What Can we Learn from Euro-Dollar Tweets? [1]

Vahid Gholampour
Bucknell University

Eric van Wincoop
University of Virginia
NBER

October 17, 2016

[1]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 and the Federal Reserve Bank of Dallas for useful comments.

Abstract

With the advent of the internet and social media, we now have real time opinions about future asset price changes by large numbers of people. This paper uses opinionated tweets about the Euro/dollar exchange rate to illustrate how information can be extracted from social media. We develop a detailed lexicon used by FX traders to translate verbal tweets into opinions that are ranked positive, negative and neutral. The methodologically novel aspect of our approach is the use of model with a precise information structure to interpret the data from opinionated FX tweets. The parameters related to the information structure are quite precisely estimated and the model is able to match a wide variety of moments involving Twitter Sentiment and the exchange rate. Based on the estimated model we are able to use daily Twitter Sentiment to predict exchange rates and compute Sharpe ratios for trading strategies. We are able to significantly outperform related results for interest differentials, which are the foundation of the large carry-trade industry.

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 potential gains from trading 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 use data on Twitter Sentiment and the exchange rate to estimate a model with dispersed information. Finally, we use

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

[2]Opinion surveys existed before, but they were infrequent (at most monthly) and limited in scope.

[3]There 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.

[4]The 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).

[5]We 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 results from the model estimation to show that a weighted average of Twitter Sentiment over a period of time is a good predictor of future exchange rate changes. The Sharpe Ratio of a strategy that exploits this predictability outperforms that based on the widely used carry-trade strategy.

As we will discuss below, related literature that uses social media to forecast asset prices is based on a data-only approach. This has significant limitations because very short data samples (often less than a year) are applied to asset prices that are volatile and hard to predict. In addition, opinions about future asset prices expressed through social media are usually directional, e.g. positive, negative and neutral. They do not specify the magnitude of the expected change or the horizon. The same is the case for our Twitter Sentiment about future exchange rate changes. We document that the *direction* of exchange rate changes is predicted by tweets in a way that is statistically significant, which does suggest that there is information content in the tweets. But we also show that Twitter Sentiment does not predict the *magnitude* of future exchange rate changes in a statistically significant way. Such predictability would be needed to develop trading strategies. This absence of predictability based on a data-only approach is not surprising as exchange rates are notoriously hard to predict, Twitter Sentiment is only directional and the data sample is only two and a half years (633 trading days).[6]

The key methodological distinction of our approach is the use of a model, with a precise information structure, to interpret opinions captured by Twitter Sentiment. By taking a stand on the information structure in the context of a specific model, we can learn much more from the data than in a data-only approach. 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. A wide variety of moments involving Twitter Sentiment and exchange rates is driven by a limited set of parameters that describe the information structure. Estimation of these parameters then sheds light on the information content of Twitter Sentiment.

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

[6]The 633 trading days is longer than most of the related literature discussed below that uses samples of no more than one year.

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.[7] 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.

We estimate the parameters of the model with the Simulated Method of Moments, using daily data on Twitter Sentiment and the exchange rate. We find that most of the parameters related to the information structure are quite precisely estimated. Moreover, the model provides a good fit of 24 moments related to Twitter Sentiment and exchange rates. We will show that many moments other than those related directly to predictability are key to the estimation of model parameters. These moments provide insight about the information quality of the agents only because of the use of a model with a precise information structure.

[7] The 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.

Since it provides a good representation of the data, we then use the model to evaluate the ability of Twitter Sentiment by informed traders to forecast future exchange rates at various horizons and to compute the Sharpe ratio of a trading strategy based on a history of the daily Twitter Sentiment index. Both predictability and Sharpe ratios outperform analogous results based on interest differentials.

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. This is why we draw a comparison to predictability based on interest differentials and Sharpe ratios from the associated carry-trade strategy as the latter is widely used in the market.

Although we are not aware of other applications to the foreign exchange market, the paper relates to a literature that has used messages from social media and the internet to predict stock prices. The main difference between this literature and what we do is that this literature has taken a data-only approach. 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.[8] 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.

Apart from the fact that it is entirely data-driven, this literature also differs from our approach in that it does not employ financial jargon used by traders to

[8] Mao et.al. (2015) uses an entire sweep of messages to search for the words "bullish" and "bearish" to classify tweets.

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 also discuss various moments based on this classification and show that this measure of Twitter Sentiment is unable to predict future exchange rate changes in a statistically significant way. In section 3 we describe the NRE model of exchange 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 Twitter Sentiment (TS) that reflects expectations about the future direction of the exchange rate. We first discuss how we use a dictionary of financial lexicon to do this. We then use the results to compute a variety of moments that will be confronted with the theory in Section 5. We also report results from regressing exchange rate changes on past Twitter Sentiment to evaluate predictability without any guidance from theory.

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 is harder to dismiss. It is possible that people are more likely to tweet if they have particularly strong beliefs about the direction of the exchange rate. This could lead to a bias in the measure of the average opinion if for example people who expect a substantial appreciation or depreciation of the Euro are more likely to tweet than those that have a more neutral opinion. In our main analysis in Sections 4 and 5 we will abstract from this bias, assuming that the decision to tweet is independent of the belief about the exchange rate itself. But in sensitivity analysis we will explicitly consider this bias. While it is present, we find that it is nonetheless small and has little effect on the results.

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 use Twitter's publicly available search tools to download the tweets and other information about them, including 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, 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

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

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. 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

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

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 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 "bullish" and "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

[11]The 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.

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 of informed individuals, the average number of tweets over the sample is larger for the informed than uninformed, 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

$$TS_t^j = \frac{1}{n_t^j} \sum_{i=1}^{n_t^j} TS_t^{j,i} \tag{1}$$

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

We use the data on $TS_t^{I,i}$, $TS_t^{U,i}$, TS_t^I, TS_t^U and s_t to ask how well Twitter Sentiment can predict future exchange rate changes. The results are reported in Tables 2, 3 and 4.

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 someone 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. 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 moment is respectively 5.7, 5.9 and as an illustration 6.4 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 for example 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 $s_{t+m} - s_t$ over the same 3 horizons ($m = 20, 40, 60$) on the Twitter Sentiment Index. 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. This suggests that the sample is too short to evaluate based on data alone the ability of Twitter Sentiment to predict the subsequent exchange rate change. It is for this reason that a model is needed to extract more information from the data. The fact that the directional moments are significant for the informed group suggests that there is information content. But the directional moments alone are not sufficient to quantify the predictive content and evaluate the returns from a trading strategy based on Twitter Sentiment.

2.7 Data Moments

We use data on $TS_t^{I,i}$, $TS_t^{U,i}$, TS_t^I, TS_t^U and s_t to compute various moments, which are reported in Table 6. 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 0.46. This suggests a significant common component in the average opinions of both groups.

The next four moments relate to the extent to which opinions differ among individuals during a particular day and are classified as disagreement moments. They are the average and variance 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

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

for which the number of tweets is 0 or 1.[13] Not surprising, 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 clearly to the right of that distribution of informed individuals.

The next six 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 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 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 0.35 for the informed and 0.26 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

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

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

$$\int_0^n b_{F,t}^{I,i} di + \int_n^1 b_{F,t}^{U,i} di = 0 \tag{2}$$

Portfolio demand by agents is

$$b_{F,t}^{I,i} = \frac{E_t^{I,i} s_{t+1} - s_t + i_t^* - i_t}{\gamma \sigma_I^2} - b_t^{I,i} \tag{3}$$

$$b_{F,t}^{U,i} = \frac{E_t^{U,i} s_{t+1} - s_t + i_t^* - i_t}{\gamma \sigma_U^2} - b_t^{U,i} \tag{4}$$

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.[14] 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.

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:

$$b_t = \int_0^n b_t^{I,i} di + \int_n^1 b_t^{U,i} di \tag{5}$$

for which we assume an AR process:

$$b_t = \rho_b b_{t-1} + \varepsilon_t^b \tag{6}$$

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. 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 Bacchetta and van Wincoop (2006) by assuming that $b_t^{j,i}$ ($j = I, U$) contains no information about the average b_t.

Standard money demand equations are assumed:

$$m_t = p_t + y_t - \alpha i_t \tag{7}$$

$$m_t^* = p_t^* + y_t^* - \alpha i_t^* \tag{8}$$

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

$$i_t - i_t^* = \frac{1}{\alpha}(s_t - f_t) \tag{9}$$

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

$$f_t - f_{t-1} = \rho(f_{t-1} - f_{t-2}) + \varepsilon_t^f \tag{10}$$

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

where $\varepsilon_t^f \sim N(0, \sigma_f^2)$. The fundamental and the 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 $d_i = 1 + \rho + ... + \rho^{i-1}$.

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 (3), (4) and (9) into the market clearing condition (2), we have

$$\omega \bar{E}_t^I s_{t+1} + (1-\omega)\bar{E}_t^U s_{t+1} - \frac{1+\alpha}{\alpha} s_t + \frac{1}{\alpha} f_t = \gamma \sigma^2 b_t \tag{11}$$

where

$$\omega = \frac{n/\sigma_I^2}{(n/\sigma_I^2) + ((1-n)/\sigma_U^2)} \tag{12}$$

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

Imposing the market clearing condition (11) allows us to solve for the equilibrium exchange rate.

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

$$v_t^{j,i} = f_{t+T} + \epsilon_t^{v,j,i} \tag{14}$$

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.

(14) 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

$$s_t = A(L)\varepsilon_{t+T}^f + B(L)\varepsilon_t^b \tag{15}$$

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. Then[15]

$$\bar{E}_t^j s_{t+1} = \theta' \bar{E}_t^j(\xi_t) + A^*(L)\varepsilon_t^f + B^*(L)\varepsilon_{t-T}^b \tag{16}$$

$$\sigma_j^2 = var_t^j(s_{t+1}) = a_1^2\sigma_f^2 + b_1^2\sigma_b^2 + \theta' var_t^j(\xi_t)\theta \tag{17}$$

where $\theta' = (a_2, a_3, ..., a_{T+1}, b_2, b_3, ..., b_{T+1})$, $\xi_t' = (\varepsilon_{t+T}^f, ..., \varepsilon_{t+1}^f, \varepsilon_t^b, ..., \varepsilon_{t-T+1}^b)$, $A^*(L) = \sum_{i=1}^{\infty} a_{T+i+1}L^{i-1}$ 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,...,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 $v_t^{j,i}$,...,$v_{t-T+1}^{j,i}$ and knowledge of the unconditional distribution of ξ_t. Solving the signal extraction problem (see BvW) yields

$$\bar{E}_t^j(\xi_t) = \mathbf{M}_j\mathbf{H}'\xi_t \tag{18}$$

$$var_t^j(\xi_t) = \tilde{\mathbf{P}} - \mathbf{M}_j\mathbf{H}'\tilde{\mathbf{P}} \tag{19}$$

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, $\tilde{P}$ is the unconditional variance of ξ_t and

$$\mathbf{H}' = \begin{bmatrix} a_1 & a_2 & ... & a_T & b_1 & b_2 & ... & b_T \\ 0 & a_1 & ... & a_{T-1} & 0 & b_1 & ... & b_{T-1} \\ ... & ... & ... & ... & ... & ... & ... & ... \\ 0 & 0 & ... & a_1 & 0 & 0 & ... & b_1 \\ d_1 & d_2 & ... & d_T & 0 & 0 & ... & 0 \\ 0 & d_1 & ... & d_{T-1} & 0 & 0 & ... & 0 \\ ... & ... & ... & ... & ... & ... & ... & ... \\ 0 & 0 & 0 & d_1 & 0 & 0 & ... & 0 \end{bmatrix} \tag{20}$$

Substituting (18) and (19) into (16) and (17) and the result into the market clearing condition (11), we have

$$\begin{aligned} &\theta'\left(\omega\mathbf{M}_I + (1-\omega)\mathbf{M}_U\right)\mathbf{H}'\xi_t - \frac{1+\alpha}{\alpha}\left(A(L)\varepsilon_{t+T}^f + B(L)\varepsilon_t^b\right) + \frac{1}{\alpha}D(L)\varepsilon_t^f \\ &+A^*(L)\varepsilon_t^f + B^*(L)\varepsilon_{t-T}^b = \gamma\sigma^2\left(1 + \rho_b L + \rho_b^2 L^2 + ...\right)\varepsilon_t^b \end{aligned} \tag{21}$$

[15]The 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.

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 $(a_1, ..., a_T, b_1, ..., b_T)$. 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

$$E_t^{j,i} s_{t+k} = \bar{E}_t^j s_{t+k} + \mathbf{z}_k' \mathbf{M}_j \mathbf{w}_t^{j,i} \tag{22}$$

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

$$\bar{E}_t^j s_{t+k} = \mathbf{z}_k' \bar{E}_t^j \xi_t + \sum_{l=0}^{\infty} a_{T+k+1+l} \varepsilon_{t-l}^f + \sum_{l=0}^{\infty} b_{T+k+1+l} \varepsilon_{t-T-l}^b \tag{23}$$

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

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.

Since no horizons are specified in the tweets, we will assume that the horizon corresponds to the period in the model over which agents have private information, which is T. From the perspective of time t agents have no information about additional fundamental and noise innovations affecting the exchange rate after time T other than that their unconditional mean is zero. So an horizon longer than T makes little sense. In sensitivity analysis we will also consider horizons 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

$$TS_t^{j,i} = \begin{cases} 1 & \text{if} \quad E_t^{j,i}(s_{t+T} - s_t) > c^j \\ 0 & \text{if} \quad -c^j \leq E_t^{j,i}(s_{t+T} - s_t) \leq c^j \\ -1 & \text{if} \quad E_t^{j,i}(s_{t+T} - s_t) < -c^j \end{cases} \tag{24}$$

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:

$$Prob\left(-c^j \leq E_t^{j,i}(s_{t+T} - s_t) \leq c^j\right) = \pi^j \tag{25}$$

Since

$$E_t^{j,i}(s_{t+T} - s_t) = \bar{E}_t^j(s_{t+T} - s_t) + \mathbf{z}_T'\mathbf{M}_j\mathbf{w}_t^{j,i} \tag{26}$$

we can compute the unconditional variance of this expectation as

$$\sigma_E^2(j) = var(E_t^{j,i}(s_{t+T} - s_t)) = var(\bar{E}_t^j(s_{t+T} - s_t)) + \mathbf{z}_T'\mathbf{M}_j\mathbf{R}_j\mathbf{M}_j'\mathbf{z}_T \tag{27}$$

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

$$Prob\left(-\frac{c^j}{\sigma_E(j)} \leq \frac{E_t^{j,i}(s_{t+T} - s_t)}{\sigma_E(j)} \leq \frac{c^j}{\sigma_E(j)}\right) = \pi^j \tag{28}$$

it must be that

$$\Phi\left(\frac{-c^j}{\sigma_E(j)}\right) = \frac{1 - \pi^j}{2} \tag{29}$$

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

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

For informed and uninformed agents, in the data we have respectively $\pi^I = 0.328$ and $\pi^U = 0.288$ (see also Figure 3).[16]

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

$$TS_t^{j,i} = TS_t^j(x_t^j) + \epsilon_t^{j,i} \tag{31}$$

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

$$TS_t^j(x_t^j) = 1 - \Psi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) - \Psi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \tag{32}$$

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

$$\epsilon_t^{j,i} = \begin{cases} 1 - TS_t^j(x_t^j) & \text{with probability} \quad 1 - \Psi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) \\ -1 - TS_t^j(x_t^j) & \text{with probability} \quad \Psi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \\ -TS_t^j(x_t^j) & \text{with probability} \quad \Psi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) - \Psi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \end{cases} \tag{33}$$

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 6. In principle the model moments 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 $\epsilon_t^{v,j,i}$. However, a limited number of simulations

[16] While 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.

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 (33). 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

$$TS_t^j = TS_t^j(x_t^j) + \frac{1}{n_t^j}\sum_{i=1}^{n_t^j} \epsilon_t^{j,i} \tag{34}$$

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. 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 (33). Doing so gives

$$var(TS_t^j(x_t^j)) + \sum_{t\in S}\frac{1}{Sn_t^j} var(\epsilon_t^{j,i}) \tag{35}$$

Here the first variance is the sample variance of $TS_t^j(x_t^j)$, while $var(\epsilon_t^{j,i})$ is computed from the distribution (33) for given x_t^j. We then finally take the mean across realizations of x_t^j across 1000 simulations of the model.

As another illustration, the average cross sectional disagreement in the sample is equal to

$$\frac{1}{S}\sum_{t\in S}\frac{n_t^j}{n_t^j-1}\left(\frac{1}{n_t^j}\sum_{i=1}^{n_t^j}\left(\epsilon_t^{j,i}\right)^2-\left(\frac{1}{n_t^j}\sum_{i=1}^{n_t^j}\epsilon_t^{j,i}\right)^2\right) \tag{36}$$

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

$$\frac{1}{S}\sum_{t\in S} var(\epsilon_t^{j,i}) \tag{37}$$

where the variance is again computed from (33) 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 moment based on a subsequent change in the exchange rate over the next k days. The sample moment is equal to

$$\frac{1}{\sum_{t\in S} n_t^j}\sum_{t\in S}\sum_{i=1}^{n_t^j} u_t^{j,i} \tag{38}$$

where

$$u_t^{j,i}=\begin{cases} 1 & \text{if} \quad \text{sign}(TS_t^{j,i})=\text{sign}(s_{t+k}-s_t) \\ -1 & \text{if} \quad \text{sign}(TS_t^{j,i})=-\text{sign}(s_{t+k}-s_t) \\ 0 & \text{if} \quad TS_t^{ji}=0 \end{cases} \tag{39}$$

The theoretical mean of the sample moment across realizations of ϵ_t^j is

$$\frac{1}{\sum_{t\in S} n_t^j}\sum_{t\in S} n_t^j TS_t^j(x_t^j)\text{sign}(s_{t+k}-s_t) \tag{40}$$

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 Technical Appendix we illustrate this approach for all the moments in the model. 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 6. The parameters are chosen in order to minimize

$$\left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right)' \mathbf{\Sigma}^{-1} \left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right) \tag{41}$$

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. $\mathbf{\Sigma}^{-1}$ is a weighting matrix. While this can in principle be any matrix, parameter estimates are efficient when $\mathbf{\Sigma}$ 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. 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 a weighted average of the squared deviations of model moments from the corresponding data moments, with the weights equal to the reciprocal of the variance of the corresponding moments.[17]

The variance covariance matrix of parameter estimates is given by

$$\frac{1}{S}\left[\left(\frac{\partial \mathbf{m}^{model}}{\partial \nu}\right)' \mathbf{\Sigma}^{-1} \left(\frac{\partial \mathbf{m}^{model}}{\partial \nu}\right)\right]^{-1} \tag{42}$$

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 fall between 0.97 and 0.98 for quarterly data. We therefore set $\alpha = 2370$ to generate a 0.975 quarterly discount rate: $(2370/2371)^{60} = 0.975$.

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 (11) it can be seen that γ enters the model multiplied

[17] Since 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.

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 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 will report results for $T = 20$, $T = 40$ and $T = 60$, corresponding to respectively a one month, two month and three month horizon. For each value of T, the other 6 parameters are chosen to minimize (41). We will compare the objective function across these values of T.

4.4 Predictability and Sharpe Ratios

In order to evaluate the usefulness of Twitter Sentiment data, we consider its ability to predict future exchange rate changes and the Sharpe ratio of a trading strategy based on Twitter Sentiment. In both cases we will draw comparisons to interest rate differentials. It is well known that interest differentials predict changes in future exchange rates and there exists a large industry of currency trade based on interest differentials, known as the carry trade.

Regarding predictability, we compute the R^2 of a forecasting regression using Twitter Sentiment:

$$s_{t+20} - s_t = \alpha + \sum_{k=1}^{l} \beta_k TS^j_{t-k+1} + \varepsilon_{t+20} \tag{43}$$

for $j = I, U$. The change in the exchange rate over 20 trading days (one month) is regressed on the most recent l values of the Twitter Sentiment index, for both informed and uninformed agents. In order to make sure that the R^2 is not upward biased due to a limited sample, we simulate the model over 200,000 trading days.

Regarding the Sharpe ratio, we consider the excess return on a strategy that goes long in Euros and short in dollars, accumulating the daily returns over 20 days (one month). The excess return is regressed on l lags of the Twitter Sentiment index. This gives

$$s_{t+20} - s_t + \sum_{k=1}^{20}(i^*_{t+k-1} - i_{t+k-1}) = \alpha + \sum_{k=1}^{l} \beta_k TS^j_{t-k+1} + \varepsilon_{t+20} \tag{44}$$

Let x_t be the amount that an agent is long in Euro denominated bonds and short in dollar bonds, both measured in dollars at time t. The trading strategy is then to go long one dollar in Euro denominated bonds ($x_t = 1$) when $\hat{\alpha}+\sum_{k=1}^{l}\hat{\beta}_k TS^j_{t-k+1} > 0$ and go short in Euro denominated bonds ($x_t = -1$) when $\hat{\alpha}+\sum_{k=1}^{l}\hat{\beta}_k TS^j_{t-k+1} < 0$. The return is x_t times the excess return on Euro denominated bonds, $s_{t+20} - s_t + \sum_{k=1}^{20}(i^*_{t+k-1} - i_{t+k-1})$. We use a simulation of the model over 200,000 days to obtain parameter estimates $\hat{\alpha}$ and $\hat{\beta}_k$ and compute the Sharpe ratio.

5 Results

Tables 5 through 11 report the results. We first discuss the estimated parameters, followed by the moments and finally consider results on predictability and Sharpe ratios.

5.1 Parameter Estimates

Table 5 reports parameter estimates for $T = 20$, $T = 40$ and $T = 60$. For comparison we use the same weighting matrix in all three cases, which is the optimal weighting matrix for $T = 40$. Even when we choose the optimal weighting matrix for $T = 20$ or $T = 60$, the objective will always be lowest for $T = 40$. A two-month horizon provides the best fit with the data and will therefore be our benchmark. 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).

We see that in all three cases $\sigma^I_v < \sigma^U_v$. Investors with a lot of followers are therefore indeed more informed. We can reject $\sigma^I_v = \sigma^U_v$ with a p-value of 0.001. We also see that σ^I_v and σ^U_v increase with the horizon T. Note that agents have three times as many signals when $T = 60$ than when $T = 20$. It is therefore natural that the quality of each signal is lower when $T = 60$ as otherwise there would be too much information.

We also see that ρ and ρ_b increase with the horizon T. A higher value of ρ implies that it takes longer for the fundamental to reach a new higher steady state level after an innovation. A higher ρ_b means that noise innovations obscure the information content of the exchange rate for a longer period of time. Both a higher ρ and a higher ρ_b have the implication that it takes longer to learn the magnitude of future fundamental innovations. This is needed when the horizon T is larger as otherwise agents would learn too soon about fundamental innovations far into the future.

This is further illustrated in Figure 6, which shows the impulse responses of average Twitter Sentiment of informed and uninformed agents ($TS_t^j(x_t^j)$) and the exchange rate in response to both fundamental and noise shocks. For $T = 60$ Twitter Sentiment changes far more gradually in response to future fundamental innovations than for $T = 20$. Related to that, for $T = 60$ Twitter Sentiment increases less and remains higher much longer in response to noise shocks. If agents would learn very quickly about fundamentals far into the future there would be too much predictability.

A final observation from Table 5 is that the estimate of n is very close to 0 under the benchmark. It is a bit larger for $T = 20$ and $T = 60$ (respectively 0.09 and 0.37). As we will see, our key results regarding predictability and Sharpe ratios do not depend much on T and are therefore not driven by the near-zero value of n under the benchmark. With a standard error of 0.01, our benchmark results are consistent with a value of n near 0.02. This may not be unrealistic. For example, Gennotte and Leland (1990) report that informed traders make up only about 2% of capital in the stock market. Similarly, Bacchetta and van Wincoop (2010) report that only about 1% of assets in the foreign exchange market are managed by informed investors such as hedge funds and currency overlay managers.[18]

[18]It also bears emphasizing that the fraction of informed agents in our Twitter data does not need to bear any relation to the fraction of informed traders in the market. Informed traders are more likely to tweet, both because they have more to say and because they wish to attract customers to their business.

5.2 Moments

5.2.1 Comparison Data and Model Moments

Table 6 reports all 24 data moments in the first column and compares them to the corresponding average model moments for the three values of T. For each T we also report the "Cost" which is the contribution of each moment to the objective function. For moment j this is

$$\frac{\left(\mathbf{m}^{data}(j) - \mathbf{m}^{model}(j)\right)^2}{\mathbf{\Sigma_{j,j}}} \tag{45}$$

where $\mathbf{\Sigma_{j,j}}$ is element (j, j) of $\mathbf{\Sigma}$, which is the variance of moment j across simulations. The sum of these "costs" across the moments is equal to the value of the objective (41), 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. For $T = 40$ we report this cost for each moment, while for $T = 20$ and $T = 60$ we report the "relative cost," which is the cost minus that under the benchmark. This allows us to quickly judge where the model performs better or worse than under the benchmark.

For $T = 40$ we can see that the difference between the data and model moments is almost always within two standard deviations and for the bulk of the moments even within one standard deviation. The only exception is the variance of disagreement. The cost is 4.07 for the informed (about 2 standard deviations) and 13.1 for the uninformed (3.6 standard deviations). Even in the latter case the deviation between the data and the model is not very large, but the standard deviation of that moment is a very small 0.0016. Note that we have not accounted for measurement error in the data, which is not included in the standard deviation of the moments.

By inspecting the relative cost of the moments for $T = 20$ and $T = 60$ we can see why $T = 40$ is preferred. A negative number means that we match the data closer than the benchmark, while a positive number implies that the model is further removed from the data than in the benchmark. The main difference relative to the benchmark is not in the moments involving the variance and disagreement of Twitter Sentiment. Some of those improve, while others deteriorate. The main difference is instead in the predictive moments and in the contemporaneous correlation between weekly Twitter Sentiment and the change in the exchange rate. Not surprisingly, for $T = 20$ the predictive moments are weaker for the longer horizons of 40 and 60 days while for $T = 60$ predictive moments are weaker at the shorter horizons, particularly 20 days. There is also a deterioration of the contemporaneous weekly correlation between Twitter Sentiment and the change in the exchange rate that is worst for $T = 60$. The more gradual change in Twitter Sentiment in response to fundamental shocks shown in Figure 6 for $T = 60$ implies a weaker contemporaneous correlation between the exchange rate and Twitter sentiment than in the data.

5.2.2 What do Moments tell us about Information Quality of the Agents?

The estimation results lead us to draw two conclusions about the information quality of the agents. First, the informed group clearly has higher quality information than the uninformed group ($\sigma_v^I < \sigma_v^U$). Second, the uninformed group have informative signals as well in that σ_v^U is not an extremely large number that would make these signals essentially useless. The information quality of the uninformed will be further confirmed in the next section when we look at predictability and Sharpe ratios.

One can ask what aspects of the data inform us about the information quality of the agents. It is not just the predictive correlations and directional moments that inform us of the higher quality information of the informed group. Table 7 reports the moments and parameter estimates when we exclude various sets of moments from the estimation (highlighted in grey). Removing both the predictive correlations and directional moments (column 3) leads to the same outcome. The estimates of σ_v^I and σ_v^U are respectively 0.0322 and 0.0618, not far from the estimates without the predictive and directional moments.

This suggests that other aspects of the data are critical to the estimation of σ_v^I and σ_v^U. We can only learn about the information quality from these other moments because of the use of a model with a precise information structure. Consider the finding that $\sigma_v^I < \sigma_v^U$, a robust result that continues to hold in Table 7 when we remove any subset of moments from the estimation. Three sets of moments are worth emphasizing in generating this result. First, the higher variance of Twitter Sentiment of informed agents in the data is matched by the model (see Table 6).

The reason for this is that more accurate signals of the informed leads to larger changes in opinion over time. Second, the higher disagreement among uninformed agents in the data is also consistent with the model (see Table 6). It is naturally the case that weaker signals imply a wider dispersion in beliefs. Finally, the higher weekly contemporaneous correlation between Twitter Sentiment and the exchange rate change for the informed group is again perfectly matched by the model. When new signals about future fundamentals lead to an appreciation of the exchange rate, the Twitter Sentiment of the informed agents rises more as they are more certain of the change in the future fundamentals. This can also be seen in Figure 6.

Next consider the information quality of the uninformed. Based on the predictability correlations and directional moments alone one might believe that these agents have no information. The sign of these moments is actually negative for the uninformed group, the opposite of what one would expect if their expectations contained information. The estimate of the information quality of this group is therefore entirely driven by other aspects of the data, where again the precise information structure of the model is critical for the interpretation.

In order to better understand what drives the estimate of σ_v^U, we have done an experiment where we raised σ_v^U while re-estimating the other parameters. This generates two inconsistencies with the data. First, less information of the uninformed agents reduces the correlation between Twitter Sentiment of the informed and uninformed to well below that in the data. In other words, the Twitter Sentiment of the uninformed group is too correlated with that of the informed group in the data to conclude that the uninformed have no information. This is also confirmed in Table 7. When we remove this correlation from the moments, the estimate of σ_v^U rises significantly to 0.0787. Second, when we raise σ_v^U the predictability moments for the informed group become too high relative to the data. Since n is close to 0, the exchange rate is almost entirely driven by the information of the uninformed. This allows the informed to predict future exchange rate changes and more so the weaker the signals of the uninformed. In other words, the informed know too much about future exchange rate changes if the uninformed know too little about future fundamentals.

5.2.3 Other Aspects of the Information Structure

The moments teach us about many aspects of the information structure other than the standard deviations of the signal errors. Consider for example the persistence parameters ρ and ρ_b. As we have discussed above, higher persistence parameters leads to slower learning, which reduces predictability. When we remove the predictive correlations and directional moments (column 3 of Table 7), these persistence parameters are lower, which leads to excessive predictability. Similarly, removing the variance of Twitter Sentiment from the moments (column 5 of Table 7) also reduces the persistence parameters. The more rapid learning leads to excessive volatility of Twitter Sentiment of both informed and uninformed agents.

The moments that are most important in the identification of the parameters are the weekly contemporaneous correlations between Twitter Sentiment and the exchange rate. Removing these moments (column 6 of Table 7) leads to the biggest changes in the estimated parameters. It leads to weaker quality signals, a considerable increase in ρ and ρ_b, a large drop in σ_b and an increase in n to almost 1. The effect of all these parameter changes is that the contemporaneous correlation between Twitter Sentiment and weekly exchange rate changes becomes much too low in the model. This naturally happens when we weaken the signals. As we have seen in Figure 6, the same happens when noise and fundamental shocks are more persistent. When the share of agents that are informed becomes very high (n close to 1), more information about future fundamentals becomes incorporated into the current exchange rate. It then becomes harder to predict further exchange rate changes, reducing the contemporaneous correlation between Twitter Sentiment and exchange rates to well below that in the data.

5.3 Predictability and Sharpe Ratios

Table 8 reports the results for predictability and Sharpe ratios as discussed in Section 4.4. To provide some perspective, it is useful to have a benchmark to compare against. For this we use the available evidence on predictability and Sharpe ratios based on interest differentials. Interest differentials are one of the few variables that have shown consistent predictive power. It is well known that high interest rate currencies tend to appreciate (Fama puzzle) and there exists an extensive industry aimed at exploiting resulting arbitrage opportunities.

Burnside et.al.(2006) report the R^2 from a standard Fama regression of the change in the exchange rate on the forward discount (interest differential). Based on monthly data for 9 currencies (relative to the British pound) over the period 1976-2005, they report an average R^2 of 0.02. This may not seem very high, but it is important to keep in mind that exchange rate changes are well known to be close to a random walk and therefore very hard to predict.[19]

Burnside et.al.(2006) report an average monthly Sharpe ratio of 0.11 for a trading strategy analogous to the one described in Section 4.4, but based on the forward discount. This is again the average for 9 currencies. We will annualize Sharpe ratios by multiplying by $\sqrt{12}$. The annualized Sharpe ratio is then 0.38. Burnside et.al.(2006) show that one can do even better by adopting an equally weighted portfolio of all currencies, yielding an annualized Sharpe ratio of 0.69. But since we will consider only one bilateral currency (Euro/dollar), the relevant comparison is the 0.38 annualized average monthly Sharpe ratio for bilateral currency strategies.

The predictability and Sharpe ratio results for the model are reported in Table 8 for the benchmark parameters ($T = 40$). We consider various lags L in (43) and (44), equal to 1 lag, 5 lags, 10 lags, 15 lags and 20 lags. As discussed in Section 4.4, these numbers are based on an average of 200,000 simulations of the model and are for monthly exchange rate changes. Sharpe ratios are annualized.

We see that for the informed agents, with at least 5 lags the R^2 is 0.06 and the Sharpe ratio is 0.68. This is considerably better than for interest differentials. It strongly suggests that there is significant information content in the tweets of agents with large numbers of followers that can be effectively used to devise a trading strategy along the lines that we have described. For the uninformed agents the R^2 is 0.03 and Sharpe ratio 0.46 with at least 10 lags. While considerably weaker than for informed traders, this still compares favorably to results based on the forward discount.

The 0.68 Sharpe ratio based on the Twitter Sentiment of the informed group is a true Sharpe ratio in the model, which would be achieved if an investor persisted in following this trading strategy and the model is a correct description of the world. Of course over a limited time period, the actual Sharpe ratio could be much larger or smaller. Figure 7 illustrates the distribution of Sharpe ratios in

[19] See for example Meese and Rogoff (1983a,b), Cheung et.al. (2005) and Engel and West (2005). Note also that this is also consistent with the near-zero autocorrelation of daily and weekly exchange rate changes in both our data and the model.

the model over the 633 day period corresponding to our sample. The numbers reported in Table 8 are at the center of this distribution. It is clear though that for such a limited sample far worse and far better Sharpe ratios are possible than the 0.68 mean for the informed and 0.46 mean for the uninformed.

If we followed the strategy implied by the model to our specific 633 day data sample, the Sharpe ratio based on TS of the informed group would be 1.07, even a bit better than the 0.68 mean implied by the theory. The Sharpe ratio based on TS of the uninformed group would be 0.47, virtually the same as the mean in the model. These are good results, but we should not make too much from this. Figure 7 shows that these numbers will vary a lot across such short samples. To illustrate this in the context of the data, if we break our data sample in two halves, the Sharpe ratio for the informed group is 2.86 in the first half and 0.50 in the second half. Similarly, the numbers are 0.75 and 0.36 for the uninformed group.

Table 9, 10 and 11 consider whether the predictability and Sharpe ratio results are robust to changes in model parameters and the set of moments upon which estimation is based. Table 9 shows that predictability and Sharpe ratios are only slightly weaker for $T = 20$ and $T = 60$. Particularly for the informed traders they remain considerably better than for interest differentials. Table 10 reports how the results for $T = 40$ are affected by lowering and raising the estimated parameters by two standard deviations. In each case we re-estimate the other parameters. Results are reported based on 20 lags. Neither predictability nor the Sharpe ratios are much affected. This suggests that the findings are robust to changing parameters within the estimated confidence intervals.

Finally, Table 11 shows how the results are affected when we exclude various subsets of moments in the estimation of parameters. This corresponds to the results of Table 7. The only set of moments that significantly affect the results are the weekly contemporaneous correlations between the exchange rate change and Twitter Sentiment. For informed traders the predictability and Sharpe ratio drops to a level that is equal to that reported above for interest differentials. This is still not bad. We have already seen that these moments are critical to the estimation of many parameters. When ignoring them in the estimation, these contemporaneous correlations are significantly lower than in the data.

While the model with the estimated parameters implies significant predictability and substantial Sharpe ratios, it is consistent with the lack of predictability in the data regressions of exchange rate changes on Twitter Sentiment reported in Tables 3 and 4. Figure 8 reports the distribution of t-stats of the regression of $s_{t+f} - s_t$ ($f = 20, 40, 60$) on TS_t^j ($j = I, U$) based on 1000 simulations of the model over 633 days. Figure 8 shows that it is more than possible, even likely, that such regressions deliver findings of statistically insignificant predictability. For example, for $f = 40$ the probability that the t-stat is less than 2 is 83% for the informed group and 94% for the uninformed group. Even though we have found that the uninformed group has significant information as well, the probability of a negative t-stat as found in the data is 43% ($f = 40$).

Figure 8 implies that it is hard to draw any conclusions about predictability from data regressions. The sample of two and a half years is simply not long enough, particularly given the volatility of exchange rates. It is for this reason that we have used a model to learn more about the information content of the tweets.

6 Conclusion

The paper has 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 simple model-free regressions of future exchange rate changes on Twitter Sentiment are not expected to deliver, and indeed do not deliver, statistically significant results. More structure is needed, particularly because exchange rates are volatile and hard to predict, Twitter opinions are only directional and we have a limited data span of 2.5 years.

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 moments generated by the model are consistent with the corresponding moments in the data.

We have used the model to evaluate whether Twitter Sentiment can be effectively used to predict future exchange rates and to determine if a profitable trading strategy can be developed based on these opinions. We have shown that Twitter Sentiment of informed traders can predict monthly exchange rate changes better than interest differentials can. Related to that, the Sharpe ratio of trade based on Twitter Sentiment is substantially better than that based on the popular carry trade.

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] Backus, David K., Allan W. Gregory and Chris I. Telmer (1993) "Accounting for Forward Rates in Markets for Foreign Currency," *Journal of Finance* 48 (5): 1887908.

[6] Bollen, Johan, Huina Mao and Xuaojun Zeng (2011), "Twitter Mood Predicts Stock Market," *Journal of Computational Science* 2, 1-8.

[7] Brunnermeier, Markus K. (2001), "Asset Pricing under Asymmetric Information," Oxford University Press, (Oxford).

[8] Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo (2006), "The Returns to Currency Speculation, "NBER Working Paper 12489.

[9] Burnside, Craig, Martin Eichenbaum, Isaac Kleshchelski and Sergio Rebelo (2011), "Do Peso Problems Explain the Carry Trade?, "*Review of Financial Studies* 24(3), 853-891.

[10] Cheung, Yin-Wong, Menzie D. Chinn and Antonio Garcia Pascual (2005), "Empirical Exchange Rate Models of the Nineties: Are they Fit to Survive?," *Journal of International Money and Finance* 24, 1150-1175.

[11] 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.

[12] 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.

[13] Dewally, Michael (2003), "Internet Investment Advice: Investing with a Rock of Salt," *Financial Analyst Journal* 59(4), 65-77.

[14] Engel, Charles and Kenneth D. West (2005), "Exchange Rates and Fundamentals," *Journal of Political Economy* 113(3), 485-517.

[15] Evans, Martin D.D. (2011), "Exchange-Rate Dynamics," Princeton University Press.

[16] Evans, D.D. and Richard K. Lyons (2002), "Order Flow and Exchange Rate Dynamics," *Journal of Political Economy* 110(1), 170-180.

[17] 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.

[18] Fama, Eugene F. (1984), "Forward and spot exchange rates," *Journal of Monetary Economics* 14, 319-338.

[19] 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.

[20] Lyons, Richard K. (2001), "The Microstructure Approach to Exchange Rates." MIT Press, Cambridge, MA.

[21] Mao, Huina, Scott Counts and Johan Bollen (2015), "Quantifying the Effects of Online Bullishness on International Financial Markets," ECB Statistics Paper 9, July 2015.

[22] 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.

[23] 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).

[24] Mittal, Anshul and Arpit Goel (2012), "Stock Prediction Using Twitter Sentiment Analysis," working paper, Stanford University.

[25] Tetlock, Paul C. (2007), "Giving content to investor sentiment: the role of media in the stock market," *The Journal of Finance* 62,1139-1168.

[26] Wang, Jiang (1994), "A Model of Competitive Stock Trading Volume," *Journal of Political Economy* 102, 127-168.

[27] Zhang, Linhao (2013), "Sentiment on Twitter with Stock Price and Significant Keyword Correlation," working paper, University of Texas at Austin.

[28] 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 *

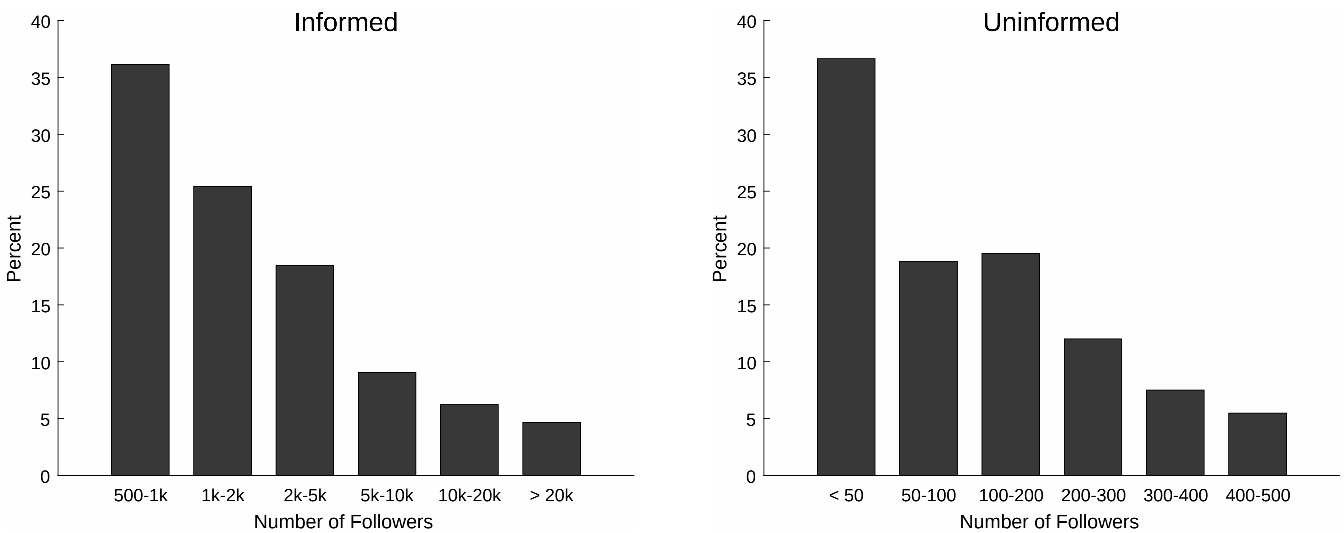

*Informed have more than 500 followers. Uninformed have fewer than 500 followers.

Figure 2: Distribution of the daily number of tweets

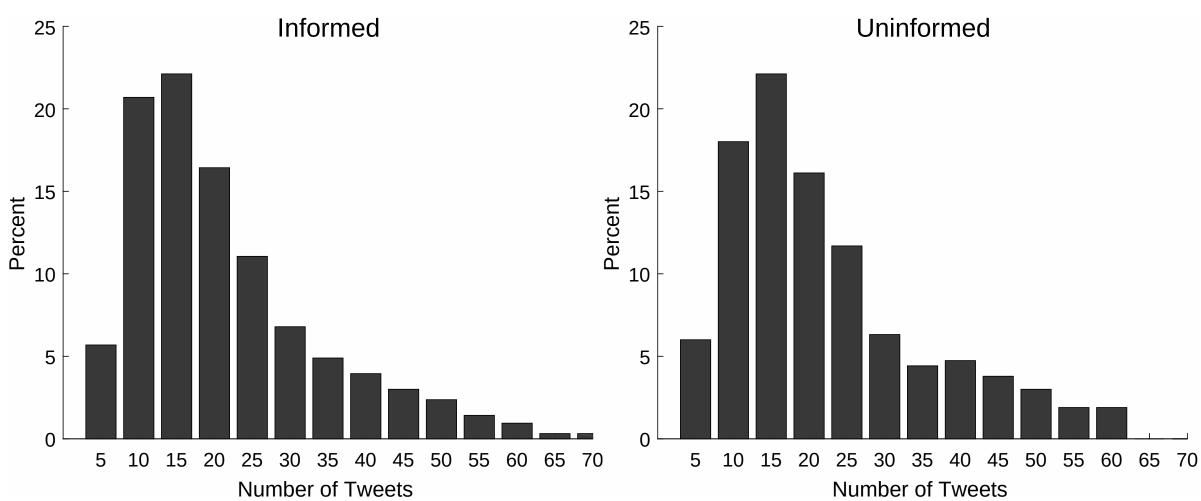

Figure 3: Distribution of Individual TS

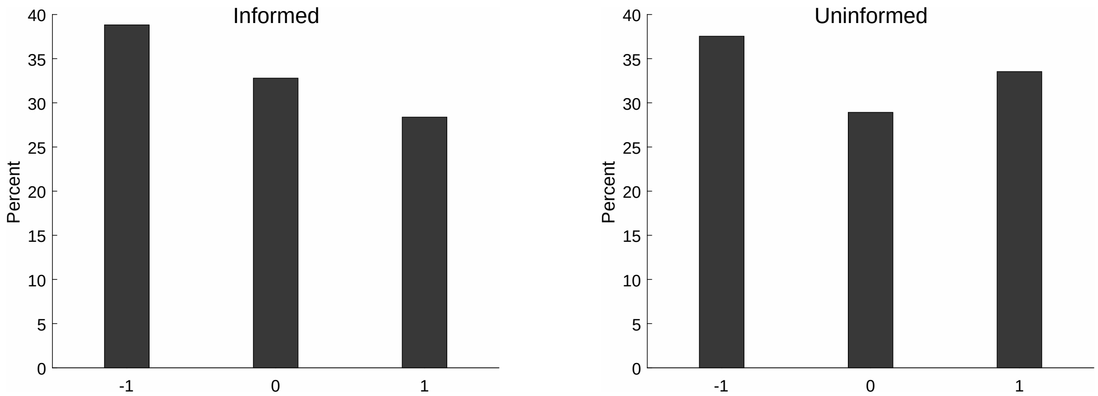

Figure 4: Distribution of daily Twitter Sentiment Index

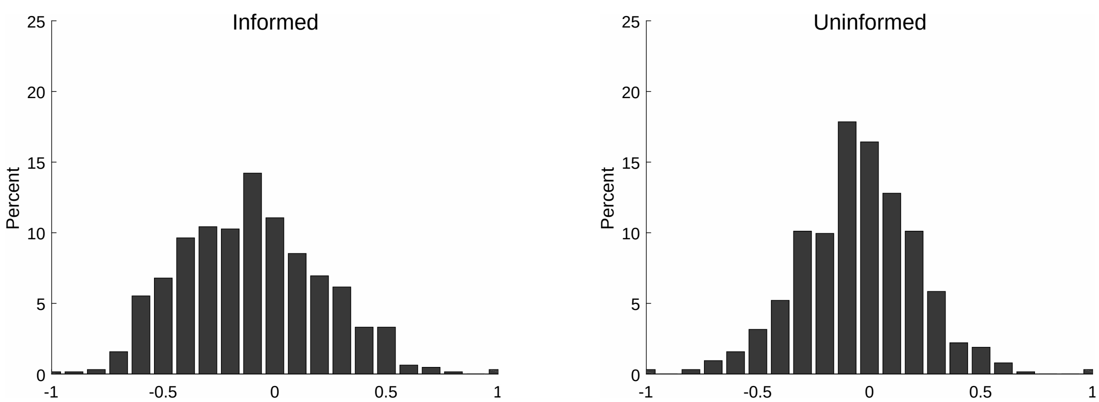

Figure 5: Distribution of daily Disagreement *

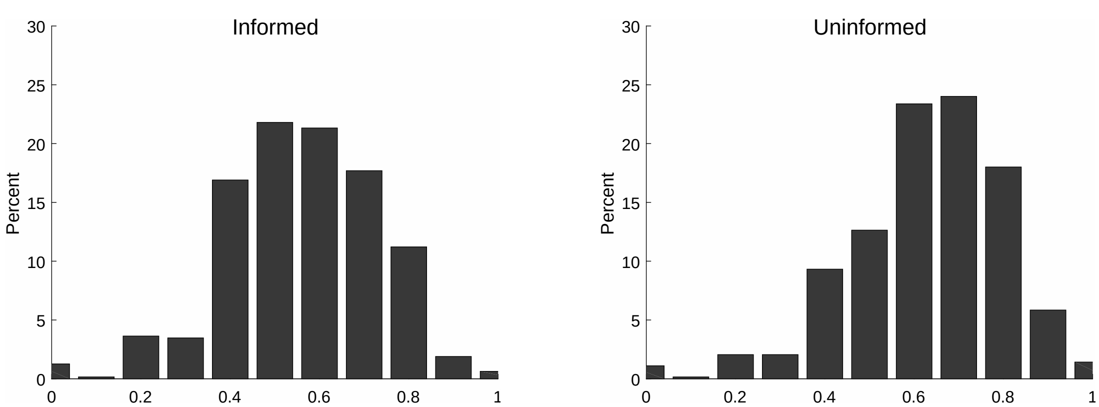

*Disagreement is defined as cross sectional variance of Twitter Sentiment across the individuals.

Figure 6: Impulse response of exchange rate and average Twitter Sentiment to fundamental and noise shocks *

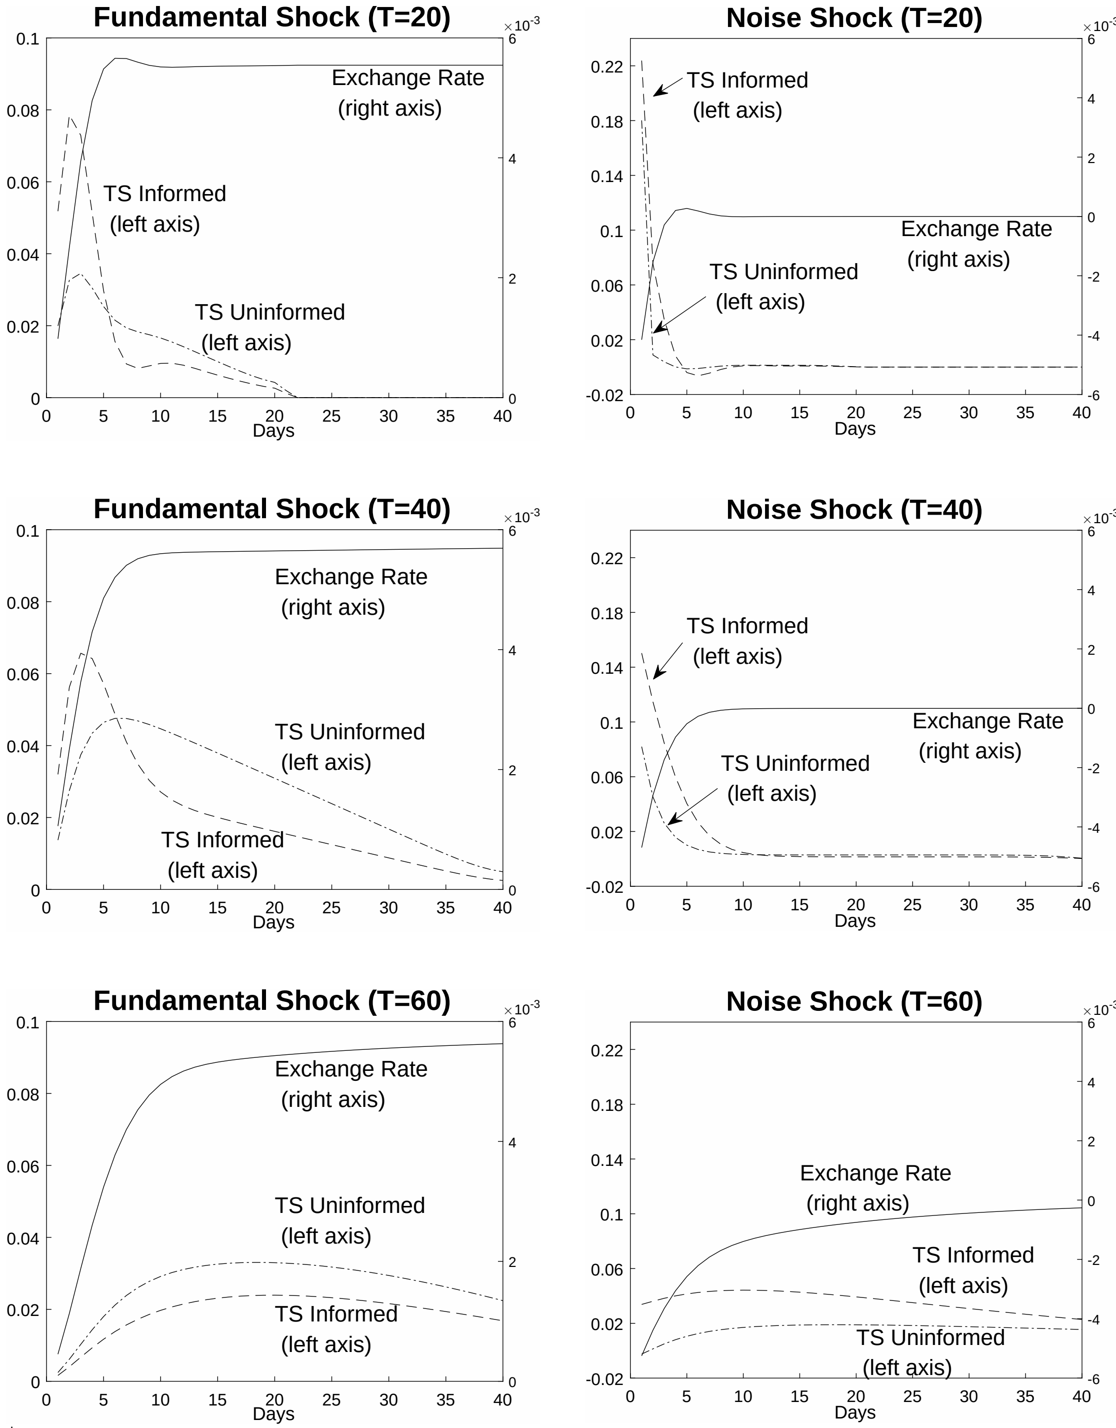

*Average Twitter Sentiment is the average of individual Twitter Sentiment if there are an infinite number of tweets.

Figure 7: Distribution of the Sharpe Ratio of the trading strategy in the model*

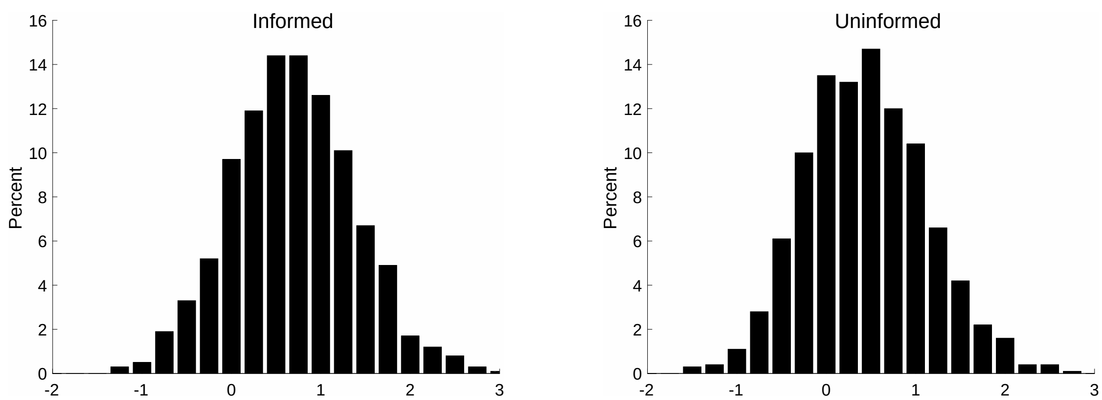

*The Sharpe Ratios are computed using 31 non-overlapping monthly observations in every 633 trading days across 1000 simulations. Twitter Sentiment and its 20 lags are used to compute the trading signal.

Figure 8: Distribution of t-stat for predictability regressions in the model*

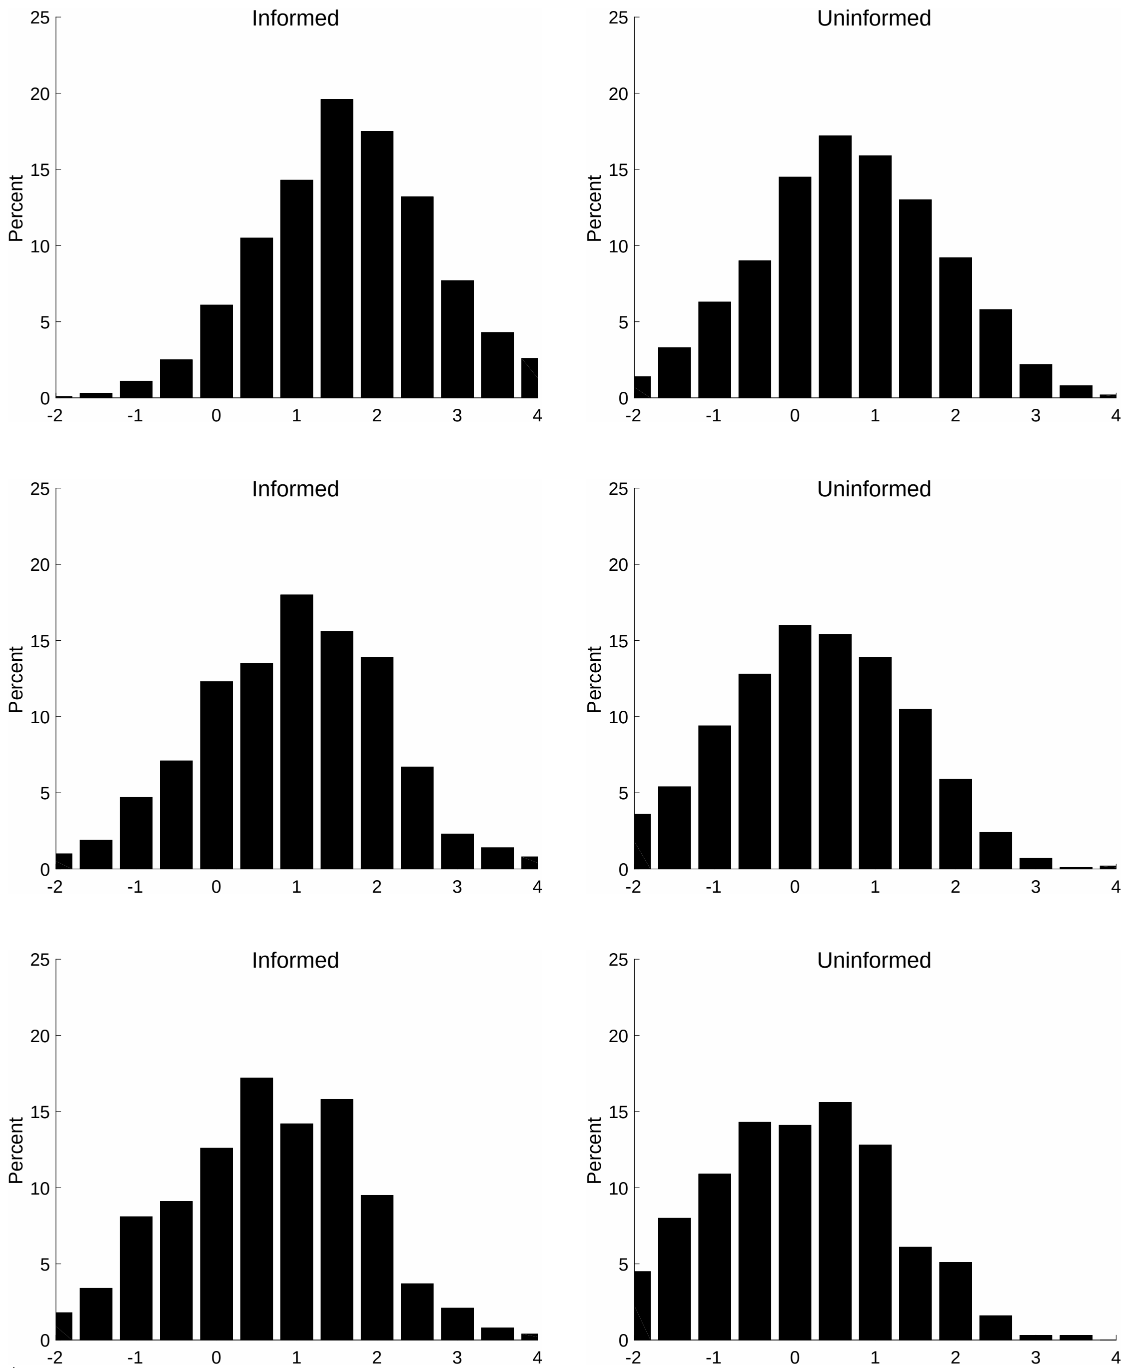

*The exchange rate and Twitter Sentiment (TS) are simulated over 633 days across 1000 simulations based on the benchmark parameters for $T = 40$. The change in the exchange rate over 20, 40 and 60 trading days is regressed on TS, for both informed and uninformed agents. The dependent variable in the regressions corresponding to top, middle and bottom figures is the subsequent 20, 40 and 60 trading days respectively.

Table 1: Examples of positive, neutral and negative tweets

Score	Category	Text
+1	Positive	$EURUSD Risks Higher High on Dovish Fed.
+1	Positive	$EURUSD: Buy dips near term.
+1	Positive	Buy eurusd market 1.2370 Stop: 1.2230 Target : 1.2600
+1	Positive	Looking to buy eurusd 1.1330
+1	Positive	Stay Long $EURUSD 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 bullish & eventually expect a rally twd 1.4045 Fibo lvl
+1	Positive	Stay calm, hold EURUSD long and USDJPY short
+1	Positive	USD Will Resume Decline; Keep EUR/USD Long For A Run Above 1.40
+1	Positive	Dollar to Face Further Losses on Dismal NFP- EURUSD to Target 1.3960
0	Neutral	I might consider selling $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 steady ahead of the Building Permits data from the United States. FOMC Meeting Minutes on focus.
0	Neutral	$EURUSD sits tight and awaits the FOMC fireworks. Levels to eye.
-1	Negative	EUR/USD Set For Further Falls With Bullish Signal Missing.
-1	Negative	$EURUSD Risks Further Losses as Growth Outlook Deteriorates.
-1	Negative	$EURUSD The pair remain bearish and looking for 1.1922 area when a 100% extension will happen .
-1	Negative	Stay Short $EURUSD, Long $USDJPY, & Resell $AUDUSD
-1	Negative	I expect $eurusd move lower, 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 Sell Signal.
-1	Negative	... said sell $EURUSD on interest rate differentials, TP 1.2800, SL 1.3700. Fair value at 1.3200. #TradersNotes #FX
-1	Negative	After ECB & Euro Squeeze, ... Adds To $EURUSD Short Exposure.
-1	Negative	We re looking to big gap at usd pairs. #eurusd will fall to the 1.23 this week.

Table 2: Directional moments in the data *

	$s_{t+20} - s_t$	$s_{t+40} - s_t$	$s_{t+60} - s_t$
Informed	0.0388	0.0398	0.0433
Uninformed	-0.0064	-0.0157	-0.0093

*Percentage of tweets that correctly forecast the direction of subsequent exchange rate changes 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)
	$s_{t+20} - s_t$	$s_{t+20} - s_t$	$s_{t+40} - s_t$	$s_{t+40} - s_t$	$s_{t+60} - s_t$	$s_{t+60} - s_t$
TS_t^I	0.342	0.205	1.196	0.748	1.418	0.975
	(0.456)	(0.310)	(0.840)	(0.495)	(1.084)	(0.686)
TS_{t-1}^I		0.082		0.501		0.698
		(0.288)		(0.428)		(0.571)
TS_{t-2}^I		0.277		0.652		0.738
		(0.268)		(0.463)		(0.544)
TS_{t-3}^I		0.185		0.771		0.531
		(0.256)		(0.499)		(0.522)
TS_{t-4}^I		0.266		0.792		0.663
		(0.286)		(0.537)		(0.568)
Observations	633	629	633	629	633	629
p-value*	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)
	$s_{t+20} - s_t$	$s_{t+20} - s_t$	$s_{t+40} - s_t$	$s_{t+40} - s_t$	$s_{t+60} - s_t$	$s_{t+60} - s_t$
TS_t^U	-0.367	-0.278	-0.766	-0.599	-0.675	-0.468
	(0.474)	(0.360)	(0.818)	(0.606)	(1.053)	(0.777)
TS_{t-1}^U		-0.319		-0.559		-0.534
		(0.315)		(0.490)		(0.656)
TS_{t-2}^U		-0.100		-0.200		-0.199
		(0.291)		(0.485)		(0.567)
TS_{t-3}^U		-0.146		-0.0643		-0.401
		(0.318)		(0.577)		(0.641)
TS_{t-4}^U		-0.304		-0.148		-0.535
		(0.344)		(0.584)		(0.678)
Observations	633	629	633	629	633	629
p-value*	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: Parameter Estimates *

	$T = 40$		$T = 20$		$T = 60$	
σ_v^I	0.0278	(0.0038)	0.0142	(0.0011)	0.0910	(0.0143)
σ_v^U	0.0572	(0.0078)	0.0241	(0.0018)	0.1820	(0.0283)
σ_b	2.7142	(0.5643)	9.2538	(1.7189)	0.4681	(0.0721)
ρ	0.4503	(0.0938)	0.5950	(0.0659)	0.7974	(0.0357)
ρ_b	0.5560	(0.0252)	0.0000	(0.0458)	0.9588	(0.0015)
n	0.0003	(0.0094)	0.0858	(0.0166)	0.3706	(0.0252)
σ_f	0.0031		0.0022		0.0012	

* Standard errors are in parentheses. The optimal weighting matrix used in the parameter estimation is the inverse of the variance of the model moments when $T = 40$. Only the diagonal elements of the weighting matrix are used. The same weighting matrix is used to estimate the parameters when $T = 20$ and $T = 60$

Table 6: Data and Model Moments*

		$T=40$		$T=20$		$T=60$	
	Data	Model	Cost	Model	Relative Cost	Model	Relative Cost
Variance of Twitter Sentiment							
TS_t^I	0.0983	0.1098	1.7144	0.1100	0.0510	0.1058	-0.9916
TS_t^U	0.0682	0.0855	3.4599	0.0812	-1.5044	0.0868	0.5376
Correlation Twitter Sentiment							
TS_t^I, TS_t^U	0.4643	0.4486	0.0823	0.4913	0.1606	0.4406	0.1054
Disagreement							
Mean Informed	0.5634	0.5640	0.0066	0.5648	0.0262	0.5636	-0.0062
Mean Uninformed	0.6356	0.6258	1.4209	0.6317	-1.1933	0.6219	1.3237
Variance Informed	0.0260	0.0230	4.0786	0.0230	0.1264	0.0229	0.4259
Variance Uninformed	0.0264	0.0205	13.1007	0.0199	2.9060	0.0209	-1.4111
Predictive Correlations							
$TS_t^I, s_{t+20}-s_t$	0.0448	0.1198	0.8391	0.0698	-0.7461	0.1504	0.8244
$TS_t^U, s_{t+20}-s_t$	-0.0400	0.0560	0.9174	0.0485	-0.1371	0.0687	0.2595
$TS_t^I, s_{t+40}-s_t$	0.1107	0.0756	0.1408	0.0449	0.3533	0.1250	-0.1174
$TS_t^U, s_{t+40}-s_t$	-0.0590	0.0219	0.4773	0.0290	0.0879	0.0463	0.3326
$TS_t^I, s_{t+60}-s_t$	0.1031	0.0513	0.2711	0.0315	0.2475	0.1025	-0.2710
$TS_t^U, s_{t+60}-s_t$	-0.0409	-0.0011	0.1017	0.0181	0.1215	0.0242	0.1707
Directional Moments							
$TS_t^I, s_{t+20}-s_t$	0.0388	0.0364	0.0069	0.0210	0.3865	0.0455	0.0479
$TS_t^U, s_{t+20}-s_t$	-0.0064	0.0197	0.7040	0.0133	-0.3039	0.0247	0.2955
$TS_t^I, s_{t+40}-s_t$	0.0398	0.0264	0.1707	0.0148	0.4277	0.0411	-0.1690
$TS_t^U, s_{t+40}-s_t$	-0.0157	0.0146	0.6814	0.0093	-0.2174	0.0232	0.4411
$TS_t^I, s_{t+60}-s_t$	0.0433	0.0216	0.3852	0.0127	0.3794	0.0369	-0.3512
$TS_t^U, s_{t+60}-s_t$	-0.0093	0.0110	0.2567	0.0079	-0.0725	0.0201	0.2821
Weekly Contemporaneous Correlations							
TS_w^I, s_t-s_{t-4}	0.3539	0.3515	0.0010	0.3892	0.2169	0.1834	5.0954
TS_w^U, s_t-s_{t-4}	0.2553	0.2558	0.0000	0.2264	0.1602	0.2207	0.2297
Exchange Rate Moments							
St. Dev. Δs	0.5718	0.5718	0.0000	0.5718	0.0000	0.5718	0.0000
Auto Corr. Δs	0.0028	0.0107	0.6088	0.0102	-0.0839	0.0110	0.0426
Auto Corr. Δs_w	0.0076	-0.0156	0.0671	-0.0166	0.0057	-0.0092	-0.0321
Objective		29.4927	29.4927	30.8908	1.3981	36.5567	7.0641

* "Cost" is the square of the difference between the model and data moment, divided by the variance of the corresponding moment. "Relative Cost" is the cost relative to that under the benchmark of $T=40$. 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 7: Excluding Sets of Moments*

	Benchmark	**Cost**	**(1)**	**Relative Cost**	**(2)**	**Relative Cost**	**(3)**	**Relative Cost**	**(4)**	**Relative Cost**
Variance of Twitter Sentiment										
TS_t^I	0.1098	1.7144	0.1096	-0.0579	0.1099	0.0288	0.1098	-0.0132	0.1095	-0.0923
TS_t^U	0.0855	3.4599	0.0854	-0.0098	0.0830	-0.9214	0.0829	-0.9387	0.0917	2.9230
Correlation Twitter Sentiment										
TS_t^I, TS_t^U	0.4486	0.0823	0.4552	-0.0548	0.4514	-0.0267	0.4624	-0.0812	0.4521	-0.0331
Disagreement										
Mean Informed	0.5640	0.0066	0.5643	0.0055	0.5640	-0.0009	0.5642	0.0035	0.5643	0.0074
Mean Uninformed	0.6258	1.4209	0.6259	-0.0237	0.6285	-0.6749	0.6287	-0.7100	0.6192	2.5465
Var Informed	0.0230	4.0786	0.0230	0.0542	0.0230	-0.0191	0.0230	0.0209	0.0230	0.0746
Var Uninformed	0.0205	13.1007	0.0205	0.0305	0.0202	1.3730	0.0202	1.4385	0.0213	-3.1495
Predictive Correlations										
$TS_t^I, s_{t+20}-s_t$	0.1198	0.8391	0.1180	-0.0414	0.1309	0.2672	0.1317	0.2866	0.0986	-0.4073
$TS_t^U, s_{t+20}-s_t$	0.0560	0.9174	0.0574	0.0270	0.0673	0.2282	0.0727	0.3464	0.0382	-0.3090
$TS_t^I, s_{t+40}-s_t$	0.0756	0.1408	0.0742	0.0111	0.0839	-0.0592	0.0846	-0.0630	0.0596	0.1571
$TS_t^U, s_{t+40}-s_t$	0.0219	0.4773	0.0232	0.0160	0.0311	0.1154	0.0358	0.1783	0.0074	-0.1552
$TS_t^I, s_{t+60}-s_t$	0.0513	0.2711	0.0501	0.0123	0.0586	-0.0707	0.0592	-0.0760	0.0373	0.1668
$TS_t^U, s_{t+60}-s_t$	-0.0011	0.1017	0.0003	0.0075	0.0077	0.0500	0.0123	0.0803	-0.0147	-0.0578
Directional Moments										
$TS_t^I, s_{t+20}-s_t$	0.0364	0.0069	0.0358	0.0040	0.0392	-0.0067	0.0394	-0.0064	0.0310	0.0675
$TS_t^U, s_{t+20}-s_t$	0.0197	0.7040	0.0198	0.0085	0.0215	0.1035	0.0226	0.1646	0.0167	-0.1524
$TS_t^I, s_{t+40}-s_t$	0.0264	0.1707	0.0261	0.0088	0.0284	-0.0458	0.0285	-0.0485	0.0226	0.1123
$TS_t^U, s_{t+40}-s_t$	0.0146	0.6814	0.0148	0.0084	0.0159	0.0597	0.0166	0.0935	0.0124	-0.0972
$TS_t^I, s_{t+60}-s_t$	0.0216	0.3852	0.0213	0.0117	0.0233	-0.0565	0.0233	-0.0561	0.0185	0.1203
$TS_t^U, s_{t+60}-s_t$	0.0110	0.2567	0.0111	0.0040	0.0121	0.0306	0.0127	0.0462	0.0091	-0.0442
Weekly Contemporaneous Moments										
TS_w^I, s_t-s_{t-4}	0.3515	0.0010	0.3518	-0.0002	0.3547	-0.0009	0.3541	-0.0010	0.3554	-0.0006
TS_w^U, s_t-s_{t-4}	0.2558	0.0000	0.2564	0.0002	0.2555	-0.0000	0.2560	0.0001	0.2584	0.0018
Exchange Rate Moments										
St. Dev. Δs	0.5718	0.0000	0.5718	0.0000	0.5718	0.0000	0.5718	0.0000	0.5718	0.0000
Auto Corr. Δs	0.0107	0.6088	0.0107	-0.0034	0.0107	-0.0086	0.0106	-0.0177	0.0108	0.0043
Auto Corr. Δs_w	-0.0156	0.0671	-0.0157	0.0004	-0.0161	0.0026	-0.0161	0.0028	-0.0148	-0.0045
ΔCost										
Included Moments				-0.0265		-0.1633		-0.2960		-0.8754
Excluded Moments				0.0454		0.5309		0.9459		2.5539
		s.e.		s.e.		s.e.		s.e.		s.e.
σ_v^I	0.0278	(0.0038)	0.0283	(0.0046)	0.0302	(0.0041)	0.0322	(0.0003)	0.0254	(0.0043)
σ_v^U	0.0572	(0.0078)	0.0566	(0.0093)	0.0605	(0.0082)	0.0618	(0.0007)	0.0534	(0.0090)
σ_b	2.7142	(0.5643)	2.9619	(0.6141)	3.1363	(0.6312)	3.8012	(0.2531)	2.7901	(0.7245)
σ_f	0.0031		0.0036		0.0036		0.0057		0.0037	
ρ	0.4503	(0.0938)	0.3671	(0.1266)	0.3753	(0.1033)	0.0000		0.3505	(0.1371)
ρ_b	0.5560	(0.0252)	0.5254	(0.0178)	0.5073	(0.0260)	0.4317	(0.0243)	0.5227	(0.0314)
n	0.0003	(0.0094)	0.0000		0.0000	(0.0108)	0.0119	(0.0121)	0.0267	(0.0124)

*The Table reports results when different sets of moments are excluded from the estimation. The excluded set of moments is shaded. ΔCost for "Included Moments" refers to the change in the overall cost of all non-excluded moments relative to the benchmark. ΔCost for "Excluded Moments" refers to the change in the overall cost of the excluded moments relative to the benchmark.

Table 7: Excluding Sets of Moments - Continued

	Benchmark	Cost	(5)	Relative Cost	(6)	Relative Cost	(7)	Relative Cost	(8)	Relative Cost
Variance of Twitter Sentiment										
TS_t^I	0.1098	1.7144	0.1169	2.7709	0.1087	-0.3166	0.1052	-1.1011	0.1096	-0.0691
TS_t^U	0.0855	3.4599	0.0922	3.2169	0.0831	-0.8989	0.0798	-1.9080	0.0853	-0.0647
Correlation Twitter Sentiment										
TS_t^I, TS_t^U	0.4486	0.0823	0.4573	-0.0658	0.4453	0.0381	0.4310	0.2876	0.4012	1.2457
Disagreement										
Mean Informed	0.5640	0.0066	0.5569	0.7259	0.5625	0.0090	0.5687	0.4778	0.5643	0.0056
Mean Uninformed	0.6258	1.4209	0.6187	2.7897	0.6273	-0.3938	0.6318	-1.2084	0.6258	0.0054
Var Informed	0.0230	4.0786	0.0237	-1.6927	0.0231	-0.1351	0.0226	1.3348	0.0230	0.0599
Var Uninformed	0.0205	13.1007	0.0214	-3.3865	0.0203	0.9517	0.0199	3.1632	0.0205	0.0203
Predictive Correlations										
$TS_t^I, s_{t+20} - s_t$	0.1198	0.8391	0.0996	-0.3912	0.0824	-0.6288	0.1470	0.7165	0.1377	0.4477
$TS_t^U, s_{t+20} - s_t$	0.0560	0.9174	0.0380	-0.3108	0.0381	-0.3095	0.0764	0.4323	0.0509	-0.0940
$TS_t^I, s_{t+40} - s_t$	0.0756	0.1408	0.0609	0.1417	0.0590	0.1640	0.0954	-0.1142	0.0887	-0.0855
$TS_t^U, s_{t+40} - s_t$	0.0219	0.4773	0.0076	-0.1538	0.0174	-0.0508	0.0380	0.2088	0.0170	-0.0560
$TS_t^I, s_{t+60} - s_t$	0.0513	0.2711	0.0390	0.1445	0.0444	0.0768	0.0683	-0.1487	0.0626	-0.1052
$TS_t^U, s_{t+60} - s_t$	-0.0011	0.1017	-0.0144	-0.0566	0.0013	0.0128	0.0137	0.0896	-0.0065	-0.0256
Directional Moments										
$TS_t^I, s_{t+20} - s_t$	0.0364	0.0069	0.0321	0.0480	0.0256	0.2076	0.0425	0.0101	0.0411	-0.0004
$TS_t^U, s_{t+20} - s_t$	0.0197	0.7040	0.0167	-0.1539	0.0144	-0.2582	0.0231	0.1943	0.0189	-0.0417
$TS_t^I, s_{t+40} - s_t$	0.0264	0.1707	0.0233	0.0903	0.0207	0.1773	0.0306	-0.0893	0.0296	-0.0719
$TS_t^U, s_{t+40} - s_t$	0.0146	0.6814	0.0123	-0.0996	0.0121	-0.1073	0.0169	0.1074	0.0140	-0.0253
$TS_t^I, s_{t+60} - s_t$	0.0216	0.3852	0.0192	0.0888	0.0181	0.1364	0.0250	-0.1102	0.0242	-0.0862
$TS_t^U, s_{t+60} - s_t$	0.0110	0.2567	0.0092	-0.0437	0.0098	-0.0275	0.0129	0.0515	0.0103	-0.0168
Weekly Contemporaneus Correlations										
$TS_w^I, s_t - s_{t-4}$	0.3515	0.0010	0.3590	0.0035	0.1289	8.8724	0.3393	0.0365	0.3531	-0.0009
$TS_w^U, s_t - s_{t-4}$	0.2558	0.0000	0.2594	0.0032	0.1265	3.1894	0.2499	0.0056	0.2524	0.0016
Exchange Rate Moments										
St. Dev. Δs	0.5718	0.0000	0.5718	0.0000	0.5718	0.0000	0.5718	0.0000	0.5718	0.0000
Auto Corr. Δs	0.0107	0.6088	0.0107	0.0015	0.0108	0.0170	0.0107	-0.0026	0.0108	0.0078
Auto Corr. Δs_w	-0.0156	0.0671	-0.0147	-0.0051	-0.0118	-0.0204	-0.0157	0.0003	-0.0145	-0.0061
ΔCost										
Included Moments				-2.3226		-1.3560		-2.0642		-0.2011
Excluded Moments				5.9878		12.0617		4.4980		1.2457
		s.e.		s.e.		s.e.		s.e.		s.e.
σ_v^I	0.0278	(0.0038)	0.0239	(0.0044)	0.0349	(0.0189)	0.0365	(0.0053)	0.0323	(0.0048)
σ_v^U	0.0572	(0.0078)	0.0528	(0.0098)	0.0711	(0.0391)	0.0733	(0.0107)	0.0787	(0.0115)
σ_b	2.7142	(0.5643)	3.0918	(0.7574)	1.1239	(0.6285)	2.5598	(0.5012)	2.2320	(0.4559)
σ_f	0.0031		0.0050		0.0015		0.0038		0.0038	
ρ	0.4503	(0.0938)	0.1253	(0.1923)	0.7374	(0.1787)	0.3331	(0.1129)	0.3428	(0.1181)
ρ_b	0.5560	(0.0252)	0.4794	(0.0259)	0.9289	(0.0030)	0.5671	(0.0210)	0.5696	(0.0206)
n	0.0003	(0.0094)	0.0287	(0.0083)	0.9990	(0.1784)	0.0000	(0.0094)	0.0402	(0.0071)

Table 8: Predictability and Sharpe Ratios (Benchmark) *

	Informed		Uninformed	
	Sharpe Ratios	R^2	*Sharpe Ratios*	R^2
$L1$	0.65	0.05	0.41	0.02
$L5$	0.68	0.06	0.44	0.02
$L10$	0.68	0.06	0.46	0.03
$L15$	0.68	0.06	0.46	0.03
$L20$	0.68	0.06	0.46	0.03

* The exchange rate and Twitter Sentiment (TS) are simulated over 200,000 days based on the benchmark parameters for $T = 40$. The change in the exchange rate over 20 trading days is regressed on TS and its lags, for both informed and uninformed agents, and the R^2 is reported in the table. The Sharpe Ratios indicate the performance of a trading strategy that takes long or short foreign currency position based on the sign of the TS signal. The TS signal is a linear combination of TS and its lags. Li in each row shows that i lags of TS are used in the predictive regression and to compute the TS trading signal.

Table 9: Predictability and Sharpe Ratios $T = 20$ and $T = 60$

	Informed		Uninformed	
	Sharpe Ratios	R^2	*Sharpe Ratios*	R^2
	$T = 20$			
$L1$	0.58	0.04	0.43	0.02
$L5$	0.59	0.04	0.44	0.02
$L10$	0.59	0.04	0.45	0.03
$L15$	0.59	0.04	0.45	0.03
$L20$	0.59	0.04	0.45	0.03
	$T = 60$			
$L1$	0.59	0.04	0.33	0.01
$L5$	0.61	0.05	0.34	0.02
$L10$	0.65	0.05	0.36	0.02
$L15$	0.65	0.05	0.36	0.02
$L20$	0.65	0.05	0.36	0.02

Table 10: Sensitivity of Predictability and Sharpe Ratios to Parameters *

		Informed		Uninformed	
		Sharpe Ratios	R^2	*Sharpe Ratios*	R^2
	Benchmark	0.68	0.06	0.46	0.03
σ_v^I	$+2s.e.$	0.71	0.06	0.48	0.03
	$-2s.e.$	0.69	0.06	0.46	0.03
σ_v^U	$+2s.e.$	0.70	0.06	0.45	0.03
	$-2s.e.$	0.69	0.06	0.46	0.03
σ_b	$+2s.e.$	0.69	0.06	0.46	0.03
	$-2s.e.$	0.70	0.06	0.48	0.03
ρ	$+2s.e.$	0.71	0.06	0.48	0.03
	$-2s.e.$	0.68	0.05	0.44	0.02
ρ_b	$+2s.e.$	0.72	0.06	0.46	0.03
	$-2s.e.$	0.72	0.06	0.48	0.03
n	$+2s.e.$	0.68	0.06	0.44	0.02

* Each parameter is increased (or decreased) by 2 standard errors from its benchmark value and other parameters are re-estimated. The resulting parameter estimates are used to simulate the exchange rate and TS. The table shows the Sharpe Ratios of the trading strategy and the R^2 of the predictive regression for each set of parameters. The benchmark is $T = 40$ as in Table 5 and 20 lags of TS are used in the predictive regression and to compute the TS trading signal.

Table 11: Predictability and Sharpe Ratios after Excluding Sets of Moments.*

	Informed		Uninformed	
	Sharpe Ratios	R^2	*Sharpe Ratios*	R^2
Excluded Moments				
Directional	0.69	0.06	0.47	0.03
Predictive Correlations	0.71	0.06	0.49	0.03
Directional and Predictive Correlations	0.73	0.06	0.52	0.03
Mean Disagreement	0.62	0.05	0.41	0.02
Variance of Twitter Sentiment	0.65	0.05	0.41	0.02
Weekly Contemporaneous Correlations	0.39	0.02	0.23	0.01
Variance of Disagreement	0.78	0.07	0.52	0.03
Correlation between TS Informed and Uninformed	0.73	0.06	0.43	0.02

* Sharpe Ratios and R^2 are shown when removing various sets of moments from the estimation for $T = 40$ as in Table 7. The results are reported for 20 lags and can be compared to the last row of Table 8.

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 ...**
”*buy??eur*” or ”*buy?eur*”	”*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*”
”*buy*lot*eur*”	”*close*buy*lot*eur*”
”*long??eur*”	”*long?term*”, ”*was?long*”, ”*close*long??eur*”, ”*close*long?eur*”, ”*exit*long??eur*”, ”*exit*long?eur*”
”*bullish*”	”*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*”
”*covered*short*”	”*short?term*”
”*buy?the?eur*”	”*never?buy?the?eur*”, ”*not?buy?the?eur*”
”*eur?usd*look?good*”	”*eur?usd*not*look?good*”
”*eur?usd*looks?good*”	”*eur?usd*not*looks?good*”
”*double*long*”	”*long?term*”
”*took*long*position*”	”*long?term*”
”*out*of*eur*short*”	”*short?term*” , ”*stop*out*of*eur*short*”
”*add*eur*long*”	”*long?term*” , ”*addict*” , ”*dadd*”

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*” ”*decrease*eur?usd?short*” ”*hold*eur\usd?short*” ”*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*” ”*increase*short?eurusd*” ”*increase*eur?usd?short*” ”*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

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