Exchange Rate Disconnect and Private Information: What Can we Learn from Euro-Dollar Tweets? [1]

Vahid Gholampour
Bucknell University

Eric van Wincoop
University of Virginia
NBER

August 27, 2018

[1]We gratefully acknowledge financial support from the Bankard Fund for Political Economy. We like to thank seminar participants at the University of Virginia, Bucknell University, the Federal Reserve Bank of Dallas, the Federal Reserve Bank of Richmond, the NBER Fall 2016 IFM Meeting, the University of Lausanne and the Hong Kong Institute for Monetary Research for useful comments. We particularly thank Laura Veldkamp for very helpful comments. Zahra Khosravi provided research assistance.

Abstract

We use Twitter opinions about the Euro-Dollar exchange rate to estimate the private information model of Bacchetta and van Wincoop (2006) and investigate the disconnect between the exchange rate and macro fundamentals over both short and long horizons. We simulate the model with the estimated parameters and replicate the methodology of three studies that document the disconnect empirically. The model is consistent with the findings of the empirical literature, though for a different reason over short than long horizons. Over short horizons private information generates a true disconnect between exchange rates and macro fundamentals that accounts for empirical findings. Over long horizons the theory shows that exchange rate changes are mostly driven by observed fundamentals, but empirical limitations in identifying this long-run relationship often lead to an appearance of disconnect in the data.

1 Introduction

The disconnect between exchange rates and observed macro fundamentals has long been a puzzle in open economy macroeconomics. In this paper we will consider the role of dispersed information in explaining the puzzle. Specifically, we will bring a model of dispersed information proposed by Bacchetta and van Wincoop (2006) to the data to see how well it can account for the observed disconnect in the data. A novelty of the paper is to use private information in the form of opinions about the future direction of the Euro-Dollar exchange rate expressed through Twitter from over four years of daily tweets. This is a natural choice as Twitter has become a widely used platform to express opinions and the importance of private information for the determination of exchange rates is well established through the FX microstructure literature.[1]

Meese and Rogoff (1983a,b) first documented the disconnect between exchange rates and fundamentals by showing that standard fundamentals do not outperform the random walk in accounting for exchange rate fluctuations. The literature since then has largely upheld this finding. Specifically, the literature considers the ratio of mean squared errors (or root mean squared errors) generated by a model versus the random walk. The errors represent the changes in the exchange rate that are not accounted for by contemporaneous changes in the fundamentals. This is compared to errors generated by random walk expectations, where any exchange rate change is an error as the expected change of the exchange rate is zero. The findings show that generally the mean squared error is as large or larger for the model as it is for the random walk. This inability of models to outperform the random walk is known as the Meese Rogoff puzzle and is the result of the very limited explanatory power of observed fundamentals (disconnect).[2] The literature has

[1]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).

[2]A substantial literature has also considered vector error correction models where future exchange rate changes are predicted based on the current deviation of the exchange rate from fundamentals. One can also produce mean squared error ratios for these models, but they tell us more about the ability of the current exchange rate and fundamentals to predict future exchange rate changes as opposed to fundamentals accounting for exchange rate fluctuations contemporaneously. It is the lack of such a contemporaneous relationship that reflects the disconnect that is the focus here.

documented the inability of contemporaneous changes in fundamentals to account for exchange rate changes both over short horizons (a month or quarter) and over longer horizons (up to 5 years).

The Bacchetta and van Wincoop (2006) model accounts for the disconnect between exchange rates and observed fundamentals in two ways. First, private information affects expectations of excess returns that determine portfolio positions and thus affect exchange rates. This private information cannot be observed by the econometrician. There is a second effect of private information that is a little more subtle. The model exhibits agent-specific portfolio shifts for reasons that are not observed, such as noise trade, liquidity trade, hedging trade, and so on. In the absence of private information these portfolio shifts generally do not have a big effect on asset prices. They affect asset prices through changes in risk-premia, which are essentially third-order. But Bacchetta and van Wincoop (2006) show that these portfolio shocks are significantly amplified in the presence of private information due to rational confusion about whether observed exchange rate changes are due to unobserved private information about the future state of the economy or unobserved portfolio shocks. They show that the amplification can be very large. However, Bacchetta and van Wincoop (2006) do not confront their model to the data. Their insights are purely theoretical and qualitative. In this paper we will make an attempt to bring data to bear by using opinions expressed through Twitter.[3]

We take several steps to connect verbal tweets about the Euro-Dollar exchange rate to the model of Bacchetta and van Wincoop (2006). We first classify tweets as positive, negative or neutral about the direction of the exchange rate. There is an existing literature related to the stock market, reviewed below, which uses classification methods that are not specific to financial markets and the language used by traders in financial markets. We instead develop a "dictionary" of word combinations based on financial lexicon used by traders in the Euro-Dollar market to automate the interpretation of verbal tweets as positive, negative or neutral.[4] This leads to a

[3]Berger et al. (2016) also confront the Bacchetta and van Wincoop (2006) to the data, using high frequency interdealer order flow data. They show that the model is consistent with the much stronger relationship between exchange rates and order flow in the short run than the long run.

[4]We do not consider other currency pairs as there are far fewer opinionated tweets and many are in different languages. But the overall method described here can certainly be applied to other languages and currency pairs.

measure of Twitter Sentiment, which is a daily measure based on all tweets during a day. To connect this to the model, we compute expectations of exchange rate changes by the agents in the model and use cutoffs for expectations to discretize them into directional beliefs: +1 (expected Euro appreciation), -1 (expected Euro depreciation) or 0 (neutral) for each agent.

We finally estimate the parameters of the model with the Simulated Method of Moments, using various moments involving exchange rates and Twitter Sentiment. These include the daily variance of Twitter sentiment, the disagreement across agents, directional moments that relate Twitter Sentiment to the direction of future exchange rate changes, predictive correlations that relate Twitter Sentiment to actual future exchange rate changes, the contemporaneous relationship between Twitter Sentiment and exchange rate changes as well as basic exchange rate moments (standard deviation and autocorrelation).

We find that the model cannot be rejected by the data and model parameters are estimated with good accuracy. This includes key parameters such as the precision of the private information and the magnitude of the portfolio shocks. It is important to note that there is nothing about the structure of the model that assures the outcome in terms of accounting for exchange rate disconnect. If the private signals in the model are weak (have a sufficiently large signal error variance), we show that the model implies that exchange rate changes are almost exclusively determined by observed fundamentals. The model is then unable to account for the Meese-Rogoff puzzle. We should also emphasize that we do not use mean squared error ratios to estimate the model. In fact, we do not use any data on macro fundamentals to estimate the model. We consider the fundamental in the model to be a summary of macro fundamentals rather than a specific variable such as relative money.

After we estimate the model with Twitter Sentiment and exchange rate data, we simulate it in order to produce mean squared error ratios in a way that is consistent with what has been done in the empirical literature. Specifically, we try to follow as close as possible the methodologies of three papers to produce these mean squared error ratios. The papers are Rossi (2013), Cheung et.al. (2005) and Cerra and Saxena (2010). We choose them as they use respectively monthly, quarterly and annual data. After we estimate the model, we can simulate it over as many years as needed (matching the respective papers). While the model is estimated with daily data, we aggregate to the frequencies that have been considered in the literature. We consider both the short-horizon relationships between exchange rates and fundamentals, which correspond to the frequencies of these papers (monthly, quarterly, annual) and 4 or 5 year long-horizon relationships between changes in the exchange rate and fundamentals. We find that the model does a very good job in accounting for the mean squared error ratios reported in the empirical literature.

This is not the first paper that aims to extract information from social media opinions to learn about asset prices. There is a literature that has used messages from social media and the internet to predict stock prices. This usually involves short data samples of no more than a year. 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.[5] 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.

This literature is purely empirical though and does not attempt to link up to models of dispersed information. In addition, this literature has not employed financial jargon used by traders to classify messages. Most of the literature uses supervised machine-based learning classifiers that are not specific to financial markets at all, such as the Naive Bayes algorithm. We will show that applying such methods to our Twitter data leads to a misleading interpretation of the messages that is of little use. 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.[6] We should also emphasize that in constrast to this literature, we do

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

[6]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.

not use the Twitter data to forecast future exchange rates. The sample is too short to do this. But we will show, even without applying a model, that Twitter Sentiment predicts the direction of exchange rate changes in a way that is statistically significant.

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 exchange rate into positive (+1), negative (-1) and neutral (0) categories. We also show that Twitter Sentiment contains information about the future direction of exchange rate changes and discuss a variety of moments based on Twitter Sentiment that will be used to estimate the moel. In Section 3 we describe the Bachetta and van Wincoop (2006) model of exchange rate determination. Section 4 discusses how to connect the model to the data and presents results from estimating the model with the Simulated Method of Moments. Section 5 applies the estimation results to compute Meese-Rogoff type mean squared error ratios and compares the results to findings from the empirical literature. Section 6 concludes.

2 Data and Methodology

The objective is to translate daily verbal tweets that express opinions about the Euro-Dollar exchange rate into a numerical measure of Twitter Sentiment (TS) that reflects expectations about the future direction of the exchange rate. After a discussion about possible biases related to the reasons for tweeting, we describe how we use a dictionary of financial lexicon to measure Twitter Sentiment. We then provide evidence from directional moments to reject the null hypothesis that the tweets are pure guesses with no relation to the exchange rate. We also discuss a broader set of moments involving Twitter Sentiment and exchange rates that will be confronted with the theory in Section 4.

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. While we acknowledge that it is impossible to be certain about the true motivation of individuals for tweeting their opinion about asset prices, we look for evidence in the data.

As we will describe below, we will identify tweets that provide an opinion about the direction of the Euro-Dollar exchange rate, which we refer to as opinionated tweets. Of these, we restrict ourselves to tweets from accounts that have at least 100 followers. Individuals who have a significant number of followers have less reason to risk their reputation by tweeting an opinion that does not match their beliefs. Most of these tweets come from accounts of individuals that are directly involved in the FX market, which again makes it less likely that they would put their reputation at stake by not revealing their true opinions through their tweets. We know from self-provided descriptions of users that most of the tweets are indeed from individuals directly involved in the FX market. We randomly selected 1000 accounts from accounts with at least 100 followers. Of these, 62 percent clearly identify themselves as connected to the FX market, with 29 percent of accounts from businesses that offer FX related products and 33 percent identifying themselves as "traders" or "experts".[7] The businesses that offer FX related products occasionally tweet their future outlook to showcase their research and gain more subscribers for their business. The "traders" may be motivated to share their views for the same reason that a sports enthusiast enjoys sharing opinions with other sports enthusiasts.

There are two potential ways in which motivations for tweeting can generate biases that 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. There is no way to know for sure whether these biases play a significant role as obviously the users do not reveal their true motivations for tweeting. However, there are reasons to believe that these biases may not be that critical, which we will now discuss.

First consider the bias associated with motivations to tweet an opinion that does not correspond to the individual's true beliefs. We have already discussed that this would be inconsistent with reputational concerns, especially since most of the tweets are from individuals directly involved in the FX market with a lot of followers. But even for FX traders one could think of motivations

[7] Of the remaining 38 percent, 11 percent left their user description blank, while 27 percent provided a description not related to the FX market or provided a description in a foreign language.

that may dominate these reputational concerns. For example, traders who believe that the Euro will appreciate in the medium run, and want to buy Euros, may post negative tweets about the Euro in the hope that others will sell the Euro. They can then buy it at a better price. Apart from reputational concerns, a counter argument to this is that the Euro-Dollar currency pair is one of the most liquid financial markets in the world. Very few individuals would therefore be able to influence the exchange rate through malicious tweets. Another possibility is that a trader is not motivated to steer the exchange rate in a certain direction, but rather aims to attract a response from others by sending out a strongly opinionated tweet with the purpose of learning from the responses by others. We find little evidence of this though. Twitter provides a variable that shows if the tweet is a reply to another Tweet. Only 2.24 percent of opinionated tweets about the Euro-Dollar exchange rate are a response to other tweets. Another way that individuals could react to a tweet is to mark it as favorite. In the data, only 6.85 percent of our opinionated tweets are favored by at least one account. Finally, it may also be possible that some are motivated to simply send random tweets, perhaps just for the sake of participating and saying something or trolling. But we will show in Section 2.5 that the average tweet is not random. It contains information about the future direction of the exchange rate. We will also show that there is a close link between the exchange rate and the average sentiment expressed through tweets. It is therefore hard to argue that most tweets are just random noise.

The second type of bias occurs if people are more motivated to tweet when they have particularly strong beliefs about the direction of the exchange rate. This can bias the average opinion. If this is the case, the percentage of neutral tweets is inversely related to the total number of tweets. To test this, we separate the days into the bottom and top 25 percent in terms of the number of tweets. For days with few tweets (bottom 25 percent), we classify 32.9 percent as neutral tweets. For days with a large number of tweets (top 25 percent), 30.0 percent are classified as neutral tweets. If there is any bias, it is clearly not very large.[8]

[8]We find the same if we separate the tweets into those with more than the average number of tweets (top 50 percent) and those with fewer than the average number of tweets (bottom 50 percent).

2.2 Overall Approach to Computing Twitter Sentiment

It is important to describe in some detail how we translate verbal tweets into a numerical Twitter Sentiment. We have used Twitter's publicly available search tools since October 9, 2013, to download the tweets in real time every half an hour. Every tweet includes the user name, the number of followers of the individual who posted the tweet, as well as the exact time and date that the tweet was posted. We start with all Twitter messages that mention EURUSD in their text and are posted between October 9, 2013 and December 31, 2017. We only consider EURUSD because it is the most important currency pair and there are far fewer tweets about other exchange rates. Moreover, many of the tweets about the Japanese yen are in languages other than English for which we do not have a dictionary. Tweets that come from individuals with less than 100 followers are excluded from the dataset.[9] As already discussed, accounts with a large number of followers have an incentive to preserve their reputation among their followers and are less likely to tweet random noise. There are on average 482 messages per day from distinct Twitter accounts that mention EURUSD with at least 100 followers, for a total of 531,272 tweets.[10] 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

[9] The data moments are little changed if we set the minimum number of followers to 200, 300, 400 or 500.

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

neutral, we search for many different word combinations. A number of recent papers, such as Tetlock (2007) and Da, Engelberg, and Gao (2015), use Harvard IV-4 dictionary and word counting to conduct text analysis. This approach is shown to be effective in analyzing the content of financial articles and Google search words. However, the dictionary is not structured to capture the vocabulary used by investors. Since opinionated tweets about the exchange rate are usually posted by individuals that are directly involved in the FX market, 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. When there are multiple tweets from one account during a day, we only use the last tweet on that day.[11]

There are on average 40.4 such opinionated tweets per day, for a total of 44,568 during our sample. Therefore only about 8.4% of all tweets with the word EURUSD are opinionated tweets. Figure 1 shows the distribution of daily tweets. It varies a lot across days. The standard deviation of the number of daily tweets is 22.4. The 44,568 opinionated tweets come from 7,567 separate accounts, implying an average of 5.9 tweets per account over the entire 1103 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 44,568 tweets all came from individuals tweeting every day, there would have been only 40 separate accounts. We are clearly capturing a far more dispersed group of people expressing opinions. Figure 2 shows the distribution of the number of followers of accounts that posted opinionated tweets during the sample period, subject

[11] On average 18.8% of tweets counted this way are from accounts from which multiple tweets were sent during a day.

to the minimum of 100 followers.

We will denote the numerical Twitter Sentiment during day t by individual i as TS_t^i. Figure 3 shows the distribution of the three values (-1, 0 and 1) that individual Twitter Sentiment takes across the entire sample. The distribution is quite even across the three values, although the percentage of negative values is somewhat larger than the percentage of positive values. This is because the Euro depreciated by 21% during this particular sample.

We also construct a daily Twitter Sentiment Index (TSI). It is computed by taking the simple average of the numerical Twitter Sentiment across tweets during a day. We denote this as TS_t on day t. So

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

where n_t is the number of opinionated tweets on day t. We set the index to 0 for trading days with no opinionated tweets.[12] Figure 4 shows the distribution of the daily TSI. Figure 5 displays the cumulative TSI, accumulated from the start of the sample, together with the EURUSD exchange rate over the sample period. It shows that the cumulative TSI is closely related to the exchange rate, which provides a first piece of evidence that the TSI is not some random noise.

2.3 Financial Lexicon

Tables B1 and B2 in Appendix B 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).[13] In order to provide some perspective, Table 1 provides examples of tweets and how they are categorized. The words in the tweet used to identify them are underlined.

In Table 1, the first tweet under the positive category is identified as positive because investors

[12] There are 42 trading days in the sample that no opinionated tweet was posted or twitter data is missing due to technical difficulties in the data collection system.

[13] It is possible that a tweet has word combinations in more than one category. We classify a tweet with both positive and neutral word combinations as positive. Similarly, a tweet with negative and neutral words combinations is classified as negative. A tweet with both positive and negative word combinations is classified as neutral.

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

It is important to note that using a dictionary that includes the lexicon of traders is crucial in correct classification of the opinionated tweets. General natural language tools are not well structured to catch opinions of traders because they are trained to measure the sentiment of conversational language. In order to show the important role of a finance dictionary, we used Google's machine learning tool to classify all opinionated tweets in our data sample. It measures the sentiment of a given text and provides a score between -1 and +1. We follow the score range guideline and place all the tweets within the score range of -1 to -0.25 in the negative category, +0.25 to +1 in the positive category and the rest in the neutral category.

As an illustration of the ineffectiveness of such a general natural language tool, consider the tweets listed in Table 1.[14] Out of 9 tweets listed under the positive category, Google's machine learning tool only identifies the seventh tweet correctly. The first six tweets are classified as neutral and the last two tweets are classified as negative. Table 1 lists three tweets under the neutral category. Machine learning identifies the second tweet as positive and the rest as neutral. Finally, there are 9 tweets listed under the negative category. Machine learning only identifies the first two and the eighth tweet as negative. The rest are incorrectly identified as neutral and in

[14]The experiment can easily be replicated by visiting Google's “Cloud Natural Language” website.

the case of the sixth tweet it is identified as positive. In the next section, which considers whether tweets get the direction of the subsequent exchange rate change correct, we will further illustrate that our classification scheme based on financial lexicon is far superior to that based on Google's machine learning tool.

2.4 Directional Moments

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 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 $TS_t^i = 1$ and the subsequent exchange rate change is positive (negative), we assign the tweet a +1 (-1). Similarly, if $TS_t^i = -1$ 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 $TS_t^i = 0$, 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 whether the tweets contain information about the subsequent direction of the exchange rate, we consider the null hypothesis that the tweets contain no information at all and therefore consist of random guesses. Under that null hypothesis one can analytically compute the standard deviation of directional moments. If the probability of posting a neutral tweet is p, then under random guessing the probability that a tweet is assigned a zero is p, while the probabilities of assigning a +1 or -1 are both $(1-p)/2$. In that case for each tweet the expectation of the directional moment is zero, with a variance of $1-p$. We have 44568 tweets in our data sample, so the variance of the directional moments over any time horizon is $\frac{1-p}{44568}$. Setting $p = 0.3$, as 30 percent of our tweets are neutral, the standard deviation of directional moments under the null of guessing is $\sqrt{\frac{1-0.3}{44568}} = 0.0039$. The directional moments reported in Table 2 are all positive and are respectively 4.2, 5.5 and 3.9 standard deviations away from zero for the one month, two month and three month subsequent exchange rate change. The null hypothesis of guessing is therefore strongly rejected.

If we used Google's machine learning tool, discussed in the previous section, the same three directional moments reported in Table 2 would be respectively 0.0027, 0.0017 and -0.0039. These are all virtually zero and we would not be able to reject the null of guessing. This illustrates again that it is critical to be familiar with the wording used by traders to properly interpret the tweets.

2.5 Moments for Model Evaluation

In Section 4 we will confront the model of Section 3 with a variety of data moments involving TS_t^i, TS_t and s_t. These moments are reported in the first column of Table 3. The first moment is the variance of TS_t. As we discuss in the Online Appendix, 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 next moment is disagreement, which captures the extent to which opinions differ among individuals during a particular day. It is the average across the 1103 days of the cross sectional variance of TS_t^i across the individuals. We again focus on the variance for easier comparison to the model. We do not include the days for which the number of tweets is 0 or 1.[15] Figure 6 shows the distribution of the daily cross sectional variance.

The next eight moments capture the relationship between Twitter Sentiment and exchange rate changes. Four of these are the correlation of the Twitter Sentiment Index with the change in the exchange rate over the next 1, 20, 40 and 60 trading days. The next three moments are the directional moments for 20, 40 and 60 day horizons, described in Section 2.4 . The next moment is the contemporaneous correlation between the weekly Twitter Sentiment Index and weekly changes in the exchange rate. The weekly Twitter Sentiment Index TS_w is defined as the average of the daily Twitter Sentiment Index over five trading days in a week.

Finally, the last two moments are the standard deviation and autocorrelation of the daily change in the exchange rate. The standard deviation of the daily change in the exchange rate is in percent, so it is 0.53%=0.0053. The daily autocorrelation is -0.003, reflecting the near random walk aspect of the exchange rate.

[15] There are 42 days during which we have less than two tweets.

3 Model Description

3.1 Some Preliminaries

We will use the model of Bacchetta and van Wincoop (2006) (from hereon BvW). We will be relatively brief in the description as BvW develop the micro foundations and provide further details. Three aspects of the model should be emphasized. First, agents rationally process all public and private information to form expectations. Second, no a priori assumption is made about the information quality. Specifically, we leave open the possibility that private signals contain no useful information about the future. The private information in the model is exogenous and takes the form of private signals. Third, we do not take a position on where the information comes from. It may for example be related to different research findings, talking to different people in the market, reading different financial news or perhaps even being exposed to different tweets or other social media opinions.

It should be emphasized that the concept of a "tweet" does not exist within the model. We interpret tweets as an expression of beliefs about the direction of the exchange rate by a subset of agents. In the next section we will relate expectations of exchange rate changes that exist in the model to directional beliefs expressed through tweets. There are several reasons to take this approach. First, unless there is reason to believe that tweets are biased for reasons discussed in Section 2.1, modeling the motivation for tweeting will not provide new insights. For example, people might tweet to advertise their investment research services or simply because they enjoy the social media interaction with others operating in the foreign exchange market. As we discussed in Section 2.1, this does not necessarily cause any bias, such as tweets that express views which differ from the investor's true opinions. Second, we do not have hard data on the relative importance of various motivations for tweeting. Third, while it may be interesting to model endogenous information acquisition with communication through tweets, this would significantly complicate the model without any clear benefit in the absence of known bias.

3.2 Model Equations

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

$$\int_0^1 b^i_{F,t} di = 0 \tag{2}$$

Portfolio demand by agents is

$$b^i_{F,t} = \frac{E^i_t s_{t+1} - s_t + i^*_t - i_t}{\gamma \sigma^2} - b^i_t \tag{3}$$

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. 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 conditional variance of s_{t+1} is σ^2, which is the same for all agents.

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 its aggregate across agents:

$$b_t = \int_0^1 b^i_t di \tag{4}$$

for which we assume an AR process:

$$b_t = \rho_b b_{t-1} + \varepsilon^b_t \tag{5}$$

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

Standard money demand equations are assumed:

$$m_t = p_t + y_t - \alpha i_t \tag{6}$$

$$m^*_t = p^*_t + y^*_t - \alpha i^*_t \tag{7}$$

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

$$i_t - i_t^* = \frac{1}{\alpha}(s_t - f_t) \tag{8}$$

where $f_t = (m_t - m_t^*) - (y_t - y_t^*)$ is the traditional fundamental.[16] Since the exchange rate is an $I(1)$ variable in the data, the fundamental is assumed to be $I(1)$ as well. We assume

$$f_t - f_{t-1} = \rho(f_{t-1} - f_{t-2}) + \varepsilon_t^f \tag{9}$$

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

Denote $\bar{E}_t s_{t+1} = \int_0^1 E_t^i s_{t+1} di$ as the average expectation across agents. Substituting (3) and (8) into the market clearing condition (2), we have

$$\bar{E}_t s_{t+1} - \frac{1+\alpha}{\alpha} s_t + \frac{1}{\alpha} f_t = \gamma\sigma^2 b_t \tag{10}$$

Imposing the market clearing condition (10) allows us to solve for the equilibrium exchange rate.

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

$$v_t^i = f_{t+T} + \epsilon_t^{v,i} \tag{11}$$

where $\epsilon_t^{v,i} \sim N(0, \sigma_v^2)$. As usual in the noisy rational expectations literature, the average of the signal errors is assumed to be zero across agents.

[16] This equation is obviously a significant simplification. PPP does not hold in the data as the real exchange rate is very volatile. Moreover, (8) implies a linear relationship between the exchange rate and the observed fundamental f_t and interest differential that we do not see in the data. More generally, there should be additional terms in (8) that are related to deviations from PPP and monetary policy shocks. This would add additional observed fundamentals beyond f_t to the model. Here we follow BvW and focus on the disconnect between the exchange rate and the single observed fundamental f_t. We do not include the interest differential as an observed fundamental as the interest differential is endogenous and (8) is obviously oversimplified.

(11) 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^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}.

We leave a discussion of the solution of the model to Appendix A as this is a bit technical. The solution takes the following form:

$$s_t = A(L)\varepsilon_{t+T}^f + B(L)\varepsilon_t^b \tag{12}$$

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. The exchange rate therefore depends on past, current and future innovations of the fundamental f_t as well as current and past noise innovations. The future innovations in the fundamental f_t enter as a result of private information.

As discussed in the Appendix A, from the solution of the exchange rate we derive expectations of future exchange rates by individual agents in the model:

$$E_t^i s_{t+k} = \bar{E}_t s_{t+k} + \mathbf{z}_k' \mathbf{M} \mathbf{w}_t^i \tag{13}$$

where $\mathbf{z}_k = (a_{k+1}, ..., a_{T+k}, b_{k+1}, ..., b_{T+k})'$, $\mathbf{w}_t^i = \left(0, ..., 0, \epsilon_t^{v,i}, ..., \epsilon_{t-T+1}^{v,i}\right)'$ and the average expectation $\bar{E}_t s_{t+k}$ can be written as a function of past and future innovations of the fundamental f_t and noise innovations. So the expectation of the future exchange rate s_{t+k} by an individual agent is equal to the average expectation of all agents plus an idiosyncratic component $\mathbf{z}_k' \mathbf{M} \mathbf{w}_t^i$ that depends on the signal errors of that agent.

4 Estimation Model Parameters

We will use data on Twitter Sentiment and exchange rates to estimate the parameters of the model. We first discuss how to compute Twitter Sentiment in the theory and then describe the estimation of model parameters through the Simulated Method of Moments.

4.1 Computing TS in the Theory

The tweets in the data express directional beliefs about changes in the exchange rate without a stated horizon. We will assume that the horizon corresponds to T in the model as this is the period over which agents have private information. Agents have no information of model innovations beyond T periods. 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, we set

$$TS_t^i = \begin{cases} 1 & \text{if} \quad E_t^i(s_{t+T} - s_t) > c \\ 0 & \text{if} \quad -c \leq E_t^i(s_{t+T} - s_t) \leq c \\ -1 & \text{if} \quad E_t^i(s_{t+T} - s_t) < -c \end{cases} \tag{14}$$

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

What remains is to identify the proper value for the cutoff c. Let π be the fraction of all observations in the data for which TS_t^i is 0. We equate this to the unconditional probability of drawing a 0 in the model:

$$Prob\left(-c \leq E_t^i(s_{t+T} - s_t) \leq c\right) = \pi \tag{15}$$

Since

$$E_t^i(s_{t+T} - s_t) = \bar{E}_t(s_{t+T} - s_t) + \mathbf{z}_T'\mathbf{M}\mathbf{w}_t^i \tag{16}$$

we can compute the unconditional variance of this expectation as

$$\sigma_E^2 = var(E_t^i(s_{t+T} - s_t)) = var(\bar{E}_t(s_{t+T} - s_t)) + \mathbf{z}_T'\mathbf{MRM}'\mathbf{z}_T \tag{17}$$

where $var(\bar{E}_t(s_{t+T} - s_t))$ is computed by first writing the average expectation as a linear function of all shocks ε_{t+T-s}^f and ε_{t-s}^b with $s \geq 0$ and then taking the unconditional variance. $\mathbf{R}$ is a $2T$ by $2T$ matrix with σ_v^2 on the last T elements of the diagonal and zeros otherwise.

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

$$Prob\left(-\frac{c}{\sigma_E} \leq \frac{E_t^i(s_{t+T}-s_t)}{\sigma_E} \leq \frac{c}{\sigma_E}\right) = \pi \tag{18}$$

it must be that

$$\Phi\left(\frac{-c}{\sigma_E}\right) = \frac{1-\pi}{2} \tag{19}$$

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

$$c = -\sigma_E \Phi^{-1}\left(\frac{1-\pi}{2}\right) \tag{20}$$

In the data we have $\pi = 0.3025$.

4.2 Simulated Method of Moments

Table 3 lists 12 moments. Eleven of these moments are used to estimate the model parameters. The standard deviation of the daily change in the exchange rate is separately used for scaling purposes as discussed below. We estimate the model using the Simulated Method of Moments. The parameters are chosen in order to minimize

$$\left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right)' \mathbf{\Sigma}^{-1} \left(\mathbf{m}^{data} - \mathbf{m}^{model}(\nu)\right) \tag{21}$$

Here $\mathbf{m}^{data}$ is the vector of 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. We compute both the model moments $\mathbf{m}^{model}$ and the variance $\mathbf{\Sigma}$ of the moments based on 1000 simulations of the model.[17] The Online Appendix provides further technical details.

Following many others, we only use the diagonal elements of the weighting matrix as the full matrix can lead to finite sample bias (e.g. Altonji and Segal (1996)). The objective function is therefore

$$\sum_{i=1}^{11}\left(\frac{\mathbf{m}^{data}(i) - \mathbf{m}^{model}(i)}{\mathbf{\Sigma}_{\mathbf{ii}}^{\mathbf{0.5}}}\right)^2 \tag{22}$$

[17] We obtain parameter estimates for a given weighting matrix, then use these parameter estimates to compute a new weighting matrix. We iterate a couple of times this way until the results no longer change.

This is equal to the sum of the of the squared t-values of all 11 moments. $\mathbf{\Sigma_{ii}^{0.5}}$ is the standard deviation of moment i across simulations of the model.

The variance covariance matrix of parameter estimates is given by

$$\frac{1}{S}\left[\left(\frac{\partial \mathbf{m}^{model}}{\partial \nu}\right)' \mathbf{\Sigma}^{-1}\left(\frac{\partial \mathbf{m}^{model}}{\partial \nu}\right)\right]^{-1} \tag{23}$$

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

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

The other parameters of the model are σ_v, σ_b, ρ, ρ_b, γ, σ_f and T. We only estimate the first 4 of these parameters. Some comments are therefore in order about γ, σ_f and T. From (10) it can be seen that γ enters the model multiplied by b_t. As a result of this we can only estimate $\gamma\sigma_b$. We therefore normalize $\gamma = 1$ and estimate σ_b. If instead one wishes to set $\gamma = 10$ the reported estimate for σ_b below simply needs to be divided by 10. We set σ_f by exploiting a scaling feature of the model. If we multiply σ_f and σ_v by a factor q, while dividing σ_b by q, the only effect is to scale up the standard deviation of the exchange rate by a factor q. None of the other moments in the model change. We can therefore choose an arbitrary σ_f and estimate the other parameters based on moments other than the standard deviation of the exchange rate. Afterwards we scale σ_f, σ_v 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. Below we will report results for $T = 20$, $T = 60$ and $T = 240$. These correspond to information about future fundamental innovations up to respectively one month, one quarter and one year into the future.

4.3 Estimation Results and Model Moments

Table 3 reports (at the bottom) the estimated parameters and standard errors for each of the three values of T. The corresponding data and model moments, and t-values, are shown on top. The objective function is also reported. The objective has a χ^2 distribution with 7 degrees of freedom (number of moments minus number of estimated parameters). At the 5 percent significance level the cutoff is 14.07. For all values of T the objective function is well below that, so that we cannot reject the null hypothesis that the model is correct. The reported t-values are all well below 2 (mostly even well below 1), with the exception for the weekly contemporaneous correlation between Twitter sentiment and the exchange rate for $T = 240$, which is 2.55. The variance of Twitter Sentiment and disagreement among agents are both matched very closely for all values of T. The predictive correlations are particularly close to the data for $T = 240$, but these moments are not very informative due to large standard errors across model simulations due to the short sample. The same is the case for the directional moments, where the numbers are also broadly in line with the data (especially for $T = 60$). The model exhibits virtually zero autocorrelation of daily exchange rate changes, as seen in the data.

The estimated standard deviation of signal errors σ_v is larger for larger T. This makes sense as agents have more private signals as T increases, so each signal naturally needs to be weaker. Importantly, all parameters are estimated with reasonable precision. For example, when $T = 240$, the point estimate of σ_v is 3.3 with a standard error of 0.19. The same applies to the other parameters.

We conduct a couple of robustness exercises. First, since the Euro depreciated by 21 percent over the sample, we also evaluate the model over the three-year period from January 1, 2015 to December 31, 2017, during which there was no trend change in the exchange rate. Based on the data moments over this period we find that the value of the objective function is respectively 2.5, 4.5 and 11.8 for the three values of T and using the parameter estimates from Table 3. We therefore are again unable to reject the model with the estimated parameters. Second, to consider sensitivity to the weighting matrix, we use the standard deviations of the moments based on the data instead of the model. These are computed using block bootstrapping with a one-month (20 trading day) block size. We find that this does not significantly change the objective function, so that we are again unable to reject the model with the estimated parameters.[18]

5 Meese-Rogoff Results

In this section we will first describe Meese-Rogoff type moments that are commonly reported in the empirical literature. They take the form of MSE (mean squared error) ratios. We then discuss various MSE ratios that can be computed based on the theory with the estimated parameters and compare them to MSE ratios that have been reported in the empirical literature. We do so separately for MSE ratios based on short-run and long-run relationships between exchange rates and fundamentals.

5.1 Results from Empirical Studies

The empirical literature considers two different types of regressions. One type, which we will focus on, considers the contemporaneous relationship between changes in the exchange rate and changes in macro fundamentals such as the money supply and output. This relates directly to the question of disconnect: how strong is the connection between changes in the exchange rate and changes in observed fundamentals. The second type of regression uses error correction specifications. These are purely predictive relationships that tell us how much the exchange rate is expected to change over a future horizon based on the current deviation of the exchange rate s_t from $\boldsymbol{\beta}'\mathbf{f}_t$, where $\mathbf{f}_t$ is a vector of fundamentals and $\boldsymbol{\beta}$ is a vector of coefficients. Since our interest here is in the disconnect question rather than exchange rate forecasting, we will focus on the contemporaneous relationship between changes in exchange rates and fundamentals.

The empirical literature usually proceeds as follows. First a relationship between the change in the exchange rate and fundamentals is estimated based on the first L data points in the sample:

$$s_t - s_{t-1} = \alpha + \boldsymbol{\beta}' \left(\mathbf{f}_t - \mathbf{f}_{t-1} \right) + \varepsilon_t \tag{24}$$

Then this relationship is used to "predict" out of sample. It is important to emphasize that the term "predict", as well as the term "forecast" used below, is in fact a misnomer as in reality the

[18] The objective function for the three values of T is now respectively 8.9, 4.5 and 10.5.

actual future fundamentals are used to "predict" what the change in the exchange rate will be. Rather than telling us something about the predictability of future exchange rates, this relationship sheds light on the explanatory power of observed macro fundamentals for changes in the exchange rate. For example, to what extent can the change in the exchange rate from January to February be explained by the change in the relative money supply from January to February?

For illustrative purposes, let the data frequency be monthly. Then the one-month out of sample prediction is

$$\hat{\alpha} + \hat{\boldsymbol{\beta}}' (\mathbf{f}_{t+1} - \mathbf{f}_t) \tag{25}$$

where $\hat{\alpha}$ and $\hat{\boldsymbol{\beta}}$ are the estimated parameters. The one-year out of sample prediction is

$$12\hat{\alpha} + \hat{\boldsymbol{\beta}}' (\mathbf{f}_{t+12} - \mathbf{f}_t) \tag{26}$$

This is then compared to the actual change in the exchange rate over this period, respectively $s_{t+1} - s_t$ and $s_{t+12} - s_t$ for the examples above. The error is recorded, which is the difference between the "predicted" and actual change in the exchange rate.

It is most common to repeat this with recursive rolling regressions. The sample each time is shifted forward one period, the regression is repeated, new forecasts and forecast errors are computed. This continues until the forecast reaches the end of the data sample. The MSE (mean squared error) is the mean of the squared forecast errors. An MSE ratio is computed by dividing the MSE to what it would be if the RW (random walk) forecast were used, in which case the forecasted change in the exchange rate is always zero.

There are various aspects of this approach that vary across papers in the literature. First, the frequency of the data varies, usually monthly, quarterly or annual. Second, the length of the data sample varies across papers. Third, the length of the sample L over which estimation takes place varies. In addition, some papers do not recursively roll the sample, but instead hold the first period of the estimation sample fixed and keep adding to the end of the estimation period. In that case the length of the estimation sample gradually increases. Fourth, the forecast period varies. When the forecast period is equal to the data frequency, we refer to it as a short-term forecast: one-month ahead forecasts with monthly data, one-quarter ahead forecasts with quarterly data. But often long-term forecasts are considered as well, where the forecast period is much longer than the data frequency. The exact forecast horizon can vary across papers. A final difference relates to the set of fundamentals included in the vector $\mathbf{f}$.

We will compare the results from the model to three papers: Rossi (2013), Cheung et.al. (2005) and Cerra and Saxena (2010). The details associated with these papers are reported in Table 4. They use respectively monthly, quarterly and annual exchange rate data. The length of the data sample is respectively 54.4 years, 27.9 years and 45 years. Regarding parameter estimation, Rossi (2013) and Cheung et.al. (2005) use recursive rolling with L equal to respectively 326 months (27.2 years) and 39 quarters (9.8 years). Cerra and Saxena (2010) start with an estimation sample of 23 years and then each time add one year to the estimation sample instead of recursive rolling. Long-horizon forecasts are over 4 years in Rossi (2013) and 5 years in the other two papers. The set of fundamentals also varies across the papers. For Cerra and Saxena (2010) we consider results where relative money and relative output are used as fundamentals (their model E). For Rossi (2013) we consider results based on their monetary model: relative interest rate, relative money and relative output. For Cheung et. al. (2005) we consider results when the fundamentals include the relative interest rate, relative money, relative output and relative inflation (their sticky price model).

Results from these three papers are reported in the last row of Tables 5 and 6. All results are in terms of MSE ratios. Rossi (2013) and Cerra and Saxena (2010) report RMSE (root mean squared error) ratios. We square these results to translate them to MSE form. Rossi (2013) and Cheung et.al. (2005) report results for various currency pairs (all relative to the dollar). In Tables 5 and 6 we show a range for these papers, which reflects the variation across these currency pairs. For Cerra and Saxena (2010) we report the findings when using the sample of all low inflation episodes. Their paper stands out for the large number of countries used, a total of 98. This includes countries with fixed exchange rate systems. They report results when combining the forecast errors of all these countries.[19] It is evident from Table 5 that the model does not outperform the random walk over relatively short horizons of one month, one quarter and one year. MSE ratios tend to be close to 1. Over the longer horizons of 4 or 5 years, we see in Table

[19]They allow the regression coefficients to vary across currency pairs. All but one of their 98 countries had low inflation episodes during their sample.

6 that in only one paper, Cerra and Saxena (2010), the model substantially outperforms the RW, with an MSE ratio of 0.56. The better long-horizon performance in Cerra and Saxena (2010) is likely due to both the much broader sample of countries and the use of annual data. The model results below shed further light on this.

5.2 MSE ratios in the Model

We have estimated the model based on daily data. Consider the change of the exchange rate over k days in the model. If $T > k$, the change of the exchange rate over the last k days is

$$\begin{aligned} s_t - s_{t-k} &= \sum_{i=1}^{k} a_i \varepsilon^f_{t+T-i+1} + \sum_{i=1}^{T-k} (a_{k+i} - a_i)\, \varepsilon^f_{t+T-k+1-i} + \\ &\quad \sum_{i=1}^{\infty} (a_{T+i} - a_{T-k+i})\, \varepsilon^f_{t+1-i} + \\ &\quad \sum_{i=1}^{k} b_i \varepsilon^b_{t+1-i} + \sum_{i=1}^{\infty} (b_{i+k} - b_i)\, \varepsilon^b_{t-k+1-i} \end{aligned} \tag{27}$$

If $T \leq k$, the change of the exchange rate over the last k days is

$$\begin{aligned} s_t - s_{t-k} &= \sum_{i=1}^{T} a_i \varepsilon^f_{t+T-i+1} + \\ &\quad \sum_{i=1}^{k-T} a_{T+i} \varepsilon^f_{t+1-i} + \sum_{i=1}^{\infty} (a_{k+i} - a_i)\, \varepsilon^f_{t-k+T+1-i} + \\ &\quad \sum_{i=1}^{k} b_i \varepsilon^b_{t+1-i} + \sum_{i=1}^{\infty} (b_{i+k} - b_i)\, \varepsilon^b_{t-k+1-i} \end{aligned} \tag{28}$$

In these expressions, the right hand side terms in the first row are future fundamental innovations, the terms in the second row are current and past fundamental innovations that can be observed, while the terms in the third row are the noise innovations. Both future fundamental innovations and noise shocks lead to a disconnect of exchange rate changes from observed fundamentals.

In the model we can compute a theoretical MSE ratio that measures the variance of exchange rate changes due to shocks unrelated to observed fundamental innovations relative to the variance of exchange changes due to all shocks. We refer to this as MSE^{theory}. This is a ratio that by construction is always between 0 and 1. Note that

$$1 - MSE^{theory}$$

measures the fraction of the variance of exchange rate changes that is explained by observed fundamental innovations. So a lower MSE^{theory} implies that observed fundamental innovations have more explanatory power for exchange rate changes.

MSE^{theory} is the best measure of the explanatory power of the observed fundamentals in the model, but it is not directly comparable to the Meese-Rogoff MSE ratios reported in the empirical literature. A first step towards the empirical Meese-Rogoff MSE ratios is to consider the explanatory power of the change in the fundamental over the same k days as the change in the exchange rate instead of the explanatory power of all past fundamental innovations. In other words, consider a regression

$$s_t - s_{t-k} = \alpha + \beta(f_t - f_{t-k}) + \varepsilon_t \tag{29}$$

We can compute the precise value of β in the model by computing $cov(s_t - s_{t-k}, f_t - f_{t-k})/var(f_t - f_{t-k})$. α is zero in the model. The corresponding MSE ratio is

$$MSE^{theory,\Delta f} = \frac{var(\varepsilon_t)}{var(s_t - s_{t-k})} \tag{30}$$

It is always the case that

$$MSE^{theory,\Delta f} > MSE^{theory}$$

as MSE^{theory} uses all past fundamental innovations as explanatory variables rather than just the fundamental change over the past k days.

This last MSE ratio still does not correspond to the Meese-Rogoff MSE ratios reported in the empirical literature. The reason is that in empirical work the parameters α and β must be estimated based on a limited finite sample of data, leading to estimation errors that increases the MSE ratio. We refer to the MSE ratio in empirical work as

$$MSE^{Meese\ Rogoff} = \frac{var(\hat{\varepsilon}_t)}{var(s_t - s_{t-k})} \tag{31}$$

where $\hat{\varepsilon}_t = s_t - s_{t-k} - \hat{\alpha} - \hat{\beta}(f_t - f_{t-k})$. Here $\hat{\alpha}$ and $\hat{\beta}$ are the estimated parameters, while the variance in the numerator and denominator of (31) are finite sample variances (mean squared errors).[20] Estimation error of the parameters raises the MSE ratio further. To summarize, we

[20] Also note that the estimated parameters are actually time-varying due to rolling regressions.

have

$$MSE^{Meese\ Rogoff} > MSE^{theory,\Delta f} > MSE^{theory}$$

Apart from estimation error of the parameters, for long-horizon forecasts (4 or 5 years in the papers above) there is one additional difference between $MSE^{Meese\ Rogoff}$ and $MSE^{theory,\Delta f}$. This is best illustrated with an example. Assume that the data frequency is monthly, as in Rossi (2013), and the forecast horizon is 4 years. Then k is equal to 20 (trading days in a month) times 12 times 4, which is 960. Instead of conducting the regression

$$s_t - s_{t-960} = \alpha_l + \beta_l(f_t - f_{t-960}) + \varepsilon_t \tag{32}$$

Rossi (2013) estimates

$$s_t - s_{t-20} = \alpha_s + \beta_s(f_t - f_{t-20}) + \varepsilon_t \tag{33}$$

Here the subscripts l and s stand for long and short. So Rossi uses monthly changes in the exchange rate and fundamentals to estimate the parameters α_s and β_s, which are then applied to the 4-year forecast horizon. The long-term forecast error is then

$$s_t - s_{t-960} - 48\hat{\alpha}_s - \hat{\beta}_s(f_t - f_{t-960}) \tag{34}$$

This creates an additional error. Not only is there estimation error of the coefficients, the incorrect short-horizon coefficients are applied to long-horizon forecasts. One can define an intermediate MSE ratio based on theory:

$$MSE^{theory,\Delta f,short} = \frac{var\left(s_t - s_{t-960} - \beta_s(f_t - f_{t-960})\right)}{var(s_t - s_{t-960})} \tag{35}$$

This theoretical MSE ratio uses the theoretical value of the short term coefficient β_s from (33) and applies it to the long-term forecast.[21] The resulting error determines the variance in the numerator. We have

$$MSE^{Meese\ Rogoff} > MSE^{theory,\Delta f,short} > MSE^{theory,\Delta f} > MSE^{theory}$$

Going from right to left, the MSE ratio rises in three steps from the ratio MSE^{theory} that measures the true explanatory power of the observed fundamentals to the ratio $MSE^{Meese\ Rogoff}$ that

[21] The theoretical intercept α_s is zero.

corresponds to empirical results. It rises first because only the change in the fundamental over the past k days is used to explain the change in the exchange rate, instead of all past changes in the fundamental. It rises further because the incorrect short-term coefficient is applied to the long-horizon relationship. Finally, it rises because the short-term coefficient is estimated with error due to finite data samples.

While for illustrative purposes we discussed a monthly data frequency and 4-year forecast horizon as in Rossi (2013), the other two papers use an analogous methodology to compute the Meese Rogoff MSE ratio for long horizons. For example Cheung et.al. (2005) regress quarterly changes in the exchange rate on quarterly changes in the fundamentals and then apply the estimated coefficients to forecasts over the next 5 years.

Below we will report all these MSE ratios in the model for short and long-horizon forecasts. When reporting $MSE^{Meese\ Rogoff}$ we apply a methodology within the model that corresponds to that in the empirical papers discussed above, as summarized in Table 4. However, there are two aspects of the papers that we are unable to capture with the model. One is that there is only one fundamental variable in the model, while there are more fundamentals in the empirical papers. The second is that the empirical results are based on multiple currencies as described above. The results in Cerra and Saxena (2010) are based on pooling these currencies, while the other papers produce a range of results that vary across currencies.

5.3 Short Horizon MSE Ratios

Tables 5 and 6 report the MSE ratios in both the model and data. Table 5 reports results for short-horizon forecasts, which is respectively 1 month, 1 quarter and 1 year for the three papers. Table 6 reports results for the long-horizon forecasts over 4 or 5 years. Results in the model are reported for $T = 20$, $T = 60$ and $T = 240$, based on the estimated parameters in Table 3.

First consider the short-horizon results. $MSE^{Meese\ Rogoff}$ in the model is computed by conducting 100,000 model simulations with sample lengths that correspond to the empirical papers. For each simulation we compute $MSE^{Meese\ Rogoff}$ following the methodology of each paper summarized in Table 4. Table 5 reports both the average $MSE^{Meese\ Rogoff}$ across these simulations as well as its standard deviation across simulations.

Comparing $MSE^{Meese\ Rogoff}$ in the model to the numbers reported in the empirical papers, we find that the theory is always consistent with the empirical papers when $T = 240$. More generally, the theory is consistent with the empirical papers as long as T is not too low. Specifically, we need $T \geq k$, where k is the short-term forecast horizon that corresponds to the data frequencies in the three papers: $k = 20$ for monthly data, $k = 60$ for quarterly data $k = 240$ for annual data. When $T \geq k$, the MSE ratios from the model are close to 1 and do not deviate in a statistically significant way from the numbers reported in the papers. When T is much smaller than the forecast horizon k, the MSE ratio is well below 1 in the model and therefore inconsistent with the data. Note that the same applies to MSE^{theory} and $MSE^{theory,\Delta f}$, so the fundamentals have little explanatory power in the theory when $T \geq k$, but too much explanatory power when T is considerably less than k.

Private information is key to the disconnect between exchange rates and fundamentals in the model. We find that setting σ_v very large, while keeping the other parameters the same, the MSE ratios in the theory are always very close to zero. In that case exchange rate changes are fully explained by observed fundamentals. The disconnect in the model is associated with the unobserved noise shocks as well as future fundamentals shocks. Private information plays a key role in their importance. We will discuss the role of these noise shocks and future fundamentals in some more detail, which also sheds light on why the model cannot explain the disconnect for values of T that are too low (well below k).

Noise shocks account for a significant share of variance of exchange rate changes over short horizons in the model. The charts on the left hand side of Figure 7 show the impulse response of exchange rate to a one standard deviation noise shock. For comparison, we also report the impulse response functions in the absence of private information, setting σ_v at an extremely high value while holding the other parameters constant. The impulse response functions show that the impact of noise shocks is much larger than without private information. For all reported T, the initial impact of a noise shock on the exchange rate is close to 25 times as large as it is in the model without private information.[22] The amplified effect of the noise shocks gradually dissipates

[22]Gennotte and Leland (1990) use some back of the envelope calculations suggesting that the amplification factor due to private information may have been as large as 250 during the October 19, 1987 stock market crash.

over time, but lasts T periods.

As emphasized by Bacchetta and van Wincoop (2006), this amplification of noise shocks due to private information is key to the disconnect between exchange rates and observed fundamentals over short horizons. They use the term "rational confusion" to refer to the phenomenon that leads to the amplification. When a noise shock raises the exchange rate, agents only observe the increase in the exchange rate. But the rise in the exchange rate may also be driven by new private information that agents have about future fundamentals. Seeing the increase in the exchange rate therefore rationally leads agents to raise their expectation of future fundamentals, which leads to a further increase in the exchange rate. The amplified noise shocks gradually dissipate with time as agents learn more by observing fundamentals and exchange rates. It dissipates entirely after T periods. When T is large, this results in a persistent disconnect between the exchange rate and observed fundamentals. This is one reason why for $T \geq k$ the model can account well for the observed disconnect in the data over an horizon k.

Disconnect also results from the private information directly, separate from its amplification of the noise shocks. Private information implies that the exchange rate depends on unobserved future fundamentals. This can be seen from the impulse response functions in Figure 7. The charts on the right hand side show the impulse response of a one standard deviation fundamentals innovation at time 0. This innovation starts to effect the exchange rate at time $-T$ because of the private information about fundamental innovations up to T periods into the future. The exchange rate response to the left of the y-axis therefore captures the effect on the exchange rate of unobserved future fundamentals. In the absence of private information these unobserved future fundamentals do not affect the exchange rate. Instead the exchange rate jumps to a level close to its long-run level at the time of a new fundamental shock. It is again clear that the larger T, the more persistent the disconnect is, again explaining why the model does well when $T \geq k$.

Consistent with these results, Figure 8 shows that a large fraction of the the variance of the exchange rate change $s_{t+k} - s_t$ is explained by the unobserved noise shocks and future fundamental innovations over horizons $k \leq T$. As k gets much larger than T, the observed fundamentals are the main drivers of the exchange rate. A larger T implies that noise shocks have a more persistent amplified effect on the exchange rate and the exchange rate depends on fundamental innovations farther into the future. This explains why the MSE ratios are close to 1 when $T \geq k$, while the MSE ratios are well below 1 (in contrast to the data) when T is well below k.

5.4 Long Horizon MSE Ratios

Next consider MSE ratios for long-horizon forecasts, reported in Table 6. The results are again consistent with the numbers reported in the literature when $T = 240$, while they are generally not consistent for $T = 20$ and $T = 60$.

First consider the case of $T = 240$. MSE^{theory} is always very low, close to 0.2. This is because a forecast horizon of 4 or 5 years is well beyond the period of 1 year over which agents have information when $T = 240$. Over 4 or 5 years the change in the exchange rate is therefore driven mostly by observed fundamentals. However, these theoretical MSE ratios cannot be compared to the empirical MSE ratios computed in the literature. There are three steps in between. The first step is to consider $MSE^{theory,\Delta f}$, the theoretical MSE ratios based on changes in fundamentals over the forecast horizon. This is only slightly higher than MSE^{theory}, between 0.2 and 0.3.

Next consider $MSE^{theory,\Delta f,short}$, where the incorrect short-horizon coefficient β_s is applied to the long-horizon change in the fundamental. This significantly raises the MSE ratio. The rise is from 0.28 to 0.68 when using the monthly data and from 0.22 to 0.66 when using quarterly data to compute the short-horizon coefficient. The increase is somewhat less, from 0.22 to 0.46, when using annual data to estimate the short-run regression coefficient, as in Cerra and Saxena (2010).

The final step considers $MSE^{Meese\ Rogoff}$, which is even higher due to parameter estimation error. It is slightly below the MSE ratios reported in the Rossi (2013), within the range of numbers reported in Cheung et al. (2005) and slightly above that of Cerra and Saxena (2010). But given the large standard errors for these long-horizon MSE ratios the difference relative to the empirical papers is never statistically significant.

Although MSE^{theory} tells us that exchange rate changes in the model are mostly driven by observed fundamentals over 4-5 years, the MSE ratios are much higher when computed in a way consistent with the empirical literature. This is the result of both parameter estimation error and applying an incorrect short-horizon regression coefficient to a long-horizon forecast. Table 7 sheds further light on the use of an incorrect short-horizon coefficient. It reports the theoretical coefficients of a regression of the change in the exchange rate on the change in the fundamentals. For $T = 240$, the coefficient is 0.85 over a 4-year horizon and 0.88 over a 5-year horizon. But over a one month horizon, the theoretical regression coefficient is only 0.21. For a quarterly horizon it is 0.22. For an annual horizon it rises a bit to 0.40. The reason that the coefficient rises with the horizon is as follows. As a result of private information the current exchange rate already incorporates some of the changes in the fundamental over the next T periods and more so for fundamental changes in the near future. This reduces the impact of fundamental changes in the near future on the exchange rate and therefore lowers the short-horizon regression coefficient.

Three more comments are worth making in this regard. First, the results imply that high MSE ratios in the data for long horizon forecasts do not mean that the observed fundamentals have little explanatory power over long horizons. Second, one may be inclined to think that the upward bias of the MSE ratios in the empirical literature due to the use of short-horizon regression coefficients can be resolved by using long-horizon regression coefficients. To some extent this is the case, but there is a tradeoff. Most datasets have only a limited number of non-overlapping 5-year intervals. This leads to a lot of noise in the estimation of long-horizon regression coefficients. The final point is related. Cerra and Saxena (2010) use a horizon of one year when regressing the change in the exchange rate on the change in the fundamentals. While this is still well short of 5 years, it is much longer than one month or one quarter and leads to a significantly higher regression coefficient. This explains why the MSE ratio in Cerra and Saxena (2010) is well below 1, in contrast to the other papers. It is 0.56 in the data and 0.65 in the theory.

So far we have focused on $T = 240$, where the model does a good job in accounting for the MSE ratios reported in the data. This is not the case for $T = 20$ and $T = 60$. Consider $T = 20$ for illustration. $MSE^{Meese\ Rogoff}$ is only 0.20 for the Cheung et al. methodology and 0.05 for the Cerra and Saxena methodology. This is well below that in the data, even allowing for standard errors. The reason for this is twofold. First, MSE^{theory} is now only 0.02, even lower than for $T = 240$. As discussed above, the lower the T relative to the forecast horizon k, the higher the explanatory power of observed fundamentals and the lower the MSE ratio. Second, there is less of an upward bias due to applying a short-horizon regression coefficient to a long-horizon forecast. $MSE^{theory,\Delta f,short}$ therefore remains low for quarterly and annual data.

Table 7 sheds light on the bias associated with using a short-horizon regression coefficient. When $T = 20$, the theoretical coefficient of a regression of the change in the exchange rate on the fundamental is only 0.22 for a monthly horizon. But it then rises significantly to 0.73 for a quarterly horizon and 0.94 for an annual horizon, while the coefficient is 0.98 for a 4-year horizon. With $T = 20$ a lot of the fundamental innovations over the next month have already been incorporated into the exchange rate at the start of the period, leading to a low regression coefficient over a monthly horizon. But most of the fundamental innovations over the next quarter or year are not yet incorporated in the current exchange rate, leading to a much higher regression coefficient for these horizons. The relatively small short-horizon bias in the regression coefficient for quarterly and annual data implies that $MSE^{theory,\Delta f,short}$ remains low, which carries over to $MSE^{Meese\ Rogoff}$.

6 Conclusion

Private information is by its nature unobservable. However, with the advent of social media many traders openly express their individual views about future asset prices. This opens up the question what can be learned from these private opinions. In this paper we have used opinionated tweets about the Euro-Dollar exchange rate to illustrate the important role of private information in the disconnect between the exchange rate and macro fundamentals.

We have developed a dictionary based on the financial lexicon used by FX traders to translate the tweets about the Euro-Dollar into beliefs about the future direction of the exchange rate. We then used these Twitter Sentiment data to estimate the parameters of a noisy rational expectations model for the foreign exchange market where agents have dispersed private information. We investigated the implications of the model, based on the estimated parameters, for the disconnect between exchange rates and macro fundamentals. We find that over both short and long horizons the model can account for the high MSE ratios that have been reported in the empirical literature and which are widely viewed as representing a disconnect between exchange rates and fundamentals.

The reason that the model can account for the high MSE ratios observed in the data is different for short than long horizons. For short horizons it reflects a true disconnect between exchange rates and fundamentals. Private information plays a key role in this regard. As a result of private information, the exchange rate depends on future fundamental innovations that the econometrician cannot yet observe. In addition, private information leads to a large amplification of the effect of unobserved noise shocks (portfolio shifts) on the exchange rate due to rational confusion. By contrast, the model tells us that over long horizons the exchange rate is mainly driven by observed fundamentals. The high MSE ratios for long horizons found in the empirical literature are instead attributed to the incorrect application of short-horizon regression coefficients to long-horizon forecasts, as well as parameter estimation error.

Appendix

A Solution of Model

In this Appendix we discuss the model solution. Agents have three sources of information. The first source consists of private signals from the last T periods. The second source consists of the past T exchange rates. They contain information about future fundamental innovations through the aggregation of the private information that agents trade on. The third source consists of current and past fundamentals f_t. We solve a signal extraction problem to compute rational expectations of the vector of innovations $\xi_t' = (\varepsilon^f_{t+T}, ..., \varepsilon^f_{t+1}, \varepsilon^b_t, ..., \varepsilon^b_{t-T+1})$, which affects expectations of future exchange rates. 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.

Start with a conjecture for the equilibrium exchange rate:

$$s_t = A(L)\varepsilon^f_{t+T} + B(L)\varepsilon^b_t \tag{A.1}$$

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

$$\bar{E}_t s_{t+1} = \theta' \bar{E}_t(\xi_t) + A^*(L)\varepsilon^f_t + B^*(L)\varepsilon^b_{t-T} \tag{A.2}$$

$$\sigma^2 = var_t(s_{t+1}) = a_1^2\sigma_f^2 + b_1^2\sigma_b^2 + \theta' var_t(\xi_t)\theta \tag{A.3}$$

where $\theta' = (a_2, a_3, ..., a_{T+1}, b_2, b_3, ..., b_{T+1})$, $A^*(L) = \sum_{i=1}^{\infty} a_{T+i+1}L^{i-1}$ and $B^*(L) = \sum_{i=1}^{\infty} b_{T+i+1}L^{i-1}$. The conditional variance $var_t(s_{t+1})$ is the same for all agents as they have the same quality information and therefore the same perceived uncertainty.

The expectation and variance of unknown innovations ξ_t is computed using signal extraction. 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^i_t,...,v^i_{t-T+1} and knowledge of the unconditional distribution of ξ_t. Solving the signal extraction problem (see BvW) yields

$$\bar{E}_t(\xi_t) = \mathbf{M}\mathbf{H}'\xi_t \tag{A.4}$$

$$var_t(\xi_t) = \tilde{\mathbf{P}} - \mathbf{M}\mathbf{H}'\tilde{\mathbf{P}} \tag{A.5}$$

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

$$\mathbf{H}' = \begin{bmatrix} a_1 & a_2 & ... & a_T & b_1 & b_2 & ... & b_T \\ 0 & a_1 & ... & a_{T-1} & 0 & b_1 & ... & b_{T-1} \\ ... & ... & ... & ... & ... & ... & ... & ... \\ 0 & 0 & ... & a_1 & 0 & 0 & ... & b_1 \\ d_1 & d_2 & ... & d_T & 0 & 0 & ... & 0 \\ 0 & d_1 & ... & d_{T-1} & 0 & 0 & ... & 0 \\ ... & ... & ... & ... & ... & ... & ... & ... \\ 0 & 0 & 0 & d_1 & 0 & 0 & ... & 0 \end{bmatrix} \tag{A.6}$$

Substituting (A.4) and (A.5) into (A.2) and (A.3) and the result into the market clearing condition (10), we have

$$\begin{aligned} &\theta'\mathbf{M}\mathbf{H}'\xi_t - \frac{1+\alpha}{\alpha}\left(A(L)\varepsilon_{t+T}^f + B(L)\varepsilon_t^b\right) + \frac{1}{\alpha}D(L)\varepsilon_t^f \\ &+A^*(L)\varepsilon_t^f + B^*(L)\varepsilon_{t-T}^b = \gamma\sigma^2\left(1 + \rho_b L + \rho_b^2 L^2 + ...\right)\varepsilon_t^b \end{aligned} \tag{A.7}$$

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 compute the expectations of future exchange rates by individual agents. In particular, we have

$$E_t^i s_{t+k} = \bar{E}_t s_{t+k} + \mathbf{z}_k'\mathbf{M}\mathbf{w}_t^i \tag{A.8}$$

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

$$\bar{E}_t s_{t+k} = \mathbf{z}_k'\bar{E}_t\xi_t + \sum_{l=0}^{\infty} a_{T+k+1+l}\varepsilon_{t-l}^f + \sum_{l=0}^{\infty} b_{T+k+1+l}\varepsilon_{t-T-l}^b \tag{A.9}$$

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

B Financial Lexicon

Tables B1 and B2 show all word combinations used to categorize tweets as positive, negative and neutral. Table B1 lists the word combinations in each category that require the explicit absence of some other words. Table B2 shows the list of word combinations whose existence in a tweet is enough to place the tweet in one of the categories.

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

Table B1: Word combinations used to identify opinionated tweets.

Positive

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

Table B1 (Continued): Word combinations used to identify opinionated tweets.

Positive

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

Table B1 (Continued): Word combinations used to identify opinionated tweets.

Negative

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

Table B1 (Continued): Word combinations used to identify opinionated tweets.

Negative

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

Neutral

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

Table B2: Word combinations used to identify opinionated tweets.

Positive

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

Table B2 (Continued): Word combinations used to identify opinionated tweets.

Negative

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

Table B2 (Continued): Word combinations used to identify opinionated tweets.

Neutral

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

References

[1] Altonji, Joseph G. and Lewis M. Segal (1996), "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business and Economic Statistics* 14 (3), 353-366.

[2] 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.

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

[4] Berger, David W., Alain P. Chaboud, Sergey V. Chernenko, Edward Howorka and Jonathan H. Wright (2008), "Order Flow and Exchange Rate Dynamics in Electronic Brokerage System Data," *Journal of International Economics* 75, 93-109.

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

[6] Cerra, Valerie, and Sweta Chaman Saxena (2010), "The Monetary Model Strikes Back: Evidence from the World," *Journal of International Economics* 81.2 : 184-196.

[7] Cheung, Yin-Wong, Menzie D. Chinn, and Antonio Garcia Pascual (2005), "Empirical Exchange Rate Models of the Nineties: Are any Fit to Survive?." *Journal of International Money and Finance* 24.7: 1150-1175.

[8] Da, Zhi, Joseph Engelberg and Pengjie Gao (2015),"The Sum of All FEARS: Investor Sentiment and Asset Prices," *Review of Financial Studies* 28(1), 1-32.

[9] Das, Sanjiv R. and Mike Y. Chen (2007), "Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web," *Management Science* 53(9), 1375-1388.

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

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

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

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

[14] Evans, D.D. and Dagfinn Rime (2012), "Micro Approaches to Foreign Exchange Determination," in Handbook Of Exchange Rates, edited by Jessica James, Ian W. Walsh and Lucio Sarno. Hoboken, New Jersey: Wiley and Sons.

[15] Gennotte, Gerard and Hayne Leland (1990), "Market Liquidity, Hedging, and Crashes," *American Economic Review* 80(5), 999-1021.

[16] 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.

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

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

[19] 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.

[20] 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).

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

[22] Rossi, Barbara 2013), "Exchange Rate Predictability." *Journal of economic literature* 51.4: 1063-1119.

[23] Tetlock, Paul C. (2007), "Giving Content to Investor Sentiment: the Role of Media in the Stock Market," *The Journal of Finance* 62,1139-1168.

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

[25] 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 Daily Number of Tweets

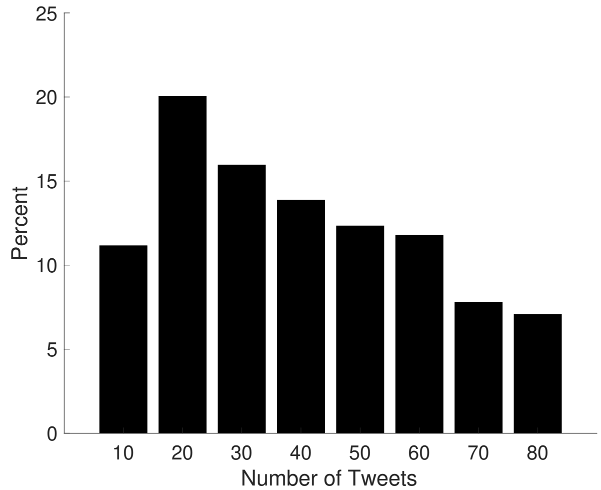

Figure 2 Distribution of the Number of Followers

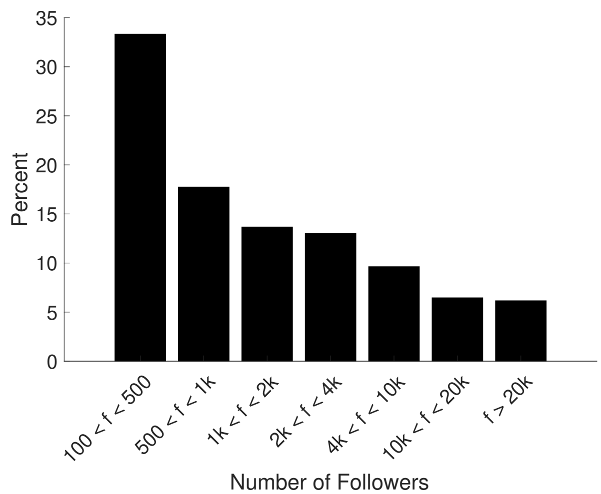

Figure 3 Distribution of Individual TS

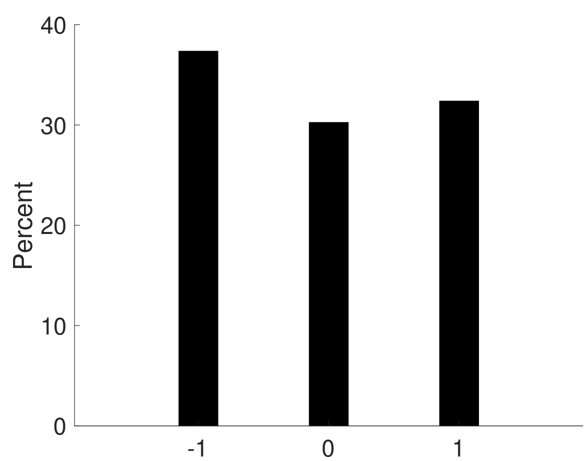

Figure 4 Distribution of Daily Twitter Sentiment Index

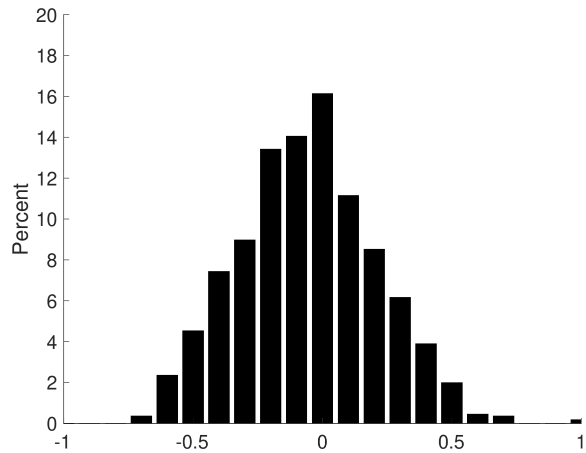

Figure 5 Daily Twitter Sentiment Index and EUR/USD Exchange Rate

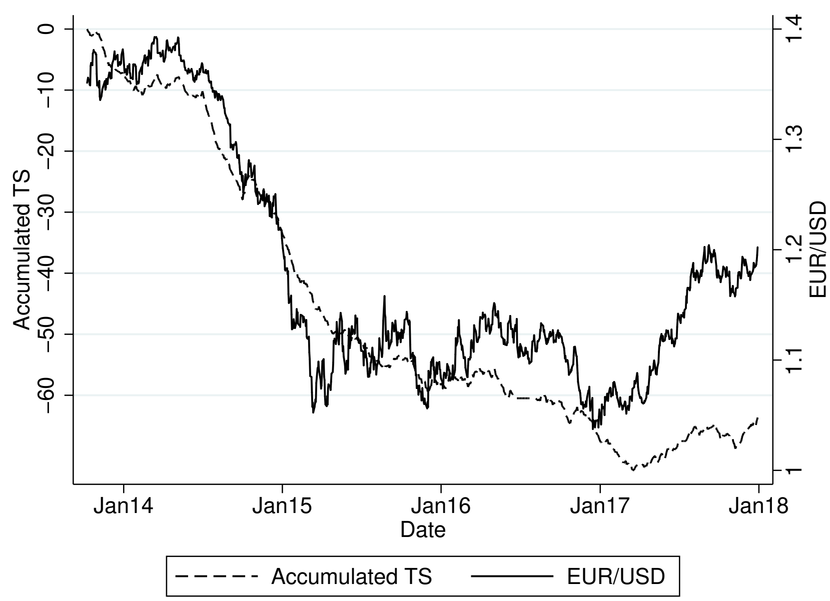

Figure 6 Distribution of Daily Disagreement *

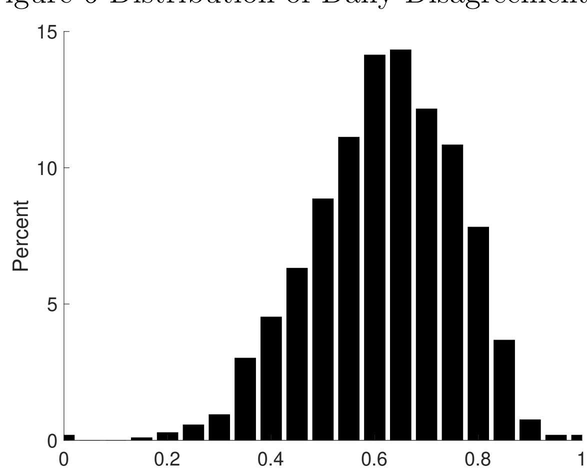

*Disagreement is defined as the daily cross sectional variance of Twitter Sentiment across the individuals.

Figure 7 Impulse Response of Exchange Rate to Fundamental and Noise Shocks*

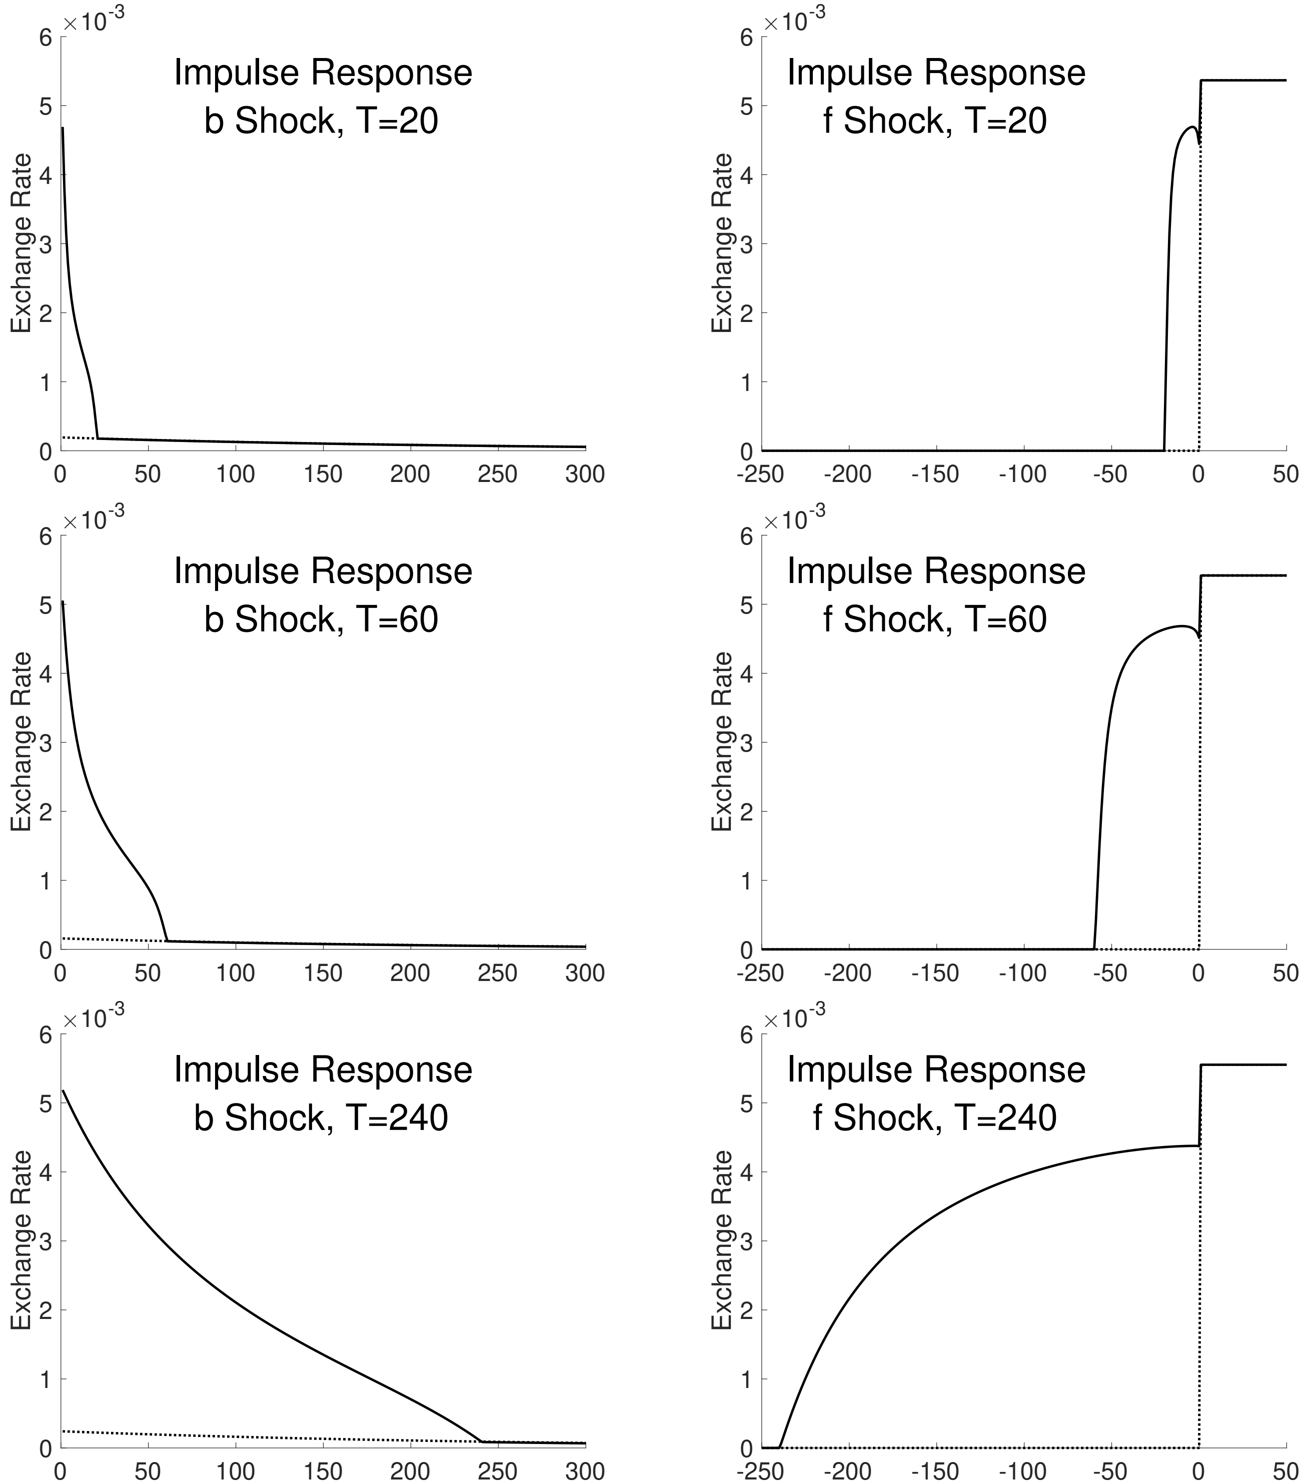

*The left side displays the impulse response to a noise shock with the magnitude of σ_b at $t = 0$ and the right side displays the impulse response to a fundamental shock with the magnitude of σ_f at $t = 0$. The dotted lines show the impulse response when $\sigma_v \to \infty$

Figure 8 Share of Variance of $s_{t+k} - s_t$ Explained by Different Shocks*

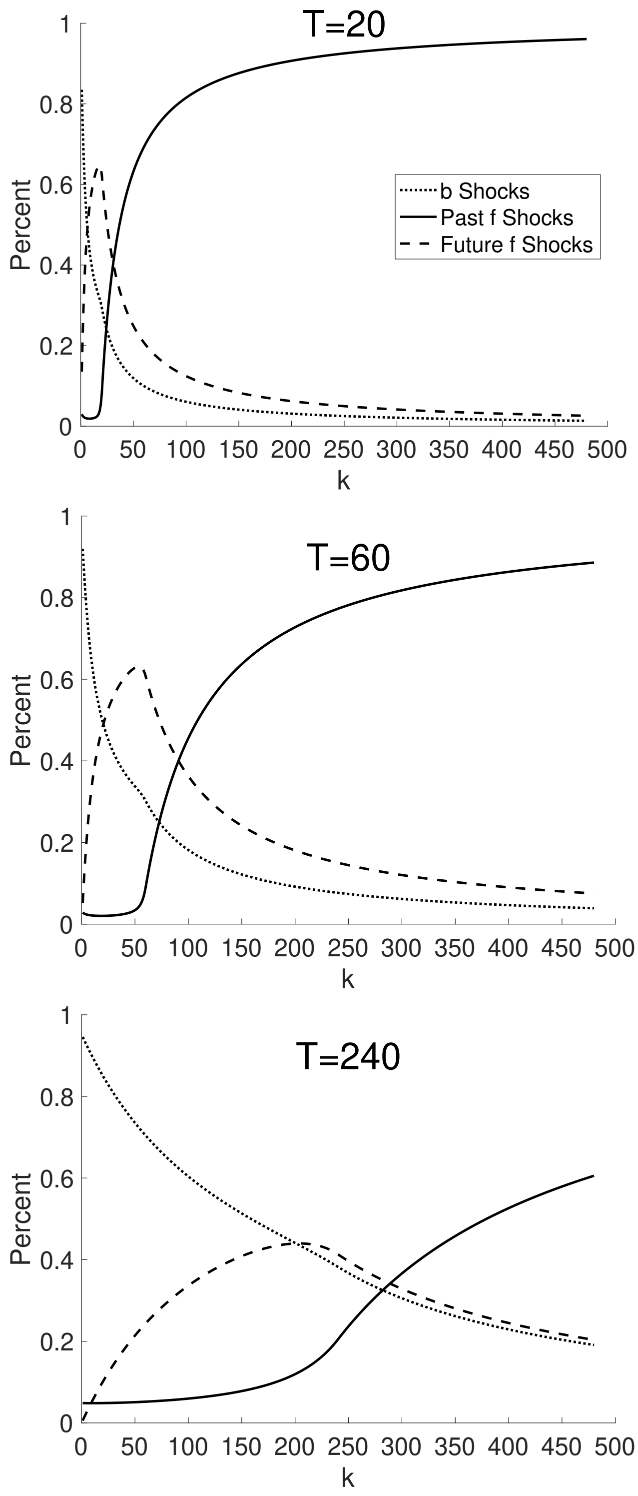

*Solid, dashed and dotted lines show the share of past fundamental shocks, future fundamental shocks and noise shocks respectively.

Table 1 Examples of Positive, Neutral and Negative Tweets

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

Table 2 Directional Moments in the Data *

	$s_{t+20} - s_t$	$s_{t+40} - s_t$	$s_{t+60} - s_t$
Directional Moment	0.0163	0.0214	0.0151
	(0.0039)	(0.0039)	(0.0039)

*Directional moments are computed as the average over all tweets of a variable that is +1 (-1) if the tweet correctly (incorrectly) predicts the direction of the subsequent exchange rate change and 0 if the tweet has a neutral opinion. Standard errors are in parentheses and are computed based on the assumption that individuals are uninformed about the exchange rate and randomly tweet their guesses about the future direction of exchange rate.

Table 3 Data and Model Moments *

	Data	Model $T=20$	t-Value	Model $T=60$	t-Value	Model $T=240$	t-Value
Variance of Twitter Sentiment							
TS_t	0.0696	0.0643	-0.2831	0.0689	-0.0481	0.0686	-0.0462
Disagreement							
Mean Informed	0.6154	0.6116	-0.1733	0.6154	-0.0055	0.6160	0.0320
Predictive Correlations							
$TS_t, s_{t+1} - s_t$	0.0349	0.0298	-0.1663	0.0230	-0.3880	0.0177	-0.5536
$TS_t, s_{t+20} - s_t$	0.0423	0.0492	0.0675	0.0542	0.1109	0.0671	0.2204
$TS_t, s_{t+40} - s_t$	0.1085	0.0300	-0.5763	0.0532	-0.3975	0.0828	-0.1694
$TS_t, s_{t+60} - s_t$	0.1148	0.0210	-0.5870	0.0407	-0.4765	0.0883	-0.1495
Directional Moments							
$TS_t, s_{t+20} - s_t$	0.0163	0.0124	-0.1251	0.0172	0.0284	0.0264	0.2967
$TS_t, s_{t+40} - s_t$	0.0214	0.0091	-0.2940	0.0188	-0.0651	0.0334	0.2935
$TS_t, s_{t+60} - s_t$	0.0151	0.0077	-0.1510	0.0177	0.0583	0.0393	0.4630
Weekly Contemporaneous Correlation							
$TS_w, s_t - s_{t-4}$	0.2874	0.2790	-0.1170	0.2685	-0.3758	0.1615	-2.5551
Exchange Rate Moments							
St. Dev. Δs	0.5341	0.5341	0.0000	0.5341	0.0000	0.5341	0.0000
Auto Corr. Δs	-0.0030	0.0188	1.4122	0.0192	1.4055	0.0194	1.3539
Objective			2.9520		2.6754		9.1594

Model Parameters **

σ_v	0.5055	(0.0538)	1.0472	(0.1058)	3.3000	(0.1936)
σ_b	0.0296	(0.0029)	0.0267	(0.0026)	0.0337	(0.0022)
ρ	0.4597	(0.0523)	0.4657	(0.0678)	0.2528	(0.0470)
ρ_b	0.9959	(0.0001)	0.9954	(0.0002)	0.9960	(0.0003)
σ_f	0.0029		0.0029		0.0041	

* "t-Value" is the difference between the model and data moment, divided by the standard deviation of the corresponding moment. St. Dev Δs is the standard deviation of the daily change in the exchange rate in percentage terms (e.g. 0.5341%=0.005341)

** Standard errors are in parentheses.

Table 4 Summary of Assumptions in the Empirical Papers

	Rossi (2013)	Cheung et. al. (2005)	Cerra and Saxena (2010)
Data Frequency	Monthly	Quarterly	Annual
Observations	1957,m02-2011,m05	1973,q2 - 2000,q4	1960-2004
Estimation Method	recursive rolling half sample	rolling 39 quarters	start with 23 years and recursively increase estimation period
Long-term Forecast Horizon	4 years	5 years	5 years
Long-term Forecast Period	second half sample	1983q1-2000q4	4 non-overlapping five years (starts 1983)

Table 5 Short Horizon MSE Ratios *

	$T = 20$			$T = 60$			$T = 240$		
	Rossi	Cheung	Cerra	Rossi	Cheung	Cerra	Rossi	Cheung	Cerra
MSE^{theory}	0.92	0.31	0.08	0.98	0.91	0.23	0.95	0.95	0.80
$MSE^{theory,\Delta f}$	0.96	0.46	0.13	0.98	0.94	0.34	0.95	0.95	0.84
$MSE^{Meese\ Rogoff}$	0.96	0.49	0.15	0.99	1.00	0.39	0.96	1.00	0.90
$std(MSE^{Meese\ Rogoff})$	(0.02)	(0.09)	(0.07)	(0.02)	(0.07)	(0.15)	(0.02)	(0.07)	(0.18)
Empirical Literature	[0.98,1.08]	[1.09,1.17]	0.96	[0.98,1.08]	[1.09,1.17]	0.96	[0.98,1.08]	[1.09,1.17]	0.96

* MSE^{theory} is the theoretical MSE ratio that measures the variance of exchange rate changes due to shocks unrelated to observed fundamental innovations relative to the variance of exchange changes due to all shocks. $MSE^{theory,\Delta f}$ is the theoretical MSE ratio based on changes in fundamentals over the forecast horizon. $MSE^{Meese\ Rogoff}$ is the average of the MSE ratio across simulations when the methodology of empirical literature is replicated.

Table 6 Long Horizon MSE Ratios *

	$T = 20$			$T = 60$			$T = 240$		
	Rossi	Cheung	Cerra	Rossi	Cheung	Cerra	Rossi	Cheung	Cerra
MSE^{theory}	0.02	0.02	0.02	0.06	0.05	0.05	0.20	0.16	0.16
$MSE^{theory,\Delta f}$	0.03	0.03	0.03	0.09	0.07	0.07	0.28	0.22	0.22
$MSE^{theory,\Delta f,short}$	0.62	0.09	0.03	0.75	0.60	0.10	0.68	0.66	0.46
$MSE^{Meese\ Rogoff}$	0.77	0.20	0.05	0.92	1.19	0.16	0.84	1.29	0.65
$std(MSE^{Meese\ Rogoff})$	(0.27)	(0.16)	(0.05)	(0.31)	(0.81)	(0.15)	(0.28)	(0.83)	(0.47)
Empirical Literature	[1.00,1.08]	[0.9 ,3.22]	0.56	[1.00,1.08]	[0.92,3.22]	0.56	[1.00,1.08]	[0.9 ,3.22]	0.56

* $MSE^{theory,\Delta f,short}$ is the theoretical MSE ratio where the incorrect short-horizon coefficient β_s is applied to the long-horizon change in the fundamental.

Table 7 Theoretical Regression Coefficient of $s_t - s_{t-k}$ on $f_t - f_{t-k}$

	$T = 20$	$T = 60$	$T = 240$
$k = 20$ (1 month)	0.2146	0.1432	0.2148
$k = 60$ (1 quarter)	0.7434	0.2365	0.2215
$k = 240$ (1 year)	0.9368	0.8114	0.3976
$k = 960$ (4 years)	0.9843	0.9530	0.8497
$k = 1200$ (5 years)	0.9874	0.9624	0.8797