Information Asymmetry and Expectations about Returns[1]

Vahid Gholampour

A Dissertation presented to the Graduate Faculty of the University of Virginia in Candidacy for the Degree of Doctor of Philosophy

Department of Economics

University of Virginia

September 2016

[1]I thank Professors Eric van Wincoop, Eric Young, Michael Gallmeyer and Zach Bethune for their extraordinary support. I gratefully acknowledge financial support from the Bankard Fund for Political Economy and the Radulovacki Fund at the University of Virginia.

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. In the first chapter, opinionated tweets about the Euro/dollar exchange rate are used 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.

The second chapter introduces a new measure of stock market investor sentiment based on the opinions shared on Twitter. The main advantage of the proposed index over existing measures of sentiment is the possibility of using the number of followers as a proxy for the quality of private signals. Moreover, the data allows for gauging sentiment directly with high frequency data. The index is used to test the implications of theories in asset pricing. The results show that (1) the follower-weighted sentiment index predicts the same day return of the stock market index, but the equal-weighted index has no predictive power for daily returns, (2) dispersion of expectations about future returns predicts volatility of the stock market returns, (3) information asymmetry is positively related to return volatility, and (4) the density of information arrival measured by the number of opinionated tweets is positively correlated with volatility and trading volume of the stock market index.

Chapter 1

What Can we Learn from Euro-Dollar Tweets ?

1.1 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 be-

[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.

come 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 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.

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. As we will discuss below, related literature that uses social media to forecast asset prices is based on a data-only approach. This has its drawbacks. 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 perhaps 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

[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.

trading days).[6]

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 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.

Several steps are taken to connect the verbal tweets in the data with expectations of future exchange rate changes 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

[6]It should be noted that 633 trading days is actually long in comparison to related literature discussed below that has looked at the predictability of asset prices by social media messages, which generally uses samples of no more than one year.

[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.

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. 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 classify messages. Most of the literature uses supervised machine-based learning classifiers that are not specific to financial markets at all. For example, the Nave 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.

1.2 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

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

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.

1.2.1 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.

1.2.2 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

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

are usually posted by investors, there is a certain type of lexicon that is found in most of these tweets. We identify this lexicon by studying large numbers of tweets. We then go through several rounds of improving our dictionary of financial lexicon by comparing the results from the automated classification to that based on manual classification. We stopped making further changes when we found only very few errors after manually checking 5000 tweets. We describe this dictionary further below.

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

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

1.2.3 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

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

them are underlined.

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

1.2.4 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

[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.

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

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

1.2.5 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

$$TS_t^j = \frac{1}{n_t^j} \sum_{i=1}^{n_t^j} TS_t^{j,i} \tag{1.1}$$

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.

1.2.6 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 and 3.

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 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 if the last 5 days before the forecasting period (TS_t^j through TS_{t-4}^j). The results are reported in Table 3 for both the informed and uninformed groups. 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. The coefficients are either insignificant or of the wrong sign when they are significant. 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.

1.2.7 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 5. 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 for which the number of tweets is 0 or 1.[12] 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 ex-

[12] 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.

change 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.

1.3 Model Description

It should be emphasized from the start that the concept of a "tweet" does not exist within the model that we are about to describe. Tweets in reality are just an expression of a belief about the direction of the exchange rate by a subset of agents. Individuals have many sources of information, including that from reading tweets by other people. In the end these beliefs still differ among individuals, reflected in a dispersion of opinions expressed through tweets. In the model the source of this dispersion is private signals, which can be thought of as related to different research findings, focusing on different pieces of information or reading different newspaper articles or different tweets. In the next section we will relate expectations of exchange rate changes that exist in the model to directional beliefs expressed through tweets.

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

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

$$\int_0^n b_{F,t}^{I,i} di + \int_n^1 b_{F,t}^{U,i} di = 0 \tag{1.2}$$

Portfolio demand by agents is

$$b_{F,t}^{I,i} = \frac{E_t^{I,i} s_{t+1} - s_t + i_t^* - i_t}{\gamma \sigma_I^2} - b_t^{I,i} \tag{1.3}$$

$$b_{F,t}^{U,i} = \frac{E_t^{U,i} s_{t+1} - s_t + i_t^* - i_t}{\gamma \sigma_U^2} - b_t^{U,i} \tag{1.4}$$

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.[13] 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:

$$b_t = \int_0^n b_t^{I,i} di + \int_n^1 b_t^{U,i} di \tag{1.5}$$

for which we assume an AR process:

$$b_t = \rho_b b_{t-1} + \varepsilon_t^b \tag{1.6}$$

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:

$$m_t = p_t + y_t - \alpha i_t \tag{1.7}$$

$$m_t^* = p_t^* + y_t^* - \alpha i_t^* \tag{1.8}$$

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

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

are denoted with a * superscript. Using PPP ($p_t = s_t + p_t^*$), subtracting these equations yields

$$i_t - i_t^* = \frac{1}{\alpha}(s_t - f_t) \tag{1.9}$$

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

$$f_t - f_{t-1} = \rho(f_{t-1} - f_{t-2}) + \varepsilon_t^f \tag{1.10}$$

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 (1.3), (1.4) and (1.9) into the market clearing condition (1.2), we have

$$\omega \bar{E}_t^I s_{t+1} + (1-\omega)\bar{E}_t^U s_{t+1} - \frac{1+\alpha}{\alpha} s_t + \frac{1}{\alpha} f_t = \gamma \sigma^2 b_t \tag{1.11}$$

where

$$\omega = \frac{n/\sigma_I^2}{(n/\sigma_I^2) + ((1-n)/\sigma_U^2)} \tag{1.12}$$

$$\sigma^2 = \frac{1}{(n/\sigma_I^2) + ((1-n)/\sigma_U^2)} \tag{1.13}$$

Imposing the market clearing condition (1.11) allows us to solve for the equilibrium exchange rate.

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

$$v_t^{j,i} = f_{t+T} + \epsilon_t^{v,j,i} \tag{1.14}$$

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.

(1.14) 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

$$s_t = A(L)\varepsilon^f_{t+T} + B(L)\varepsilon^b_t \tag{1.15}$$

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[14]

$$\bar{E}^j_t s_{t+1} = \theta' \bar{E}^j_t(\xi_t) + A^*(L)\varepsilon^f_t + B^*(L)\varepsilon^b_{t-T} \tag{1.16}$$

$$\sigma^2_j = var^j_t(s_{t+1}) = a_1^2\sigma_f^2 + b_1^2\sigma_b^2 + \theta' var^j_t(\xi_t)\theta \tag{1.17}$$

where $\theta' = (a_2, a_3, ..., a_{T+1}, b_2, b_3, ..., b_{T+1})$, $\xi'_t = (\varepsilon^f_{t+T}, ..., \varepsilon^f_{t+1}, \varepsilon^b_t, ..., \varepsilon^b_{t-T+1})$, $A^*(L) = \sum_{i=1}^{\infty} a_{T+i+1}L^{i-1}$ and $B^*(L) = \sum_{i=1}^{\infty} b_{T+i+1}L^{i-1}$. The conditional variance $var^j_t(s_{t+1})$ 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 $v^{j,i}_t$,...,$v^{j,i}_{t-T+1}$ and knowledge of the unconditional distribution of ξ_t. Solving the signal extraction problem (see BvW) yields

$$\bar{E}^j_t(\xi_t) = \mathbf{M}_j\mathbf{H}'\xi_t \tag{1.18}$$

$$var^j_t(\xi_t) = \tilde{\mathbf{P}} - \mathbf{M}_j\mathbf{H}'\tilde{\mathbf{P}} \tag{1.19}$$

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 $(\sigma^j_v)^2$ on the last T elements of the diagonal and zeros otherwise, $\tilde{P}$ is the unconditional variance of ξ_t and

$$\mathbf{H}' = \begin{bmatrix} a_1 & a_2 & ... & a_T & b_1 & b_2 & ... & b_T \\ 0 & a_1 & ... & a_{T-1} & 0 & b_1 & ... & b_{T-1} \\ ... & ... & ... & ... & ... & ... & ... & ... \\ 0 & 0 & ... & a_1 & 0 & 0 & ... & b_1 \\ d_1 & d_2 & ... & d_T & 0 & 0 & ... & 0 \\ 0 & d_1 & ... & d_{T-1} & 0 & 0 & ... & 0 \\ ... & ... & ... & ... & ... & ... & ... & ... \\ 0 & 0 & 0 & d_1 & 0 & 0 & ... & 0 \end{bmatrix} \tag{1.20}$$

[14] The innovations ε^f_{t-s} are known at t for $s \geq 0$. The innovations ε^b_{t-T-s} 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.

Substituting (1.18) and (1.19) into (1.16) and (1.17) and the result into the market clearing condition (1.11), we have

$$\theta' \left(\omega \mathbf{M}_I + (1-\omega)\mathbf{M}_U\right) \mathbf{H}'\xi_t - \frac{1+\alpha}{\alpha}\left(A(L)\varepsilon^f_{t+T} + B(L)\varepsilon^b_t\right) + \frac{1}{\alpha}D(L)\varepsilon^f_t$$
$$+A^*(L)\varepsilon^f_t + B^*(L)\varepsilon^b_{t-T} = \gamma\sigma^2\left(1+\rho_b L + \rho_b^2 L^2 + ...\right)\varepsilon^b_t \tag{1.21}$$

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

$$E^{j,i}_t s_{t+k} = \bar{E}^j_t s_{t+k} + \mathbf{z}'_k \mathbf{M}_j \mathbf{w}^{j,i}_t \tag{1.22}$$

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

$$\bar{E}^j_t s_{t+k} = \mathbf{z}'_k \bar{E}^j_t \xi_t + \sum_{l=0}^{\infty} a_{T+k+1+l}\varepsilon^f_{t-l} + \sum_{l=0}^{\infty} b_{T+k+1+l}\varepsilon^b_{t-T-l} \tag{1.23}$$

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 $\mathbf{z}'_k \mathbf{M}_j \mathbf{w}^{j,i}_t$ that depends on the signal errors of that agent.

1.4 Empirical Methodology

1.4.1 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

$$TS_t^{j,i} = \begin{cases} 1 & \text{if } \; E_t^{j,i}(s_{t+T} - s_t) > c^j \\ 0 & \text{if } \; -c^j \leq E_t^{j,i}(s_{t+T} - s_t) \leq c^j \\ -1 & \text{if } \; E_t^{j,i}(s_{t+T} - s_t) < -c^j \end{cases} \tag{1.24}$$

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:

$$Prob\left(-c^j \leq E_t^{j,i}(s_{t+T} - s_t) \leq c^j\right) = \pi^j \tag{1.25}$$

Since

$$E_t^{j,i}(s_{t+T} - s_t) = \bar{E}_t^j(s_{t+T} - s_t) + \mathbf{z}_T'\mathbf{M}_j\mathbf{w}_t^{j,i} \tag{1.26}$$

we can compute the unconditional variance of this expectation as

$$\sigma_E^2(j) = var(E_t^{j,i}(s_{t+T} - s_t)) = var(\bar{E}_t^j(s_{t+T} - s_t)) + \mathbf{z}_T'\mathbf{M}_j\mathbf{R}_j\mathbf{M}_j'\mathbf{z}_T \tag{1.27}$$

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

$$Prob\left(-\frac{c^j}{\sigma_E(j)} \leq \frac{E_t^{j,i}(s_{t+T} - s_t)}{\sigma_E(j)} \leq \frac{c^j}{\sigma_E(j)}\right) = \pi^j \tag{1.28}$$

it must be that

$$\Phi\left(\frac{-c^j}{\sigma_E(j)}\right) = \frac{1-\pi^j}{2} \tag{1.29}$$

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

$$c^j = -\sigma_E(j)\Phi^{-1}\left(\frac{1-\pi^j}{2}\right) \tag{1.30}$$

For informed and uninformed agents, in the data we have respectively $\pi^I = 0.328$ and $\pi^I = 0.288$ (see also Figure 3).[15]

For what follows, it is useful to characterize the distribution of TS_t^{ji} 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

$$TS_t^{i,j} = TS_t^j(x_t^j) + \epsilon_t^{j,i} \tag{1.31}$$

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

$$TS_t^j(x_t^j) = 1 - \Psi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) - \Psi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \tag{1.32}$$

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

$$\epsilon_t^{j,i} = \begin{cases} 1 - TS_t^j(x_t^j) & \text{with probability} \quad 1 - \Psi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) \\ -1 - TS_t^j(x_t^j) & \text{with probability} \quad \Psi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \\ -TS_t^j(x_t^j) & \text{with probability} \quad \Psi\left(\frac{c^j - x_t^j}{\sigma_w^j}\right) - \Psi\left(\frac{-c^j - x_t^j}{\sigma_w^j}\right) \end{cases} \tag{1.33}$$

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.

1.4.2 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

[15] 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.

5. 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 $\epsilon_t^{v,ji}$. However, a limited number of simulations 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^{ji} 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 $\varepsilon_t^{v,ji}$ translate into realizations of ϵ_t^{ji}, whose distribution is given by (1.33). We can then write a specific sample moment as $m = m(\mathbf{e}, \mathbf{x})$, where $\mathbf{e}$ consists of the realizations of ϵ_t^{ji} 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

$$TS_t^{ji} = TS_t^j(x_t^j) + \frac{1}{n_t^j} \sum_{i=1}^{n_t^j} \epsilon_t^{ji} \tag{1.34}$$

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

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

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 ϵ_t^{ji} given in (1.33). Doing so gives

$$var(TS_t^j(x_t^j)) + \sum_{t\in S} \frac{1}{Sn_t^j} var(\epsilon_t^{ji}) \tag{1.35}$$

Here the first variance is the sample variance of $TS_t^j(x_t^j)$, while $var(\epsilon_t^{ji})$ is computed from the distribution (1.33) 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

$$\frac{1}{S}\sum_{t\in S} \frac{n_t^j}{n_t^j - 1} \left(\frac{1}{n_t^j} \sum_{i=1}^{n_t^j} \left(\epsilon_t^{ji}\right)^2 - \left(\frac{1}{n_t^j} \sum_{i=1}^{n_t^j} \epsilon_t^{ji} \right)^2 \right) \tag{1.36}$$

The mean across the distribution of ϵ_t^{ji} is

$$\frac{1}{S}\sum_{t\in S} var(\epsilon_t^{ji}) \tag{1.37}$$

where the variance is again computed from (1.33) 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

$$\frac{1}{\sum_{t\in S} n_t^j} \sum_{t\in S} \sum_{i=1}^{n_t^j} u_t^{ji} \tag{1.38}$$

where

$$u_t^{ji} = \begin{cases} 1 & \text{if} \quad \text{sign}(TS_t^{ji}) = \text{sign}(s_{t+k} - s_t) \\ -1 & \text{if} \quad \text{sign}(TS_t^{ji}) = -\text{sign}(s_{t+k} - s_t) \\ 0 & \text{if} \quad TS_t^{ji} = 0 \end{cases} \tag{1.39}$$

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

$$\frac{1}{\sum_{t\in S} n_t^j} \sum_{t\in S} n_t^j TS_t^j(x_t^j) \text{sign}(s_{t+k} - s_t) \tag{1.40}$$

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 $\varepsilon_t^{v,ji}$.

1.4.3 Estimation of Model Parameters

We estimate the model using the simulated method of moments, based on the 24 moments in Table 5. The parameters are chosen in order to minimize

$$\left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right)' \mathbf{\Sigma}^{-1} \left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right) \tag{1.41}$$

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.[16]

The variance covariance matrix of parameter estimates is given by

$$\frac{1}{S}\left[\left(\frac{\partial \mathbf{m}^{model}}{\partial \nu}\right)' \mathbf{\Sigma}^{-1} \left(\frac{\partial \mathbf{m}^{model}}{\partial \nu}\right)\right]^{-1} \tag{1.42}$$

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

[16] 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.

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_u and T. We only estimate the first 6 of these parameters. Some comments are therefore in order about γ, σ_u and T. From (1.11) it can be seen that γ enters the model multiplied by b_t. As a result of this we can only estimate $\gamma\sigma_b$. We therefore normalize $\gamma = 1$ and estimate σ_b. If instead one wishes to set $\gamma = 10$ the reported estimate for σ_b below simply needs to divided by 10. We set σ_u by exploiting a scaling feature of the model. If we multiply σ_u, 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_u and estimate the other parameters based on moments other than the standard deviation of the exchange rate. Afterwards we scale σ_u, 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 (1.41). We will compare the objective function across these values of T.

1.4.4 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:

$$s_{t+20} - s_t = \alpha + \sum_{k=1}^{l} \beta_k TS^j_{t-i+1} + \varepsilon_{t+f} \tag{1.43}$$

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

$$s_{t+20} - s_t + \sum_{k=1}^{20}(i^*_{t+k-1} - i_{t+k-1}) = \alpha + \sum_{k=1}^{l} \beta_k TS^j_{t-k+1} + \varepsilon_{t+f} \tag{1.44}$$

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.

1.5 Results

Tables 4 through 10 report the results. We first discuss the estimated parameters, followed by the moments and finally consider results on predictability and Sharpe ratios.

1.5.1 Parameter Estimates

Table 4 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_u 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 $\sigma_v^I < \sigma_v^U$. Investors with a lot of followers are therefore indeed more informed. We can reject $\sigma_v^I = \sigma_v^U$ with a p-value of less than 0.001. We also see that σ_v^I and σ_v^U 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.

1.5.2 Moments

Table 5 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

$$\left(\frac{\mathbf{m}^{data}(j) - \mathbf{m}^{model}(j)}{\mathbf{\Sigma_{j,j}^{0.5}}}\right)^2 \tag{1.45}$$

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 (1.41), 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.

Table 5 also sheds light on why σ_v^I is estimated to be smaller than σ_v^U. The finding that investors with more followers are indeed more informed is supported by many of the moments. First, the higher variance of Twitter Sentiment by informed agents in the data is reflected in the model as more accurate signals lead to larger changes in opinion over time. Second, the higher disagreement among uninformed agents in the data is also consistent with the model as weaker signals imply a wider dispersion in beliefs. Third, the higher quality information of the informed group is naturally reflected in the predictability moments. The predictive correlations and directional moments are positive for the informed group in the data, while negative for the uninformed. Finally, the higher weekly contemporaneous correlation between Twitter Sentiment and the exchange rate change for the informed group is also consistent with more precise signals in 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.

While the uninformed have less information than the informed, one may wonder what aspects of the data allow us to determine that they have any information at all. Based on the predictability moments alone there is no reason to believe that the uninformed group has any information. In order to better understand this, we have looked at the effects on the moments when we increase σ_v^U while re-estimating the other parameters (not reported). 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. Second, the predictability moments of the informed group become too high relative to the data. Since n is close to 0, the exchange rate does not capture the additional information that the informed have. 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.

Table 6 reports the results of both model moments and parameter estimates after we exclude various subsets of moments. This provides insight into how important moments are to the results. We consider 8 subsets of excluded moments, which are highlighted in shaded grey. This means that the shaded moments, and the associated cost, are ignored in estimating the parameters. Note first of all that $\sigma_v^I < \sigma_v^U$ remains the case in all cases. As already discussed, the higher information quality of the informed group is reflected in a broad set of moments.

Table 6 provides insight into the role of various moments in the value of parameter estimates. For example, removing the predictability moments (predictive correlations and directional moments) raises the estimates of both σ_v^I and σ_v^U. This suggests that the quality of information is at least partially attributed to the predictability moments. The predictability moments also play a role in the estimates of ρ and ρ_b. They become lower when the predictability moments are removed. As discussed already, this implies faster learning about the fundamentals, which leads to excessive predictability.

If we remove the variance of Twitter Sentiment, σ_v^I, σ_v^U, ρ and ρ_b are all lower. This suggests that the estimates of these parameters are at least partially driven by the variance of Twitter Sentiment. Lower values of σ_v^I and σ_v^U improve the information quality, which leads to too much variation in Twitter Sentiment. Similarly, a lower ρ and ρ_b implies faster learning, which also raises the variance of Twitter Sentiment to levels that are too high.

Removing the weekly contemporaneous correlations between Twitter Sentiment and the exchange rate change leads to the biggest changes in estimated parameters, suggesting that these moments are most critical to the estimation. It leads to weaker quality signals, a considerable increase in ρ and ρ_b and an increase in n to almost 1. The weaker signals reduce the contemporaneous link between Twitter Sentiment and exchange rates. As we have seen in Figure 6, the same happens when noise and fundamental shocks are more persistent. Perhaps most surprising is the increase in n to almost 1. In that case almost all traders are informed, so that more information about future fundamentals is incorporated into the current exchange rate and less in Twitter Sentiment, reducing the contemporaneous correlation.

As a final example, when we remove the correlation between Twitter Sentiment and the informed and uninformed agents, there is a significant increase in σ_v^U and drop in the correlation. This moment therefore plays an important role in estimating the information quality of the uninformed group. The more correlated the Twitter Sentiment of the uninformed is with the informed group, the less far behind one expects the information quality of the uninformed to be.

1.5.3 Predictability and Sharpe Ratios

Table 7 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 then 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.[17]

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 7 for the benchmark parameters ($T = 40$). We consider various lags L in (1.43) and (1.44), 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

[17] 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 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.

Table 8 shows that predictability and Sharpe ratios are only slightly weaker for $T = 20$ and $T = 60$, but for the informed traders they remain considerably better than for interest differentials. Table 9 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 results are robust to changing parameters within the estimated confidence intervals.

Finally, Table 10 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 6. 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.

1.6 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.

Chapter 2

Quality of Private Signals and Expectations about Returns

2.1 Introduction

In this paper, Twitter messages about U.S. stock market are used as the data source to construct daily measures of investor sentiment, dispersion of sentiment, and density of information arrival. These high frequency measures are used to shed light on some of the theories in asset pricing and test their implications. More specifically, the paper is focused on three questions. Does heterogeneity in quality of private signals matter for constructing a measure of sentiment? Is dispersion of sentiment associated with more return volatility? Is the information arrival rate correlated with return volatility and trading volume?

The first question is whether the quality of private signals should be taken into account when aggregating individual opinions to measure investor sentiment. Surveys assign equal weight to individual opinions when aggregating the responses. In this study, the number of followers is used as an indication for the quality of information shared by individuals. In other words, the marketplace for information determines the weight of each opinionated tweet in the daily measure of sentiment. The regression results show that the follower-weighted sentiment index predicts the daily open-to-close return of the S&P500 index. The equal-weighted index, however, shows no predictive power for the daily stock market returns. An increase in the level of pre-market bullishness in Twitter messages on a given day also predicts lower volatility and negative returns for volatility futures contracts on that day.

The relation between disagreement among investors and return volatility is the focus of the second question. Several models of dynamic equilibrium asset pricing, such as Gallmeyer and Hollifield (2008) and Buraschi and Jiltsov (2006), show that under heterogeneous beliefs, disagreement between agents is positively correlated with return volatility. Estimating the parameters of a GARCH(1,1) model with additional control variables shows that an increase in the pre-market disagreement about the short-term returns predicts more volatility during the trading day. In addition, the daily number of opinionated tweets posted by high follower individuals is used as a proxy for information asymmetry to test an implication of Wang (1993). The paper shows that in a dynamic model with information heterogeneity, information asymmetry can cause more volatility in stock prices. Consistent with Wang's model, there is a positive relation between the volatility and empirical measure of information asymmetry constructed from Twitter data.

The third question is about the relation between the number of opinionated tweets and trading volume. The literature has documented a positive correlation between return volatility and trading volume. The empirical model developed by Andersen (1996) attributes the relation to an underlying process. The assumption is that information arrival governs both volatility and trading volume. An increase in the density of information arrival leads to more trading volume and return volatility. In this study, the daily number of opinionated tweets is used as a proxy for the density of information arrival to study the explanatory power of information arrival for trading volume and return volatility. The results show that the daily number of opinionated tweets is positively correlated with both trading volume and volatility of the S&P500 index.

In recent years, Twitter has become a major source of information and an effective communication tool for traders and investors in the financial markets. Stock market participants use Twitter to share their information with others and receive real-time information about the stock market and individual companies. There are several stories similar to the ones mentioned in Chapter 1 that highlight the role of Twitter for active participants in the financial markets.

The contribution of this paper is threefold. First, this is the first study that systematically aggregates Twitter's stock market related messages to measure investor sentiment. Second, the text analysis method used in this paper allows for capturing the opinion of stock market traders who express their opinion using certain phrases or by announcing their recent trades. Third, the proposed measure of investor sentiment can be constructed at a daily frequency and be used to study the role of sentiment in the daily returns of the stock market index. Moreover, the daily measure of disagreement among investors and the daily number of opinionated tweets are valuable pieces of information to test the predictions of theory about return volatility and trading volume.

A number of papers examined the ability of internet messages to predict stock returns. Using a large number of messages posted during 2000 on Yahoo Finance and Raging Bull about 45 companies in the Dow Jones Industrial Average and the Dow Jones Internet Index, Antweiler and Frank (2004) report three major findings. First, an increase in the number of messages predicts negative return on the next day. Second, disagreement is associated with more trades. Third, the number of messages is positively correlated with contemporaneous volatility and predicts volatility on the next day. Dewally (2000) studies the effect of two newsgroups and finds no predictive ability in the forecasts posted on the newsgroups. Das and Chen (2001) examine the impact of stock messages on nine firms and find no evidence that messages forecast stock returns.

Text analysis is an important element of extracting useful information from internet message boards and social media networks. Antweiler and Frank (2004) use Naive Bayes as the main algorithm to classify messages. Tetlock (2007) and Da, Engelberg, and Gao (2015) rely on the Harvard IV-4 Dictionary to count the number of words in different categories. Both methods are shown to be effective in practice given the data sources used in the studies. In this paper, a different approach is used to classify the tweets. A number of word combinations are defined as indicators for bullish, bearish and neutral tweets. If a tweet contains one of the bullish word combinations and does not mention any of the bearish or neutral word combinations, it is placed in the most bullish category and is associated with the numeric score of +1. The details of message classification is provided in subsection 2.2. Given that many traders express their opinion by announcing their current option or ETF positions and the Harvard dictionary is not structured for the vocabulary of traders, using word combinations is more effective than alternative methods in measuring the short-term investor sentiment.

The rest of the paper is organized as follows. Section 2 describes the details of Twitter data and the methodology of constructing the Twitter sentiment (TSI) and Dispersion of sentiment (DS) indices. Section 3 investigates the relation between sentiment and short-term returns of the U.S. stock market. The relation between sentiment and return volatility is examined in section 4. The focus of section 5 is on the link between daily trading volume and information arrival rate measured by the number of opinionated tweets. Section 6 concludes the paper.

2.2 Data and Methodology

Public tweets are used as the main data source for the empirical tests. This section discusses the details of data and the methodology of constructing the Twitter sentiment (TSI) and Dispersion of sentiment (DS) indices.

2.2.1 Data

The dataset includes all the tweets that mention a major U.S. equity index or an Exchange Traded Fund (ETF) that tracks the daily return of the index from September 2013 to August 2015. For instance, tweets that mention the S&P500 index or ticker symbols "SPY", "SSO", "SDS", "UPRO", or "SPXU" are included in the dataset. These tickers represent the ETFs that correspond to unleveraged and leveraged long and short positions on the S&P500 index. The full list of search words and ticker symbols is provided in the Appendix B. Tweets that provide a positive, neutral or negative outlook for the future return of the equity indices are included in the construction of sentiment and dispersion of opinion index. Marketing tweets or those with no obvious opinion are excluded because they do not provide any useful information about sentiment. More details about the identification of opinionated tweets is provided in subsection 2.2.

The dataset contains a daily average of 347 opinionated tweets about the index of U.S. large cap stocks (S&P500), technology stocks (NASDAQ), Dow Jones industrial average, small cap stocks (Russell 2000), and CBOE volatility index (VIX). Figure 7 shows the distribution of opinionated tweets during the hours of a trading day. It is no surprise that a large portion of opinionated tweets are posted during the U.S. trading hours. Figure 8 shows that there is little variation between the average number of tweets posted on the weekdays and the number drops significantly over the weekends.

2.2.2 Twitter sentiment (TSI) index

The Twitter sentiment index measures the sentiment of traders and investors by aggregating the opinion of all individuals that express their forward looking outlook of the stock market in a tweet. The idea is to (1) collect all the relevant tweets posted in a time period, (2) examine each one for a positive, neutral or negative view on the stock market, (3) assign a weight to each tweet based on the number of followers of the individual that posted the tweet and then (4) aggregate them to generate the Twitter sentiment index.

Step 1 involves collecting the tweets that mention a major U.S. equity index or an ETF that tracks the daily return of the index. Twitter's publicly available search tools are used to download the tweets that mention any of the search words listed in Appendix B.

In step 2, positive, negative, and neutral tweets are identified and tweets with no opinion are filtered out. Tweets that express a bullish outlook are put in the positive category. Neutral tweets indicate that a trader is indecisive or waiting for more information before taking a position in the stock market. Finally, bearish comments in a tweet put it in the negative category. Since the objective is to construct a measure for sentiment of traders and investors, the methodology adopted in this paper captures the type of words that traders and investors use to describe their current position or outlook for the markets. Opinionated tweets are identified using 1148 word combinations. For instance, an optimistic trader could buy call option on the S&P index ETF (SPY) so a tweet that contains the words "bought, spy, call" in this order is placed in positive category. A tweet that contains the words "increase, spy, short" in this order is identified as negative because it indicates that a trader expects further drop in equity prices and is willing to increase the size of an existing short position. The words "will, buy, if" put a tweet in neutral category because they indicate a decision to buy the equity index conditional on some event. Table **??** provides examples of actual positive, neutral, and negative tweets. The placement of the tweets in the categories is independent of daily returns in the stock market and is determined only by the content of tweets. Several recent studies, such as Tetlock (2007) and Da, Engelberg, and Gao (2015), use the Harvard IV-4 Dictionary to carry out text analysis. The dictionary provides tools to place words in categories such as "Positive", "Strong", and "Weak". In this study, the Harvard dictionary is not used because certain words that traders use to express their opinion convey a different message according to the Harvard dictionary. For example, traders often use the word "Bullish" to express a positive outlook. The closest word to bullish in the Harvard dictionary is "Bull" which is placed in "Male" category.

In step 3, a numerical value is assigned to each category. All positive tweets are coded as +1, neutral tweets as 0 and negative tweets as −1. Tweets that mention both positive and neutral word combinations are coded as +0.5 and those that mention negative and neutral combinations are coded as −0.5. The numerical score of a tweet is then multiplied by the number of followers of the account that posted the tweet. If we assume that the number of followers of each account is the equilibrium value given by the information marketplace to the opinions posted by that account, weighting individual opinions based on their followers takes into account the differences between the quality of private signals. Although the number of followers might not always reflect the quality of information shared by an individual, it is a useful observable variable for extracting information from Twitter messages.

Finally step 4 involves aggregating the views expressed in the tweets. The daily Twitter sentiment index (denoted by TSI_t), is the weighted average of the score of all the opinionated tweets posted during a day:

$$TSI_t = \frac{\sum_{i=1}^{n_t} w_i s_i}{\sum_{i=1}^{n_t} w_i},$$

where s_i is the numeric score associated with the tweet i and $s_i \in \{-1, -0.5, 0, +0.5, +1\}$, w_i is the number of the followers of the account that posted tweet i and n_t is the number of tweets with positive, negative or neutral view posted during the time period t.

By definition, TSI index can take any value between −1 and +1. If all the tweets in a given day express a negative outlook for the stock market, TSI index would be −1, which is the most bearish value for the sentiment. Conversely, TSI index takes the most bullish value of +1 if every tweet contains a positive outlook. The TSI index by construction gives higher weight to the opinion of individuals with more followers. Accounts that post tweets about stock market and have a large number of followers are typically controlled by professional traders or money managers. As the result, the TSI index is a measure of optimism among stock market traders and investors. However, The TSI index is not dominated by few individuals as evidenced by the distribution of weight of tweets over the number of followers of the accounts. Figure 9 shows that the opinion of accounts with more than 500,000 followers has an average weight of 0.13 in the daily TSI index. Table **??** shows the summary statistics of the daily TSI index over 476 trading days from September 18th 2013 to August 8th 2015.

2.2.3 Dispersion of Sentiment (DS) index

The daily dispersion of opinion is measured by aggregating the deviation of the numeric score of each tweet from the daily sentiment index. Dispersion of sentiment index (denoted by DS_t) is the weighted standard deviation of the numeric value associated with opinionated tweets posted during a day:

$$DS_t = \sqrt{\frac{\sum_{i=1}^{n_t} w_i (s_i - TSI_t)^2}{\sum_{i=1}^{n_t} w_i}}.$$

The measure of dispersion of opinion could take any value between 0 and 1. Dispersion of opinion will be zero if all the tweets posted on a trading day express the same view about the direction of change in the price of U.S. equity indices. The other extreme case is when the view expressed by tweets weighted by their followers is equally split between positive and negative categories. In this case, dispersion of opinion will be 1. Table **??** provides the summary statistics of the daily DS index.

2.3 Other Measures of Sentiment

There are several individuals and organizations that send a questionnaire to a group of investors or households and report the summary of their findings as a measure of sentiment or confidence. In order to understand the benefits of using the TSI index, this section provides a brief summary of existing surveys including their methodology, frequency of data collection, and their target participants.

Robert Shiller has been conducting an investor attitude survey every month since 1989. The survey questions a sample of U.S. wealthy individual investors and institutional investors about their confidence that the stock market will go up in the succeeding year. The participants are also questioned about their confidence on some aspects of the stock market and the results are published as stock market confidence indices separately for individual and institutional investors[1]. Another example is the UBS/Gallup survey that has been carried out every month since 1996. UBS index of investor optimism is usually published on the fourth Monday of each month. The Gallup organization randomly selects 800 private investors across the U.S. who have more than 10000 USD in investable assets and conducts phone interviews during the first two weeks of every month. The survey questions cover the personal financial aspect and the macroeconomic dimension of investments in the United States[2]. Investors Intelligence is another entity that measures investor sentiment by aggregating the opinion of market newsletter writers and is released every week on Wednesdays. They study over a hundred independent weekly and monthly market newsletters and assess each author's current stance on the market. Advisors' Sentiment Report provides a summary of the assessments and was originated in 1963[3]. Finally, the AAII sentiment survey conducted by American Association of Individual Investors (AAII) measures the percentage of individual investors who are bullish, bearish, and neutral on the stock market for the next six months. Every individual who subscribes to the AAII services can submit a vote in the weekly surveys. The proportion of investors in each group is measured based on the data received each week by Wednesday and the results are reported on Thursdays. The survey was first conducted in July 1987[4].

Another approach to measure investor sentiment is to combine multiple imperfect proxies. Baker and Wurgler (2007) show that their six sentiment proxies have a common sentiment component. The proxies are trading volume measured by NYSE turnover, the dividend premium, the closed-end fund discount, the number of IPOs, first day return on IPOs and equity share in new issues. They remove the major macroeconomic influences from the proxies and construct an investor

[1]Source:http://som.yale.edu/faculty-research/our-centers-initiatives/international-center-finance/data/stock-market-confidence

[2]Source:http://extranet.datastream.com/ContentUpdate/detail.asp?MainID=2321

[3]Source:www.investorsintelligence.com

[4]Source:www.aaii.com/sentimentsurvey

sentiment index using the principal component of the proxies. They show that their index captures the major anecdotal accounts of bubbles and crashes between 1966 and 2000.

Monthly and weekly surveys are valuable sources of information and capture the long term trends of sentiment but are unable to show the short term changes due to their low frequency. By contrast, the TSI index provides a high frequency measure of sentiment that captures daily changes in investor sentiment. Moreover, individuals have the incentive to express their true opinion about the stock market in social media networks because they care about their reputation and would like to attract more followers by posting their best analysis about asset prices. However, there is no such incentive for those who participate in a survey because the opinion of each survey participant is equally weighted in the results and is not individually disclosed.

2.4 TSI and Stock Market Returns

This section examines the relationship between the daily TSI index and daily return of the stock market index for the same day and future days.

2.4.1 Return Predictability

In this subsection, the TSI index is used to predict the same day return of the S&P500 index. In order to avoid the feedback problem caused by the influence of price changes on the sentiment of traders, a daily pre-market TSI index is constructed by using only the tweets posted between midnight and 9:30am EST on a trading day. The pre-market TSI index on a given day is created using the exact same method described in section 2.2 but its data universe is limited to tweets posted before the market open on that day. Daily open-to-close price change of the S&P500 index is regressed on the pre-market TSI index and some control variables. The regression results are reported in Table **??**. In order to see the effect of weighting individual opinions on the predictive power of the sentiment index, follower-weighted sentiment index is used as the independent variable in the first two regressions and the last two regressions of Table **??** use equal-weighted sentiment index.

The daily open-to-close return of the S&P500 is the dependent variable in the regressions of columns (1) and (3). Since the open price of the index does not reflect the price changes during the U.S. market close time overnight, the daily open-to-close return of a liquid ETF, SPDR S&P500 NYSEARCA:SPY, is used as the dependent variable in the regressions of columns (2) and (4). Investors can take a long or short position on the index at the market open through SPY.

In order to control for the effect of changes in macroeconomic activities on the stock market, a daily measure of macroeconomic conditions is included in the regressions. Aruoba, Diebold, and Scoitti (2009) use a number of macroeconomic variables of different frequencies and construct the ADS index as a daily measure of business conditions. The ADS index is obtained from the website of the Federal Reserve Bank of Philadelphia and includes seasonally adjusted value of quarterly real GDP, industrial production, manufacturing and trade sales, personal income minus transfer payments, monthly payroll employment, and weekly jobless claims.

Uncertainty about economic policy could influence investor sentiment and asset prices. Baker, Bloom, and Davis (2013) define three categories of words associated with economic policy uncertainty (EPU) and count the number of U.S. newspaper articles that mention any of the words in these categories. Baker, Bloom, and Davis (2013) show that their measure of EPU offers a good proxy for uncertainty related to economic policy over time. The EPU index is the control variable that accounts for economic policy uncertainty in the regressions.

One day lagged volatility index (VIX) and lagged return of the corresponding stock market index are included as additional control variables in all the regressions. According to the daily data over 476 trading days, one standard deviation increase in the follower-weighted pre-market TSI index predicts a return of 12.8 basis points for the S&P500 index[5]. As indicated in the second column of Table **??**, the return drops to 6.8 basis points if we control for the practical issues of trading the S&P500 index. The last two columns of Table **??** indicate that the predictive power of equal-weighted sentiment index disappears once the overnight price changes are excluded from the daily return of the stock market index.

The regression results highlight the role of assigned weights to individual opinions when aggregating the expectation of investors. The equal-weighted sentiment index does not capture the differences between the quality of private signals given by the individuals and consequently shows no predictive power for the daily return

[5]one standard deviation change in the pre-market TSI index corresponds to 0.3937.

of the stock market index. The predictive power of follower-weighted sentiment index for the daily open-to-close return of the stock market index indicates that the information marketplace is effective in identifying the individuals with higher quality private signals.

2.4.2 Robustness Checks

This section provides a number of robustness tests for return predictability of the follower-weighted pre-market TSI index. Table **??** reports the result of regressions that include more controls and use an alternative method to construct the sentiment index.

There are a number of market-based measures that are often used as a gauge for investor sentiment. Brown and Cliff (2004) show that investor sentiment measured by surveys is related to the market measures such as the ratio of advancing to declining stocks and the change in short interest. One might argue that the TSI index is influenced by these market indicators and contains no additional information. To ensure that the index is not a reflection of lagged stock market indicators, some of these indicators are included as additional control variables in the first column of Table **??**. More specifically, one day lagged the number of advancing stocks divided by the number of declining stocks in NYSE, ARMS index[6], the number of stocks at their 52 weeks high price divided by the number of stocks at their 52 weeks low price, the number of put options on the S&P500 index divided by the number of call options, and the difference between the interest rate of 10 years and 3 months U.S. Treasury bonds are included in the regression. According to the results, including the market indicators does little to change the effect of pre-market TSI on the daily returns.

Since the information flow is continuous and news could influence the prices through futures market during the U.S. market close time, there is a concern that the TSI index reflects the price changes of the futures market. In order to control for the effect of price changes in the futures market, the difference between the open price of SPY and its close price on the prior day is included as a control variable in the regression. The second column of Table **??** shows that including the overnight price changes does not alter the main result.

[6]ARMS index is the ratio of the number of advancing to declining stocks divided by their trading volume.

Several empirical studies, such as Keim (1983) and Gultekin and Gultekin (1983), provide evidence for seasonality in the stock market returns. To ensure that the results are not driven by known seasonal effects, month of the year dummies are included in the regression and the results are reported in the third column of Table **??**. Including the month dummies increases the effect of the pre-market TSI index.

The news-based EPU index is included in the regressions to control for the effect of economic policy uncertainty that could influence the sentiment and the stock market returns. Since the twitter data used in this study is related to the time period that Federal Reserve was communicating its exit strategy from the unconventional monetary policy, one might still be concerned that the results are driven by a few sizable returns following unexpected announcements in the FOMC meetings. Column (4) of Table **??** shows that including a dummy variable for the FOMC rate decision days does not change the sign and significance of the pre-market TSI index.

Given that major economic data such as GDP and unemployment are reported at 8:30am EST, the pre-market TSI index and the stock market index could reflect the surprises in the economic data on the data release dates. To ensure that the predictability of the TSI index is not driven by a few data announcement days, an economic news day dummy is included in the regression. The dummy variable is one on the days that GDP, unemployment or inflation data is released and zero otherwise. The regression result is reported in column (5) of Table **??** and shows that the effect of pre-market TSI index is not driven by economic news announcement days.

The pre-market TSI index is constructed by using the opinionated tweets posted between midnight and 9:30am. A potential concern is that some public information released after the market close on the prior day could be missed as the result of this choice for the time interval of tweets included in the construction of the index. The last column of Table **??** shows the results of a regression similar to the second column of Table **??** except that the daily pre-market TSI index uses the opinionated tweets posted from market close (4:00pm EST) on the prior day up to market open (9:30am EST). The coefficient of the pre-market TSI index remains positive and statistically significant in the regression following this alteration in the construction of the index.

2.5 Volatility

In this section, the relationship between the sentiment and volatility of stock market returns is investigated. We start with testing the link between the disagreement among investors and the level of return volatility. Then, the prediction of Wang (1993) about the positive relation between information asymmetry and volatility is tested.

2.5.1 Volatility and disagreement among investors

The literature offers many ways to model volatility. Hansen and Lunde (2005) compare 330 ARCH models and find no evidence that sophisticated models can predict the conditional variance better than a simple GARCH(1,1) model. In this study, a GARCH(1,1) with additional independent variables is used to model the stock market volatility. More specifically, the volatility is given by:

$$r_t = \sqrt{h_t}\epsilon_t$$
$$h_t = \alpha_0 + \alpha_1 r_{t-1}^2 + \alpha_2 h_{t-1} + \big(\beta_1 TSI_t + \beta_2 DS_t + \beta_3 logDOT_t\big) + \gamma Z_t,$$

where r_t is the daily return of the S&P500 index, and h_t represents the return volatility, TSI_t is the Twitter sentiment, DS_t is the Dispersion of sentiment, and $logDot_t$ is the daily log number of opinionated tweets. Z_t is the vector of control variable that includes the ADS business conditions index and the Economic Policy Uncertainty (EPU) index.

In the first test, opinionated tweets posted before the market open on each trading day are used to predict the return volatility of that day. The left panel of Table **??** shows that an increase in the disagreement about near-term returns, measured by the DS indx, predicts higher daily volatility. This observation is consistent with dynamic models of heterogeneous beliefs, such as Gallmeyer and Hollifield (2008) and Buraschi and Jiltsov (2006), that predict a positive correlation between dispersion of opinions and return volatility. The table also shows that an increase in the pre-market sentiment predicts lower volatility. The right panel of table **??** shows the parameter estimates of the same model except that the Twitter indicators are constructed using the tweets posted on a trading day. Disagreement remains positive and statistically significant.

The daily log number of opinionated tweets is used as a proxy for density of information arrival and is included to control for the effect of information arrival rate on the volatility[7]. As mentioned earlier in section 2.1, the dataset includes the tweets that provide a forward looking opinion about the direction of stock market. Individuals usually express their opinion in a tweet by directly pointing out their outlook or announcing their current position or recent trades. Opinionated tweets are typically posted when an individual acquires new information. For instance, a trader sends a bullish tweet following the public announcement of some economic data or based on her conclusion after analyzing technical indicators. Therefore, the number of opinionated tweets is considered a proxy for density of public and private information. Positive and statistically significant coefficient of $logDOT_t$ is consistent with the model of Andersen (1996) that assumes a positive relation between the return volatility and information flow rate.

To examine the relation between the sentiment and changes in volatility, we consider the daily returns of a tradable security based on volatility, iPath S&P 500 VIX Short Term Futures ETN (NYSEARCA:VXX). VXX provides investors with exposure to a daily rolling long position in 30 days VIX futures contracts and the daily returns of VXX tracks the changes in the price of first and second month VIX futures contracts. The daily return of VXX is regressed on the daily sentiment, dispersion of opinions, log number of opinionated tweets, and control variables. The regression results are reported in Table **??**. The daily TSI index is negatively related to contemporaneous return of VXX. More specifically, one standard deviation increase in the TSI index is associated with 50 basis points decrease in return of VIX futures contracts. The regression also highlight the contemporaneous positive correlation between the daily number of opinionated tweets and return of volatility futures contracts.

2.5.2 Volatility and information asymmetry

Wang (1993) develops a model of intertemporal asset prices in which investors are different in terms of their information about dividend growth. The model implies that the information asymmetry can cause price to be more volatile. In this section, the daily number of opinionated tweets posted by individuals with a large number of followers is used as a measure of information asymmetry to test the implication of the model. Wang (1993) assumes that informed investors have private information about stock's dividend growth and uninformed investors extract information from

[7] See Andersen (2001)

public signals, prices and realized dividends. If we define informed investors as those who have more than 500 followers and uninformed investors as those with less than 500 followers, we can follow Wang (1993) and assume that uninformed investors tweet the information that can be extracted from other resources such as past prices or realized earning data. The tweets posted by informed investors, however, contain their private signals. We use the daily number of tweets posted by informed traders as the empirical measure of information asymmetry and examine the link between information asymmetry among investors and return volatility. More specifically, we are interested in the coefficient of information asymmetry (denoted by β_3) in the GARCH model of volatility given by:

$$
\begin{aligned}
& r_t = \sqrt{h_t}\epsilon_t \\
& h_t = \alpha_0 + \alpha_1 r_{t-1}^2 + \alpha_2 h_{t-1} + \left(\beta_1 TSI_t + \beta_2 DS_t + \beta_3 DOT_t^I + \beta_4 DOT_t^U\right) + \gamma Z_t,
\end{aligned}
$$

where DOT_t^I is the daily number of opinionated tweets posted by informed investors and DOT_t^U is the same measure for uninformed investors. DOT_t^U is included to control for the effect of public signals on volatility. The first column of Table **??** shows the coefficient estimates of the model. The next two columns of the table report the regression results using different thresholds for the number of followers that separates informed and uninformed agents. Consistent with Wang's model, information asymmetry is positively related to volatility as evidenced by the positive and statistically significant coefficient of the daily number of tweets posted by informed agents.

2.6 Trading Volume and Information Flow

The literature has documented a positive correlation between daily trading volume and return volatility. Several studies, such as Epps and Epps (1976) and Tauchen and Pitts (1983), attribute the joint dependence of volume and volatility to information flow. The intuition is that an increase in the intensity of information arrival leads to more transactions by informed traders and results an increase in the trading volume and return volatility in the stock market. Andersen (1996) presents a model based on the assumption that a random information arrival process drives both return volatility and trading volume. The paper estimates the parameters of the model using the moments of volatility and volume. In this section, the daily number of opinionated tweets is used as an empirical measure of information arrival rate and the positive correlation between information arrival rate and trading volume is investigated. Table **??** provides the summary statistics of the daily number of opinionated tweets in the dataset. Figure 10 shows the daily log trading volume of the S&P500 index ETF (NYSEARCA:SPY) and number of opinionated tweets from September 2013 to August 2015. The direct relation between the two time series is clearly evident in Figure 10 . In order to investigate the relationship more rigorously, log trading volume of SPY is regressed on the log daily number of opinionated tweets and control variables. There are many empirical papers in the literature that highlight the link between trading volume and other market measures. For instance, Statman, Thorley, and Vorkink (2006) document the positive correlation between trading volume and past returns. Chen, Firsth, and Rui (2001) points out the direct relation between trading volume and absolute value of contemporaneous return. Epps (1975) and several other studies show that trading volume and return volatility are positively related. In order to control for the effect of these variables on the trading volume, the absolute price changes of the stock market index and its past returns as well as the volatility index are included in the regression.

In summary, the coefficient of log Daily Opinionated Tweets (denoted by $logDOT_t$) in the following model is estimated using an OLS regression with Newey-West variance. Since the daily trading volume is an autocorrelated variable, Newey-West variance estimator produces consistent estimation of standard errors for coefficient estimates[8]. In the regression, it is assumed that the errors are correlated for up to 20 days.

$$logVol_t = \beta_0 + \beta_1 DOT_t + \gamma Z_t,$$

$logVol_t$ is the daily log trading volume of the S&P500 index ETF (SPY) and Z_t is the vector of control variables that includes the absolute price changes of the S&P500 index, five days lagged returns of the index, logVIX and its five days lags, , five days lagged number of opinionated tweets, the business conditions index (ADS), economic policy uncertainty index (EPU), and the weekday and month dummies. Table **??** reports the regression results.

The first column of Table **??** shows that the coefficient of information flow rate measured by the number of opinionated tweets is positive and statistically

[8]See Newey and West (1987).

significant after controlling for other known factors that influence trading volume. The result is consistent with the assumption of Andersen (1996) that the density of information arrival is directly related to trading volume in the stock market. Week of the day and month of the year dummies are included to capture seasonal effects on trading volume but one might argue that the regression result reflect simultaneous year over year increase in trading volume and number of tweets. To address this concern, the daily number of tweets and log trading volume are tested for unit root using Dickey-Fuller test and MacKinnon approximate p-values reject the null hypothesis at the 0.01 level.

Chae (2005) documents unusual patterns in the trading volume of stocks before scheduled announcements due to changes in the behavior of informed and liquidity traders. Extending the logic to the broad market, there might be a concern that the result is influenced by the unusual patterns of trading volume on the release days of scheduled economic news. The second column of Table **??** reports the result of the regression that includes dummy variables for FOMC policy announcement and monthly payroll data release days. The number of opinionated tweets remains positive and statistically significant after including the dummies.

Option expiry days are usually accompanied with large trading volume in the stock market. To ensure that the result is not driven by large trading volume on a few option expiry days, a dummy variable that is one on monthly option expiry days and zero otherwise is included in the regression. The last column of Table **??** shows that controlling for the effect of larger than usual trading volume on option expiry days does not change the positive relation between the trading volume and the number of opinionated tweets.

2.7 Conclusion

The role of information heterogeneity in the asset prices is highlighted in a number of papers, such as Bacchetta and van Wincoop (2006) and Evans and Lyons (2002). In this paper, the expectations about stock market returns expressed through Twitter messages are aggregated to construct a measure of sentiment. In order to account for the heterogeneity of quality of private signals, opinionated messages are weighted by the number of followers of the individuals. The predictive power of the follower-weighted sentiment index for the daily return of the stock market index suggests that the marketplace for information is efficient in giving more weight to individuals with higher quality private signals.

The opinionated tweets are used to construct a daily measure of dispersion of sentiment. Consistent with the theory, the volatility of daily returns is positively correlated with disagreement among investors about future returns. It is also shown that the empirical measure of information asymmetry, constructed using the Twitter data, is positively related to return volatility.

Given that individuals often use the Twitter platform to share their information with others, the number of opinionated tweets is used as a proxy for density of information arrival to test the relation between information arrival rate and trading volume. Positive correlation between the information arrival rate and trading volume is consistent with the theory that assumes the information flow as a common factor affecting both trading volume and volatility.

Bibliography

[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] Andersen, Torben G. (1996), Return volatility and trading volume: An information flow interpretation of stochastic volatility, *The Journal of Finance*, 51(1), 169-204.

[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] Aruoba, S. Bora?an, Francis X. Diebold, and Chiara Scotti (2009), Real-time measurement of business conditions, *Journal of Business & Economic Statistics* 27:417?27.

[5] Bacchetta, Philippe and Eric van Wincoop (2006), "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?," *American Economic Review* 96(3), 522-576.

[6] 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.

[7] Baker, Scott R., Nicholas Bloom, and Steven J. Davis (2013), Measuring Economic Policy Uncertainty, Chicago Booth Research Paper No. 13-02.

[8] Baker, Malcolm, and Jeffrey Wurgler (2007), Investor sentiment in the stock market. *Journal of Economic Perspectives* 21(2): 129-151.

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

[10] Brown, Gregory W., and Michael T. Cliff (2004), Investor sentiment and the near-term stock market, *Journal of Empirical Finance* 11: 1-27.

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

[12] Buraschi, Andrea, and Alexei Jiltsov (2006), Model uncertainty and option markets with heterogeneous beliefs, *The Journal of Finance*, 61(6), 2841-2897.

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

[14] 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.

[15] Chae, Joon (2005), Trading volume, information asymmetry, and timing information, *The Journal of Finance*, 60(1), 413-442.

[16] Chen, Gong?meng, Michael Firth, and Oliver M. Rui (2001), The dynamic relation between stock returns, trading volume, and volatility, *Financial Review*, 36(3), 153-174.

[17] 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.

[18] 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.

[19] 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.

[20] De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann (1990), Noise Trader Risk in Financial Markets, *Journal of Political Economy* 98(4): 703-38.

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

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

[23] Epps, Thomas W. (1975), Security Price Changes and Transaction Volumes: Theory and Evidence, *American Economic Review*, 65 586-597.

[24] Epps, Thomas W., and Mary Lee Epps (1976), The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis, *Econometrica: Journal of the Econometric Society*, 305-321.

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

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

[27] 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.

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

[29] Gallmeyer) Gallmeyer, Michael, and Burton Hollifield (2008), An Examination of Heterogeneous Beliefs with a Short-Sale Constraint in a Dynamic Economy, *Review of Finance* 12(2): 323-364.

[30] Gultekin, Mustafa N., and N. Bulent Gultekin (1983), Stock market seasonality: International evidence, *Journal of Financial Economics* 12(4): 469-481.

[31] Hansen, Peter R., and Asger Lunde (2005), A forecast comparison of volatility models: does anything beat a GARCH (1,1)? *Journal of applied econometrics*, 20(7), 873-889.

[32] Keim, Donald B. (1983), Size-Related Anomalies and Stock Return Seasonality: Further Empirical Evidence, *Journal of Financial Economics* 12,13-32.

[33] 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.

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

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

[36] 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.

[37] 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).

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

[39] Newey, Whitney K., and Kenneth D. West (1987), A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55: 703?708.

[40] Statman, Meir, Steven Thorley, and Keith Vorkink (2006), Investor overconfidence and trading volume, *Review of Financial Studies* 19, 1531?1565.

[41] Tauchen, George E., and Mark Pitts (1983), The price variability-volume relationship on speculative markets, *Econometrica: Journal of the Econometric Society*, 485-505.

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

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

[44] Wang, J. (1993). A model of intertemporal asset prices under asymmetric information. *The Review of Economic Studies*, 60(2), 249-282.

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

[46] 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 *

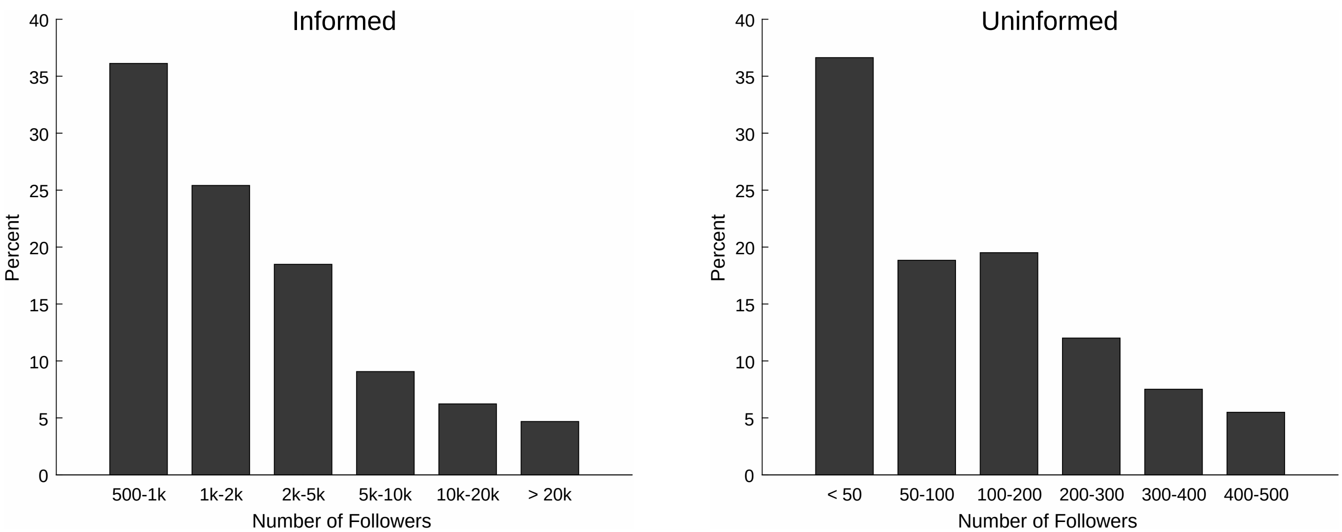

*Informed have more than 500 followers. Uninformed have fewer than 500 followers.

Figure 2: Distribution of the daily number of tweets

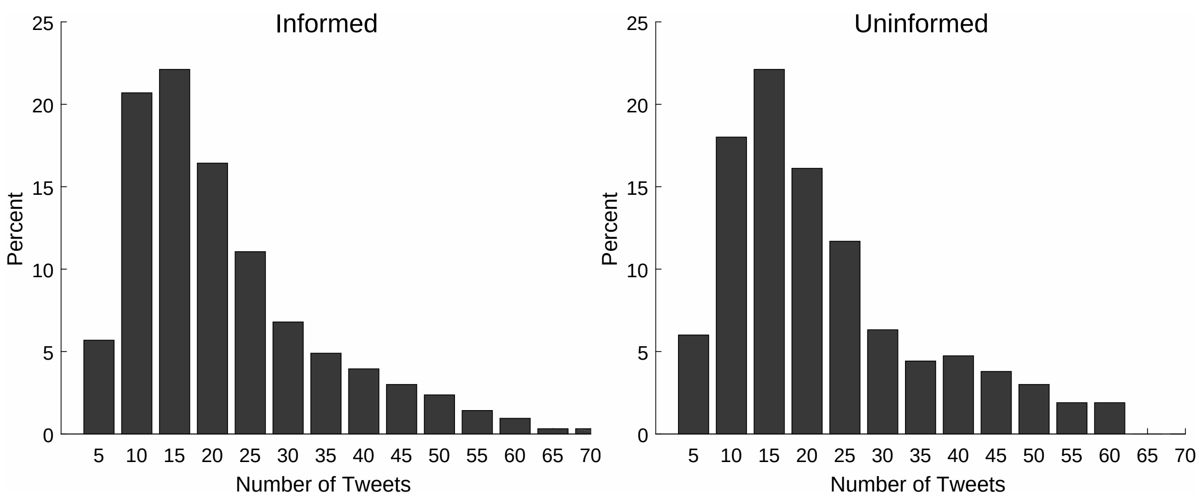

Figure 3: Distribution of Individual TS

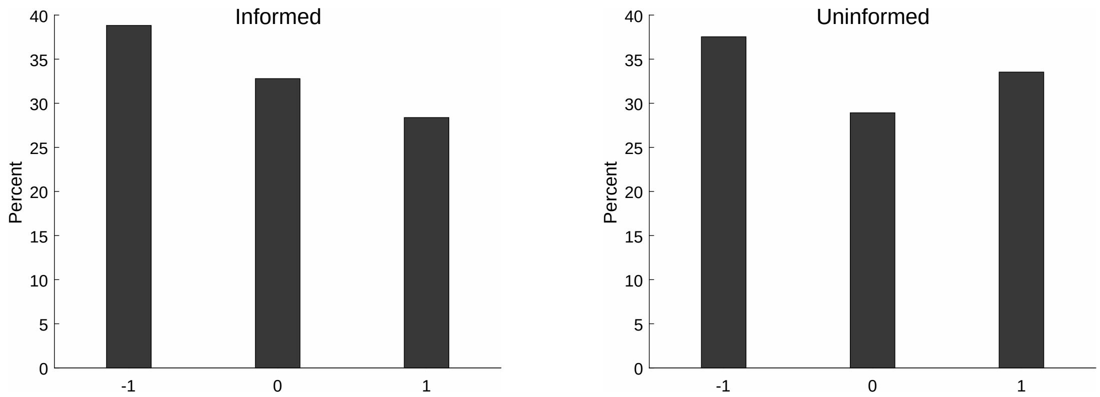

Figure 4: Distribution of daily Twitter Sentiment Index

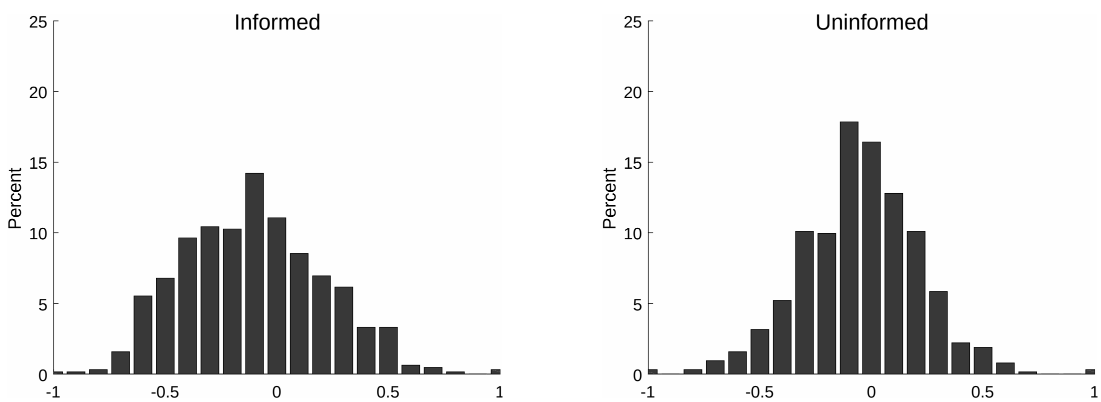

Figure 5: Distribution of daily Disagreement *

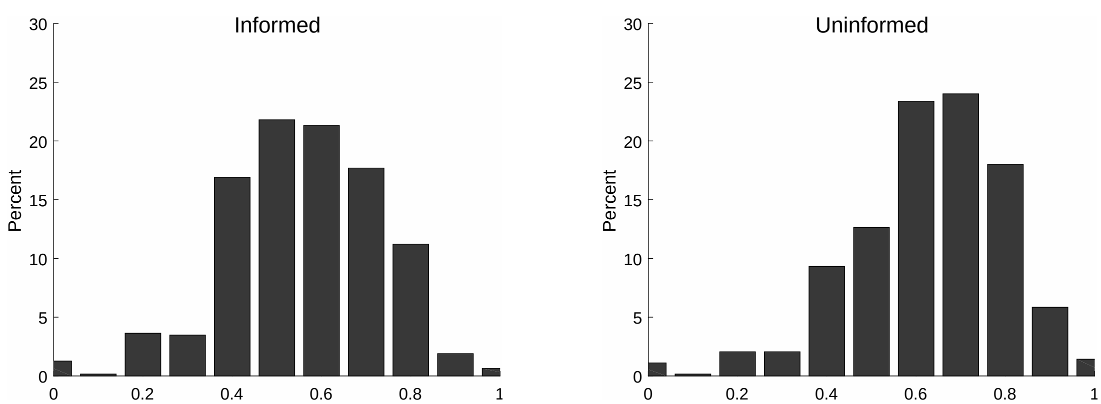

*Disagreement is defined as cross sectional variance of Twitter Sentiment across the individuals.

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

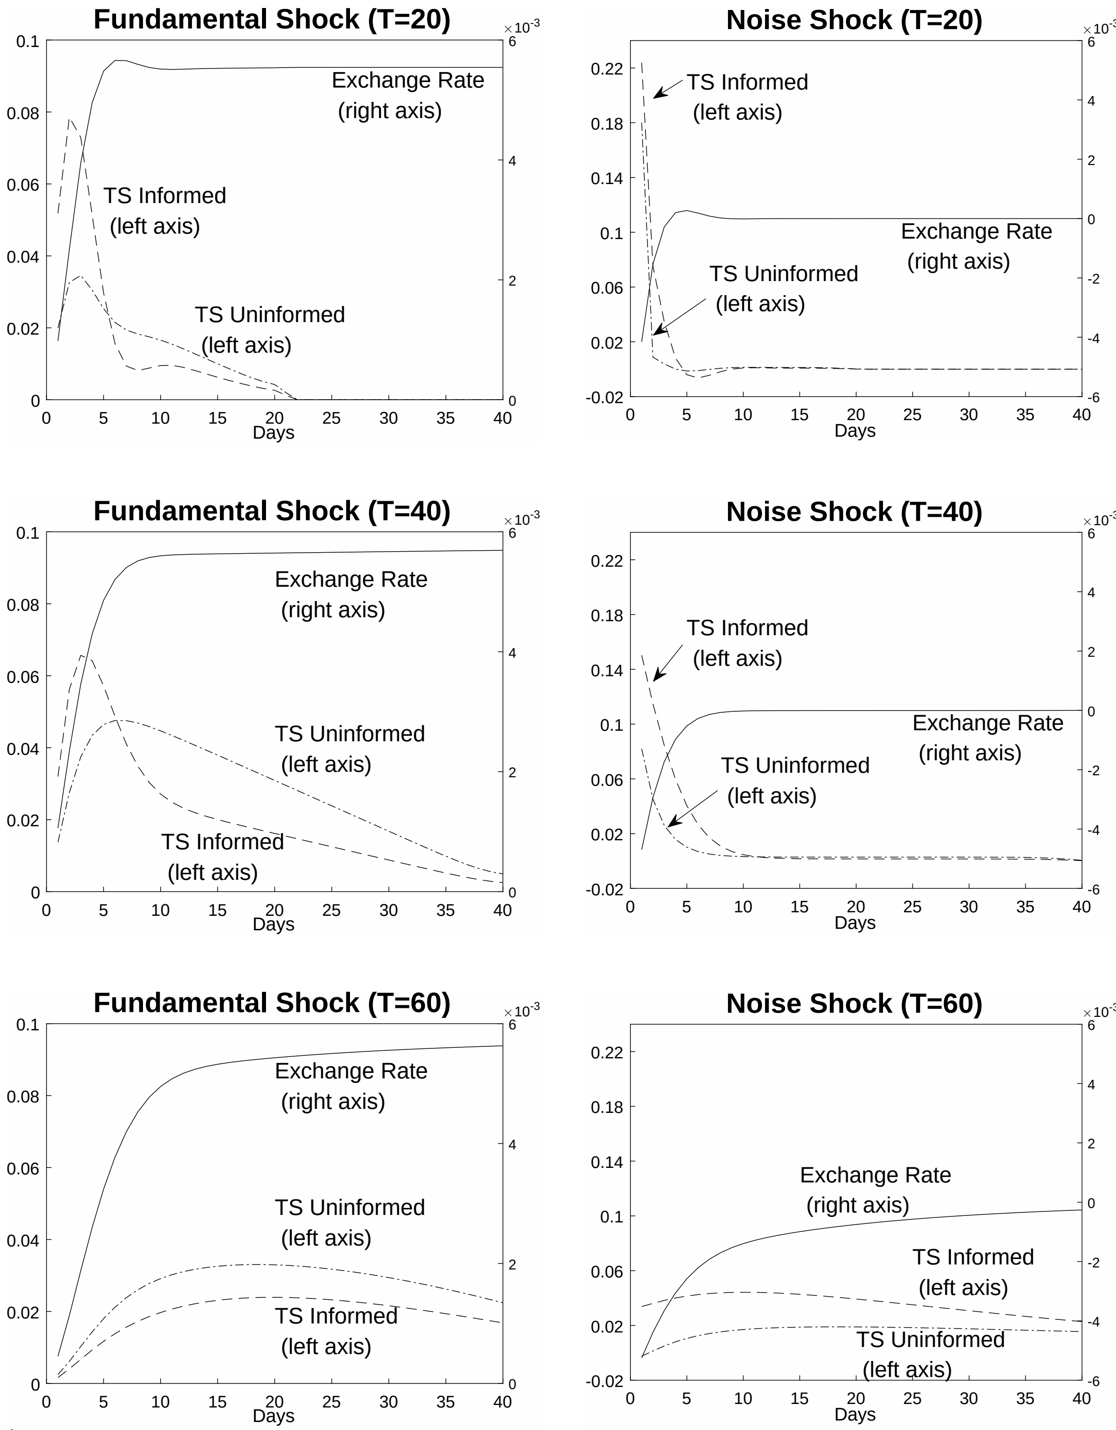

*Average Twitter Sentiment is the average of individual Twitter Sentiment if there are an infinite number of tweets.

Figure 7: Distribution of opinionated tweets about U.S. stock market indices posted during the hours of a day

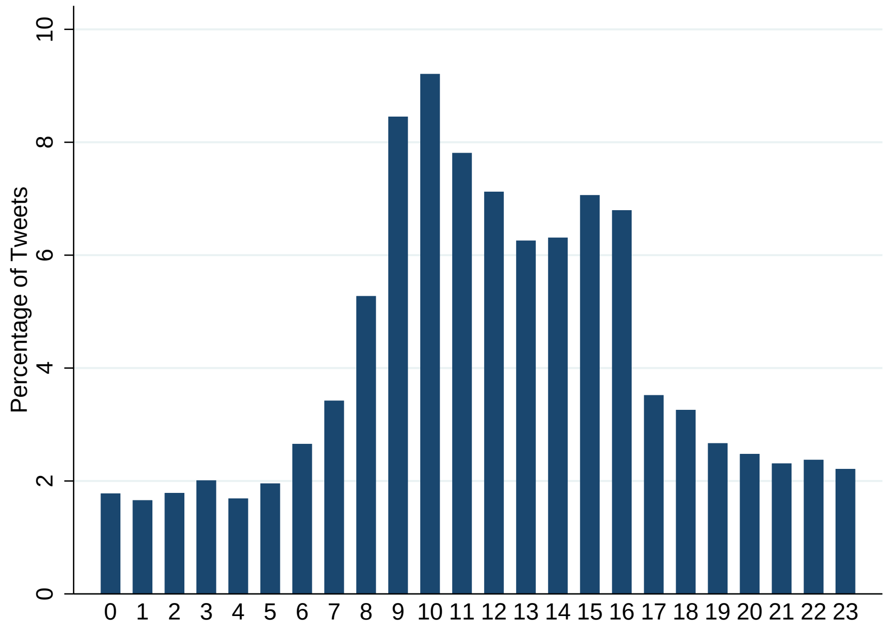

Figure 8: Average number of opinionated tweets about U.S. stock market indices posted during the days of a week

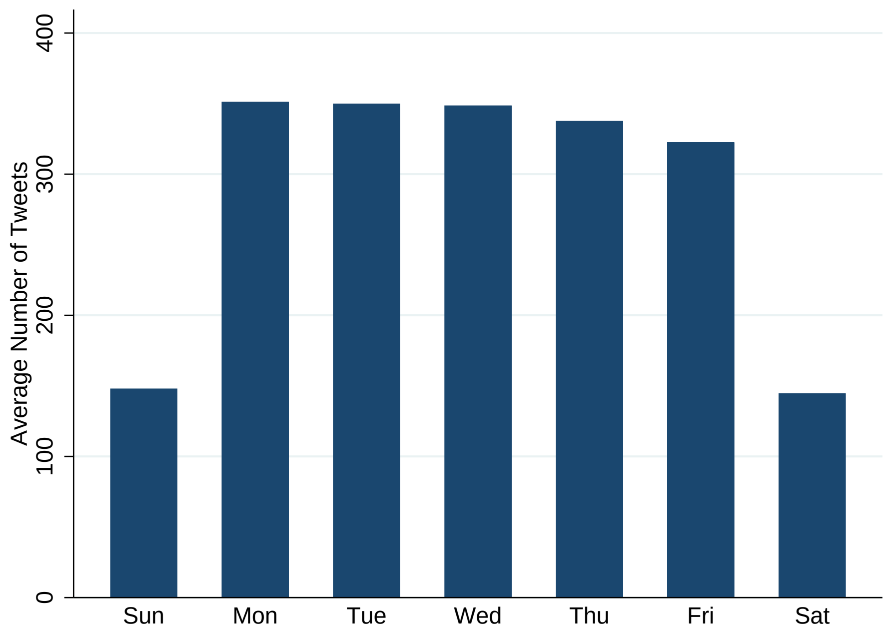

Figure 9: Average daily weight of tweets in the sentiment index over the number of followers of the accounts.

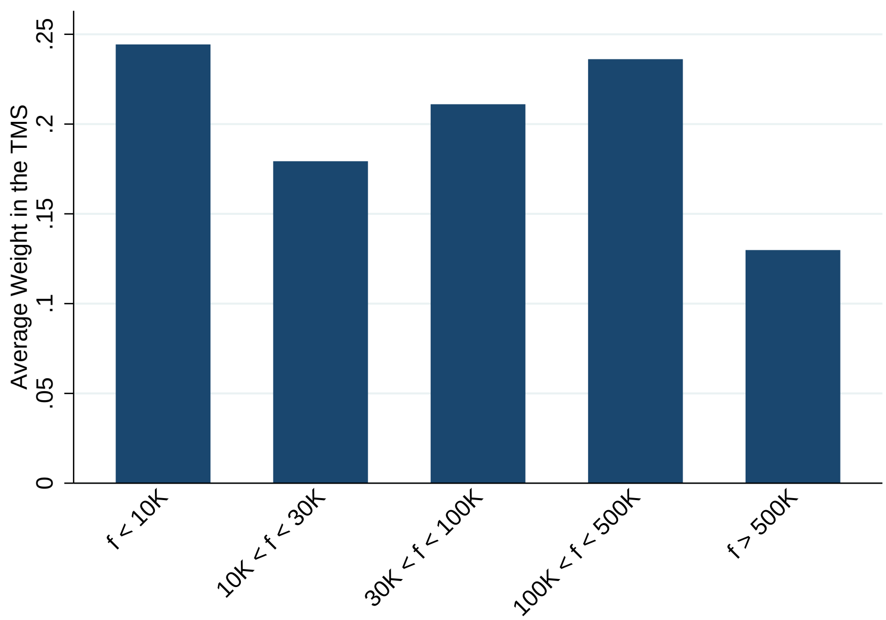

Figure 10: Daily log trading volume of the S&P500 index (SPY) and number of opinionated tweets from September 2013 to August 2015