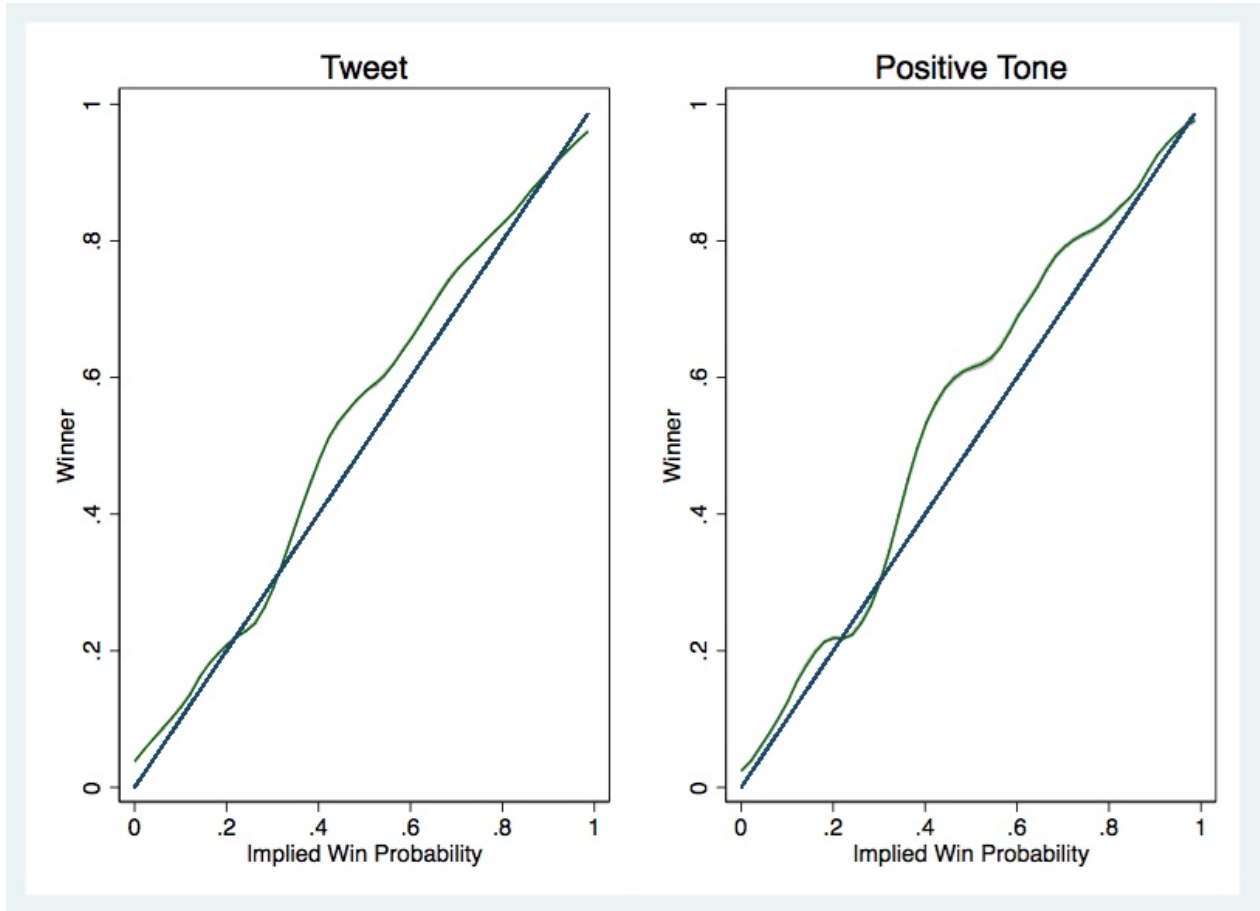


Figure 3: **Information Contained in Tweets.** Local polynomial regression estimations (of degree 3) of the relationship between the win indicator variable (equalling 1 if the team won the match), and the implied win probability (inferred from the midpoint of the bid-ask spread). The two panels include 1) seconds where a Tweet for that team was observed, and 2) just seconds where aggregate tone for that team was positive. A 45 degree line is added in blue for comparison, and 95% confidence intervals are displayed. These intervals are very narrow because of the high number of observations.

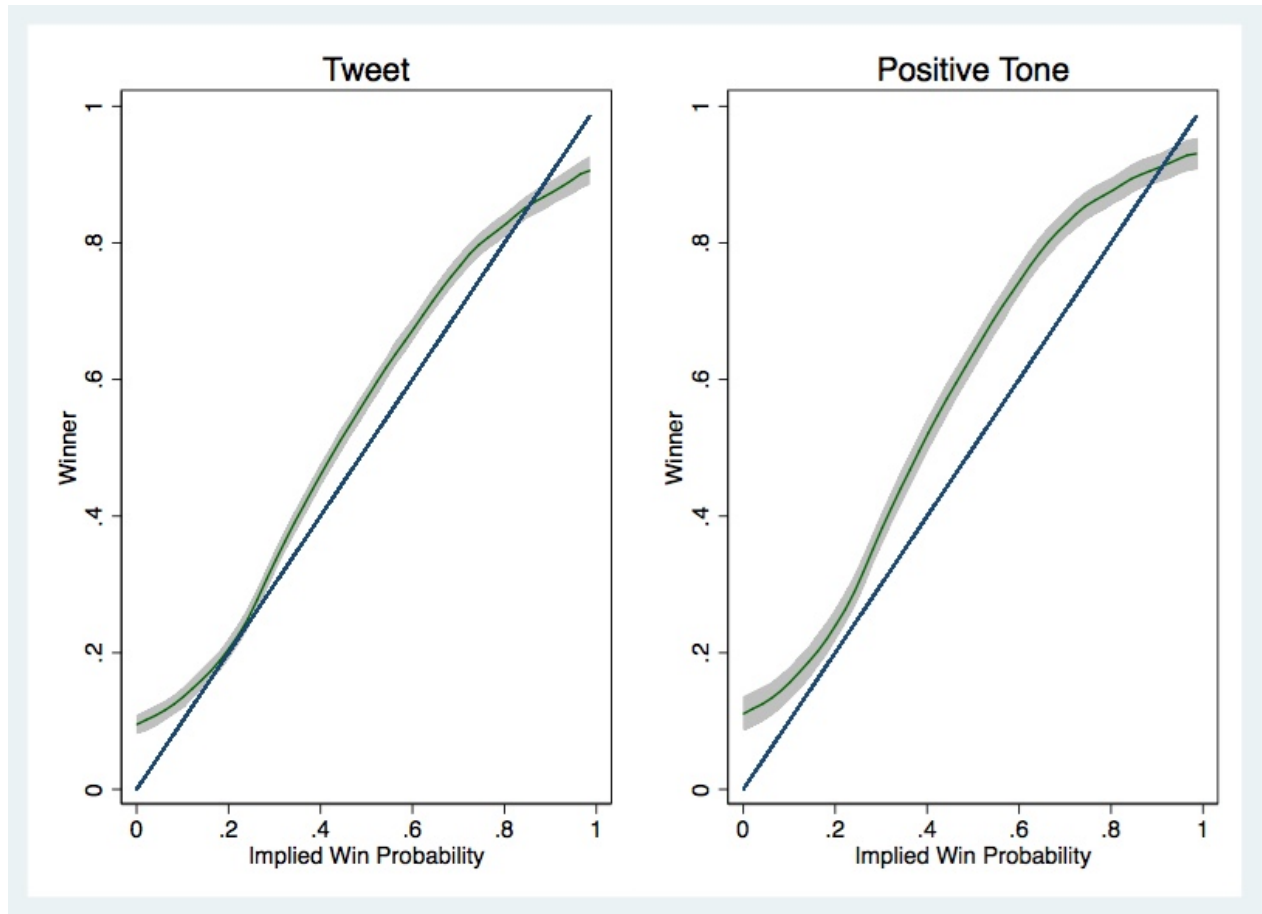


Figure 4: **Information Contained in Tweets after Events.** Local polynomial regression estimations (of degree 3) of the relationship between the win indicator variable (equalling 1 if the team won the match), and the implied win probability (inferred from the midpoint of the bid-ask spread). In this figure, data are only taken from the 5 seconds immediately after a market suspension (for a goal/red card etc.). The two panels include 1) seconds where a Tweet for that team was observed, and 2) just seconds where aggregate tone for that team was positive. A 45 degree line is added in blue for comparison, and 95% confidence intervals are displayed.

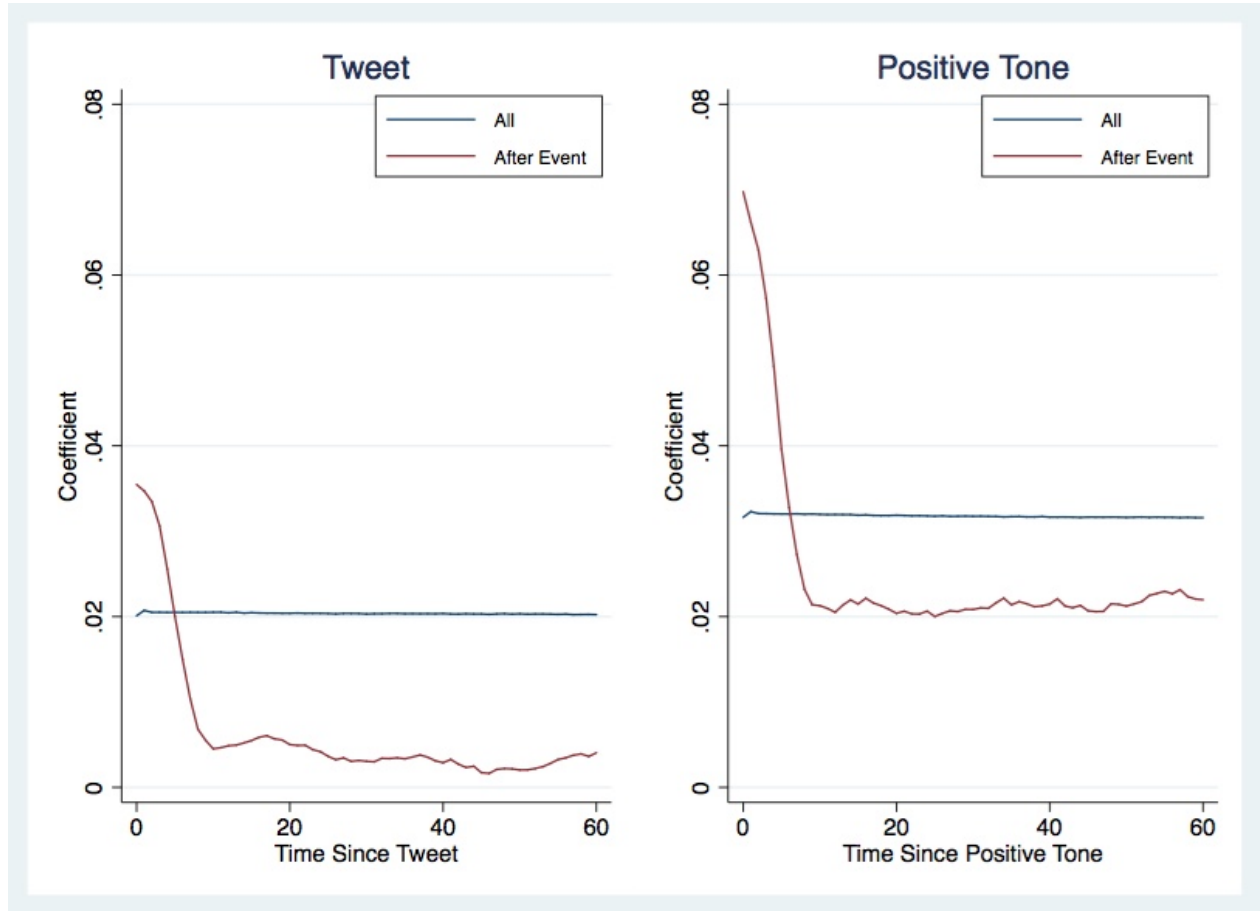


Figure 5: **Decay of Information in Tweets.** Estimates of the β_2 coefficient from Equation (1), this time varying the time (in seconds) since the Tweet(s). These plots indicate the decay in social media information as the Tweets become more stale. Results for all Tweets and aggregate positive tone tweeting are displayed in the left and right panels respectively. Separate lines are displayed for all inplay Tweets, and for Tweets in the aftermath of market suspensions/events (e.g. goals, red cards etc.)

Table 1: Summary Statistics

	(1)	(2)	(3)	(4)	(5)
Variables	N	mean	sd	min	max
No. of Tweets	5,249,502	2.608	7.346	0	264
No. of Original Tweets	5,249,502	1.377	3.878	0	163
No. of Retweets	5,249,502	1.231	3.942	0	181
Aggregate Tone	5,249,502	1.398	7.629	-144	471
Tweet	5,249,502	0.441	0.496	0	1
Original Tweet	5,249,502	0.369	0.483	0	1
Retweet	5,249,502	0.296	0.457	0	1
Positive Tone	5,249,502	0.209	0.406	0	1
Implied Win Probability	4,447,057	0.377	0.305	0.00101	0.985
Win	5,249,502	0.399	0.490	0	1

In the top panel, we display summary statistics on the number of Tweets, the number of original Tweets, the number of retweets, and the aggregate tone of all Tweets, each second. In the middle panel we display indicator variables equalling 1 if, each second, there is at least one Tweet, equalling 1 if there is at least one original Tweet, equalling 1 if there is at least one retweet, and equalling 1 if aggregate tone is positive. In the bottom panel, we display summary statistics on implied win probability – inferred from betting prices – and a win indicator variable. Draws are not included (Tweets are seldom tagged #draw).

Table 2: Main Analysis				
	(1)	(2)	(3)	(4)
Variables	Win	Win	Win	Win
Implied Win Probability	1.028*** (0.0280)	1.028*** (0.0280)	1.026*** (0.0281)	1.025*** (0.0281)
Tweet	0.0218 (0.0149)			
Original Tweet		0.0219 (0.0147)		
Retweet			0.0233 (0.0172)	
Positive Tone				0.0339** (0.0148)
Constant	-0.00355 (0.0161)	-0.00214 (0.0157)	2.10e-05 (0.0155)	0.000169 (0.0151)
Observations	4,447,057	4,447,057	4,447,057	4,447,057
R-squared	0.416	0.416	0.416	0.416

Robust standard errors – clustered at the selection level – in parentheses

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

The main analysis in the paper. An indicator variable equalling 1 if the bet/team won, is regressed on the implied win probability inferred from the betting prices, and four different indicator variables in four different regressions. These indicators equal 1 if, in that second, there is at least one Tweet, or equal 1 if there is at least one original Tweet, or equal 1 if there is at least one retweet, or equal 1 if aggregate tone is positive, respectively.

Table 3: Nearest Neighbour Matching

Variables	(1) Win	(2) Win
Tweet	0.0198*** (0.0011)	
Positive Tone		0.0284*** (0.0014)
Observations	446,543	446,543

Robust standard errors in parentheses

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

A repetition of the main Table 2 analysis, this time using nearest-neighbour matching methods. ‘Treated’ observations (e.g. Positive Tone==1) are matched with ‘non-treated’ observations (e.g. Positive Tone==0) according to their Implied Win Probabilities. A random sample, approximately equal to only a 1/10 of the full sample, is used for computational reasons.

Table 4: Wisdom of Crowds

	(1)	(2)
Variables	Win	Win
Implied Win Probability	1.025*** (0.0281)	1.031*** (0.0279)
Positive Sentiment	0.0339** (0.0148)	
Positive Sentiment: One Random Tweet		0.0188* (0.0109)
Constant	0.000169 (0.0151)	0.00258 (0.0150)
Observations	4,447,057	4,447,057
R-squared	0.416	0.415

Robust standard errors – clustered at the selection level – in parentheses

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

Analysis to examine whether social media forecasts match outcomes due to the aggregation of crowd wisdom. An indicator variable equalling 1 if the bet/team won, is regressed on the implied win probability inferred from the betting prices, and two different indicator variables in two different regressions. These indicators equal 1 if, in that second, aggregate tone is positive, or equal 1 if the tone of a randomly selected Tweet is positive, respectively.

Table 5: Predicting Positive Events

	(1)	(2)	(3)	(4)
Variables	+Event+5	+Event+60	+Event+5	+Event+60
Tweet	-4.09e-05 (5.67e-05)	-0.000429 (0.000335)		
Positive Tone			-0.000137** (5.55e-05)	-0.00116*** (0.000312)
Constant	0.00116*** (4.85e-05)	0.00724*** (0.000302)	0.00117*** (4.10e-05)	0.00729*** (0.000263)
Observations	5,294,202	5,294,202	5,294,202	5,294,202
R-squared	0.000	0.000	0.000	0.000

Robust standard errors – clustered at the selection level – in parentheses

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

Analysis to examine whether Tweets, for a particular team, predict positive events, either in the next 5 seconds (+Event+5) or in the next 60 seconds (+Event+60). Both '+Event+5' and '+Event+60' are indicator variables. These indicator variables are regressed on two other indicator variables, which equal 1 if, in that second, there is at least one Tweet, or equal 1 if aggregate tone is positive, respectively. The classification of positive events is based on the use of betting data, so if Tweets predict events in the next 5 seconds this implies that the event has already occurred and social media is actually 'breaking news' faster than the betting market.

Sample	Event-5	Event-60	Event-5	Event-60
	(1)	(2)	(3)	(4)
Variables	Win	Win	Win	Win
Implied Win Probability	1.056***	1.032***	1.049***	1.029***
	(0.0413)	(0.0307)	(0.0414)	(0.0311)
Tweet	0.0427**	0.0167		
	(0.0217)	(0.0198)		
Positive Tone			0.0812***	0.0300*
			(0.0195)	(0.0170)
Constant	-0.0244	-0.0104	-0.0177	-0.00714
	(0.0238)	(0.0195)	(0.0214)	(0.0167)
Observations	17,339	198,577	17,339	198,577
R-squared	0.343	0.462	0.346	0.462

Robust standard errors – clustered at the selection level – in parentheses

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

Analysis to examine whether Tweets aid in the interpretation of information. Subsamples of the data are examined, either up to 5 seconds after a market event/suspension (Event-5), or up to 60 seconds after a market event/suspension (Event-60). An indicator variable equalling 1 if the bet/team won, is regressed on the implied win probability inferred from the betting prices, and two different indicator variables in two different regressions. These indicators equal 1 if, in that second, there is at least one Tweet, or equal 1 if aggregate tone is positive, respectively.

Table 7: Processing +/- Information

Sample	+Event-5	-Event-5	+Event-5	-Event-5
	(1)	(2)	(3)	(4)
Variables	Win	Win	Win	Win
Implied Win Probability	1.062*** (0.0508)	1.039*** (0.0461)	1.055*** (0.0512)	1.035*** (0.0465)
Tweet	0.0487* (0.0268)	0.0531** (0.0247)		
Positive Tone			0.0957*** (0.0243)	0.0733*** (0.0229)
Constant	-0.0631** (0.0285)	-0.00188 (0.0263)	-0.0564** (0.0251)	0.00950 (0.0238)
Observations	6,230	7,609	6,230	7,609
R-squared	0.301	0.382	0.305	0.383

Robust standard errors – clustered at the selection level – in parentheses

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

Analysis to examine whether Tweets aid in the interpretation of positive or negative information, or both. Subsamples of the data are examined, either up to 5 seconds after a positive market event (+Event-5), or up to 5 seconds after a negative market event (-Event-5). An indicator variable equalling 1 if the bet/team won, is regressed on the implied win probability inferred from the betting prices, and two different indicator variables in two different regressions. These indicators equal 1 if, in that second, there is at least one Tweet, or equal 1 if aggregate tone is positive, respectively.

Table 8: Executable Tweets

	(1)	(2)
Variables	Win	Win
Implied Win Probability	1.028*** (0.0280)	1.025*** (0.0281)
Executable Tweet	0.0218 (0.0149)	
Executable Tone		0.0338** (0.0147)
Constant	-0.00356 (0.0161)	0.000166 (0.0151)
Observations	4,447,057	4,447,057
R-squared	0.416	0.416

Robust standard errors – clustered at the selection level – in parentheses

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

Analysis to examine whether information contained in Tweets still has predictive power even if we take the Betfair speed bump into account. The speed bump varies from 5-9 seconds. If the speed bump is X seconds, we link Tweets to betting prices X+1 seconds into the future. This allows time to cross the speed bump, and an extra 1 second to source Twitter data, calculate tone if applicable, and place a bet. An indicator variable equalling 1 if the bet/team won, is regressed on the implied win probability inferred from the betting prices, and two different indicator variables in two different regressions. These indicators equal 1 if there is at least one Tweet, or equal 1 if aggregate tone is positive, respectively.

Table 9: Betting Returns (%)					
	(1)	(2)	(3)	(4)	(5)
Variables	N	mean	sd	min	max
0% Commission					
All Returns	4,447,057	-2.504	368.1	-100	64,900
Tweet Returns	1,970,960	0.740	209.5	-100	7,900
Positive Tone Returns	903,821	4.908	195.9	-100	7,400
2% Commission					
All Returns	4,447,057	-3.667	361.1	-100	63,602
Tweet Returns	1,970,960	-0.365	205.8	-100	7,742
Positive Tone Returns	903,821	3.858	192.5	-100	7,252
5% Commission					
All Returns	4,447,057	-5.411	350.5	-100	61,655
Tweet Returns	1,970,960	-2.023	200.3	-100	7,505
Positive Tone Returns	903,821	2.282	187.4	-100	7,030

Summary statistics on betting returns. Returns are for *back bets* for three different strategies: bet on all teams at all times (every second), bet when there is a Tweet for that team, and bet when aggregate tone for that team is positive. The strategies allow $X+1$ seconds for a X second speed bump; the extra second is to allow for extraction of Twitter data and execution of the bet. Returns are displayed for 0% commission (a hypothetical rate), 2% commission (the minimum rate), and 5% commission (the maximum rate).