

What Can we Learn from Euro-Dollar Tweets? *

Vahid Gholampour

Eric van Wincoop

Bucknell University

University of Virginia

NBER

February 27, 2017

Abstract

We use 633 days of tweets about the Euro/dollar exchange rate to determine their information content and the profitability of trading based on Twitter Sentiment. We develop a detailed lexicon used by FX traders to translate verbal tweets into positive, negative and neutral opinions. The methodologically novel aspect of our approach is the use of a model with heterogeneous private information to interpret the data from FX tweets. After estimating model parameters, we compute the Sharpe ratio from a trading strategy based on Twitter Sentiment. The Sharpe ratio outperforms that based on the well-known carry trade and is precisely estimated.

*We gratefully acknowledge financial support from the Bankard Fund for Political Economy. We like to thank seminar participants at the University of Virginia, Bucknell University, the Federal Reserve Bank of Dallas, the Federal Reserve Bank of Richmond and the NBER Fall 2016 IFM Meeting for useful comments. We particularly thank Laura Veldkamp for very helpful comments.

1 Introduction

Asset pricing models with asymmetric information commonly assume that information is widely dispersed among traders.¹ Traders have different information about future events or may interpret the same information differently. All this information affects asset prices, which therefore in turn provide a (noisy) looking glass into the dispersed private information in the market. With the advent of the internet and social media, large numbers of people now go online to directly express their opinions about the direction of asset prices.² This leads to questions about the information content of these online opinions and the profitability of trading based on this information.³ In this paper we investigate what can be learned from Twitter by considering two and a half years of tweets that express opinions about the Euro/dollar exchange rate. This is a natural choice as Twitter has become a widely used platform to express opinions and the importance of private information for the determination of exchange rates is well established through the FX microstructure literature.⁴

The paper makes several contributions. First, we develop a “dictionary” based on financial lexicon used by traders in the Euro/dollar market to automate the interpretation of verbal tweets as positive, negative or neutral.⁵ This leads to a measure of Twitter Sentiment, which we consider separately for individuals with a lot of followers and few followers. Second, we document that little can be concluded about the information quality of Twitter Sentiment without imposing the structure of a model. While there is evidence of predictability of the *direction* of the exchange rate change from accounts with a lot of followers, Twitter Sentiment does not predict the *magnitude* of

¹See Brunnermeier (2001) for a review of the literature.

²Opinion surveys existed before, but they were infrequent (usually monthly) and limited in scope.

³There exists lots of anecdotal information suggesting that such information can be important. For example, on August 13, 2013, Carl Icahn, an activist investor, tweeted about his large position in Apple. As a result, the stock surged by over four percent in a few seconds. Almost two years later, on April 28, 2015, a data mining company obtained Twitter’s quarterly earnings and posted it on Twitter before the scheduled release time. Twitter’s stock plummeted by twenty percent and trading was halted by the NYSE.

⁴The seminal contribution by Evans and Lyons (2002) established a close relationship between exchange rates and order flow, with the latter seen as aggregating private information. Reviews of the FX microstructure literature can be found in Evans (2011), Evans and Rime (2012), King, Osler and Rime (2013) and Lyons (2001).

⁵We do not consider other currency pairs as a lot of the tweets are in different languages. But the overall method described here can certainly be applied to other languages and currency pairs.

the future exchange rate changes and implies Sharpe ratios from trading strategies that are very imprecise (could be large or small). Third, we then estimate a model with dispersed information in the foreign exchange market and use the results to compute the implied Sharpe ratio. We find the Sharpe ratio to be high, outperforming the popular carry trade strategy, and very precisely measured.

The key contribution of the paper is in the use of a model with a precise information structure. Without the use of a model, our data sample of 633 days is not long enough to draw strong conclusions about predictability and Sharpe ratios. Since exchange rates are notoriously hard to predict (are close to a random walk) and Twitter Sentiment is only directional, much longer data samples would be needed to draw conclusions about the information content of the tweets. By taking a stand on the information structure in the context of a specific model, we can learn much more from our short data sample. The model allows us to interpret many aspects of the Twitter Sentiment data, such as sentiment volatility, disagreement among agents, the relationship with current and future exchange rates and the different information quality of different groups of agents.

We therefore leverage the short sample by imposing structure to interpret the data. There remains uncertainty about the parameters of the model, which are estimated. This translates into uncertainty about the Sharpe ratio from a Twitter Sentiment trading strategy. But by using a wide variety of data moments, which are driven by a limited set of parameters that describe the information structure, we can obtain quite precise estimates of model parameters and therefore of the Sharpe ratio. Of course this all assumes that the model is correct. Our conclusions about the information content in Twitter Sentiment are critically dependent on this assumption. No model is exactly correct, but we show that the model cannot be rejected by the data and does a good job in matching moments in the data.

The model we use is an extension of a noisy rational expectation (NRE) model for exchange rate determination developed by Bacchetta and van Wincoop (2006), from here on BvW. Each period (day) agents receive new private signals about future fundamentals. As a result of noise trade the exchange rate does not reveal the aggregate of the private information, a common feature of NRE models. We extend the model of BvW to allow for two categories of agents, referred to as

informed and uninformed traders. They both receive private signals, but informed traders receive higher quality private signals.⁶ The information structure in the model is defined by the precision of the signals of both groups of agents, the horizon of future fundamentals over which they receive signals, the relative size of the informed group and the known processes of observed fundamentals and unobserved noise shocks.

There are no tweets in the model. Tweets are interpreted as the expression of an opinion by a subset of agents. Several steps are taken to connect the verbal tweets in the data with expectations of future exchange rate changes by individual agents in the model. First, we develop a large set of word combinations to classify tweets as positive (+1), negative (-1) or neutral (0) about the outlook for the Euro/dollar exchange rate. The word combinations are based on language typically used in the Euro/dollar market by traders. Second, in line with the theory, we separate the tweets into two groups, those with more than 500 followers and those with fewer than 500 followers. While we make no assumption in the estimation of the model about which group is the informed group, the results show that those with more followers have much more precise signals. Third, we use cutoffs for expectations in the model to obtain a theoretical Twitter Sentiment of +1, -1 or 0 for each agent. The cutoffs are such that the unconditional distribution across the three values in the model corresponds to that in the data.

After estimation with the Simulated Method of Moments, the model is used to evaluate predictability and Sharpe ratios. It should be emphasized that predicting exchange rates is no easy matter. It is well known since the results by Meese and Rogoff (1983a,b) that the exchange rate is close to a random walk. Engel and West (2005) show that reasonable estimates of the discount rate of future fundamentals in exchange rate models (close to 1) indeed imply a near-random walk behavior. The same will be the case in the model in this paper. Predictability will therefore always be limited, no matter the quality of the private information. But very limited predictability is sufficient to generate attractive Sharpe ratios.

Although we are not aware of other applications to the foreign exchange market, the paper

⁶The distinction between informed and uninformed agents is actually quite common in NRE models. A good example is Wang (1994). But in those models it is assumed that uninformed agents do not receive any private signals. We instead assume lower quality private signals, with the precision of the signals to be estimated.

relates to a literature that has used messages from social media and the internet to predict stock prices. Results are based on regressions of stock price returns on either “mood” states (like hope, happy, fear, worry, nervous, upset, anxious, positive, negative) or an opinion about the direction of stock price changes (along the line of positive, negative or neutral). Predictability is considered at most a couple of days into the future. Papers focusing on mood states, like Bollen et.al. (2011), Zhang et.al. (2011), Mittal and Goel (2012) and Zhang (2013), use an entire sweep of all Twitter messages, or random sets of messages, rather than messages specifically related to financial markets.⁷ Some of the literature prior to Twitter did focus specifically on financial messages. These include Antweiler and Frank (2004) and Das and Chen (2007), who use message boards like Yahoo!Finance, and Dewally (2003), who uses messages from newsgroups about US stocks. Evidence of predictability in most of these papers is limited at best, which is not surprising as they are based on short data samples of no more than a year.

The main difference in our approach is the use of a model to leverage the short sample and learn more about the information content of the messages. But there is another important difference as well. The literature has not employed financial jargon used by traders to classify messages. Most of the literature uses supervised machine-based learning classifiers that are not specific to financial markets at all. For example, the Naive Bayes algorithm is a popular classifier, which uses the words of a message to update the probabilities of various classification categories, based on a pre-classified training set. Tetlock (2007) has used a dictionary approach to consider the ability of verbal text to predict stock prices. But it is based on the Harvard IV dictionary that is not specifically related to financial news. Moreover, it is applied to WSJ articles as opposed to the diverse opinions expressed by a broad set of individuals on message boards and social media.

The remainder of the paper is organized as follows. In Section 2 we describe the Twitter data and methodology used to translate opinionated tweets about the Euro/dollar into positive (+1), negative (-1) and neutral (0) categories. We show that this measure of Twitter Sentiment is unable to predict future exchange rate changes in a statistically significant way and implies a Sharpe ratio that is very imprecisely measured. In Section 3 we describe the NRE model of exchange

⁷Mao et.al. (2015) uses an entire sweep of messages to search for the words “bullish” and “bearish” to classify tweets.

rate determination used to interpret the data. Section 4 discusses the empirical methodology and Section 5 presents the results. Section 6 concludes.

2 Data and Methodology

The objective is to translate daily verbal tweets that express opinions about the dollar/Euro exchange rate into a numerical measure of Twitter Sentiment (TS) that reflects expectations about the future direction of the exchange rate. After a discussion about possible biases related to the reasons for tweeting, we describe how we use a dictionary of financial lexicon to measure Twitter Sentiment. We then use the results to provide evidence on predictability and Sharpe ratios without any guidance from theory. We also compute a variety of moments that will be confronted with the theory in Section 5.

2.1 Why Individuals Tweet

Before we describe Twitter data and the steps of constructing Twitter Sentiment, a brief discussion of potential motivations by individuals for tweeting their outlook is in order. There are two potential ways in which such motivations can generate biases that can affect the analysis. The first bias occurs when individuals are motivated to tweet something that does not correspond to their actual beliefs. The second bias occurs when individuals are more or less likely to tweet in a way that is correlated with their outlook for the exchange rate.

The first bias is not likely to be much of a concern, for various reasons. First, it is hard to think of a reason to tweet the opposite of one's belief. Even if the objective of a tweet is to steer the market in a certain direction, there is little reason to steer it in a direction opposite to one's beliefs, especially if the individual has a stake in the outcome. Second, the market for the Euro-dollar currency pair is one of the most liquid financial markets in the world, so few individuals would be able to influence the exchange rate through malicious tweets. Finally, the self-provided user descriptions provide some information about the motivation for the tweets. A significant fraction of accounts with a lot of followers are controlled by individuals or businesses that provide investment research services. They occasionally tweet their future outlook to showcase their research and

gain more subscribers for their business. Businesses have no incentive to tweet an opinion that is in contradiction with their internal research because misleading the followers could hurt their reputation.

The second type of bias occurs if people are more likely to tweet when they have particularly strong beliefs about the direction of the exchange rate. This can bias the average opinion. If this is the case, the percentage of neutral tweets is inversely related to the total number of tweets. To test this, we separate the days into the bottom and top 25 percent in terms of the number of tweets. For days with few tweets (bottom 25 percent), we classify 31.7 percent as neutral tweets. For days with a large number of tweets (top 25 percent), 29.7 percent are classified as neutral tweets. If there is any bias, it is clearly not very large.⁸

2.2 Overall Approach to Computing Twitter Sentiment

It is important to describe in some detail how we translate verbal tweets into a numerical Twitter Sentiment. We have used Twitter's publicly available search tools since October 9, 2013, to download the tweets in real time every half an hour. Every tweet includes the user name, the number of followers of the individual who posted the tweet, as well as the exact time and date that the tweet was posted. We start with all Twitter messages that mention EURUSD in their text and are posted between October 9, 2013 and March 11, 2016. There are on average 578 such messages coming from distinct Twitter accounts per day, for a total of 268,770 tweets.⁹ However, the bulk of these messages do not include an opinion about the future direction in which the exchange rate will move. For example, many mention changes in the Euro/dollar exchange rate that have already happened or advertise a link to a web site discussing the Euro/dollar exchange rate.

The next step then is to look for opinionated tweets that express a positive, negative, or neutral outlook about the direction of the exchange rate. The exchange rate is dollars per Euro,

⁸We find the same if we separate the tweets into those with more than the average number of tweets (top 50 percent) and those with fewer than the average number of tweets (bottom 50 percent). We also get similar results when we separately consider tweets from those with a lot of followers and those with few followers, as we will do below.

⁹Here we count multiple tweets from the same account during a day as one EURUSD tweet.

denoted s_t in logs. A positive sentiment therefore means an expected Euro appreciation, while a negative sentiment indicates an expected Euro depreciation. A neutral outlook indicates a lack of conviction or dependency of the outlook on the outcome of a future event. Numerically we measure a positive outlook as +1, a negative outlook as -1 and a neutral outlook as 0. Unfortunately the tweets are not sufficiently precise to capture further gradations. The tweets are also not precise about the horizon of the expectation, an issue to which we return in Section 4 when discussing the connection to the theory.

In order to identify such opinionated tweets, and categorize them as positive, negative or neutral, we search for many different word combinations. A number of recent papers, such as Tetlock (2007) and Da, Engelberg, and Gao (2015), use Harvard IV-4 dictionary and word counting to conduct text analysis. This approach is shown to be effective in analyzing the content of financial articles and Google search words. However, the dictionary is not structured to capture the vocabulary used by investors. Since opinionated tweets about the exchange rate are usually posted by investors, there is a certain type of lexicon that is found in most of these tweets. We identify this lexicon by studying large numbers of tweets. We then go through several rounds of improving our dictionary of financial lexicon by comparing the results from the automated classification to that based on manual classification. We stopped making further changes when we found only very few errors after manually checking 5000 tweets. We describe this dictionary further below.

A day is defined as the 24 hour period that ends 12 noon EST. This corresponds well to our data on exchange rates as the Federal Reserve reports daily spot exchange rates at 12PM in New York. We allow only one opinion for each Twitter account on any given day to ensure that the measure of sentiment is not dominated by few individuals who express their opinion multiple times. We are not interested in intra-day price fluctuations. When there are multiple tweets from one account during a day, we only use the last tweet on that day.¹⁰

There are on average 43.5 such opinionated tweets per day, for a total of 27,557 during our sample. Therefore only about 8.5% of all tweets with the word EURUSD are opinionated tweets.

¹⁰On average 16.9% of tweets counted this way are from accounts from which multiple tweets were sent during a day.

The 27,557 opinionated tweets come from 6,236 separate accounts, implying an average of 4.4 tweets per account over the entire 633 day period of our sample. The opinions are therefore from a very diverse set of individuals as opposed to the same individuals repeating their opinions day after day. If the 27,557 tweets all came from individuals tweeting every day, there would have been only 43 separate accounts. We are clearly capturing a far more dispersed group of people expressing opinions.

2.3 Financial Lexicon

Tables A1 and A2 in Appendix A provide the list of all word combinations used to identify tweets as positive, negative or neutral. As can be seen, there are various ways that a tweet can be identified to be in one of the three categories. It might involve simply the combination of certain words, or the combination of some words together with the explicit absence of other words (positive and negative word combinations). In order to provide some perspective, Table 1 provides examples of tweets and how they are categorized. The words in the tweet used to identify them are underlined.

In Table 1, the first tweet under the positive category is identified as positive because investors use “higher high” to describe an uptrend in the price charts. In this example, using the individual words to extract the opinion could be misleading because the word “risk” might be interpreted as a negative word and the word “high” by itself is not enough to identify a positive opinion because investors use the word combination “lower high” to describe a downtrend. The first tweet under the neutral category is placed in this category because the words “might” and “sell” indicate lack of a definitive decision. Finally, the first tweet under the negative category is classified as bearish because the words “further” and “fall” indicate that the individual expects Euro to depreciate further against the dollar. We should note that the tweet mentions the word “bullish” which is a positive word. However, as mentioned earlier, we require the existence of certain words in the absence of other words to place a tweet in a category. In this example, the tweet is not identified as positive because a tweet should mention “bullish” and not mention “missing” to be placed in the positive category. This tweet is another example that highlights the significance of using word combinations instead of words to classify the opinionated tweets.

2.4 Separation by Number of Followers

We separate the opinionated tweets into those coming from individuals with at least 500 followers from those that have fewer than 500 followers. The idea is that those with more followers may be better informed investors. There are 4496 accounts with less than 500 followers and 2007 accounts with more than 500 followers.¹¹ Figure 1 shows the distribution of the number of followers, separately for accounts with more and less than 500 followers. For those with less than 500 followers, a large number has fewer than 50 followers. Of those that have more than 500 followers, 725 accounts have between 500 and 1000 followers, while 1282 accounts have more than 1000 followers.

When in Section 5 we confront the data to the theory developed in Section 3, we will see that the evidence strongly bears out the suspicion that individuals with a lot of followers are more informed. It should be noted that those with at least 500 followers are not famous people outside of the financial world, like movie stars who happen to tweet about the Euro/dollar exchange rate. Typical examples are brokers, technical analysts, financial commentators and people with research websites. One would expect these individuals to be well informed. From hereon we will simply refer to these two groups as informed and uninformed investors. The extent of the information difference will be documented in Section 5.

With this split, the daily average of opinionated tweets posted by informed and uninformed investors is respectively 21 and 22, so that we have a similar number of tweets in both groups. It may be the case that for example individuals with 1000 followers are even more informed than those with 500 followers, but splitting the data into more than 2 groups based on followers has the disadvantage of lowering the number of daily tweets per group. Figure 2 shows the distribution of daily tweets for both groups. It varies a lot across days. The standard deviation of the number of daily tweets is respectively 12.7 and 13.2 for the informed and uninformed. Since the average number of daily tweets of both groups is about the same, and there are fewer accounts from informed individuals, the average number of tweets per account during the sample is larger for

¹¹The total of these accounts is a bit larger than the 6236 mentioned above. This is because 267 individuals switch between both groups during the sample. We categorize the tweets each day based on the number of followers on that day.

informed than uninformed individuals, respectively 6.6 and 3.1.

We will denote the numerical Twitter Sentiment during day t by individual i from the informed group as $TS_t^{I,i}$. Analogously, when the individual is from the uninformed group it is denoted as $TS_t^{U,i}$. Figure 3 shows the distribution of the three values (-1, 0 and 1) that individual Twitter Sentiment of both groups takes across the entire sample. Especially for the informed group the percentage of negative values is a bit larger than the percentage of positive values. This is because the Euro depreciated by 21% during this particular sample.

2.5 Twitter Sentiment Index

For each of the two groups (informed and uninformed) we construct a daily Twitter Sentiment Index by taking the simple average of the numerical Twitter Sentiment across individuals during a day. We denote this as TS_t^I and TS_t^U for respectively informed and uninformed investors on day t . So

$$TS_t^j = \frac{1}{n_t^j} \sum_{i=1}^{n_t^j} TS_t^{j,i} \quad (1)$$

where $j = I, U$ is the group and n_t^j is the number of opinionated tweets on day t in group j . There are no tweets during 2 days in the sample for the informed group and 3 days for the uninformed. We set the index to 0 for those days. Figure 4 shows the distribution of the daily Twitter Sentiment Index for both groups.

2.6 Predictability of Exchange Rate by Twitter Sentiment

Tables 2, 3 and 4 report what we can learn from data on $TS_t^{I,i}$, $TS_t^{U,i}$, TS_t^I , TS_t^U and s_t about exchange rate predictability without any guidance from theory.

Table 2 reports directional moments, which capture how well tweets predict the subsequent direction of the exchange rate change. These moments are computed as follows. Consider a tweet by agent i of group j on day t . We look at how well it can forecast the direction of the exchange rate change over the next month, two months and three months. For example, $s_{t+40} - s_t$ is the change in the exchange rate over the next two months as there are about 20 trading days in a month. If $TS_t^{ji} = 1$ and the subsequent exchange rate change is positive (negative), we assign

the tweet a +1 (-1). Similarly, if $TS_t^{ji} = -1$ and the subsequent exchange rate change is positive (negative), we assign the tweet a -1 (+1). So +1 will be assigned if the direction is consistent with the Twitter Sentiment and -1 if the direction is inconsistent with the sentiment. A zero is assigned if $TS_t^{ji} = 0$, so that there is no directional opinion. We then take the average across all the tweets in the sample. A positive number suggests that the direction was more often correct than wrong, while a negative number suggests the opposite.

In order to evaluate if there is any information content in the tweets, we need to compare to what the moment would be if the agents guessed. To this end we conducted 1000 simulations over 633 days. The simulations are constructed such that on average the fraction of tweets that are zero corresponds to the average for each group, while on average the number of +1 tweets corresponds to the number of -1 tweets. The subsequent exchange rate change is unrelated to the number of +1 or -1 tweets as we consider the effect of guessing. In this case the mean of the moment is obviously 0. The standard error of the mean across the 1000 simulations is 0.0068 and does not depend on the horizon of the subsequent exchange rate change.

Based on this, the moments for the informed group are highly significant. For the one month, two month and three month subsequent exchange rate change the directional moment is respectively 3.3, 7.6 and 4.0 standard errors away from zero. This is strong evidence that there is valuable information content in the tweets of the informed group. The same cannot be said of the uninformed group, where the moments are slightly negative.

But directional moments themselves do not tell us if Twitter Sentiment is a good predictor of the actual magnitude of subsequent exchange rate changes. This is important if we wish to use Twitter Sentiment for trading purposes, where actual returns matter rather than the sign of the return. To this end we regress the exchange rate change $s_{t+m} - s_t$ on the Twitter Sentiment Index. We do this again for one month, two month and three month horizons ($m = 20$, $m = 40$ and $m = 60$), but also for a one-day horizon ($m = 1$). The one-day horizon is useful as we will report Sharpe ratios from a strategy based on Twitter Sentiment that is based on daily returns. We regress either on Twitter Sentiment TS_t^j at the start of the forecasting period or on Twitter Sentiment on each of the last 5 days before the forecasting period (TS_t^j through TS_{t-4}^j). The results are reported in Tables 3 and 4 for respectively the informed and uninformed groups. The regressions

are based on overlapping data intervals with Newey-West standard errors in parenthesis.¹² The bottom of the table reports the p-value associated with the F-test of zero coefficients on all lags of the Twitter Sentiment index.

There is no evidence of predictability as the coefficients are insignificant. The only exception is for the one-day horizon for uninformed agents when we use 5 days of Twitter Sentiment. But in that case the significant coefficients have the wrong sign. This lack of predictability suggests that the sample is too short to evaluate the ability of Twitter Sentiment to predict the subsequent exchange rate change without imposing structure through a model. The fact that the directional moments are significant for the informed group suggests that there is information content. But profitable trade based on Twitter Sentiment requires predictability of not just the direction, but also the magnitude of the exchange rate change.

2.7 Sharpe Ratio Trading Strategy Based on Twitter Sentiment

Table 5 reports the Sharpe ratio of a trading strategy based on Twitter Sentiment of both groups. For group i we take a position of TS_t^i dollars in Euro denominated bonds and $-TS_t^i$ dollars in dollar denominated bonds. We therefore go long in Euros when Twitter Sentiment is positive (expected Euro appreciation) and short in Euros when Twitter Sentiment is negative. It is important that positions are proportional to Twitter Sentiment as a stronger Twitter Sentiment implies more confidence that the exchange rate will go in a certain direction. The return from this strategy is

$$TS_t^i (s_{t+1} - s_t + i_t^* - i_t) \quad (2)$$

Here i_t and i_t^* are overnight dollar and Euro interest rates. We use overnight Libor rates for both, obtained from the St. Louis Fed.¹³

Table 5 shows that the Sharpe ratio is 1.09 for the informed group and -0.19 for the uninformed group. These are annualized Sharpe ratios based on daily returns. We annualize by multiplying by the square root of 250, the number of trading days in a year. The Sharpe ratio of 1.09 for the informed group is attractive, but it has a large standard error of 0.6. The 95% confidence

¹²Following Andrews (1991), we use $m + 1$ lags for the Newey-West standard errors.

¹³The dollar overnight Libor rate is the series USDONTD156N. The Euro overnight Libor rate is the series EURONTD156N.

interval is therefore very wide, ranging from -0.09 to 2.27. Based on the short data sample alone, without using a model, we can therefore say very little about the merits of trading based on Twitter Sentiment. This is consistent with the results we reported for exchange rate predictability in Tables 3 and 4. For the uninformed group the 95% confidence interval is [-1.41,1.03]. While the point estimate of -0.19 suggests that there is no information content in Twitter Sentiment of the uninformed, with the wide confidence interval we cannot rule out that the true Sharpe ratio is near 1. We clearly need more precision, which is the purpose of the model in the next section.

2.8 Data Moments

In Section 5 we will confront the model of Section 3 with a variety of data moments involving $TS_t^{I,i}$, $TS_t^{U,i}$, TS_t^I , TS_t^U and s_t . These moments are reported in the first column of Table 7. The first 3 moments relate to the Twitter Sentiment indices. The first and second moment are the variance of TS_t^I and TS_t^U . As we will discuss in Section 4, in the model the average variance of Twitter Sentiment is easier to compute than the average standard deviation. That is why we use the variance in the data as well. The variance is a bit higher for informed individuals (0.098 versus 0.068). This is not surprising as new information leads to changes in expectations. Figure 4 illustrates this graphically. Average opinions in the uninformed group are more centered towards the neutral 0, while the informed group shows a wider distribution. The third moment is the correlation between the TS index of informed and uninformed agents, which is 0.44. This suggests a significant common component in the average opinions of both groups.

The next two moments are disagreement moments, which relate to the extent to which opinions differ among individuals during a particular day. They are the average across the 633 days of the cross sectional variance of $TS_t^{I,i}$ and $TS_t^{U,i}$ across the individuals in that group. We again focus on the variance for easier comparison to the model. We do not include the few days for which the number of tweets is 0 or 1.¹⁴ Unsurprisingly, the average difference in opinion is a bit larger for uninformed individuals. This is also illustrated in Figure 5, which shows the distribution of the daily cross sectional variance for both groups. The distribution of uninformed individuals is

¹⁴There are 6 days during which the number of tweets is 0 or 1 for the informed group and 5 days for the uninformed group.

clearly to the right of that distribution of informed individuals.

The next eight moments capture the correlation between Twitter Sentiment and future exchange rate changes. We consider the correlation of the Twitter Sentiment index with the change in the exchange rate over the next 1, 20, 40 and 60 trading days for both informed and uninformed groups. The correlations are positive for the informed group, but negative for the uninformed group. As we will see in the model, these moments can vary a lot across different 633-day samples even when there is economically significant predictability. These correlations by themselves therefore only provide limited information.

The next six moments are the directional moments described in Section 2.6 over 20, 40 and 60 days for the informed and uninformed groups. As discussed, these moments suggest significant information content for the informed group.

The next two moments are the contemporaneous correlation between the weekly Twitter Sentiment index and weekly changes in the exchange rate. The weekly Twitter Sentiment Index TS_w^j is defined as the average of the daily Twitter Sentiment Index over five trading days in a week. The correlation is 0.26 for the informed and -0.01 for the uninformed group.

Finally, the last three moments are the standard deviation and autocorrelation of the daily change in the exchange rate and the autocorrelation of the weekly change in exchange rate. The standard deviation of the daily change in the exchange rate is in percent, so it is $0.57\%=0.0057$. The daily and weekly autocorrelation are 0.003 and 0.008 respectively, reflecting the near random walk aspect of the exchange rate.

3 Model Description

It should be emphasized from the start that the concept of a “tweet” does not exist within the model that we are about to describe. Tweets in reality are just an expression of a belief about the direction of the exchange rate by a subset of agents. Individuals have many sources of information, including that from reading tweets by other people. In the end these beliefs still differ among individuals, reflected in a dispersion of opinions expressed through tweets. In the model the source of this dispersion is private signals, which can be thought of as related to different

research findings, focusing on different pieces of information or reading different newspaper articles or different tweets. In the next section we will relate expectations of exchange rate changes that exist in the model to directional beliefs expressed through tweets.

The model used to shed light on the data is an extension of BvW. They develop a noisy rational expectations exchange rate model in which all agents have private signals about future fundamentals. We will extend the BvW model in only one dimension. In BvW all agents receive different signals, but the quality of these signals is identical in that the variance of the signal errors is equal across all agents. In the extension developed here we assume that there are two groups of agents, which we refer to as informed and uninformed. They are modeled in the same way, except that the informed agents have higher quality private signals. The variance of signal errors is smaller for informed agents. We will be relatively brief in the description of the model as BvW develop the micro foundations and provide further details.

The model focuses on the demand for Foreign bonds. Let $b_{F,t}^{I,i}$ and $b_{F,t}^{U,i}$ be the demand for Foreign bonds by informed and uninformed agent i . There is a continuum of such agents, with $i \in [0, n]$ for informed agents and $i \in [n, 1]$ for uninformed agents. Since Foreign bonds are in zero net supply, the market clearing condition is

$$\int_0^n b_{F,t}^{I,i} di + \int_n^1 b_{F,t}^{U,i} di = 0 \quad (3)$$

Portfolio demand by agents is

$$b_{F,t}^{I,i} = \frac{E_t^{I,i} s_{t+1} - s_t + i_t^* - i_t}{\gamma \sigma_I^2} - b_t^{I,i} \quad (4)$$

$$b_{F,t}^{U,i} = \frac{E_t^{U,i} s_{t+1} - s_t + i_t^* - i_t}{\gamma \sigma_U^2} - b_t^{U,i} \quad (5)$$

Portfolio demand has two components. The first depends on the expected excess return on the Foreign bonds, divided by the product of absolute risk aversion γ and the variance of the excess return.¹⁵ s_t is the log exchange rate (Home currency per unit of Foreign currency), i_t and i_t^* are the Home and Foreign nominal interest rates. The variance of s_{t+1} is respectively σ_I^2 and σ_U^2 for informed and uninformed agents. The computation of first and second moments of s_{t+1} is discussed below.

¹⁵The effect of allowing for different rates of risk-aversion of the two groups is analogous to changing n .

The second term in the portfolio is unrelated to expected returns. In BvW it represents a hedge against non-asset income. In the literature it has alternatively been modeled as noise trade or liquidity trade. What matters is their aggregate across agents:

$$b_t = \int_0^n b_t^{I,i} di + \int_n^1 b_t^{U,i} di \quad (6)$$

for which we assume an AR process:

$$b_t = \rho_b b_{t-1} + \varepsilon_t^b \quad (7)$$

where $\varepsilon_t^b \sim N(0, \sigma_b^2)$. b_t represents the noise that is present in all noisy rational expectations models. In equilibrium the exchange rate will be affected by both shocks to b_t and private information about future fundamentals. By assuming that b_t is unobservable (only its AR process is known), the equilibrium exchange rate will not reveal the aggregate of private information in the market. We also follow BvW by assuming that $b_t^{j,i}$ ($j = I, U$) contains no information about the average b_t .

Standard money demand equations are assumed:

$$m_t = p_t + y_t - \alpha i_t \quad (8)$$

$$m_t^* = p_t^* + y_t^* - \alpha i_t^* \quad (9)$$

m_t is the log money demand, which must equal the log of money supply. y_t is log output. p_t is the log price level. The analogous variables for the Foreign country are denoted with a * superscript. Using PPP ($p_t = s_t + p_t^*$), subtracting these equations yields

$$i_t - i_t^* = \frac{1}{\alpha}(s_t - f_t) \quad (10)$$

where $f_t = (m_t - m_t^*) - (y_t - y_t^*)$ is the traditional fundamental. Since the exchange rate is an $I(1)$ variable in the data, the fundamental is assumed to be $I(1)$ as well. We assume

$$f_t - f_{t-1} = \rho(f_{t-1} - f_{t-2}) + \varepsilon_t^f \quad (11)$$

where $\varepsilon_t^f \sim N(0, \sigma_f^2)$. The fundamental and its process are known to all agents. We will also write the fundamental as $f_t = D(L)\varepsilon_t^f$, where $D(L) = \sum_{i=1}^{\infty} d_i L^{i-1}$ is an infinite order polynomial in the lag operator L , with $d_i = 1 + \rho + \dots + \rho^{i-1}$.

Denote $\bar{E}_t^I s_{t+1} = \int_0^n E_t^{I,i} s_{t+1} di/n$ as the average expectation across informed agents and analogously $\bar{E}_t^U s_{t+1}$ for uninformed agents. Substituting (4), (5) and (10) into the market clearing condition (3), we have

$$\omega \bar{E}_t^I s_{t+1} + (1 - \omega) \bar{E}_t^U s_{t+1} - \frac{1 + \alpha}{\alpha} s_t + \frac{1}{\alpha} f_t = \gamma \sigma^2 b_t \quad (12)$$

where

$$\omega = \frac{n/\sigma_I^2}{(n/\sigma_I^2) + ((1 - n)/\sigma_U^2)} \quad (13)$$

$$\sigma^2 = \frac{1}{(n/\sigma_I^2) + ((1 - n)/\sigma_U^2)} \quad (14)$$

Imposing the market clearing condition (12) allows us to solve for the equilibrium exchange rate.

Finally, agents receive private signals about future values of the fundamental:

$$v_t^{j,i} = f_{t+T} + \epsilon_t^{v,j,i} \quad (15)$$

where $\epsilon_t^{v,j,i} \sim N(0, (\sigma_v^j)^2)$ for $j = I, U$. We assume that $\sigma_v^I < \sigma_v^U$, so that informed agents receive more precise signals than uninformed agents. As usual in the noisy rational expectations literature, the average of the signal errors is assumed to be zero across agents.

(15) says that each period each agent receives a signal about the value of the fundamental T periods later. This is equivalent to assuming that agents receive a signal about the growth rate $f_{t+T} - f_t$ of the fundamental over the next T periods. At time t agents will not just use their latest signal $v_t^{j,i}$ to forecast future fundamentals, but all signals from the last T periods. The signal at time $t - T + 1$ remains informative about f_{t+1} .

The equilibrium exchange rate is solved as follows. Start with the conjecture

$$s_t = A(L)\varepsilon_{t+T}^f + B(L)\varepsilon_t^b \quad (16)$$

where $A(L) = \sum_{i=1}^{\infty} a_i L^{i-1}$ and $B(L) = \sum_{i=1}^{\infty} b_i L^{i-1}$ are polynomials in the lag operator L . Then¹⁶

$$\bar{E}_t^j s_{t+1} = \theta' \bar{E}_t^j(\xi_t) + A^*(L)\varepsilon_t^f + B^*(L)\varepsilon_{t-T}^b \quad (17)$$

$$\sigma_j^2 = \text{var}_t^j(s_{t+1}) = a_1^2 \sigma_f^2 + b_1^2 \sigma_b^2 + \theta' \text{var}_t^j(\xi_t) \theta \quad (18)$$

¹⁶The innovations ε_{t-s}^f are known at t for $s \geq 0$. The innovations ε_{t-T-s}^b are known as well at time t for $s \geq 0$ as they can be extracted from the equilibrium exchange rate at time $t - T$ and earlier.

where $\theta' = (a_2, a_3, \dots, a_{T+1}, b_2, b_3, \dots, b_{T+1})$, $\xi'_t = (\varepsilon_{t+T}^f, \dots, \varepsilon_{t+1}^f, \varepsilon_t^b, \dots, \varepsilon_{t-T+1}^b)$, $A^*(L) = \sum_{i=1}^{\infty} a_{T+i+1} L^{i-1}$ and $B^*(L) = \sum_{i=1}^{\infty} b_{T+i+1} L^{i-1}$. The conditional variance $var_t^j(s_{t+1})$ only has a superscript $j = I, U$ for the group. All agents within the same group have the same quality information and therefore the same perceived uncertainty.

The expectation and variance of unknown innovations ξ_t are computed using a signal extraction problem. Agents have exchange rate signals s_t, \dots, s_{t-T+1} , which all depend on at least some of the unknown innovations of the vector ξ_t . They also have the private signals $v_t^{j,i}, \dots, v_{t-T+1}^{j,i}$ and knowledge of the unconditional distribution of ξ_t . Solving the signal extraction problem (see BvW) yields

$$\bar{E}_t^j(\xi_t) = \mathbf{M}_j \mathbf{H}' \xi_t \quad (19)$$

$$var_t^j(\xi_t) = \tilde{\mathbf{P}} - \mathbf{M}_j \mathbf{H}' \tilde{\mathbf{P}} \quad (20)$$

where $\mathbf{M}_j = \tilde{\mathbf{P}} \mathbf{H} [\mathbf{H}' \tilde{\mathbf{P}} \mathbf{H} + \mathbf{R}_j]^{-1}$, \mathbf{R}_j is a $2T$ by $2T$ matrix with $(\sigma_v^j)^2$ on the last T elements of the diagonal and zeros otherwise, $\tilde{\mathbf{P}}$ is the unconditional variance of ξ_t and

$$\mathbf{H}' = \begin{bmatrix} a_1 & a_2 & \dots & a_T & b_1 & b_2 & \dots & b_T \\ 0 & a_1 & \dots & a_{T-1} & 0 & b_1 & \dots & b_{T-1} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & a_1 & 0 & 0 & \dots & b_1 \\ d_1 & d_2 & \dots & d_T & 0 & 0 & \dots & 0 \\ 0 & d_1 & \dots & d_{T-1} & 0 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & d_1 & 0 & 0 & \dots & 0 \end{bmatrix} \quad (21)$$

Substituting (19) and (20) into (17) and (18) and the result into the market clearing condition (12), we have

$$\begin{aligned} & \theta' (\omega \mathbf{M}_I + (1 - \omega) \mathbf{M}_U) \mathbf{H}' \xi_t - \frac{1 + \alpha}{\alpha} \left(A(L) \varepsilon_{t+T}^f + B(L) \varepsilon_t^b \right) + \frac{1}{\alpha} D(L) \varepsilon_t^f \\ & + A^*(L) \varepsilon_t^f + B^*(L) \varepsilon_{t-T}^b = \gamma \sigma^2 (1 + \rho_b L + \rho_b^2 L^2 + \dots) \varepsilon_t^b \end{aligned} \quad (22)$$

Equating coefficients on the various innovations on both sides yields analytical expressions for a_{T+s} and b_{T+s} for $s \geq 1$ and a set of $2T$ non-linear equations in the remaining parameters $(a_1, \dots, a_T, b_1, \dots, b_T)$. The latter are solved numerically.

Once the equilibrium exchange rate is computed, we can also compute the expectations of future exchange rates by individual agents. In particular, we have

$$E_t^{j,i} s_{t+k} = \bar{E}_t^j s_{t+k} + \mathbf{z}'_k \mathbf{M}_j \mathbf{w}_t^{j,i} \quad (23)$$

where $\mathbf{z}_k = (a_{k+1}, \dots, a_{T+k}, b_{k+1}, \dots, b_{T+k})'$, $\mathbf{w}_t^{j,i} = (0, \dots, 0, \epsilon_t^{v,j,i}, \dots, \epsilon_{t-T+1}^{v,j,i})'$ and

$$\bar{E}_t^j s_{t+k} = \mathbf{z}'_k \bar{E}_t^j \xi_t + \sum_{l=0}^{\infty} a_{T+k+1+l} \epsilon_{t-l}^f + \sum_{l=0}^{\infty} b_{T+k+1+l} \epsilon_{t-T-l}^b \quad (24)$$

So the expectation of the future exchange rate s_{t+k} is equal to the average expectation of all agents in that group (informed or uninformed) plus an idiosyncratic component $\mathbf{z}'_k \mathbf{M}_j \mathbf{w}_t^{j,i}$ that depends on the signal errors of that agent.

4 Empirical Methodology

In this section we discuss the approach we adopt in confronting the model to the data. We first discuss how we compute Twitter Sentiment in the model. After that we discuss the computation of moments in the model and the estimation of model parameters through the Simulated Method of Moments.

4.1 Computing TS in the Theory

The tweets in the data express directional beliefs about the exchange rate without a specific horizon. In connecting the theory to these data, there are two issues that we need to confront. The first is how to think about the horizon of opinions expressed through the tweets. The second is how to relate expected exchange rate changes by individual agents in the model to the directional beliefs in the tweets that can take on the numeric values -1, 0 and 1.

The opinions expressed through tweets do not specify the horizon over which the exchange rate is expected to change. In most of the analysis we will assume that the horizon corresponds to the period in the model over which agents have private information, which is T . Agents have no information about fundamental and noise innovations more than T periods into the future other than that their unconditional mean is zero. So an horizon longer than T makes little sense. In

sensitivity analysis in Section 5.4 we will also consider results when opinions expressed through tweets are based on expected exchange rate changes over an horizon shorter than T .

Regarding the second issue, the model provides no guidance in how to translate expectations of $s_{t+T} - s_t$ into the numeric values -1, 0 and 1. But it is natural that sufficiently large positive (negative) expectations of $s_{t+T} - s_t$ are interpreted as a positive (negative) sentiment, while intermediate expectations are neutral. We will therefore use the following measure of Twitter Sentiment in the theory. For agent i from group j ($j = I, U$), we set

$$TS_t^{j,i} = \begin{cases} 1 & \text{if } E_t^{j,i}(s_{t+T} - s_t) > c^j \\ 0 & \text{if } -c^j \leq E_t^{j,i}(s_{t+T} - s_t) \leq c^j \\ -1 & \text{if } E_t^{j,i}(s_{t+T} - s_t) < -c^j \end{cases} \quad (25)$$

We therefore assign an opinion of +1 if the expected change in the exchange rate is above the cutoff c^j , so that agents are sufficiently confident that the Euro will appreciate. Analogously, we assign a -1 if the expected change is below $-c^j$ and 0 otherwise.

What remains is to identify the proper value for the cutoff c^j . Let π^j be the fraction of all observations in the data for group j ($j = I, U$) for which $TS_t^{j,i}$ is 0. We equate this to the unconditional probability of drawing a 0 in the model:

$$Prob(-c^j \leq E_t^{j,i}(s_{t+T} - s_t) \leq c^j) = \pi^j \quad (26)$$

Since

$$E_t^{j,i}(s_{t+T} - s_t) = \bar{E}_t^j(s_{t+T} - s_t) + \mathbf{z}'_T \mathbf{M}_j \mathbf{w}_t^{j,i} \quad (27)$$

we can compute the unconditional variance of this expectation as

$$\sigma_E^2(j) = var(E_t^{j,i}(s_{t+T} - s_t)) = var(\bar{E}_t^j(s_{t+T} - s_t)) + \mathbf{z}'_T \mathbf{M}_j \mathbf{R}_j \mathbf{M}'_j \mathbf{z}_T \quad (28)$$

where $var(\bar{E}_t^j(s_{t+T} - s_t))$ is computed by first writing the average expectation as a linear function of all ε_{t+T-s} and ε_{t-s}^b with $s \geq 0$ and then taking the unconditional variance.

Using that $E_t^{j,i}(s_{t+T} - s_t)/\sigma_E(j)$ has a $N(0, 1)$ unconditional distribution, and that

$$Prob\left(-\frac{c^j}{\sigma_E(j)} \leq \frac{E_t^{j,i}(s_{t+T} - s_t)}{\sigma_E(j)} \leq \frac{c^j}{\sigma_E(j)}\right) = \pi^j \quad (29)$$

it must be that

$$\Phi\left(\frac{-c^j}{\sigma_E(j)}\right) = \frac{1 - \pi^j}{2} \quad (30)$$

where $\Phi(\cdot)$ is the cumulative normal distribution. Therefore

$$c^j = -\sigma_E(j)\Phi^{-1}\left(\frac{1 - \pi^j}{2}\right) \quad (31)$$

For informed and uninformed agents, in the data we have respectively $\pi^I = 0.328$ and $\pi^U = 0.288$ (see also Figure 3).¹⁷

For what follows, it is useful to characterize the distribution of $TS_t^{j,i}$ conditional on the average expectation, which we will denote $x_t^j = \bar{E}_t^j(s_{t+T} - s_t)$. Then $E_t^{j,i}(s_{t+T} - s_t) = x_t^j + \mathbf{z}'_T \mathbf{M}_j \mathbf{w}_t^{j,i}$. Let σ_w^j be the standard deviation of the second term, associated with signal errors. We can write

$$TS_t^{j,i} = TS_t^j(x_t^j) + \epsilon_t^{j,i} \quad (32)$$

Here $TS_t^j(x_t^j)$ is the mean value of $TS_t^{j,i}$ conditional on x_t^j . This is equal to the average Twitter Sentiment if there were an infinite number of tweets that day. We have

$$TS_t^j(x_t^j) = 1 - \Phi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) - \Phi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \quad (33)$$

where $\Phi(\cdot)$ is the cumulative normal distribution. It follows that

$$\epsilon_t^{j,i} = \begin{cases} 1 - TS_t^j(x_t^j) & \text{with probability } 1 - \Phi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) \\ -1 - TS_t^j(x_t^j) & \text{with probability } \Phi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \\ -TS_t^j(x_t^j) & \text{with probability } \Phi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) - \Phi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \end{cases} \quad (34)$$

We now know the distribution of Twitter Sentiment of individual agents conditional on x_t^j . Below we will use in particular the variance $var(\epsilon_t^{j,i})$ conditional on x_t^j .

4.2 Computing Model Moments

In order to estimate the model parameters, discussed in Section 4.3, we need to compute the model moments. We focus on the 24 moments listed in Table 7. In principle the model moments

¹⁷While it is possible that these percentages are affected by the Euro depreciation over the sample, the values of π^j remain virtually identical for the last 270 days of the sample during which the exchange rate remains almost unchanged.

correspond to the average across an infinite number of simulations of the model over the 633 days for which we have data. In practice model moments are usually computed as the average over a finite number of simulations, like 1000. When considering different sets of model parameters, the model moments are computed using the same set of shocks for the simulations. In our case the shocks are $\varepsilon_t^f, \varepsilon_t^b$ and $\varepsilon_t^{v,j,i}$. However, 1000 simulations, or any other reasonably finite number, creates too much inaccuracy in the context of our application. The reason is that Twitter Sentiment is a discrete variable, so that for given set of shocks a tiny change in model parameters can lead to a discrete change in $TS_t^{j,i}$ for some days and agents, which leads to a discrete change in various moments. Such discontinuities create problems in estimating the parameters as moments are not smooth functions of parameters.

We resolve this as follows. Realizations of the signal error shocks $\varepsilon_t^{v,j,i}$ translate into realizations of $\epsilon_t^{j,i}$, whose distribution is given by (34). We can then write a specific sample moment as $m = m(\mathbf{e}, \mathbf{x})$, where \mathbf{e} consists of the realizations of $\epsilon_t^{j,i}$ and \mathbf{x} consists of the realizations of the shocks ε_t^f and ε_t^b . We need to compute the mean of $m(\mathbf{e}, \mathbf{x})$ across all possible outcomes for \mathbf{e} and \mathbf{x} . We do so by first computing a theoretical expression for the mean across all possible outcomes for \mathbf{e} . This theoretical expression is for one particular set of values of \mathbf{x} . We next simulate the model 1000 times by drawing the shocks ε_t^f and ε_t^b in order to approximate the mean of the moment across all values of \mathbf{x} .

As an illustration, consider the sample variance of Twitter sentiment for group j . We can write

$$TS_t^j = TS_t^j(x_t^j) + \frac{1}{n_t^j} \sum_{i=1}^{n_t^j} \epsilon_t^{j,i} \quad (35)$$

Let S stand for the number of days in the sample (here 633) as well as the set of days in the sample. Then the sample variance is equal to

$$\frac{1}{S-1} \sum_{t \in S} \left(TS_t^j(x_t^j) + \frac{1}{n_t^j} \sum_{i=1}^{n_t^j} \epsilon_t^{j,i} \right)^2 - \frac{S}{S-1} \left[\frac{1}{S} \sum_{t \in S} \left(TS_t^j(x_t^j) + \frac{1}{n_t^j} \sum_{i=1}^{n_t^j} \epsilon_t^{j,i} \right) \right]^2$$

In this case \mathbf{x} consists of the values of x_t^j in the sample, which only depend on the fundamental and noise shocks. We first compute the theoretical mean of this expression for given values of x_t^j

over the distribution of the $\epsilon_t^{j,i}$ given in (34). Doing so gives

$$\text{var}(TS_t^j(x_t^j)) + \sum_{t \in S} \frac{1}{Sn_t^j} \text{var}(\epsilon_t^{j,i}) \quad (36)$$

Here the first variance is the sample variance of $TS_t^j(x_t^j)$, while $\text{var}(\epsilon_t^{j,i})$ is computed from the distribution (34) for given x_t^j . We then finally take the mean across realizations of x_t^j across 1000 simulations of the model. When simulating the model over 633 days, we always assume that the number of tweets each day, n_t^j , corresponds exactly to that in the data for that day.

As another illustration, disagreement is measured as the average cross-sectional variance of Twitter Sentiment across the days in the sample, which is equal to

$$\frac{1}{S} \sum_{t \in S} \frac{n_t^j}{n_t^j - 1} \left(\frac{1}{n_t^j} \sum_{i=1}^{n_t^j} (\epsilon_t^{j,i})^2 - \left(\frac{1}{n_t^j} \sum_{i=1}^{n_t^j} \epsilon_t^{j,i} \right)^2 \right) \quad (37)$$

The mean across the distribution of $\epsilon_t^{j,i}$ is

$$\frac{1}{S} \sum_{t \in S} \text{var}(\epsilon_t^{j,i}) \quad (38)$$

where the variance is again computed from (34) as a function of x_t^j . We finally take the mean across 1000 simulations of the model that lead to different values of x_t^j .

As a final illustration consider the directional moments based on a subsequent change in the exchange rate over the next m days. The sample moment is equal to

$$\frac{1}{\sum_{t \in S} n_t^j} \sum_{t \in S} \sum_{i=1}^{n_t^j} u_t^{j,i} \quad (39)$$

where

$$u_t^{j,i} = \begin{cases} 1 & \text{if } \text{sign}(TS_t^{j,i}) = \text{sign}(s_{t+m} - s_t) \\ -1 & \text{if } \text{sign}(TS_t^{j,i}) = -\text{sign}(s_{t+m} - s_t) \\ 0 & \text{if } TS_t^{j,i} = 0 \end{cases} \quad (40)$$

The theoretical mean of the sample moment across realizations of $\epsilon_t^{j,i}$ is

$$\frac{1}{\sum_{t \in S} n_t^j} \sum_{t \in S} n_t^j TS_t^j(x_t^j) \text{sign}(s_{t+m} - s_t) \quad (41)$$

We again take the average across \mathbf{x} through 1000 simulations. Note that the exchange rate change is part of \mathbf{x} as it depends on the shocks ε_t^f and ε_t^b .

In the Appendix we illustrate this approach for the remaining moments. We double check that the model moments obtained this way are the same as obtained by simulating across all shocks, including the $\varepsilon_t^{v,ji}$. We have done this for 100,000 simulations for a particular parameterization. While it is possible to do this for one set of parameters, it is extremely time-consuming (it takes 8 hours) and therefore runs into computational constraints when estimating parameters. In addition, even for such a large number of simulations the moments are still not completely smooth functions of the parameters when simulating across all shocks, including the $\varepsilon_t^{v,ji}$.

4.3 Estimation of Model Parameters

We estimate the model using the Simulated Method of Moments, based on the 24 moments in Table 7. The parameters are chosen in order to minimize

$$\left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right)' \Sigma^{-1} \left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right) \quad (42)$$

Here \mathbf{m}^{data} is the vector of 24 data moments and $\mathbf{m}^{model}(\nu)$ are the corresponding moments in the model. The latter are a function of the vector ν of model parameters and computed as described in Section 4.2. Σ^{-1} is a weighting matrix. While this can in principle be any matrix, parameter estimates are efficient when Σ corresponds to the variance of the vector of moments. There are different ways this can be approximated. We compute the variance of the moments based on 1000 simulations of the model.¹⁸ Following many others, we only use the diagonal elements of the weighting matrix as the full matrix can lead to finite sample bias (e.g. Altonji and Segal (1996)). The objective function is therefore

$$\sum_{i=1}^{24} \left(\frac{\mathbf{m}^{data}(i) - \mathbf{m}^{model}(i)}{\Sigma_{ii}^{0.5}} \right)^2 \quad (43)$$

¹⁸For this purpose we draw all shocks ε_t^f , ε_t^b and $\varepsilon_t^{v,j,i}$ over 1000 simulations. We obtain parameter estimates for a given weighting matrix, then use these parameter estimates to compute a new weighting matrix. We iterate a couple of times this way until the results no longer change.

This is equal to the sum of the squared t-values of all 24 moments.¹⁹ $\Sigma_{ii}^{0.5}$ is the standard deviation of moment i across simulations of the model.

The variance covariance matrix of parameter estimates is given by

$$\frac{1}{S} \left[\left(\frac{\partial \mathbf{m}^{model}}{\partial \nu} \right)' \Sigma^{-1} \left(\frac{\partial \mathbf{m}^{model}}{\partial \nu} \right) \right]^{-1} \quad (44)$$

where S is the sample length and the derivatives $\partial \mathbf{m}^{model} / \partial \nu$ are evaluated at the estimated parameter vector $\hat{\nu}$.

There is one parameter that we set without estimation, which is the interest elasticity of money demand α . As shown in BvW, we can write the exchange rate as the present discounted value of current and future fundamentals f and noise b . The discount rate in this present value equation is $\alpha / (1 + \alpha)$. Engel and West (2005) report a variety of estimates of this discount rate, which are close to 0.98 for quarterly data. We therefore set $\alpha = 2969$ to generate a 0.98 quarterly discount rate: $(2969/2970)^{60} = 0.98$.

The other parameters of the model are σ_v^I , σ_v^U , σ_b , ρ , ρ_b , n , γ , σ_f and T . We only estimate the first 6 of these parameters. Some comments are therefore in order about γ , σ_f and T . From (12) it can be seen that γ enters the model multiplied by b_t . As a result of this we can only estimate $\gamma \sigma_b$. We therefore normalize $\gamma = 1$ and estimate σ_b . If instead one wishes to set $\gamma = 10$ the reported estimate for σ_b below simply needs to be divided by 10. We set σ_f by exploiting a scaling feature of the model. If we multiply σ_f , σ_v^I and σ_v^U by a factor q , while dividing σ_b by q , the only effect is to scale up the standard deviation of the exchange rate by a factor q . None of the other moments in the model change. We can therefore choose an arbitrary σ_f and estimate the other parameters based on moments other than the standard deviation of the exchange rate. Afterwards we scale σ_f , σ_v^I , σ_v^U and σ_b to match the standard deviation of the daily change in the exchange rate.

The last parameter, T , is different from the others in that it is discrete. We have considered $T = 20$, $T = 40$ and $T = 60$. It is hard to compare them based on the value of the objective function as in each case the estimated weighting matrix is different. We set $T = 40$ as the benchmark, but will report results for $T = 20$ and $T = 60$ as well in sensitivity analysis. The key

¹⁹Since the variance of the moments depends on the parameters themselves, we iterate a couple of times on the optimal weighting matrix and the estimated parameters.

results do not depend on which of these values of T we choose. We set $T = 40$ as it leads to the closest match with the predictive correlations in the data (in terms of average absolute difference, average squared difference, average t-value as well as average squared t-value).

5 Results

We first discuss the estimated parameters. After that we discuss the fit of the model in terms of the moments and examine the role that the moments play for the estimation of key model parameters. We then turn to our ultimate objective, the implications for predictability and Sharpe ratios. We finally perform some robustness analysis.

5.1 Parameter Estimates

Table 6 reports parameter estimates. The standard errors are generally relatively small, so that the data is very informative about the values of our parameters. The standard error is not reported for σ_f as it is simply scaled to match the standard deviation of the change in the exchange rate (see Section 4.3).

The important parameters σ_v^I and σ_v^U relate to the quality of the signals of the informed and uninformed. The two key findings are:

Result 1 *The informed group has higher quality signals than the uninformed group.*

Result 2 *The uninformed group has significant information as well.*

Result 1 says that $\sigma_v^I < \sigma_v^U$. We can reject $\sigma_v^I = \sigma_v^U$ with a p-value of 0.001. One would have expected that the informed have higher quality signals based on the directional moments and predictive correlations. But as we will see, other data moments are also consistent with this higher information quality of the informed. Result 2 is based on the finding that $\sigma_v^U = 0.1$. If they had no private information, σ_v^U would be infinity or a very large number. While the negative directional moments for the uninformed suggest that they have no information at all, other aspects of the data strongly suggest that they do carry information.

A final observation we can make from Table 6 is that the estimate of n is very close to 1, with a standard error of 0.022. This suggests that a large fraction of the wealth is managed by what we call informed traders. Agents with few followers that express opinions through Twitter do not appear to play a big role in the foreign exchange market according to these estimates. However, this is a conclusion that is not robust and not critical to the results. For example, in sensitivity analysis we consider horizons for the tweets that are smaller than T , which generates very similar results for Sharpe ratios but lower values of n of about 0.8. This is discussed in Section 5.4.

5.2 Moments

Table 7 reports all 24 data moments in the first column and compares them to the corresponding average model moments based on the 1000 simulations. The “Model” column reports average model moments. The next column reports the standard deviations of the moments across 1000 simulations. The “Cost” reported in the last column is the contribution of each moment to the objective function. For moment i this is

$$\left(\frac{\mathbf{m}^{data}(i) - \mathbf{m}^{model}(i)}{\Sigma_{ii}^{0.5}} \right)^2 \quad (45)$$

It is the squared t-value of the moment. The sum of these “costs” across the moments is equal to the value of the objective (42), which is shown at the bottom of the table. When the difference between the data and model moments is within two standard deviations of the moment, the cost is less than 4. If the cost is less than 1, the model moment deviates less than one standard deviation from the data.

The model matches various features of the data well: the variance of Twitter Sentiment is higher for the informed than the uninformed, there is a modestly positive correlation between Twitter Sentiment of both groups, disagreement is larger among the uninformed, predictive correlations and directional moments are higher for the informed and the exchange rate is close to a random walk. One point of weakness is that the weekly contemporaneous correlation between Twitter Sentiment and the exchange rate change is higher for the informed in the data, while it is only barely higher in the model.

Overall the model cannot be rejected by the data. The value of the objective function is

16.08. It has a χ^2 distribution with 17 degrees of freedom (number of moments minus number of estimated parameters). Under the null hypothesis that the model is correct, the probability that the objective function is at least 16.08 is 52%, so the null cannot be rejected.

We would like to know what role various moments play in Results 1 and 2. Based on the directional moments and predictive correlations alone, without the use of a model, we would have expected that the uninformed group has no information. After all, the uninformed get the direction of the exchange rate change more often wrong than correct. However, the standard deviations of the directional moments in Table 7 are very large. It is therefore quite possible that the true directional moments are significantly positive for the uninformed and they have substantial information.

One way to see how the moments affect the estimated parameters is to see what happens when they are removed. The directional moments have little effect on the parameters. Removing them leaves parameter estimates virtually unchanged. Table 8 reports the results when the predictive correlations are removed from the estimation. Their values in the model are reported, but there are no entries under the columns for the standard deviation and the “Cost” as these moments do not enter the objective function that we minimize. Even without these moments we see that Results 1 and 2 still hold. We also see that the predictive correlations now become too high relative to what we see in the data. This implies that the predictive correlations, when used in the benchmark estimation, actually reduce the ability of agents to predict future exchange rates. In other words, the strongest evidence of information content in Twitter Sentiment comes from moments other than the predictive correlations and directional moments. Without a model it would be impossible to exploit the information in these moments.

Figure 6 sheds light on the role that some of these other moments play. It shows the variance of Twitter Sentiment, disagreement and the correlation between Twitter Sentiment of the informed and uninformed, all as a function of σ_v^U .²⁰ In doing so we keep all other parameters equal to their estimated values reported in Table 6. It should be kept in mind that the estimated value of σ_v^U is

²⁰For the purpose of Figure 6 we set $\pi^I = \pi^U = 0.3$ and assume a constant number of daily tweets that is equal for both groups. On average this is close to the data. In that case the only difference between the informed and uninformed is associated with the difference between σ_v^U and σ_v^I .

0.1 and the estimated value of σ_v^I is 0.036.

It is useful to consider the first and second charts of Figure 6 simultaneously as they are related. The way that we set the cutoff c^j , discussed in Section 4.1, implies that the unconditional distribution of individual Twitter Sentiment $TS_t^{j,i}$ is always close to that the data. Using (32) and (35) we have

$$TS_t^{j,i} = TS_t^j + \left(\epsilon_t^{ji} - \frac{1}{n_t^j} \sum_{m=1}^{n_t^j} \epsilon_t^{j,m} \right) \quad (46)$$

As we hold constant the variance of individual Twitter Sentiment, the variance of average Twitter Sentiment TS_t^j falls (rises) when the variance of the last term rises (falls). The last term captures the idiosyncratic component of Twitter Sentiment. When its variance rises, agents disagree more. There is therefore a close relationship between the variance of the Twitter Sentiment index TS_t^j and disagreement. We indeed see in the first two charts of Figure 6 that the variance of Twitter Sentiment and disagreement of the uninformed move in opposite directions as we change σ_v^U .

Focusing on disagreement, we can see that only for an intermediate range of values of σ_v^U between $0.036(= \sigma_v^I)$ and 0.18 is disagreement larger among the uninformed than the informed. Since in the data disagreement is larger among the uninformed, this explains why the estimate of σ_v^U is larger than σ_v^I , but not too large. Intuitively, if $\sigma_v^U < \sigma_v^I$, the private signals of the uninformed would have very little noise, which leads to the counterfactual that disagreement is less among the uninformed. At the other extreme, if σ_v^U were very large, the uninformed would give little weight to their poor signals. In the extreme where $\sigma_v^U = \infty$, they only pay attention to public signals. Then there is no disagreement at all, again counterfactual. Only in the intermediate range for σ_v^U is disagreement higher for the uninformed. In that case their signals still contain significant information, so that the agents pay attention to them, but since the signals are noisier than for the informed it causes agents to disagree more than the informed as we see in the data.

The last chart of Figure 6 shows that the correlation between Twitter Sentiment of the informed and uninformed drops uniformly as σ_v^U rises and $\sigma_v^U > \sigma_v^I$. As the uninformed have less and less information about future fundamentals, while the information quality of the informed does not change, the expectations of the uninformed will deviate more from those of the informed. This causes a drop in the correlation. This again explains why the uninformed cannot have very poor

signals. It would lead to a correlation between Twitter Sentiment of both groups that is too low relative to the data.

5.3 Predictability and Sharpe Ratios

We can now turn to the key question of the paper. Is Twitter Sentiment a useful piece of information for predictability and trading purposes? As we discussed in Section 2.7, the Sharpe ratio of a strategy based on Twitter Sentiment leaves inconclusive results based on the limited 633 days of our sample. The 95 percent confidence interval is $[-0.09, 2.27]$ for the informed and $[-1.41, 1.03]$ for the uninformed. We would need a much longer sample, which we do not have. That is why we have used a model to obtain more precision.

In the model we compute the Sharpe ratio in the same way as in the data, as described in Section 2.7. The difference of course is that we are not constrained in the model in terms of sample length. We will therefore compute the Sharpe ratio by simulating the model over 200 years (50,000 trading days).²¹ However, there remains uncertainty about the precise level of the Sharpe ratio as there is uncertainty about the exact model parameters (which in turn is associated with the short data sample). We therefore take 1000 draws from the distribution of the estimated parameters and for each draw compute the Sharpe ratio by simulating the model over 200 years. The findings are reported in Table 9 and Figure 7.

Table 9 shows that the Sharpe ratio for the informed is now on average 1.68, with a 95 percent confidence interval of $[1.59, 1.78]$. The Sharpe ratio is measured with great precision as the model parameters have relatively small standard errors. By taking a stand on the model we have therefore significantly reduced the uncertainty about the level of the Sharpe ratio. Moreover, a Sharpe ratio in the range $[1.59, 1.78]$ is very high by any reasonable standard.

We can compare to Sharpe ratios for the popular currency carry trade strategy based on interest differentials. Burnside et.al (2010) report an average annualized Sharpe ratio of 0.44 for 20 currencies against the dollar based on a carry trade strategy. This is based on data from February 1976 to July 2009. When using an equally weighted strategy that uses all 20 currencies, they find

²¹In doing so we assume that for both groups the number of tweets each day is drawn from the distribution of tweets per day from our 633 day sample.

a Sharpe ratio of 0.91. Since these are annualized Sharpe ratios based on monthly returns, for comparison we have computed the Sharpe ratio based on Twitter Sentiment for monthly returns as well. The strategy remains the same, but the returns are accumulated over one month (20 trading days). We do not report the results as the annualized Sharpe ratios based on monthly returns are virtually identical to those based on daily returns in Table 9. These Sharpe ratios clearly compare very favorably to the popular carry trade.

For the uninformed the annualized Sharpe ratio for daily returns has a 95 percent confidence interval of [0.64,0.90]. Results are again similar for monthly returns. Figure 7 shows the distribution of the Sharpe ratios for both the informed and uninformed. While as expected the Sharpe ratio for the uninformed is well below that for the informed, a mean Sharpe ratio of 0.77 for the uninformed is not bad, particularly in comparison to carry trade Sharpe ratios. As discussed in Section 5.2, while the directional moments and predictive correlations suggest that the uninformed have little information, several other moments strongly suggest that they do have good quality signals.

Figure 8 illustrates why a 633 day sample of data, without any model guidance, will tell us little about the Sharpe ratio. It reports the distribution of the Sharpe ratio for both the informed and uninformed based on 1000 simulations of the model over 633 days, using the estimated parameters. The range of Sharpe ratios from these simulations is very wide for both groups of investors. This gets back to the key point that one either needs a much longer sample than we have or use a model to leverage the small sample and obtain more precision.

Moving on to predictability, Table 10 reports the R^2 from predictability regressions. We regress $s_{t+m} - s_t$ on TS_t^j for $m = 1, 20, 40, 60$ for both informed and uninformed. The reported R^2 numbers are based on simulating the model over 100,000 days based on the estimated parameters. They therefore measure true predictability, not a random small sample relationship. For the informed the R^2 is largest for the one month (20 day) horizon, but even then it is only 0.0144. For the uninformed the highest R^2 is only 0.0022. Such limited predictability is a result of the near-random

walk behavior of exchange rates, but it is sufficient to achieve strong Sharpe ratios.²²

We can also bring the predictability results in Tables 3 and 4 for the data into perspective with the model. Figure 9 reports the distribution of the t-statistic of regressions of $s_{t+m} - s_t$ on TS_t^j , with $m = 1$ in the upper chart, $m = 20$ in the middle chart and $m = 40$ in the bottom chart. These distributions are based on 1000 simulations of the model over 633 days, using the estimated parameters. Figure 9 shows that the probability that the t-stat is less than 2 is in all cases very high. Even though the agents have high quality information, as we know from the Sharpe ratios, the sample is too short to expect statistically significant predictability. This emphasizes once again that data regressions with such a short sample are not sufficient to learn about the information content from Twitter Sentiment.

5.4 Robustness Checks

Table 11 reports results from some robustness checks. It only reports estimates of the most critical parameters, σ_v^I and σ_v^U and Sharpe ratios for both groups of agents, together with standard errors.²³

We have assumed that the horizon of the agents implicit in Twitter Sentiment corresponds to the horizon T in the model over which agents have private signals. This is the horizon over which expected changes in exchange rates are computed in the model in order to compute Twitter Sentiment. In the three rows below the benchmark, we keep $T = 40$, but lower the horizon k implicit in Twitter Sentiment to three values less than 40 (10, 20 and 30). It remains the case that the information quality of the informed agents is well above that of the uninformed. Even more importantly, the Sharpe ratio results remain similar. Sharpe ratios are high and precisely measured with a low standard error. This is the case even though n is now well below that under the benchmark, respectively 0.80 and 0.82 for $k = 10$ and $k = 20$, suggesting that n near 1 is not

²²When we compute the R^2 by simulating the model 1000 times over 633 days, the result is very similar when $m = 1$, but we obtain a higher average R^2 of respectively 0.0274, 0.0285 and 0.0305 when $m = 20, 40, 60$. These are for the informed. These higher R^2 numbers are a result of upward bias of the R^2 in small samples.

²³To compute the Sharpe ratio and its standard error we take 100 draws from the distribution of the estimated parameters, for each draw simulating the model over 200 years. For the benchmark Sharpe ratio results reported in Table 9 we took 1000 draws from the distribution of parameters. We limited it to 100 here for computational reasons as the results are virtually the same with 100 draws as with 1000 draws.

critical to the results.

The next two rows report results for $T = 20$ and $T = 60$. The standard deviation of signal errors drops when $T = 20$ and rises when $T = 60$. But this is because there are fewer signals when $T = 20$ (agents form expectations based on 20 private signals) and more signals when $T = 60$. So it does not represent an overall change in the quality of private information and it remains the case that the informed have better quality signals. Again the Sharpe ratio results remain very similar.

Finally, in the last two rows we report results when we remove days where the number of tweets is less than 5 and less than 10. Such days may not be very informative because of the low number of tweets. For the informed (uninformed) this reduces the number of days from 633 to 616(623) with at least 5 tweets and 561(554) with at least 10 tweets.²⁴ We find that this has very little effect on our estimates of σ_v^j and Sharpe ratios.²⁵

6 Conclusion

Private information is by its nature unobservable. However, with the advent of social media many traders openly express their individual views about future asset prices. This opens up the question what can be learned from these private opinions. In this paper we have used opinionated tweets about the Euro/dollar exchange rate to illustrate how information can be extracted from social media. We have developed a detailed lexicon used by FX traders to translate verbal tweets into opinions that are ranked positive, negative and neutral. Our approach is methodologically different from a related literature that has used social media and the internet to predict future stock price changes. We have aimed to learn from data on Twitter Sentiment and exchange rates through the lens of a model with a precise information structure. This is necessary as in the absence of a model the sample is too short to draw conclusions about predictability and Sharpe ratios.

²⁴In order to compute the weekly contemporaneous correlations we only use weeks for which there are tweets each day of the week. For the informed (uninformed) this reduces the number of weeks from 127 to 115 (120) with at least 5 tweets and 77 (74) with at least 10 tweets.

²⁵The estimate for n drops to 0.89 with at least 10 tweets, again suggesting that the benchmark value of n near 1 is not critical to the results.

The information structure in the model we have used is rich enough to encompass many aspects. There is dispersed heterogeneous information. There are two groups of agents that have different information quality. The horizon over which agents have information is specified. There are unobserved noise shocks as well as observed fundamentals. The parameters of this information structure are mapped into a wide range of moments involving Twitter Sentiment and the exchange rate, allowing us to estimate them with great precision.

The model does a good job in matching the moments in the data and the null that the model is correct cannot be rejected. Estimation provides parameter estimates with small standard errors, which translates into precise estimates of the Sharpe ratio from a trading strategy based on Twitter Sentiment. The large Sharpe ratios that we have reported suggest that there are significant gains from trading strategies based on Twitter Sentiment. The methodology developed here can easily be applied to other currencies or portfolios of currencies, as well as other financial markets such as the stock market.

Appendix: Computation of Moments

As discussed in Section 4.2, we compute the model moments by first taking the theoretical mean of the moments across realizations of ϵ_t^{ji} for given realizations of the shocks ϵ_t^f and ϵ_t^b , and then taking the average across the latter shocks through 1000 simulations of the model. In Section 4.2 we illustrate how to compute the theoretical mean across realizations of ϵ_t^{ji} for the variance of Twitter Sentiment, the cross-sectional disagreement of Twitter Sentiment and the directional moments. Here we consider the remaining moments: the correlation between Twitter Sentiment of the informed and uninformed and the moments involving correlations between Twitter Sentiment and exchange rate changes.

First a general comment is in order about all correlation moments. In principle we should first write down an expression for the sample correlation and then take the mean over realizations of ϵ_t^{ji} . However, the resulting expression is a complicated non-linear function of the ϵ_t^{ji} for which we cannot compute the theoretical mean. We therefore proceed slightly differently. We first compute the mean over realizations of ϵ_t^{ji} for the covariance and the variance of both variables. We then use this result to compute the correlation (covariance divided by product of the standard deviations). It turns out that this delivers virtually identical results. We have verified this by numerically computing the correlation both ways for given realizations of ϵ_t^f and ϵ_t^b . In addition, we have compared the results based on 100,000 simulations of all the shocks over 633 days ($\epsilon_t^b, \epsilon_t^f, \epsilon_t^{v,ji}$), computing the correlation as the average across the simulations, to the moments obtained with our approach. The results are again virtually identical. The reason it makes little difference is that across different draws for the ϵ_t^{ji} across simulations, the variance of variables entering correlations changes very little.

In what follows we will use that for two variables x_t and y_t , we can write the sample covariance as

$$cov(x, y) = \frac{1}{S-1} \sum_{t \in S} (x_t - \bar{x})(y_t - \bar{y}) = \frac{1}{S} \sum_{t \in S} x_t y_t - \frac{1}{S(S-1)} \sum_{t_1 \neq t_2} x_{t_1} y_{t_2} \quad (47)$$

Here \bar{x} and \bar{y} are the sample means of x and y .

Next apply this to the covariance between TS_t^I and TS_t^U . Let S be the number of days, as

well as set of days, for which we have data on both variables. The sample covariance is

$$\begin{aligned} & \frac{1}{S} \sum_{t \in S} \left(TS_t^I(x_t^I) + \frac{1}{n_t^I} \sum_{i=1}^{n_t^I} \epsilon_t^{Ii} \right) \left(TS_t^U(x_t^U) + \frac{1}{n_t^U} \sum_{i=1}^{n_t^U} \epsilon_t^{Ui} \right) \\ & - \frac{1}{S(S-1)} \sum_{t_1 \neq t_2} \left(TS_{t_1}^I(x_{t_1}^I) + \frac{1}{n_{t_1}^I} \sum_{i=1}^{n_{t_1}^I} \epsilon_{t_1}^{Ii} \right) \left(TS_{t_2}^U(x_{t_2}^U) + \frac{1}{n_{t_2}^U} \sum_{i=1}^{n_{t_2}^U} \epsilon_{t_2}^{Ui} \right) \end{aligned} \quad (48)$$

Taking the mean across the ϵ_t^{ji} gives

$$\frac{1}{S} \sum_{t \in S} TS_t^I(x_t^I) TS_t^U(x_t^U) - \frac{1}{S(S-1)} \sum_{t_1 \neq t_2} TS_{t_1}^I(x_{t_1}^I) TS_{t_2}^U(x_{t_2}^U) \quad (49)$$

This is simply the sample covariance between $TS_t^I(x_t^I)$ and $TS_t^U(x_t^U)$. We divide this by the product of the square root of the variance of TS_t^I and TS_t^U , whose computation is discussed in Section 4.2. This results in an expression that depends on the values of x_t^I and x_t^U over the sample, for which we take the average across 1000 simulations.

Next consider predictive correlations. These are the correlation between the Twitter Sentiment Index and the change in the exchange rate over the next m days. We first consider the sample covariance between TS_t^j and $s_{t+m} - s_t$. Since $s_{t+m} - s_t$ only depends on realizations of ϵ_t^f and ϵ_t^b , the average of this sample covariance across realizations of ϵ_t^{ji} is equal to the sample covariance between $TS_t^j(x_t^j)$ and $s_{t+m} - s_t$. This is because $TS_t^j(x_t^j)$ is the mean of TS_t^j across realizations of ϵ_t^{ji} . We then divide the sample covariance between $TS_t^j(x_t^j)$ and $s_{t+m} - s_t$ by the product of the square root of the variance of TS_t^j and $s_{t+m} - s_t$. We discussed the variance of TS_t^j in Section 4.2. The variance of $s_{t+m} - s_t$ only depends on realizations of ϵ_t^f and ϵ_t^b . We finally again take the mean of the resulting correlation across values of ϵ_t^f and ϵ_t^b over 1000 simulations.

The final moments involving the relationship between Twitter Sentiment and exchange rate changes are the contemporaneous weekly correlation between Twitter Sentiment and the exchange rate change. We only consider non-overlapping weeks. Let w be a particular week (5 days in the model). The average Twitter Sentiment in week w is

$$\frac{1}{5} \sum_{m=1}^5 TS_{5*(w-1)+m}^j \quad (50)$$

The exchange rate change in the corresponding week is $s_{5w} - s_{5w-4}$. Since again the exchange rate change does not depend on the ϵ_t^{ji} , the mean of the sample covariance is equal to the sample

covariance of

$$\frac{1}{5} \sum_{m=1}^5 TS_{5*(w-1)+m}^j(x_{5*(w-1)+m}^j) \quad (51)$$

and $s_{5w}-s_{5w-4}$.

In order to compute the correlation we also need the variance of (50). This can be computed analogous to the variance of TS_t^j . After writing down the sample variance and taking the mean across the distribution of the ϵ_t^{ji} , we get

$$\text{var} \left(\frac{1}{5} \sum_{m=1}^5 TS_{5*(w-1)+m}^j(x_{5*(w-1)+m}^j) \right) + \sum_{t \in S} \frac{1}{5Sn_t^j} \text{var}(\epsilon_t^{ji}) \quad (52)$$


References

- [1] Altonji, Joseph G. and Lewis M. Segal (1996), “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business and Economic Statistics* 14 (3), 353-366.
- [2] Andrews, Donald W.K. (1991), “Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation,” *Econometrica* 59(3), 817-858.
- [3] Antweiler, Werner and Murray Z. Frank (2004), “Is All that Talk Just Noise? The Information Content of Internet Stock Message Boards,” *The Journal of Finance* 59 (3), 1259-1294.
- [4] Bacchetta, Philippe and Eric van Wincoop (2006), “Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?,” *American Economic Review* 96(3), 522-576.
- [5] Bollen, Johan, Huina Mao and Xuaojun Zeng (2011), “Twitter Mood Predicts Stock Market,” *Journal of Computational Science* 2, 1-8.
- [6] Brunnermeier, Markus K. (2001), “Asset Pricing under Asymmetric Information,” Oxford University Press, (Oxford).
- [7] Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo (2010), “Do Peso Problems Explain the Returns to Carry Trade?,” *Review of Financial Studies* 24(3), 853-891.
- [8] Da, Zhi, Joseph Engelberg and Pengjie Gao (2015), “The Sum of All FEARS: Investor Sentiment and Asset Prices,” *Review of Financial Studies* 28(1), 1-32.
- [9] Das, Sanjiv R. and Mike Y. Chen (2007), “Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web,” *Management Science* 53(9), 1375-1388.
- [10] Dewally, Michael (2003), “Internet Investment Advice: Investing with a Rock of Salt,” *Financial Analyst Journal* 59(4), 65-77.
- [11] Engel, Charles and Kenneth D. West (2005), “Exchange Rates and Fundamentals,” *Journal of Political Economy* 113(3), 485-517.
- [12] Evans, Martin D.D. (2011), “Exchange-Rate Dynamics,” Princeton University Press.

- [13] Evans, D.D. and Richard K. Lyons (2002), “Order Flow and Exchange Rate Dynamics,” *Journal of Political Economy* 110(1), 170-180.
- [14] Evans, D.D. and Dagfinn Rime (2012), “Micro Approaches to Foreign Exchange Determination,” in *Handbook Of Exchange Rates*, edited by Jessica James, Ian W. Walsh and Lucio Sarno. Hoboken, New Jersey: Wiley and Sons.
- [15] King, Michael R., Carol L. Osler and Dagfinn Rime (2013), “The Market Microstructure Approach to Foreign Exchange: Looking Back and Looking Forward,” *Journal of International Money and Finance* 38, 95-119.
- [16] Lyons, Richard K. (2001), “The Microstructure Approach to Exchange Rates.” MIT Press, Cambridge, MA.
- [17] Mao, Huina, Scott Counts and Johan Bollen (2015), “Quantifying the Effects of Online Bullishness on International Financial Markets,” ECB Statistics Paper 9, July 2015.
- [18] Meese, Richard A. and Kenneth Rogoff (1983a), “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?,” *Journal of International Economics* 14, 345-373.
- [19] Meese, Richard A. and Kenneth Rogoff (1983b), “The Out of Sample Failure of Empirical Exchange Models,” in *Exchange Rates and International Macroeconomics*, edited by Jacob A. Frenkel. Chicago: Univ. Chicago Press (for NBER).
- [20] Mittal, Anshul and Arpit Goel (2012), “Stock Prediction Using Twitter Sentiment Analysis,” working paper, Stanford University.
- [21] Tetlock, Paul C. (2007), “Giving content to investor sentiment: the role of media in the stock market,” *The Journal of Finance* 62,1139-1168.
- [22] Wang, Jiang (1994), “A Model of Competitive Stock Trading Volume,” *Journal of Political Economy* 102, 127-168.
- [23] Zhang, Linhao (2013), “Sentiment on Twitter with Stock Price and Significant Keyword Correlation,” working paper, University of Texas at Austin.

- [24] Zhang, Xue, Hauke Fuehres and Peter A. Gloor (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.

Figure 1 Distribution of the Number of Followers *



*Informed have more than 500 followers. Uninformed have fewer than 500 followers.

Figure 2 Distribution of the Daily Number of Tweets




Figure 3 Distribution of Individual TS




Figure 4 Distribution of Daily Twitter Sentiment Index





Figure 5 Distribution of Daily Disagreement *



*Disagreement is defined as cross sectional variance of Twitter Sentiment across the individuals.

Figure 6 Moments of the Informed and Uninformed Agents as a Function of σ_v^U .





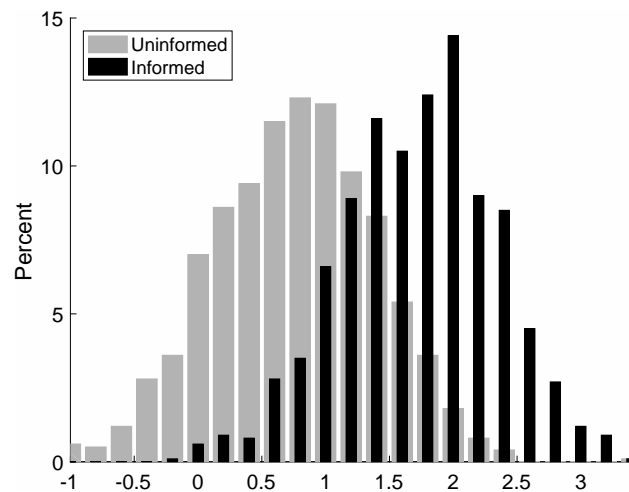


Figure 7 Distribution of Sharpe Ratio*




*The Sharpe ratios are computed using daily observations over 200 years across 1000 draws from the distribution of estimated parameters.

Figure 8 Distribution of Sharpe Ratio due to Small Sample Variation*



*The Sharpe ratios are computed using daily observations over 633 trading days across 1000 simulations.

Figure 9 Distribution t-stat in Predictability Regressions*



*The distribution of t-stats of regressions of $s_{t+m} - s_t$ on TS_t^j are reported based on 1000 simulations of the model over 633 days. The horizon m is 1, 20 and 40 days, respectively, in the top, middle and bottom figure.

Table 1 Examples of Positive, Neutral and Negative Tweets

Score	Category	Text
+1	Positive	\$EURUSD Risks <u>Higher High</u> on Dovish Fed.
+1	Positive	\$EURUSD: <u>Buy dips</u> near term.
+1	Positive	<u>Buy eurUSD</u> market 1.2370 Stop: 1.2230 Target : 1.2600
+1	Positive	Looking to <u>buy eurUSD</u> 1.1330
+1	Positive	Stay <u>Long \$EURUSD</u> For 1.3700; Add At 1.3474/64
+1	Positive	\$EURUSD is right between the two Fibonacci pivot points: 1.3520 and 1.3720. I remain <u>bullish</u> & eventually expect a rally twd 1.4045 Fibo lvl
+1	Positive	Stay calm, <u>hold EURUSD long</u> and USDJPY short
+1	Positive	USD Will Resume Decline; <u>Keep EUR/USD Long</u> For A Run Above 1.40
+1	Positive	<u>Dollar</u> to Face <u>Further Losses</u> on Dismal NFP- EURUSD to Target 1.3960
0	Neutral	I <u>might consider selling</u> \$EURUSD at 1.36 if we spike up.1st probe the market with a small position, and add if we decide to plunge aftrwrds.
0	Neutral	EURUSD trading <u>steady ahead</u> of the Building Permits data from the United States. FOMC Meeting Minutes on focus.
0	Neutral	\$EURUSD sits tight and <u>awaits</u> the FOMC fireworks. Levels to eye.
-1	Negative	EUR/USD Set For <u>Further Falls</u> With Bullish Signal Missing.
-1	Negative	\$EURUSD Risks <u>Further Losses</u> as Growth Outlook Deteriorates.
-1	Negative	\$EURUSD The pair remain <u>bearish</u> and looking for 1.1922 area when a 100% extension will happen .
-1	Negative	Stay <u>Short \$EURUSD</u> , Long \$USDJPY, & Resell \$AUDUSD
-1	Negative	I <u>expect</u> \$eurUSD <u>move lower</u> , just not yet. Daily SRC approaching 5% mark and FT already below -3.530 A Short near 1.3480 makes sense.
-1	Negative	EURUSD Downtrend Intact, Waiting for <u>Sell Signal</u> .
-1	Negative	... said <u>sell \$EURUSD</u> on interest rate differentials, TP 1.2800, SL 1.3700. Fair value at 1.3200. #TradersNotes #FX
-1	Negative	After ECB & Euro Squeeze, ... <u>Adds To \$EURUSD Short</u> Exposure.
-1	Negative	We re looking to big gap at usd pairs. <u>#eurUSD will fall</u> to the 1.23 this week. 46

Table 2 Directional Moments in the Data *

	$s_{t+20} - s_t$	$s_{t+40} - s_t$	$s_{t+60} - s_t$
Informed	0.0224	0.0519	0.0273
Uninformed	-0.0078	-0.0089	-0.0147

*Percentage of tweets that correctly forecast the direction of subsequent exchange rate change minus the percentage of incorrect directional forecasts by individuals in group j . Neutral tweets are counted as neither correct nor incorrect forecasts.

Table 3 Predictability in the Data for the Informed Agents.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$s_{t+1} - s_t$	$s_{t+1} - s_t$	$s_{t+20} - s_t$	$s_{t+20} - s_t$	$s_{t+40} - s_t$	$s_{t+40} - s_t$	$s_{t+60} - s_t$	$s_{t+60} - s_t$
TS_t^I	0.0998 (0.0725)	0.118 (0.076)	0.342 (0.456)	0.205 (0.310)	1.196 (0.840)	0.748 (0.495)	1.418 (1.084)	0.975 (0.686)
TS_{t-1}^I		-0.065 (0.077)		0.082 (0.288)		0.501 (0.428)		0.698 (0.571)
TS_{t-2}^I		0.068 (0.077)		0.277 (0.268)		0.652 (0.463)		0.738 (0.544)
TS_{t-3}^I		0.004 (0.077)		0.185 (0.256)		0.771 (0.499)		0.531 (0.522)
TS_{t-4}^I		-0.123 (0.076)		0.266 (0.286)		0.792 (0.537)		0.663 (0.568)
Observations	633	629	633	629	633	629	633	629
R^2	0.003	0.009	0.002	0.006	0.012	0.035	0.011	0.024
p-value*	0.169	0.318	0.453	0.603	0.155	0.402	0.191	0.612

Newey West Standard errors are in parentheses

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

* p-values are the probability associated with the F-test that all coefficients are zero.

Table 4 Predictability in the Data for the Uninformed Agents.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$s_{t+1} - s_t$	$s_{t+1} - s_t$	$s_{t+20} - s_t$	$s_{t+20} - s_t$	$s_{t+40} - s_t$	$s_{t+40} - s_t$	$s_{t+60} - s_t$	$s_{t+60} - s_t$
TS_t^U	-0.0481 (0.0872)	-0.002 (0.090)	-0.367 (0.474)	-0.278 (0.360)	-0.766 (0.818)	-0.599 (0.606)	-0.675 (1.053)	-0.468 (0.777)
TS_{t-1}^U		-0.274*** (0.0899)		-0.319 (0.315)		-0.559 (0.490)		-0.534 (0.656)
TS_{t-2}^U		0.0119 (0.091)		-0.100 (0.291)		-0.200 (0.485)		-0.199 (0.567)
TS_{t-3}^U		0.125 (0.090)		-0.146 (0.318)		-0.0643 (0.577)		-0.401 (0.641)
TS_{t-4}^U		-0.211** (0.090)		-0.304 (0.344)		-0.148 (0.584)		-0.535 (0.678)
Observations	633	629	633	629	633	629	633	629
R^2	0.000	0.025	0.002	0.005	0.003	0.006	0.002	0.005
p-value*	0.582	0.007	0.439	0.694	0.349	0.395	0.522	0.673

Newey West Standard errors are in parentheses

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

* p-values are the probability associated with the F-test that all coefficients are zero.

Table 5 Annualized Sharpe Ratio of Daily Returns in the Data

	Estimate	S.E.	95% Conf. Int.
Informed	1.09	0.60	[-0.09 , 2.27]
Uninformed	-0.19	0.62	[-1.41 , 1.03]

Table 6 Parameter Estimates *

σ_v^I	0.0362	(0.0067)
σ_v^U	0.1031	(0.0189)
σ_b	0.9906	(0.1877)
ρ	0.0001	(0.1584)
ρ_b	0.9263	(0.0029)
n	0.9999	(0.0221)
σ_f	0.0057	

* Standard errors are in parentheses.

Table 7 Data and Model Moments*

	Data	Model	Standard Deviation	Cost
Variance of Twitter Sentiment				
TS_t^I	0.0983	0.1074	(0.0150)	0.3660
TS_t^U	0.0682	0.0740	(0.0086)	0.4584
Correlation Twitter Sentiment				
TS_t^I, TS_t^U	0.4424	0.3385	(0.0846)	1.5091
Disagreement				
Mean Informed	0.5634	0.5643	(0.0111)	0.0058
Mean Uninformed	0.6356	0.6372	(0.0079)	0.0395
Predictive Correlations				
$TS_t^I, s_{t+1} - s_t$	0.0547	0.1083	(0.0378)	2.0146
$TS_t^U, s_{t+1} - s_t$	-0.0219	0.0440	(0.0396)	2.7720
$TS_t^I, s_{t+20} - s_t$	0.0448	0.1158	(0.1184)	0.3596
$TS_t^U, s_{t+20} - s_t$	-0.0400	0.0393	(0.1029)	0.5937
$TS_t^I, s_{t+40} - s_t$	0.1107	0.0859	(0.1452)	0.0292
$TS_t^U, s_{t+40} - s_t$	-0.0590	0.0180	(0.1281)	0.3615
$TS_t^I, s_{t+60} - s_t$	0.1031	0.0689	(0.1604)	0.0455
$TS_t^U, s_{t+60} - s_t$	-0.0409	0.0018	(0.1367)	0.0974
Directional Moments				
$TS_t^I, s_{t+20} - s_t$	0.0224	0.0333	(0.0392)	0.0777
$TS_t^U, s_{t+20} - s_t$	-0.0078	0.0135	(0.0293)	0.5286
$TS_t^I, s_{t+40} - s_t$	0.0519	0.0263	(0.0487)	0.2762
$TS_t^U, s_{t+40} - s_t$	-0.0089	0.0111	(0.0367)	0.2982
$TS_t^I, s_{t+60} - s_t$	0.0273	0.0224	(0.0548)	0.0081
$TS_t^U, s_{t+60} - s_t$	-0.0147	0.0087	(0.0400)	0.3433
Weekly Contemporaneous Correlations				
$TS_w^I, s_t - s_{t-4}$	0.2659	0.1018	(0.0844)	3.7752
$TS_w^U, s_t - s_{t-4}$	-0.0113	0.0926	(0.0857)	1.4684
Exchange Rate Moments				
St. Dev. Δs	0.5718	0.5718		0.0000
Auto Corr. Δs	0.0028	0.0108	(0.0103)	0.6097
Auto Corr. Δs_w	0.0076	-0.0120	(0.0865)	0.0513
<i>Objective</i>				16.0889

*“Cost” is the square of the difference between the model and data moment, divided by the variance of the corresponding moment. St. Dev Δs is the standard deviation of the daily change in the exchange rate in percentage terms (e.g. 0.5718%=0.005718)

Table 8 Excluding Predictability Moments*

	Data	Model	Standard Deviation	Cost
Variance of Twitter Sentiment				
TS_t^I	0.0983	0.1062	(0.0150)	0.2773
TS_t^U	0.0682	0.0751	(0.0086)	0.6517
Correlation Twitter Sentiment				
TS_t^I, TS_t^U	0.4424	0.3912	(0.0846)	0.3659
Disagreement				
Mean Informed	0.5634	0.5676	(0.0111)	0.1439
Mean Uninformed	0.6356	0.6369	(0.0079)	0.0259
Predictive Correlations				
$TS_t^I, s_{t+1} - s_t$	0.0547	0.3004		
$TS_t^U, s_{t+1} - s_t$	-0.0219	0.1596		
$TS_t^I, s_{t+20} - s_t$	0.0448	0.2013		
$TS_t^U, s_{t+20} - s_t$	-0.0400	0.1002		
$TS_t^I, s_{t+40} - s_t$	0.1107	0.1377		
$TS_t^U, s_{t+40} - s_t$	-0.0590	0.0582		
$TS_t^I, s_{t+60} - s_t$	0.1031	0.1074		
$TS_t^U, s_{t+60} - s_t$	-0.0409	0.0343		
Directional Moments				
$TS_t^I, s_{t+20} - s_t$	0.0224	0.0556	(0.0392)	0.7147
$TS_t^U, s_{t+20} - s_t$	-0.0078	0.0262	(0.0293)	1.3463
$TS_t^I, s_{t+40} - s_t$	0.0519	0.0395	(0.0487)	0.0652
$TS_t^U, s_{t+40} - s_t$	-0.0089	0.0187	(0.0367)	0.5686
$TS_t^I, s_{t+60} - s_t$	0.0273	0.0325	(0.0548)	0.0090
$TS_t^U, s_{t+60} - s_t$	-0.0147	0.0144	(0.0400)	0.5312
Weekly Contemporaneous Moments				
$TS_w^I, s_t - s_{t-4}$	0.2659	0.2153	(0.0844)	0.3593
$TS_w^U, s_t - s_{t-4}$	-0.0113	0.0126	(0.0857)	0.0777
Exchange Rate Moments				
St. Dev. Δs	0.5718	0.5718		0.0000
Auto Corr. Δs	0.0028	0.0106	(0.0103)	0.5743
Auto Corr. Δs_w	0.0076	-0.0183	(0.0865)	0.0896
			s.e.	
σ_v^I		0.0291	(0.0057)	
σ_v^U		0.0680	(0.0131)	
σ_b		2.2986	(0.6936)	
σ_f		0.0019		
ρ		0.6579	(0.0865)	
ρ_b		0.7718	(0.0197)	
n		0.2533	(0.0089)	

*The Table reports results when predictive correlations are excluded from the estimation. The excluded set of moments is shaded.

Table 9 Annualized Sharpe Ratio of Daily Returns in the Model

	Estimate	S.E.	95% Conf. Int.
Informed	1.68	0.05	[1.59 , 1.78]
Uninformed	0.77	0.06	[0.64 , 0.90]

Table 10 R^2 from Predictability Regressions*

	Informed	Uninformed
$s_{t+1} - s_t$	0.0130	0.0022
$s_{t+20} - s_t$	0.0144	0.0015
$s_{t+40} - s_t$	0.0073	0.0006
$s_{t+60} - s_t$	0.0042	0.0003

* The exchange rate and Twitter Sentiment (TS) are simulated over 100,000 days. The change in the future exchange rate is regressed on TS_t , for both informed and uninformed agents, and the R^2 is reported in the table.

Table 11 Sensitivity Tests *

	σ_v^I	σ_v^U	Sharpe Ratio (I)	Sharpe Ratio (U)
Benchmark	0.0362 (0.0067)	0.1031 (0.0189)	1.6852 (0.0502)	0.7704 (0.0642)
$k = 10$	0.0313 (0.0037)	0.0757 (0.0088)	1.3524 (0.0437)	0.6952 (0.0548)
$k = 20$	0.0410 (0.0076)	0.1001 (0.0184)	1.4867 (0.0525)	0.7666 (0.0549)
$k = 30$	0.0400 (0.0057)	0.0996 (0.0142)	1.5543 (0.0497)	0.7839 (0.0574)
$T = 20$	0.0176 (0.0012)	0.0412 (0.0028)	1.4648 (0.0522)	0.8122 (0.0519)
$T = 60$	0.0740 (0.0212)	0.2120 (0.0602)	1.6269 (0.1116)	0.7451 (0.0913)
Number of Tweets ≥ 5	0.0362 (0.0066)	0.1031 (0.0187)	1.7187 (0.0493)	0.7670 (0.0632)
Number of Tweets ≥ 10	0.0396 (0.0047)	0.1029 (0.0123)	1.6180 (0.0427)	0.7702 (0.0482)

* Standard errors are in parentheses.

Appendix A

Tables A1 and A2 show all word combinations used to categorize tweets as positive, negative and neutral. Table A1 lists the word combinations in each category that require the explicit absence of some other words. Table A2 shows the list of word combinations whose existence in a tweet is enough to place the tweet in one of the categories.

“*” and “?” are wildcard characters. “*” represents one or more characters and “?” represents one character. For instance, “*buy??eur*” is a match with any tweet that contains the words “buy” and “eur” in this order and with exactly two characters between them. “*” before “buy” and after “eur” means that there could be any number of characters in a tweet before “buy” or after “eur”. This word combination is intended to identify positive tweets that contain expressions such as “buy \$eurusd” or “buy #euro”. In both cases, all the criteria of a match with “*buy??eur*” are satisfied. There are exactly two characters between “buy” and “eur”. In the case of “buy \$eurusd”, there are three characters after “eur” and in “buy #euro” there is only one character after “eur”. Both are acceptable replacements for the wildcard character “*”. In both examples, there is no character before “buy”. Since “*” could be replaced with zero or any number of characters, no character before “buy” is considered a match with “*buy??eur*”.

Table A1: Word combinations used to identify opinionated tweets.

Positive

Include ...	and not include ...
<p>"*buy??eur*" or "*buy?eur*"</p>	<p>"*close*buy*eur*" , "*exit*buy*eur*", "*close*buy?eurusd*" , "*close*buy?eurusd*", "*close*buy??eur\usd*" , "*buy*,*eur*" , "*buy*:*eur*" , "*buy*fade*" , "*close*buy??eur\usd*" "*never*buy*eur*"</p>
<p>"*buy*lot*eur*"</p>	<p>"*close*buy*lot*eur*"</p>
<p>"*long??eur*"</p>	<p>"*long?term*" , "*was?long*" , "*close*long??eur*" , "*close*long?eur*" , "*exit*long??eur*" , "*exit*long?eur*"</p>
<p>"*bullish*"</p>	<p>"*absent*" , "*absence*" , "*void*" , "*lack*" , "*bullish*fail*" , "*fail*bullish*" , "*bullish*invalid*" , "*bullish*break*" , "*nothing*bullish*" , "*missing*" , "*were?bullish*" , "*was?bullish*" , "*no?bullish*" , "*not?bullish*" , "*market is bullish*"</p>
<p>"*covered*short*"</p>	<p>"*short?term*"</p>
<p>"*buy?the?eur*"</p>	<p>"*never?buy?the?eur*" , "*not?buy?the?eur*"</p>
<p>"*eur?usd*look?good*"</p>	<p>"*eur?usd*not*look?good*"</p>
<p>"*eur?usd*looks?good*"</p>	<p>"*eur?usd*not*looks?good*"</p>
<p>"*double*long*"</p>	<p>"*long?term*"</p>
<p>"*took*long*position*"</p>	<p>"*long?term*"</p>
<p>"*out*of*eur*short*"</p>	<p>"*short?term*" , "*stop*out*of*eur*short*"</p>
<p>"*add*eur*long*"</p>	<p>"*long?term*" , "*addict*" , "*dadd*"</p>

Table A1 (Continued): Word combinations used to identify opinionated tweets.

Positive

Include ...	and not include ...
”*increase*eur*long*”	”*long?term*” , ”*long?off*”
”*up*accelerate*trend*”	”*update*”
”*signals?buy*eur*”	”*forexsignals*”
”*long?signal*”	”*long?term*” , ”*wait*for*long?signal*”
”*higher?high*”	”*if*higher?high*”
”*take*eur?usd*long*”	”*took*profit*”
”*took*eur?usd*long*”	”*took*profit*” , ”*took*opportunity*”
”*further*buying*”	”*buying*usd*”
”*further*eur*gain*”	”*against*”
”*dip*buy*” or ”*buy*dip*”	”*dip*;*eurusd*” , ”*dip*;*eur?usd*” , ”*buy*dips?in?cable*” , ”*buy*dip?in?cable*” ”*sell*rall*”
”*look*to*buy*”	”*looks?like*” , ”*look*to*buy*put*”
”*buying?the?eur*”	”*buying?the?eur*was*” , ”*about*buying?the?eur*” , ”*buying?the?eur*tomorrow*”
”*trigger*further*eurusd*gain*” or ”*trigger*further*eur?usd*gain*”	”*against*”
”*offer*long*entr*”	”*long?term*”
”*look*to*long*”	”*long?term*” , ”*looks*”
”*eur?usd*may*extend*gain*” or ”*eurusd*may*extend*gain*” or ”*eur?usd*will*extend*gain*” or ”*eurusd*will*extend*gain*” or ”*eur?usd*set*extend*gain*” or ”*eurusd*set*extend*gain*”	”*against*”
”eurusd*targets?higher*” or ”*eur?usd*target?higher*”	”*higher?low*”

Table A1 (Continued): Word combinations used to identify opinionated tweets.

Negative

Include ...	and not include ...
""*bearish*"	""*absent*"", ""*bearish*void*"", ""*bearish*lack*"", ""*missing*"", ""*bearish*fail*"", ""*void*bearish*"", ""*lack*bearish*"", ""*fail*bearish*"", ""*bearish*break*"", ""*were?bearish*"", ""*was?bearish*"", ""*not?bearish*"", ""*bearish*invalid*"", ""*nothing*bearish*"", ""*market is bearish*"", ""*no?bearish*"
""*short?eurusd*"" or ""*short??eurusd*"" or ""*short?eur?usd*"" or ""*short??eur?usd*"" or ""*short?euro*"	""*covered*short*"", ""*exit*short*"", ""*stop*short*eur*"", ""*close*short*"
""*took*short*position*"	""*short?term*"
""*short?signal*"	""*short?term*"
""*sell?signal*"	""*buy*signal*"
""*shorted??euro*"" or ""*shorted??eurusd*"" or ""*shorted??eur?usd*"	""*short?term*"
""*sell?eurusd*"" or ""*sell??eurusd*"" or ""*sell?eur?usd*"" or ""*sell??eur?usd*"	""*close*sell*eur*"", ""*exit*sell*eur*"", ""*stop*sell*eur*"", ""*if*sell*eur*"", ""*where*sell*eur*"", ""*no?reason*sell*eur*"

Table A1 (Continued): Word combinations used to identify opinionated tweets.

Negative

Include ...	and not include ...
”*sell the eur*”	”*where*sell the eur*”
”*short the eur*”	”*was*short the eur*”
”*add*eur*short*”	”*short?term*” , ”*addict*” , ”*dadd*”
”*sold*rally*”	”*oversold*”
”*sold*bounce*”	”*oversold*bounce*”
”*eurusd*toppy*” or ”*eurusd*topping*” or ”*eur?usd*toppy*” or ”*eur?usd*topping*”	”*stopp*” , ”*dollar?topp*” , ”*audusd??topp*”
”*bounce*sold*”	”*oversold*”
”*good?short*”	”*short?term*”
”*take*eur*short*”	”*take*profit*eur*short*” , ”*take*out*eur*short*” , ”*take*rest*eur*short*”
”*took*eur*short*”	”*took*profit*eur*short*” , ”*took*out*eur*short*” , ”*took*rest*eur*short*”
”*further*loss*”	”*dollar*further*loss*”
”*further?fall*”	”*dollar*further?fall*”
”*next*leg*lower*”	”*long?term*”

Neutral

”*watch*”	”*video*” , ”*marketwatch*” , ”*watchlist*” , ”*iwatch*”
”*out*eur*long*”	”*break*out*”

Table A2: Word combinations used to identify opinionated tweets.

Positive

""*buy??fxe*""	""*long??fxe*""	""*buy?signal*""
""*upside*breakout*""	""*eur?usd*bull*intact*""	""*expect*move*higher*""
""*oversold*eur*""	""*eur?usd*oversold*""	""*ascending*triangle*""
""*increase*bullish*bet*""	""*bought*rebound*""	""*will*move*higher*today*""
""*bought*dip*""	""*bought*bounce*""	""*will*higher*today*""
""*should?buy*dip*""	""*rally*has*leg*""	""*buy*above*moving*average*""
""*tradable*bottom*""	""*eur?usd*good?buy*""	""*raise*eur*exposure*""
""*will*see*higher*"" or ""*going?to*see*higher*""	""*eur?usd*bias*upside*"" or ""*eur?usd*bias*positive*""	""*eur?usd?will?rise*"" or ""*eur?usd?will?continue?to?rise*""
""*euro?bottoming*"" or ""*eur?usd?bottoming*""	""*staying?long?eur?usd*"" or ""*staying?long??eur?usd*""	""*eur?usd*heads*higher*"" or ""*eur?usd*heading*higher*""
""*bottom?is?in*""	""*oversold*bounce*""	""*sticking*with*long*""
""*potential?buy*""	""*resume*bull*trend*""	""*dollar*further*loss*""
""*long?favor*""	""*suggest*bull*control*""	""*suggest*advance*continue*""
""*spark*eur*buy*"" or ""*initiat*eur*buy*""	""*eurusd?could?bottom*"" or ""*euro?could?bottom*""	""*further*rise*ahead*"" or ""*further*advance*ahead*""
""*stay??eurusd?long*"" or ""*stay?eurusd?long*"" or ""*stay??eur?usd?long*"" or ""*stay?eur?usd?long*""	""*currently?long??eurusd*"" or ""*currently?long?eurusd*"" or ""*currently?long??eur?usd*"" or ""*currently?long?eur?usd*""	""*increase*eurusd?long*"" ""*increase*long?eurusd*"" ""*increase*eur\usd?long*"" ""*increase*long?eur\usd*""
""*increase*eurusd?long*"" , ""*decrease*eurusd?long*"" , ""*hold*eurusd?long*"" , ""*keep*eurusd?long*"" , ""*increase*eur?usd?long*"" , ""*decrease*eur?usd?long*"" , ""*hold*eur?usd?long*"" , ""*keep*eur?usd?long*""		

Table A2 (Continued): Word combinations used to identify opinionated tweets.

Negative

"*short?fxe*" or "*short??fxe*"	"*buy?uup*" or "*buy??uup*"	"*eurusd*buying?put*" or "*eur?usd*buying?put*"
"*eur*overbought*"	"*expect*move*lower*"	"*descending*triangle*"
"*sell?resistance*"	"*selling?resistance*"	"*down*accelerate*trend*"
"*buy?euo*" or "*buy??euo*"	"*eurusd?will?fall*" or "*eur?usd?will?fall*"	"*staying?short*eur?usd*" or "*staying?short*eurusd*"
"*fade*rally*"	"*eur*overpriced*"	"*signals?sell*eur*"
"*bias*down*"	"*eur*overvalued*"	"*stall*retrace*"
"*top?is?in*"	"*deeper?correction*"	"*sticking*with*short*"
"*sell*bounce*"	"*will*see*lower*"	"*prepare*eur*downturn*"
"*recovery*fail*"	"*bear*intact*"	"*eyes*downside*target*"
"*eur*look?bad*" or "*eur*looks?bad*"	"*eur?usd*has*topped*" or "*eurusd*has*topped*"	"*buy the u.s. dollar*" or "*buy?the?dollar*"
"*potential?sell*"	"*downside*remain*"	"*eur*eyes*downside*"
"*stay??eurusd?short*" or "*stay?eurusd?short*" or "*stay??eur?usd?short*" or "*stay?eur?usd?short*"	"*eur?usd*bias*downside*" or "*eur?usd*bias*negative*" or "*eurusd*bias*downside*" or "*eurusd*bias*negative*" or	"*increase*eur?usd?short*" or "*decrease*eur?usd?short*" or "*hold*eur\usd?short*" or "*keep*eur\usd?short*"
"*eurusd*over?bought*" or "*euro*over?bought*" or "*eur?usd*over?bought*" or "*eurusd*overbought*" or "*euro*overbought*" or "*eur?usd*overbought*"	"*further?selling*" or "*further?eurusd?selling*" or "*further??eurusd?selling*" or "*further?eur?usd?selling*" or "*further??eur?usd?selling*" or	"*increase*eurusd?short*" or "*increase*short?eurusd*" or "*increase*eur?usd?short*" or "*increase*short?eur?usd*"
"*look*to*sell*" or "*look*to*buy*put*"	"*will*selling?the?eur*" or "*am?selling?the?eur*"	"*will*head*lower*" or "*heads*lower*" or "*heading*lower*"
"*increase*eurusd?short*" or "*decrease*eurusd?short*" or "*hold*eurusd?short*" or "*keep*eurusd?short*"	"*currently?short??eurusd*" or "*currently?short?eurusd*" or "*currently?short??eur?usd*" or "*currently?short?eur?usd*"	

Table A2 (Continued): Word combinations used to identify opinionated tweets.

Neutral

<p>”*were?bearish*” or ”*were?bullish*” or ”*was?bearish*” or ”*was?bullish*”</p>	<p>”*no?bearish*” or ”*no?bullish*” or ”*not?bearish*” or ”*not?bullish*”</p>	<p>”*not*expect*move*higher*” or ”*not*expect*move*lower*”</p>
<p>”*bullish*absent*” or ”*bearish*absent*” or ”*bullish*void*” or ”*bearish*void*”</p>	<p>”*bullish*lack*” or ”*lack*bullish*” or ”*bearish*lack*” or ”*lack*bearish*”</p>	<p>”*bullish*missing*” or ”*missing*bullish*” or ”*bearish*missing*” or ”*missing*bearish*”</p>
<p>”*bought*sold*” or ”*sold*bought*”</p>	<p>”*will*buy*if*” or ”*will*sell*if*”</p>	<p>”*might*buy*eur*” ”*might*sell*eur*”</p>
<p>”*could?go?higher*” or ”*could?go?lower*” or ”*could?move?higher*” or ”*could?move?lower*”</p>	<p>”*needs?confirm*” or ”*need?to?see*” or ”*needs?to?hold*” or ”*need?to?hold*”</p>	<p>”*bullish*fail*” or ”*fail*bullish*” or ”*bearish*fail*” or ”*fail*bearish*”</p>
<p>”*must?close*” or ”*should?close*”</p>	<p>”*buy*signal*watch*” or ”*sell*signal*watch*”</p>	<p>”*bullish*decline*” or ”*bullishness*decrease*”</p>
<p>”*bull*lose*steam*” or ”*bear*lose*steam*”</p>	<p>”*neutral?on??eur*” or ”*neutral?on?eur*”</p>	<p>”*eur*need*go*lower*” or ”*eur*need*go*higher*”</p>
<p>”*wait*”</p>	<p>”*not*trading*”</p>	<p>”*staying?in?cash*”</p>
<p>”*rally*weak*”</p>	<p>”*patience*”</p>	<p>”*no?need*do?anything*”</p>
<p>”*not?yet*”</p>	<p>”*looking?for*”</p>	<p>”*not?doing?much*”</p>
<p>”*staying?flat*”</p>	<p>”*bounce?possible*”</p>	<p>”*will?be?telling*”</p>
<p>”*all?eyes*on*”</p>	<p>”*steady*ahead*”</p>	<p>”*could*accelerate*”</p>
<p>”*bull*doubt*” or ”*bear*doubt*”</p>	<p>”*no*new*trade*”</p>	<p>”*out*eur*short*” or ”*out*eur*long*”</p>
<p>”*no?trend*”</p>	<p>”*range?in?focus*”</p>	<p>”*may*hold*range*”</p>
<p>”*indecision*”</p>	<p>”*look?to?see*”</p>	<p>”*bias*remain*neutral*”</p>
<p>”*hesitation*”</p>		