Master's Thesis Economics and Business

News Sentiment and Foreign Exchange Markets

Quantifying investors' beliefs about future fundamentals through quantitative content analysis and its influence on the GBP/USD exchange rate

Author
Yannick Brouwer, 433638

Supervisor
Dr. Agnieszka Markiewicz

1st Reader
Prof. Dr. Casper de Vries

Erasmus University Rotterdam
Erasmus School of Economics

August 27, 2017
Rotterdam, the Netherlands

ABSTRACT

This thesis aims at quantifying foreign exchange investors' beliefs about future fundamentals through extracting a news sentiment score from economic news articles using quantitative content analysis. It aims at contributing to a better understanding of the FX market, thereby closing the gap between exchange rate economics and what actually determines, or drives, exchange rates. OLS regressions of the news sentiment measurement on the GBP/USD exchange rate show that news about future fundamentals are incorporated in the exchange rate within the same day of the news' arrival. Incorporating investors' beliefs in existing exchange rate models yields a better understanding of the FX market and possibly more accurate forecasting ability of these models.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

The foreign exchange (FX) market is, unlike the stock market, a fairly opaque market with little regulations and low transparency. This makes explaining — and even more so forecasting — exchange rates a difficult task. Since the end of the Bretton-Woods in the 1970s, and therefore the beginning of a system of floating exchange rates, a vast literature on exchange rate models has emerged. Many existing exchange rate models have a documented failure to outperform a simple random walk model in their forecasting performance. With that, a large body of literature has evolved trying to explain this failure. Most explanations are, like the models themselves, based and focused on macroeconomics. There could be, however, other explanations. Lyons (2000), for example, states that there is a gap between exchange rate economics and the actual determination of exchange rates. This thesis aims at contributing to a better understanding of the FX market, thereby closing the gap between exchange rate economics and what actually determines, or drives, exchange rates.

Most exchange rate models express the exchange rate as the present discounted value of current and expected future macroeconomic fundamentals (Meese & Rogoff, 1983). But, as state Engel and West (2005), these models take expectations of market participants (investors) of future fundamentals too little into account. In their view, exchange rates are, at least in the short run, determined far more by expectations than by current fundamentals. There is a need, then, for a more effective way of incorporating expectations about future fundamentals into exchange rate models.

Some efforts have been made to incorporate expectations about future fundamentals in exchange rate models. For example, Lyons (2000) acknowledges the importance of investors' expectations — or as he calls it, beliefs — about future fundamentals when trying to explain or forecast exchange rates. This is called the 'micro-structure approach to exchange rates'. This approach is based on micro-structure, or behavioural, finance. The micro-structure approach to exchange rates emphasizes that exchange rates are determined in the trading room, by the traders. Those traders carry certain beliefs about the future direction of the exchange rates.

Investors' beliefs can be quantified, or measured, in multiple ways. One way is to look at private information that is contained in the order flow (Fratzscher, Rime, Sarno, & Zinna, 2015). Order flow contains information about the net buying or selling pressure of a certain currency. A net buying pressure would make investors believe that the exchange rate will appreciate. Another way beliefs can be measured is through media sentiment. Tetlock (2007) uses quantitative content analysis to extract a media sentiment measure for the stock market from a recurring *Wall Street Journal* column. He calls this measure the 'media pessimism factor'. The results of Tetlock's research indicate that movements in broad indicators of stock market prices and volume can be predicted from changes in the pessimism factor. For the FX market, there does not yet exist a comprehensive quantitative measure or model for investors' beliefs, even though Hopper (1997) finds that, in the short to medium term, exchange rates are driven by these beliefs. Hopper refers to the way investors interpret the values of fundamentals, based on news about these fundamentals, as 'market sentiment'.

Market sentiment can be defined in different ways. Brown and Cliff (2004), for example, define market sentiment in the stock market as market participants' beliefs about future earnings relative to a norm. It can also be seen as periods of excessive optimism or pessimism in the market. Barberis, Shleifer, and Vishny (1998) also link market sentiment to investors' beliefs. Their model is based on the premise that investors can overreact to a string of good news about a specific company or market, and therefore adjust their beliefs about future earnings too optimistically. This could apply to the FX market, with news about economic fundamentals.

News about economic fundamentals should, according to the Efficient Market Hypothesis (EMH), be compounded in the exchange rate directly. The EMH, in its strong form, states that prices (the exchange rate in this case) cannot be predicted based on new information, both publicly available and private (Fama, 1970). Presence of lasting optimism or pessimism in the market affecting the exchange rate over the short to medium term means that not all information is compounded in the exchange rate directly. The EMH then, does not hold. This would mean that the exchange rate can be predicted based on investors' private information. This private information, or private beliefs, are in this case the lasting optimism or pessimism.

The results of a survey of FX traders in the United States (US) conducted by Cheung and Chinn (2001) indicate that news about macroeconomic variables (fundamentals) is incorporated into exchange rates within minutes or hours after the news' arrival. This corroborates the EMH. To empirically test if the EMH holds for the GBP/USD exchange rate, I construct a measure of news sentiment from economic news articles posted on the *Reuters* news website about the United Kingdom (UK) and the US. This measure allows me to quantify investors' beliefs about future fundamentals and test whether these beliefs are incorporated in the exchange rate directly (the same day). Any effect of the news sentiment measure on the exchange rate lasting more than a day violates the EMH.

I construct the news sentiment measure using quantitative content analysis, also known as text sentiment analysis. As far as I know, incorporating investors' beliefs about future fundamentals by extracting a news sentiment measure from news articles has not been done before. Understanding the relationship between investors' beliefs about future fundamentals and the exchange rate will help close the gap between exchange rate economics and actual exchange rate determination.

The question that is answered in this thesis is: Can news content be used to quantify investors' beliefs about future fundamentals? After answering this question, other questions can be answered. For example: Are investors' beliefs compounded in the exchange rate directly: does the EMH hold? If news content can be used to incorporate investors' beliefs in exchange rate models, can this improve these models? If these short to medium term dynamics can be incorporated in exchange rate models, this will hopefully improve the explanatory capabilities and forecasting abilities of these models and yield a better understanding of the FX market.

The remaining text is structured as follows: Section 2 outlines the existing literature on exchange rate modelling and market sentiment and lays the foundation for the methodology. The methodology on news sentiment extraction and descriptive statistics of the final news sentiment score are explored in Section 3. Section 4 shows the empirical results from the regressions of the exchange rates on the news sentiment score. In Section 5, empirical results are discussed in relation to the theory, conclusions are drawn, and suggestions for future research are made.

2. LITERATURE REVIEW

In this section, I lay out the foundation of exchange rate modelling, starting at the earliest standard present value models. I also explore their documented failure to outperform simple random walk models, and possible explanation for this failure. The rest of the section introduces microstructure finance and different definitions and measures of investors' beliefs in the stock market and the FX market.

2.1 Standard Present Value Models

A considerable amount of research has gone into finding out what determines, or drives, exchange rates. Many models have been developed for this purpose. The earliest exchange rate models assume that money is the only asset whose supply and demand plays a role in the determination of exchange rates (Leventakis, 1987). There are two main principles that separate three of the most famous models: flexible prices and sticky prices. For example, the flexible-price monetary (Frenkel-Bilson) model assumes purchasing power parity (PPP), implying a constant real exchange rate. The sticky price monetary (Dornbusch-Frankel) model and the sticky-price asset (Hooper-Morton) model assume that prices adjust slowly to the goods market equilibrium.

These models are the so-called standard present-value, or 'asset', models. This is what stock market return forecasting models are also based on (French, Schwert, & Stambaugh, 1987). The models express the exchange rate as the present discounted value of current and expected future economic fundamentals (Engel, Mark, & West, 2006). A mathematical representation of these models would look like:

$$s_t = a_0 + a_1(m_t - m_t^*) + a_2(y_t - y_t^*) + a_3(r_t - r_t^*) + a_4(\pi_t^e - \pi_t^{e*}) + a_5 TB_t + a_6 TB_t^* + \epsilon_t$$

where s is the exchange rate, m the money supply, y the real income (all in logarithms), r the short-term interest rate, π^e the long-term inflation rate, TB the trade balance, and ϵ an error term. The subscript t represent the time period. The asterisk indicates the value of the variable in the foreign country. The difference in the models lies in the value of the coefficients ($= 0$ or $\neq 0$). The GBP/USD exchange rate here (and in the rest of the text) is defined as the amount of USD per GBP. Consequently, a GBP/USD rate of 1.611 would mean that one GBP yields 1.611 USD. In this situation, the UK is the home country and the US is the foreign country (indicated by the asterisk).

As an example, look at the effect of a change in the interest rate differential between the UK and the US on the GBP/USD exchange rate in the above models. A rise in the short-term interest rate in the UK, relative to the US, would increase the interest rate differential between the two countries. According to the uncovered interest parity (UIP), which all three models assume (Frankel & Rose, 1995; Mark, 1995), the GBP would depreciate in value, therefore lowering the value of the GBP/USD exchange rate.[1] The UIP is a no arbitrage condition stating that the interest rate differential between two countries should be equal to the expected change in exchange rate between the two countries (Isard, 2006).

There are several ways to assess the performance of exchange rate models. Meese and Rogoff (1983) do this by comparing the out-of-sample forecasting accuracy of the Frenkel-Bilson, Dornbusch-Frankel, and Hooper-Morton model. Their results show that, for different specifications of these models, the forecasting performance does not outperform a (drift-less) random walk. The random walk model of exchange rates simply states that tomorrow's exchange rate equals the exchange rate today. Their results also suggest that changes in the exchange rate are largely unpredictable, and that the spot exchange rate approximately follows a random walk. Cheung, Chinn, and Garcia Pascual (2005) repeat this exercise for the newer models, and their conclusions are a little more optimistic than those of Meese and Rogoff (1983). Still, not one model outperforms a random walk for multiple currencies and different

[1] The coefficient of a_3 is negative.

time periods. This means that there still is no comprehensive exchange rate model that can be applied to different currencies and broad time periods.

Others have suggested that exchange rate models could possess more explanatory value than one would draw from the above results. Engel, Mark, and West (2007) state that exchange rates are "not as bad as you think". They build on their earlier theorem — which is introduced in Engel and West (2005) — stating that the models actually imply that the exchange rate should nearly follow a random walk and are not influenced by fundamentals in the short run. This means that it is not surprising that the models do not render better forecasts than a random walk. The main insight of the Engel and Mark theorem is that current fundamentals play a small role in determining the exchange rate. Expectations about future fundamentals play a far bigger role. In the short run, much more weight is put on these expectations when determining the exchange rate.

2.2 Microstructure Finance

Instead of looking for an explanation within macroeconomics, others have suggested taking from the field of microstructure finance (or behavioural finance). For example, Lyons (2000) finds that there is a gap between exchange rate economics and microstructure finance, and that closing this gap could help improve the quality of exchange rate models. Lyons realised that this gap exists after spending a day shadowing an actual currency trader. The microstructure approach is sometimes referred to as 'trading room perspective', as the field claims that exchange rates are determined in the trading room, by the traders. Therefore, understanding the psychology of these traders is key in understanding exchange rates.

The hallmark of the micro-structure approach is order flow. Order flow is defined as the net of buyer-initiated and seller-initiated orders (Evans & Lyons, 2002). It conveys information about the net buying pressure of exchange rates. For example, a large, positive order flow would signal buying pressure and, as with any price, then signals an expected rise in the exchange rate. A direct application of order flow is shown by Fratzscher et al. (2015). They use order flow as a proxy for unobservable exchange rate determinants and design a model that incorporates traders' expectations about fundamentals. Their model shows increased out-of-sample forecasting capabilities over a random walk. Another application is found in Rime, Sarno, and Sojli (2010), where the authors find that order flow is related to current and expected future fundamentals, and that order flow is a powerful predictor of daily movements in exchange rates. This does not mean that order flow is a main determinant of exchange rates, but that it can be used to proxy investors' expectations or beliefs about future fundamentals. It also shows that it is important to incorporate these expectations into exchange rate models.

The problem with order flow data is that it is very hard and costly to get. This means that there is need for a different measure, or proxy, of investors' beliefs. The idea of incorporating investor's beliefs, or expectations, in price modelling is more widely applied in the stock market. One way this has been done is through investor, or market sentiment.[2] Investor sentiment can be defined in multiple ways, and one of the definitions is the way investors form beliefs about future earnings of stocks (Barberis et al., 1998). Barberis et al. construct a parsimonious model of investor sentiment, based on overreaction to a string of good or bad news about stocks and underreaction to news like earnings announcements. Their results show that security prices can overreact to consistent patterns of good or bad news over periods of up to five years.

Another definition of sentiment is that it represents market participants' beliefs relative to a norm

[2] Investor sentiment and market sentiment are used interchangeably throughout this text.

(Brown & Cliff, 2004).[3] This is called speculator's bias, which can be seen as excessive optimism or pessimism. These are periods of lasting bullishness or bearishness. The sentiment measures that are used here are quantitative measures like the advance-decline ratio, short interest, and closed-end fund discounts. These measures are used to empirically test the relationship between sentiment and future returns. Their results show that sentiment and contemporaneous returns are positively related, returns predict future sentiment, but sentiment does not predict future returns.

A different way of measuring sentiment is through quantifying media content. Tetlock (2007) uses quantitative content analysis to construct a measure of media pessimism based on a daily *Wall Street Journal* column about the stock market. His results show that changes in the pessimism (news sentiment) factor predict movements in broad indicators of stock market activity like prices and volume. The above two researches show different ways to measure sentiment, and different outcomes of sentiment being useful for stock market return forecasting. Sentiment is found to be correlated with contemporaneous and future returns, but results are mixed regarding the direction of causality. What can be concluded though, is that sentiment does play an important role in the stock market.

The influence, and potential forecasting capabilities, of sentiment in the FX market has been less widely researched. Hopper (1997) was one of the first to attribute exchange rate movements to sentiment in the market. Hopper finds that, in the short run, exchange rates seem to be influenced by market sentiment, rather than by fundamentals. It is therefore necessary to have an understanding of the psychology of the FX market (and its traders) in order to understand what drives exchange rates. A few attempts have been carried out to construct a measure of sentiment in the FX market. The main focus has been on Twitter and financial message boards like *Yahoo Finance*, *Google Finance* or *StockTwits*. These outlets are used to extract investors' beliefs about the FX market. Financial message boards are preferred over Twitter as they contain solely financial information.

Papaioannou, Russo, Papaioannou, and Siettos (2013), for example, filter Twitter for messages about the EUR/USD exchange rate. They find that each Tweet containing EUR/USD also contains information about the type of order that the user made. Most Tweets contain limit orders with prices included. Since the orders have already been placed at a broker, they reflect the trader's belief about the future direction of the exchange rate. Plakandaras, Papadimitriou, Gogas, and Diamantaras (2015) use econometric and machine learning algorithms to extract a measure of sentiment from financial message boards. Their algorithms classify messages into two categories: expectation for a decline ('Bearish') and expectation for an increase ('Bullish'). Using these classifications, they try to forecast the direction of four currency pairs.

2.3 Efficient Market Hypothesis

Both researches laid out above challenge the EMH. The EMH exists in three forms: the weak, the semi-strong and the strong. The weak form states that future prices cannot be predicted based on historic prices. The semi-strong form states that prices cannot be predicted based on new, publicly available information such as macroeconomic shocks. The strong form expands the semi-strong to all sorts of information (e.g. people's private beliefs and non-public information) (Fama, 1970). Papaioannou et al. (2013) challenge the EMH by concluding that, in the short run, private information distilled from social microblogs can be used to outperform the random walk hypothesis and other exchange rate models for at-level forecasting. Plakandaras et al. (2015) reject the strong form of the EMH in concluding that

[3] The norm is seen as the average returns of the stock (or stock market).

exchange rates are driven by market expectations formed by private information.

If the EMH holds, macro news is compounded in the exchange rate instantaneously. According to Cheung and Chinn (2001), who survey US FX brokers, 70% of currency traders say that macro news is compounded in the exchange rate within one minute. This would confirm the EMH: news about (future) macro fundamentals changes traders' expectations about these fundamentals. If, after reading the news, traders expect the UK short term interest rate to rise in the future, they will start buying GBP now, in expectation of a GBP appreciation. This in turn results in a direct appreciation of the GBP.

Evans and Lyons (2005) empirically test the above made conclusion through the effect of macro news on traders' subsequent trading behaviour. They conclude that macro news is absorbed in the exchange rate at least on the same day but trading behaviour can stay affected days after the news arrival. Evans and Lyons (2008) repeat this exercise and conclude that the arrival of macro news can account for 30 percent of daily price variance. Both researches, though, say exchange rates are also affected indirectly through order flow. News tends to affect the order flow, which subsequently affects the exchange rate through the (non-public) information contained in the order flow.

The direct effect of news content on the exchange rate has not yet been empirically tested. Through quantitative content analysis, I construct a measure for news sentiment, hereby capturing how investors form beliefs about future fundamentals. Herewith I can test if the different forms of the EMH hold in the FX market, and if all information is compounded in the exchange rate directly.

3. METHODOLOGY AND DATA

In the previous section, different ways to define and measure investors' beliefs are introduced. The way I have chosen to do this, is through quantifying media content. I measure the sentiment conveyed by an economic news article to represent the beliefs of investors about economic fundamentals. For example, a news article about the negative effects of a growing trade imbalance would carry a negative sentiment, suggesting investors updating their beliefs and consequently a future GBP depreciation. If all of this new information is directly compounded in the exchange rate, i.e. investors have updated their beliefs and acted upon this, the EMH holds. If periods of media pessimism or optimism affect the exchange rate, the EMH is violated. This could mean that a change in the exchange rate is not necessarily caused by changing fundamentals but due to certain media coverage which leads investors to be overly optimistic or pessimistic for a sustained period of time.

3.1 Quantitative Content Analysis

I extract a news sentiment measure from media content through quantitative content analysis. Merriam-Webster's Dictionary (2017), defines content analysis as "analysis of the manifest and latent content of a body of communicated material through a classification, tabulation, and evaluation of its key symbols and themes in order to ascertain its meaning and probable effect." In this case, the body of communicated material is made up of news articles. The meaning (and probable effect) of the content of these articles will carry the sentiment and will be quantified. Krippendorff (2012) states that content analysis is unique in its combinations of features. It allows the researcher to analyse unstructured data and extract the meanings, symbolic qualities, and expressive contents they possess.

Creating a large enough sample of news sentiment scores means analysing and quantifying a large amount of texts. This is called 'text mining'. Mining is a metaphor for the fact that the needed information is buried in a big stack of (irrelevant) information. The pile of news articles contains a large amount of relevant information. What needs to be 'mined' is the sentiment that it carries: is it good or bad news? Will expectations about future fundamentals be adjusted upwards or downwards? (Moniz & de Jong, 2014) use text mining techniques to analyse the financial media and its effect on investors' behaviour. They state that, due to limitations in time cognitive processing abilities, investors use the financial media to assess the quality of news when making their investing decisions. They find a significant relationship between financial media content and investor behaviour, highlighting the possible use of text mining techniques in investment analysis.

3.2 Text Mining

This research is focussed on the GBP/USD exchange rate. To, in first instance, restrict the text mining exercise to economic news, articles that appear under the *Economy* tab from *Reuters*[1] are used. For the UK, 4623 articles are extracted over a six year period (April 2011 through mid-June 2017). The *Economy* tab is only there for the UK, so I filter news about the US for the same economic subjects as the UK. This results in the extraction of 912 articles between December 2013 and mid-June 2017 for the US.

Apart from news about economic fundamentals affecting investors' beliefs, news about other subjects might also have an impact on the exchange rate. To test whether news sentiment in other news affects the exchange rate, I also extract *Headline* news articles for both the UK and the US. This results in the extraction of 13,932 and 6505 articles for the UK and US respectively, over almost 4 years (December 2013 through mid-June 2017).

I use the R programming language to extract URLs to all news articles within the above specified time periods. An example of the code that is used for the UK economic news articles is presented in Appendix B.1. Subsequently, the HTML code of all URLs is scraped. R provides an HTML-scraping tool, which is part of the *BoilerpipeR* package. All uninformative HTML code is removed from the scraped text. The code for this step is presented in Appendix B.2. The below text is an example of one scraped UK economic news article:

> x
> 1 Economy — Mon Apr 4, 2011 — 9:46am BST
> LONDON Britain's financial services sector saw growth in business volumes during the first quarter of 2011, the Confederation of British Industry said in a survey conducted with Price-WaterhouseCoopers on Monday.
> The survey added that profitability was increasing in the UK banking sector, which is dominated by the Big Four of Lloyds (LLOY.L), Royal Bank of Scotland (RBS.L), Barclays (BARC.L) and HSBC (HSBA.L), due to lower losses on loans.
> A third quarter of strong volume growth shows the financial services recovery is building strength, said CBI chief economic adviser Ian McCafferty.
> While business with private individuals has again shown the fastest growth, business with companies also shows some signs of improvement. Firms' profitability has improved due to higher incomes and another big drop in the value of non-performing loans, he added.
> Financial services represent some 8 percent of Britain's gross domestic product but the sector

[1] http://uk.reuters.com/business/economy

has been hit by the credit crunch and debt crises in Ireland and other European countries. Britain's Independent Commission on Banking is also considering reforms which could see retail banks split off from investment banks, although analysts expect it to recommend ring-fencing retail and trading operations rather than a full break-up between the two.
The CBI/PWC survey was conducted between February 24 and March 10, polling 94 parties from banks, building societies, insurers, trading houses and fund managers.

All texts are combined into four corpora.[2] There is one corpus per group of news articles: The UK economic, US economic, UK headline, and US headlines news. Through the corpora, all news articles in the specific groups can be manipulated and analysed at once.

Before analysing, the corpora need to be cleaned up. Cleaning the corpora means removing unnecessary information like the time stamp, converting all letters to lower case, removing stop words,[3] punctuation, numbers, and empty white spaces. After the corpora are cleaned, all words in the corpora are stemmed. Stemming reduces words to their root morphological form, leaving just the stem. Stemming words makes sure that, among other things, different tenses are not treated differently (Das, 2012). So the words 'stagnate' and 'stagnated' will become 'stagnat' and are treated equally. The cleaned and stemmed example text looks like this:

london britain financi servic sector saw growth busi volum first quarter confeder british industri said survey conduct pricewaterhousecoop monday
survey ad profit increas uk bank sector domin big four lloyd lloyl royal bank scotland rbsl barclay barcl hsbc hsbal due lower loss loan
third quarter strong volum growth show financi servic recoveri build strength said cbi chief econom advis ian mccafferti
busi privat individu shown fastest growth busi compani also show sign improv firm profit improv due higher incom anoth big drop valu nonperform loan ad
financi servic repres percent britain gross domest product sector hit credit crunch debt crise ireland european countri
britain independ commiss bank also consid reform see retail bank split invest bank although analyst expect recommend ringfenc retail trade oper rather full break two
cbipwc survey conduct februari march poll parti bank build societi insur trade hous fund manag

3.3 News Sentiment analysis

The news sentiment analysis is run on the cleaned and stemmed texts. Different sentiment analysis techniques can be categorised into two groups: the dictionary-based, or lexicon-based, approach and the statistical, or machine learning, approach (Kearney & Liu, 2014). The dictionary-based approach is based on the relative amount of negative and positive words in the text. A pre-classified library (or lexicon) is used to mark a word as negative or positive. The relative amount of positive and negative words will make up the news sentiment score of the text. Machine learning techniques assign texts into categories using statistical classifiers.

[2] In text mining, a corpus is a collection of a large number of texts.
[3] Stop words are non-contextual words like 'the', 'a' and 'are'. Removing them improves the quality of the text analysis (Das, 2014).

The majority of machine learning techniques categorizes texts based on a set of pre-classified texts, called the training data. This means that, in order to get training data, a large part of the total sample has to be classified by hand. One classifier that is widely used is the Bayesian classifier (Das, 2012). This classifier uses Bayes' theorem to establish 'prior' probabilities. These probabilities are then used to classify the remainder of the sample. Another classifying algorithm that has gained popularity is called Support Vector Machine (SVM) (Antweiler & Frank, 2004). The SVM is a discriminant classifier that produces an optimal hyperplane that categorizes other texts based on training data.

Machine learning techniques require great technological knowledge, and pre-classifying a training data set is very time consuming. Furthermore, creating the training data requires personal judgement about the texts and can results in biased training data, potentially biassing the final news sentiment scores. I therefore employ the dictionary-based approach. It is seen as a valid method of sentiment analysis and requires less technical knowledge than the machine learning approach (Kearney & Liu, 2014; Taboada, Brooke, Tofiloski, Voll, & Stede, 2011; Tetlock, 2007).

The most used library in natural language processing is the Harvard General Inquirer[4] (GI) dictionary (Kearney & Liu, 2014). Loughran and McDonald (2011), though, find a high degree of misclassification in the GI dictionary when applied to financial texts. They create a new dictionary: the Loughran and McDonald Sentiment Word Lists[5] (LM), which is more suited for financial texts. Others also find that the GI dictionary yields misclassification for financial and economic texts (Das, 2012; Huang, Zang, & Zheng, 2012).

To make the dictionary even more suitable for the analysis of economic news texts, I read 100+ articles in order to find words that are used in relation to economics and have clear economic meaning. Examples of negative words that I found are: 'darkened', 'squeeze', 'fragile', and 'contraction'. And positive words: 'recover', 'boom', 'accelerate', and 'thrive'. All words in the dictionaries are stemmed to match the stemmed words in the corpora. As a means of robustness check, I run the text analysis on all three versions of the dictionary: the GI, the LM and the amended LM.

The dictionary-based method is also called the 'bag-of-words' approach, since all texts are represented as lists (columns) of words. All the words in the 'bag' can be matched to the negative and positive word lists from the library. The net of the number of positive and negative words is divided by the total number of positive and negative words in the text (Huang et al., 2012). The score can then also be seen as the difference in the percentage of positive and negative words in the article. More formally, the sentiment score then looks like this:

$$score = \frac{pos - neg}{pos + neg} = \%pos - \%neg$$

where pos is the number of positive matches and neg the number of negative matches. The R code for the cleaning of the UK economic news corpus and creating of the UK economic news sentiment score is presented in Appendix B.3. On some days, multiple articles are posted. To construct a daily sentiment measure, all scores on one day are averaged. The whole process described above is summarised in Figure 3.1.

When using the amended LM library, the text displayed above has seven positive matches: 'profit' (twice), 'strong', 'recovery', 'strength', 'improvement', and six negative matches: 'loss', 'drop', 'non-performing', 'hit'. 'crunch', 'crises'. This yields a sentiment score of 0.0769. This means that this article

[4] http://www.wjh.harvard.edu/~inquirer
[5] http://sraf.nd.edu/textual-analysis/resources

Fig. 3.1: Sentiment Score Creation

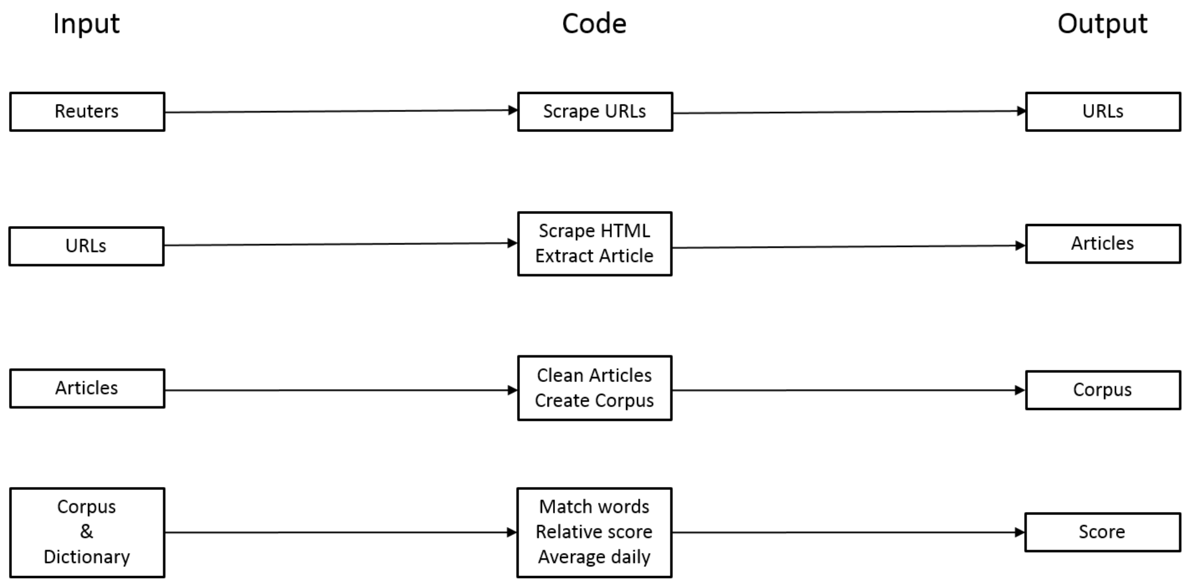

has a small, positive sentiment. The sentiment score does depend on the library that is used, as visible from Table 3.1. The GI dictionary yields a more positive score than the LM and the amended LM

Tab. 3.1: Sentiment Scores Dependent on Library

Dictionary	posmatch	negmatch	score	overall average
GI	16	2	0.7778	0.3329
LM	6	5	0.0909	-0.3747
LM*	7	6	0.0769	-0.2183

Notes: LM* is the amended LM dictionary. The overall average is the average of the sentiment scores for UK economic news.

dictionary. The GI dictionary counts positive words like 'confederation', 'economic', and 'independent'. As negative words, the GI dictionary only counts 'dominated' and 'loss'. The miscounting of positive and negative words results in an upwards biased (overoptimistic) sentiment score. This bias is also seen in the overall average sentiment scores. This corroborates Loughran and McDonald's results that the GI dictionary misclassifies financial and economic texts. The amended LM is, overall, slightly less pessimistic than the LM dictionary. I choose only to use the amened LM dictionary for the text analysis.

3.4 News Sentiment Score

The relative amount of positive and negative words in an article should reflect the sort of information contained in the article. The example text shown above, for example, expresses the growing profitability of the financial services sector in the UK. This positive signal about the financial services sector would lead to investors adjusting their beliefs about the strength of the financial system upwards. In theory, they would believe that the GBP will appreciate in the future, therefore leading to a rise in GBP demand and an appreciation of the GBP today. The opposite would be true with an article containing more negative than positive words like, for example, an article about the growing government deficit in the UK with a news sentiment score of -0.59.

The final news sentiment score that is used for regression analysis is the average of the news sentiment scores of all articles posted one a specific day. A string of days with a negative or positive news sentiment score could result in periods of optimism or pessimism (bullishness or bearishness) in the market. The month after the Brexit vote in the UK on the 23rd of June 2016 has an average (negative) economic news sentiment score of -0.45. This would suggest investors adjusting their beliefs about the economy downwards, even though it was not clear how fundamentals would change in the future, or that fundamentals changed at all. This period is therefore a period with an (excessive) pessimistic sentiment.

Next to the analysis of the news sentiment score, I look if what Tetlock (2007) and Brown and Cliff (2004) identify as periods of (excessive) optimism or pessimism are present. To identify this, I use the Hodrick-Prescott (HP) filter to decompose the news sentiment score time series into a cyclical and a trend component. The filtered trend component is more sensitive to long-term fluctuations than to short-term fluctuations. The cyclical component represent regular or periodic fluctuations around the trend. One cycle could represent a period of optimism or pessimism.

The equation that is used in this filter looks as follows (Hodrick & Prescott, 1997):

$$\min_{\tau}\left(\sum_{t=1}^{T}(y_t-\tau_t)^2+\lambda\sum_{t=2}^{T-1}[(\tau_{t+1}-\tau_t)-(\tau_t-\tau_{t-1})]^2\right)$$

where y_t for $t = 1, 2, ..., T$ is the logarithm of the news sentiment score and τ is the trend component. The news sentiment score (y_t) is composed of the trend τ and cyclical component c so that: $y_t = \tau_t + c_t + \epsilon_t$. With a λ that is dependent on the frequency of observations,[6] there exists a trend component that solves the above equation.

The downside of using the HP filter is the fact that, as can be seen from above equation, it is not solely backward looking. This makes the filtered cyclical and trend components not useful for forecasting. In this research this does not pose a problem, as the cycle and trend components are only used to identify periods of media optimism or pessimism that possibly affect the exchange rate. The news sentiment score itself is the main focus of this research.

A graphic representation of the final news sentiment score, filtered trend and cycle for the UK economic and headline news over the entire sample period is visible in Figure 3.2. The blue lines in Figure 3.2a and 3.2c represent the news sentiment score itself, the red lines represent the trend component. Figure 3.2b and 3.2d show the news sentiment score's cyclical components. All observations are at daily frequency.

Table 3.2 shows descriptive statistics of all news sentiment scores. What stands out, is that all scores have a negative mean. When it is assumed that the amended LM dictionary does not result in misclassification, the negative means would suggest that the overall sentiment in news is negative. Also, headline news has a more negative sentiment than economic news. Looking at the standard deviation, economic news seems to be more volatile than headline news. This can also be seen from Figure 3.2a and 3.2c. The difference in observations is firstly the time sample, and secondly the amount of news articles per day, with some days having no news articles at all.

Tab. 3.2: Descriptive Statistics of News Sentiment Score

		Mean	Max.	Min.	Std. Dev.	Skweness	Kurtosis	Obs.
Economic News	UK	-0.2183	1	-1	0.2495	0.6539	4.2612	1437
	US	-0.1830	1	-1	0.2739	0.3615	3.5006	518
Headline News	UK	-0.3357	0.7778	-1	0.1633	0.5283	5.8834	925
	US	-0.3874	0.5	-1	0.1950	0.4201	4.1086	883

[6] Daily frequency: $\lambda = 100 * 365^2 = 13,322,500$

Fig. 3.2: UK News Sentiment Score, Trend and Cycle

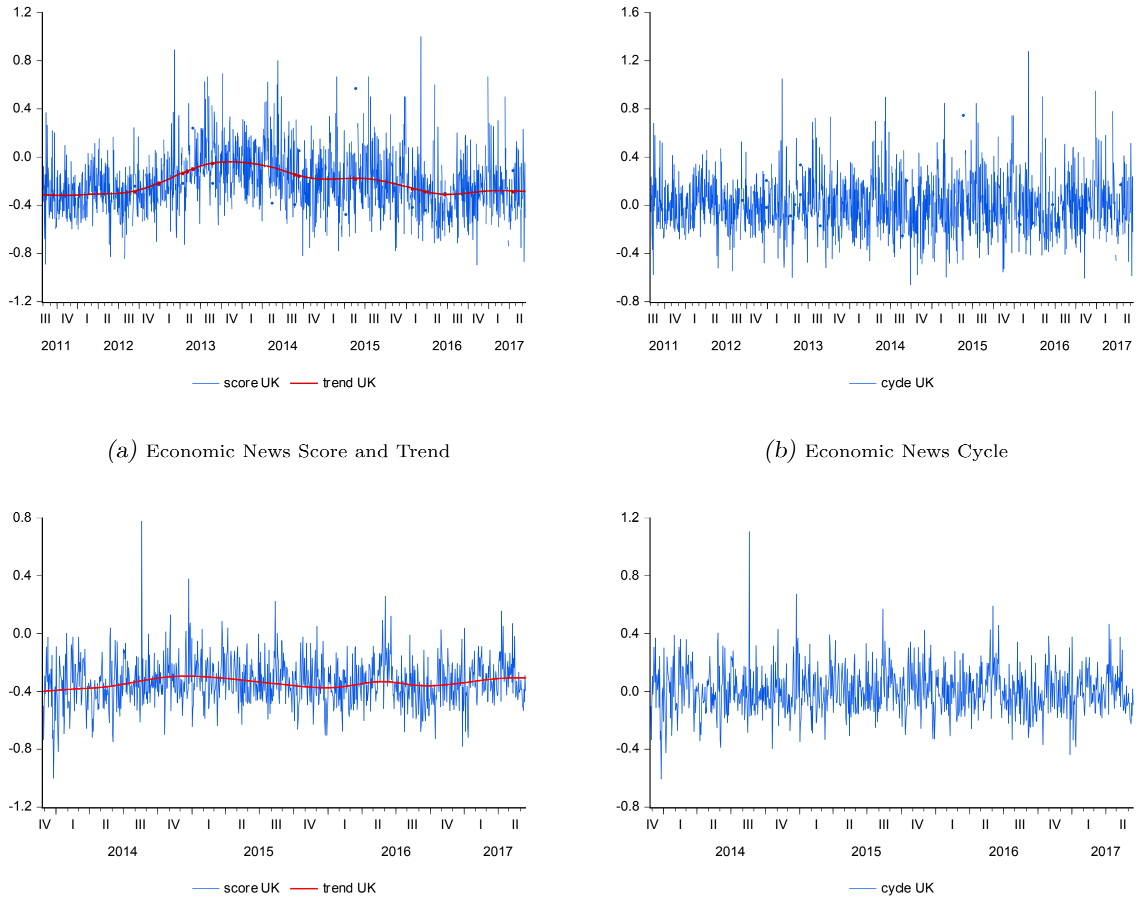

(a) Economic News Score and Trend

(b) Economic News Cycle

(c) Headline News Score and Trend

(d) Headline News Cycle

For illustrative purposes, Figure 3.3 shows the filtered trend components of the news sentiment scores as differentials between the UK and the US, compared to the GBP/USD exchange rate.[7] The blue line in Figure 3.3 represents the GBP/USD exchange rate. The red line represents the differential news sentiment score between the UK and the US. All observations are at daily frequency.

The differential in economic news sentiment in Figure 3.3a shows, like the GBP/USD exchange rate, a clear downward trend. This would suggest that the exchange rate and the economic sentiment score move together over the long run. In terms of investors' beliefs, this could mean that the premise of economic news affecting investors' beliefs, and therefore the exchange rate, is true. This inference does depend on the direction of causality between the two variables and cannot be made without caution. A matching trend between the GBP/USD exchange rate and the differential in headline news is less clear. In Figure 3.3b, one can see that the two lines do not move in similar fashion, suggesting that the exchange rate is more related to information contained in economic news than in headline news, which was deducted from theory and hypothesized in the beginning of this text.

Both Figure 3.3a and 3.3b show interesting developments of the sentiment score and the exchange rate around the two biggest economic and political events of 2016. The Brexit vote in 2016Q2 and Donald Trump's presidential election win in 2016Q4 have no visible effect on the economic news sentiment trend

[7] Data on the end-of-day Bloomberg Generic Composite (BGN) exchange rates are gathered from Bloomberg (2017a, 2017d).

Fig. 3.3: Difference in Trends Between UK and US

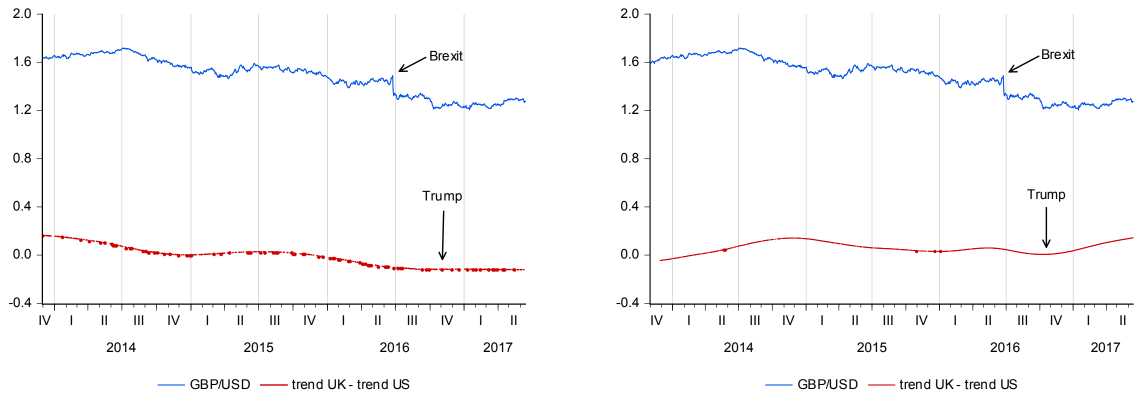

(a) Economic news

(b) Headline news

differential. The headline news sentiment trend, however, does move in accordance to expectations. After the Brexit vote, there is a dip in the sentiment trend differential. This dip is soon turned around due to the large negative media coverage of the Trump election win. Both developments, though, do not have an impact on the exchange rate. To confirm these statements deducted from the graphs, the impact of the sentiment score, cycle, and trend need to be tested empirically.

3.5 Regressions

To measure the short-run relationship between the news sentiment score and the exchange rate, and simultaneously if the EMH holds, I run simple OLS regressions of the first-differenced exchange rate on the first-differenced news sentiment score. I run regressions for the news sentiment measure for economic news and headline separately. To test if there are periods of news pessimism or optimism affecting the exchange rate present, I regress the first-differenced exchange rate on the first-differenced sentiment cycles. These cycles represent short to medium term fluctuations around the trend. For long-term relationships, I regress the first-differenced exchange rate on the filtered trend.

I add the short term interest differential between the UK and the US to the regression equations to control for the UIP, like the present value exchange rate models. The short term interest rate is measured by the GBP and USD overnight LIBOR rate.[8] Other variables used in the asset models like inflation or money supply are not available in daily frequency and can therefore need be used to control for here.

4. EMPIRICAL RESULTS

4.1 News Sentiment Regressions

Table 4.1 lists the results of the OLS regressions of the first-differenced GBP/USD exchange rate on the news sentiment score, cycle, and trend differentials with the UK as the home country. Panel (a) in Table 4.1 shows the values for economic news, panel (b) shows the values for headline news. Both the economic news and the headline news sample run from December 2013 through mid-June 2017. The headline news

[8] Data on GBP and USD LIBOR overnight rates are gathered from (Bloomberg, 2017b, 2017c).

sample has more observations as there was more news available, with less dates of no news. Gaps in the data result in a lower amount of observations with every lag added.

The validity of all regressions is assessed through the Durbin-Watson (DW) statistic and other residual diagnostics. The DW statistics show no sign of autocorrelation in the residuals in any of the regressions as it falls in the range between 1.5 and 2.5. Any DW statistic lower or higher than these values is a sign of negative or positive autocorrelation respectively (Durbin & Watson, 1951). The regression of the lagged economic news cycle has a DW statistic of 1.5087. Even though this still falls within the accepted bounds, it would suggest slight negative autocorrelation. Other residual diagnostics of the non-lagged regressions in Table 4.1 are presented in Table A.1 in Appendix A. The residuals of all regressions have a zero mean. There is no correlation between the residuals and the explanatory variable visible in the regressions.

The non-lagged regressions for the economic sentiment score and trend show signs of heteroskedasticity. The Breusch-Pagan-Godfrey statistic rejects the hypothesis of homoskedasticity at the 5% and 1% level for the score and trend respectively. This means that the regressions' standard errors are biased, which can lead to biased test statistics. I therefore re-run these regressions with robust standard errors. These regressions are represented in Table 4.1 by the ‡ superscript. Using robust standard errors means that there is no DW statistic.

Based on multiple researches explored in the literature review (Cheung & Chinn, 2001; Evans & Lyons, 2005, 2008) and EMH theory, one would expect that news about fundamentals is incorporated in the exchange rate directly (within minutes, hours or at least the same day). To check this premise, I run the regressions with one and two lags of the independent variable to see if the news sentiment measure has an effect on the exchange rate one and two days after its arrival. I also run lagged regressions for the sentiment cycles. The trended news sentiment measure is not lagged. Ideally, the optimal lag length would be assessed with the Akaike Information Criterion (AIC) and the Bayesian, or Schwartz, Information Criterion (BIC). This approach is not valid here, however, as the observations in the sample differ too much due to gaps on the data. The values of the AIC and the BIC depend on the sample size, so regressions with different sample sizes cannot be compared.

The first row of panel (a) in Table 4.1 shows that the economic news sentiment score has a positive and significant effect on the exchange rate. An increase in the economic news sentiment score of 1 means a rise in the GBP/USD exchange rate (an appreciation of the GBP) of 0.0022. Adding one lag to the regression lowers the significance of the contemporaneous value. The effect on the exchange rate becomes slightly larger, namely 0.0026. The first lag itself is not significant, indicating that the lagged value of the sentiment score does not have any effect on the exchange rate. The economic news sentiment score therefore does not affect the exchange rate beyond the day of arrival. This means that, in line with the EMH, all information contained in the news sentiment score is absorbed in the exchange rate on the same day of the news' arrival.

The cyclical component shows the same pattern as the sentiment score. The contemporaneous value has a positive, significant effect on the exchange rate. A rise in the news sentiment cycle of 1, results in an increase of the GBP/USD exchange rate of 0.0021. The significance of the cyclical component points to the existence of periods (cycles) of optimism or pessimism in the media that affect the exchange rate, like was found by Tetlock (2007) and Brown and Cliff (2004). Adding a lagged value of the sentiment cycle lowers the significance and increases the effect to 0.0025. The lagged cyclical component itself is not significant and therefore does not have an effect on the exchange rate. Looking at the outcome of the trend component, there does not seem to be a long-run trend in the economic news sentiment.

Almost all regressions in panel (b) of Table 4.1 point to the outcome that headline news does not affect

Tab. 4.1: Regression Outcomes UK and US News Sentiment Differential

	Lag	Var.	Coeff.	Std. Error	Prob.	Adj. R^2	DW†	AIC	BIC	Obs.
(a) Economic News	t†	score_t	0.0022	0.0010	0.0232**	0.0072	n/a	521.9151	530.6864	472
	t-1	score_t	0.0026	0.0014	0.0629*	0.0139	1.5176	-6.9333	-6.9168	260
		score_{t-1}	0.0016	0.0013	0.2186					
	t-2	score_t	0.0014	0.0019	0.7745	-0.0125	1.7688	-6.9458	-6.8602	136
		score_{t-1}	0.0025	0.0020	0.2006					
		score_{t-2}	0.0010	0.0018	0.5744					
	t	cycle_t	0.0021	0.0010	0.0504*	0.0059	1.9296	-6.8082	-6.7909	482
	t-1	cycle_t	0.0025	0.0014	0.0860*	0.0099	1.5087	-6.9391	-6.8990	269
		cycle_{t-1}	0.0017	0.0013	0.2050					
	t-2	cycle_t	0.0013	0.0019	0.5016	-0.0152	1.7573	-6.9631	-6.8786	139
		cycle_{t-1}	0.0013	0.0020	0.5192					
		cycle_{t-2}	0.0003	0.0019	0.8591					
	t†	trend_t	0.0041	0.0039	0.2961	-0.0001	n/a	531.8693	540.7182	482
(b) Headline News	t	score_t	0.0007	0.0012	0.5706	-0.0008	1.9917	-6.6100	-6.5992	882
	t-1	score_t	0.0009	0.0012	0.4933	-0.0018	1.9395	-6.5893	-6.5724	842
		score_{t-1}	0.0000	0.0012	0.9790					
	t-2	score_t	0.0015	0.0013	0.2526	0.0014	1.9536	-6.5947	-6.5714	807
		score_{t-1}	0.0007	0.0013	0.5588					
		score_{t-2}	-0.0022	0.0013	0.0815*					
	t	cycle_t	0.0008	0.0012	0.5102	-0.0006	1.9917	-6.6101	-6.5993	882
	t-1	cycle_t	0.0009	0.0013	0.4740	-0.0018	1.9394	-6.5894	-6.5725	842
		cycle_{t-1}	0.0000	0.0013	0.9871					
	t-2	cycle_t	0.0015	0.0013	0.2555	0.0014	1.9536	-6.5946	-6.5714	807
		cycle_{t-1}	0.0008	0.0013	0.5608					
		cycle_{t-2}	-0.0022	0.0013	0.0864*					
	t	trend_t	-0.0028	0.0065	0.6619	-0.0009	1.9935	-6.6098	-6.5990	882

Notes: *, ** and *** denote significance at the 10% ($p < 0.10$), 5% ($p < 0.05$), and 1% ($p < 0.01$) level, respectively.
† DW is the Durbin-Watson statistic on autocorrelation.
‡ These regressions are run with robust standard errors as the Breusch-Pagan-Godfrey tests show heteroskedasticity. They therefore do not have a DW statistic.

the exchange rate. This would mean that investors are not affected by non-economic news sentiment. Only the second lag of the news sentiment score and cycle are significant at the 10% level. The sign of the coefficients for these two regressions, however, is negative. A negative sign would mean that a negative news sentiment has a positive impact on the exchange rate. This is counter-intuitive and goes against theory.

4.2 News Sentiment and LIBOR Regressions

To test if the above results still hold when I control for UIP, as is done in the present value exchange rate models, I add the short term interest rate differential between the UK and the US to the regression equations at different lags. The short term interest rate is represented by the overnight GBP and USD LIBOR rate. I only run these regressions for the variables that had significant effects on the exchange rate in the previous Table (i.e. the economic news sentiment score and cycle) to check if the significance of the results hold with the LIBOR included. If this analysis yields the same or similar results, it will increase the robustness of the earlier analysis.

The results of the first-differenced GBP/USD exchange on both the economic news sentiment score and cycle and the LIBOR rate are presented in Table 4.2. The sample period again runs from December 2013 through mid-June 2017. Adding lags to the LIBOR rate results in loss of observations due to gaps in the data.

Again, the validity of all regressions is assessed through the Durbin-Watson (DW) statistic and other residual diagnostics. The DW statistic falls in the 1.5 - 2.5 range, which points to no sign of autocorrelation in the residuals. Other residual diagnostics are presented in Table A.2 in Appendix A. The residuals of all regressions have a zero mean. There is no correlation between the residuals and the explanatory variables (at all lags) visible in the regressions.

Tab. 4.2: OLS Regression Outcomes with LIBOR included

		Coeff.	Std. Error	Prob.	Adj. R^2	DW†	AIC	BIC	Obs.
Economic News	score_t	0.0021	0.0010	0.0461**	0.0089	1.9219	-6.7853	-6.7587	467
	libor_t	0.0000	0.0011	0.9727					
	score_t	0.0020	0.0010	0.0546*	0.0075	1.9211	-6.9056	-6.8691	451
	libor_t	0.0255	0.0175	0.1457					
	libor_{t-1}	-0.0252	0.0176	0.1527					
	score_t	0.0020	0.0011	0.0627*	0.0138	1.9303	-6.8864	-6.8388	426
	libor_t	0.0264	0.0178	0.1383					
	libor_{t-1}	-0.0279	0.0331	0.3997					
	libor_{t-2}	0.0016	0.0262	0.9522					
	cycle_t	0.0021	0.0011	0.0478**	0.0045	1.9369	-6.8022	-6.7760	477
	libor_t	0.0005	0.0010	0.6327					
	cycle_t	0.0020	0.0010	0.0485**	0.0081	1.9440	-6.9233	-6.8875	461
	libor_t	0.0260	0.0173	0.1350					
	libor_{t-1}	-0.0252	0.0174	0.1485					
	cycle_t	0.0020	0.0011	0.0666*	0.0043	1.9359	-6.9037	-6.8569	435
	libor_t	0.0268	0.0176	0.1289					
	libor_{t-1}	-0.0284	0.0328	0.3881					
	libor_{t-2}	0.0020	0.0259	0.9380					

Notes: *, ** and *** denote significance at the 10% ($p < 0.10$), 5% ($p < 0.05$), and 1% ($p < 0.01$) level, respectively.
† DW is the Durbin-Watson statistic on autocorrelation.

Adding the LIBOR to the regressions equations of Table 4.1 strengthens the results that were described above. With different lags of the LIBOR added, the contemporaneous sentiment score and cycle stay significant and positive. Only the magnitude of the effect becomes somewhat smaller. With the contemporaneous value of the LIBOR included in the regression, a rise in the economic news sentiment score of 1 raise the GBP/USD exchange rate by 0.0021 against 0.0022 in the regression without the LIBOR. None of the lags of the LIBOR have a significant effect on the exchange rate, which suggests that the short term interest rate does not have an effect on the exchange rate. This goes against the UIP, which is assumed by all three 'asset' exchange rate models described earlier in the text.

5. DISCUSSION

There exists a gap between exchange rate economics and what actually drives exchange rates. This has been documented by outcomes of various studies about the (in)effectiveness of exchange rate forecasting. But it is also documented in the literature on micro-structure finance. At least in the short to medium term, exchange rates are not determined by fundamentals, but by FX traders and their beliefs about current and future fundamentals. Even though macro exchange rate models acknowledge the importance of investors' beliefs, these beliefs are not captured by the models. The aim of this thesis was to contribute to a better understanding of the FX market through quantifying investors' beliefs about future fundamentals and incorporating them in exchange rate models.

The main question that I sought to answer in this thesis was: Can news content be used to quantify investors' beliefs about future fundamentals? After answering this question, other questions can be answered. For example: Are investors' beliefs compounded in the exchange rate directly: does the EMH hold? Can incorporating investors' beliefs in exchange rate models improve these models? The answers to these questions will be given below.

The first question is answered through the literature review. Existing literature showed that measuring and quantifying investors' beliefs is possible. Different models and ways to measure investors' beliefs and expectations in both the stock market and the FX market were explored. In this research, I followed Tetlock (2007) and measured beliefs using quantitative content analysis to construct a measure of news sentiment contained in *Reuters* economic news articles in the UK and the US. This news sentiment score represents the way investors form their beliefs. For example, an article with a positive news sentiment signals a positive change in future fundamentals. After reading a positive news article about the UK economy, for example, investors should update their beliefs and trade based on them.

The new information contained in the news articles should be reflected in the exchange rate immediately. This relates to the second question: Does the EMH hold? This was empirically tested through regressing the economic news sentiment score differential between the UK and the US on the GBP/USD exchange rate. The outcomes of the regressions have shown that the economic news sentiment score positively affects the exchange rate. Adding several lags of the score confirms that the EMH holds in that new information contained in economic news is compounded in the exchange rate directly since the first and second lag of the sentiment score are both not significant.

To test the strength of these outcomes, I also ran the above regressions with the short term interest (GBP and USD LIBOR) rate included. This means including the UIP, which the three asset models introduced in the beginning of this text assume. The results of the regressions strengthen the conclusions made above. The significant effect of the economic news sentiment score is still present, and is not attributed to UIP.

As a separate test, the effect of news sentiment contained in all kinds of news was tested through regressing the news sentiment score differential based on headline news on the GBP/USD exchange rate. None of the results point to the fact that the exchange rate is affected by sentiment contained in all sorts of news. This result was already deducted from looking at the sentiment trend differentials compared to the GBP/USD exchange rate in Figure 3.3b. The flow of bad news about Brexit did not negatively impact the GBP/USD exchange rate. Neither did the USD depreciate against the GBP after the election of Donald Trump, which clearly started a downward trend in the headline news sentiment differential.

To test the premise that there are periods of media pessimism and optimism present in the market affecting the exchange rate, the news sentiment scores were filtered for a cyclical and trend component using the HP filter. The regressions of these components on the exchange rate yielded that for economic news, there are periods of (excessive) optimism or pessimism present that affect the exchange rate. This means that investors' beliefs, and thereby the exchange rate, could be affected by certain media coverage, and not necessarily by changing fundamentals.

The above conclusions are subject to some limitations in the research. Firstly, the dictionary approach in the form used here has some flaws. The representation of a text as lists of words does not take interaction among words into account. A phrase like 'moderate growth' would be valued as much as ‘rapid growth’, whereas the latter phrase carries a more positive sentiment. Also, a word gets a different meaning if it has the word ‘not’ or ‘no’ in front of it. Not capturing these dynamics could misclassify a news text. Secondly, the HP filter uses future values of the time series to produce the cyclical and the trend components. If these components are used for forecasting, they will need to be filtered using a method that is strictly backward looking. Thirdly, the naive form of the regressions omits other factors that could also affect the exchange rate in the short to medium term like, for example, order flow.

To solve the above issues one would first need to apply more advanced methods for text sentiment analysis. Techniques called tagging ‘negations’ and ‘intensifications’ could capture word relationships within sentences. Also, ‘term weighting’ accounts for strengths of certain phrases when assigning a sentiment score. In order to further improve the quality of the news sentiment score, more news articles from different outlets over a longer horizon should be analysed. The gaps in the data for the news sentiment score results in loss of observations when lagging the variables.

Furthermore, to make the regressions less naive and suitable for forecasting, other variables need to be added to the regressions. Order flow data would be a good first example. This could lead to the ‘modelling’ of investors’ beliefs with multiple variables. This ‘model’ could then be added to existing exchange rate models, which can then be assessed on their forecasting capabilities.

With more advanced sentiment analysis techniques, sentiments in news articles can be more accurately captured. This would lead to more accurately capturing investors’ beliefs. And, together with other factors like order flow, the understanding of what drives exchange rates can be improved. Closing the gap between exchange rate economics on a macro level and exchange rate economics on a micro-finance level, could mean a more comprehensive understanding of the FX market and more accurate forecasting abilities of existing exchange rate models.

BIBLIOGRAPHY

Antweiler, W., & Frank, M. Z. (2004). Is all that talk just noise? The information content of Internet stock message boards. *Journal of Finance*, *59*(3), 1259–1294.

Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, *49*, 307–343.

Bloomberg. (2017a). *GBPUSD BGN Curncy.* Retrieved June 2, 2017 from Bloomberg Terminal.

Bloomberg. (2017b). *ICE LIBOR GBP Overnight.* Retrieved August 12, 2017 from Bloomberg Terminal.

Bloomberg. (2017c). *ICE LIBOR USD Overnight.* Retrieved August 12, 2017 from Bloomberg Terminal.

Bloomberg. (2017d). *USDGBP BGN Curncy.* Retrieved June 2, 2017 from Bloomberg Terminal.

Brown, G. W., & Cliff, M. T. (2004). Investor sentiment and the near-term stock market. *Journal of Empirical Finance*, *11*(1), 1–27.

Cheung, Y. W., & Chinn, M. D. (2001). *Currency traders and exchange rate dynamics: A survey of the US market* (Vol. 20) (No. 4).

Cheung, Y. W., Chinn, M. D., & Garcia Pascual, A. (2005). Empirical exchange rate models of the nineties: Are any fit to survive? *Journal of International Money and Finance*, *24*(7), 1150–1175.

Das, S. R. (2012). News Analytics: Framework, Techniques, and Metrics. *The Handbook of News Analytics in Finance*, 41–71.

Das, S. R. (2014). Text and Context: Language Analytics in Finance. *Foundations and Trends® in Finance*, *8*(3), 145–261.

Durbin, J., & Watson, G. (1951). Testing for Serial Correlation in Least Squares Regression . I. *Biometrika*, *38*(1), 159–177.

Engel, C., Mark, N. C., & West, K. (2006). Exchange-rate models. *NBER Reporter Fall*, *2006*, 17–34.

Engel, C., Mark, N. C., & West, K. D. (2007). Exchange Rate Models Are Not as Bad as You Think. *NBER Macroeconomics Annual*, *22*, 381–473.

Engel, C., & West, K. D. (2005). Exchange Rates and Fundamentals. *Journal of Political Economy*, *113*(3), 485–517. doi: 10.1086/429137

Evans, M. D. D., & Lyons, R. K. (2002). Order Flow and Exchange Rate Dynamics. *Journal of Political Economy*, *110*(1).

Evans, M. D. D., & Lyons, R. K. (2005). Do currency markets absorb news quickly? *Journal of International Money and Finance*, *24*(2), 197–217.

Evans, M. D. D., & Lyons, R. K. (2008). How is macro news transmitted to exchange rates? *Journal of Financial Economics*, *88*(1), 26–50.

Fama, E. (1970). American Finance Association Efficient Capital Markets : A Review of Theory and Empirical. *The Journal of Finance*, *25*(2), 383–417.

Frankel, J. A., & Rose, A. K. (1995). Empirical research on nominal exchange rates. *Handbook of International Economics*, *3*(C), 1689–1729.

Fratzscher, M., Rime, D., Sarno, L., & Zinna, G. (2015). The scapegoat theory of exchange rates: The first tests. *Journal of Monetary Economics*, *70*, 1–21.

French, K. R., Schwert, G. W., & Stambaugh, R. F. (1987). Expected Stock Returns and Volatility. *Journal of Financial Economics*, *19*, 3–29.

Hodrick, R. J. ., & Prescott, E. C. (1997). Postwar U . S . Business Cycles : An Empirical Investigation. , *29*(1), 1–16.

Hopper, G. P. (1997). What determines the exchange rate: economic factors or market sentiment? *Business Review Federal Reserve Bank of Philadelphia*, 17–29.

Huang, A., Zang, A. Y., & Zheng, R. (2012). Large Sample Evidence on the Informativeness of Text in Analyst Reports. *SSRN Electronic Journal*(March 2017).

Isard, P. (2006). Uncovered Interest Parity. *IMF Working Paper*.

Kearney, C., & Liu, S. (2014). Textual sentiment in finance: A survey of methods and models. *International Review of Financial Analysis*, *33*(Cc), 171–185.

Krippendorff, K. (2012). *Content Analysis: An Introduction to Its Methodology* (Third Edit ed., Vol. 79). SAGE.

Leventakis, J. A. (1987). Exchange Rate Models: Do They Work? *Weltwirtschaftliches Archiv*, *Bd. 123*(H. 2), 363–376.

Loughran, T., & McDonald, B. (2011). When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *Journal of Finance*, *66*(1), 35–65.

Lyons, R. K. (2000). *The Microstructure Approach to Exchange Rates.* Cambridge: MIT Press.

Mark, N. C. (1995). Exchange Rates and Fundamentals: Evidence on Long-Horizon Predictability. *The American Economic Review*, *85*(1), 201–218.

Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies. *Journal of International Economics*, *14*(1-2), 3–24.

Merriam-Webster's Dictionary. (2017). *Content analysis.* Retrieved 2017-06-13, from https://www.merriam-webster.com/dictionary/contentanalysis

Moniz, A., & de Jong, F. (2014). Classifying the influence of negative affect expressed by the financial media on investor behavior. In *5th information interaction in context symposium, iiix 2014* (pp. 275–278).

Papaioannou, P., Russo, L., Papaioannou, G., & Siettos, C. I. (2013). Can social microblogging be used to forecast intraday exchange rates? *NETNOMICS: Economic Research and Electronic Networking*, *14*(1-2), 47–68.

Plakandaras, V., Papadimitriou, T., Gogas, P., & Diamantaras, K. (2015). Market sentiment and exchange rate directional forecasting. *Algorithmic Finance*, *4*(1-2), 69–79.

Rime, D., Sarno, L., & Sojli, E. (2010). Exchange rate forecasting, order flow and macroeconomic information. *Journal of International Economics*, *80*(1), 72–88.

Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-Based Methods for Sentiment Analysis. *Computational Linguistics*, *37*(2), 267–307.

Tetlock, P. C. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. *The Journal of Finance*, *62*(3), 1139–1168.

APPENDIX

Appendix A

RESIDUAL DIAGNOSTICS

Tab. A.1: Residual Diagnostics to Regressions in Table 4.1

			Mean	Std. Dev.	Prob. $\chi^{2\dagger}$	Corr. t‡	Corr. t-1‡	Corr. t-2‡
(a) Economic News	t	score	-2.94^{-20}	0.0081	0.0341*	-1.62^{-17}		
	t-1	score	1.87^{-19}	0.0075	0.4807	-4.33^{-17}	-1.77^{-17}	
	t-2	score	-2.28^{-19}	0.0073	0.5375	-2.74^{-17}	-2.27^{-17}	0.0000
	t	cycle	3.17^{-19}	0.0080	0.1172	1.03^{-17}		
	t-1	cycle	-4.00^{-19}	0.0075	0.8499	-4.92^{-20}	-1.14^{-17}	
	t-2	cycle	-2.31^{-19}	0.0073	0.6799	-3.13^{-17}	-1.13^{-17}	2.59^{-34}
	t	trend	-4.46^{-19}	0.0080	0.0011**	-5.07^{-18}		
(b) Headline News	t	score	6.74^{-19}	0.0089	0.5367	3.50^{-18}		
	t-1	score	4.17^{-19}	0.0089	0.8295	-2.56^{-17}	-1.64^{-17}	
	t-2	score	-3.46^{-19}	0.0089	0.9073	5.81^{-18}	-1.34^{-17}	2.29^{-17}
	t	cycle	1.2^{-19}	0.0089	0.5587	7.13^{-18}		
	t-1	cycle	3.74^{-19}	0.0089	0.8462	1.49^{-17}	3.76^{-18}	
	t-2	cycle	5.9^{-19}	0.0089	0.9133	1.97^{-18}	-1.95^{-17}	5.42^{-17}
	t	trend	1.14^{-19}	0.0089	0.8155	-6.39^{-17}		

Notes: This table shows residual diagnostics for the non-lagged regressions in Table 4.1.
† This is the test statistic for the Breusch-Pagan-Godfrey test for heteroskedasticity. * and ** represents a rejection of the null of homoskedasticity at the 5% and 1% level respectively.
‡ These values represent the correlation between the residuals and the explanatory variable in the regression with the different lags.

Tab. A.2: Residual Diagnostics to Regressions in Table 4.2

	LIBOR lag	Variables	Mean	Std. Dev.	Prob. $\chi^{2\dagger}$	Corr. Var. 1‡	Corr. LIBOR t‡	Corr. LIBOR t-1‡	Corr. LIBOR t-2‡
Economic News	t	score & libor	-3.71^{-19}	0.0081	0.0939	-1.63^{-17}	1.90^{-33}		
	t-1	score & libor	-9.69^{-19}	0.0076	0.3950	8.50^{-18}	7.04^{-16}	-2.91^{-16}	
	t-2	score & libor	5.13^{-19}	0.0077	0.5104	-5.96^{-18}	6.32^{-16}	-5.17^{-16}	1.73^{-16}
	t	cycle & libor	-3.93^{-19}	0.0080	0.1917	-2.85^{-17}	-1.03^{-17}		
	t-1	cycle & libor	-6.02^{-20}	0.0075	0.5137	8.17^{-18}	5.82^{-16}	-5.14^{-16}	
	t-2	cycle & libor	4.94^{-19}	0.0076	0.6330	2.45^{-17}	3.46^{-16}	-2.23^{-16}	-2.20^{-16}

Notes: This table shows residual diagnostics for the non-lagged regressions in Table 4.2. Panel (a) represents Economic News and panel (b) represent All Headline News.
† This is the test statistic for the Breusch-Pagan-Godfrey test for heteroskedasticity. * and ** represents a rejection of the null of homoskedasticity at the 5% and 1% level respectively.
‡ These values represent the correlation between the residual and the explanatory variables in the regression.

Appendix B

R CODE

B.1 Scraping Reuters URLs

```
# parse command line arguments
suppressMessages({
        library(argparser)
        library(magrittr)
})
argv <- arg_parser("Scrape www.reuters.com") %>%
        add_argument("--from",  help="page number (start)",    default=1L) %>%
        add_argument("--to",    help="page number (end)",      default=463L) %>%
        add_argument("--cache", help="number of pages cached", default=100L) %>%
        add_argument("--db",    help="sqlite database path",   default="") %>%
        add_argument("--csv",   help="output csv file",        default="") %>%
        parse_args()

# Packages
        suppressMessages({
        library(progress)
        library(foreach)
        library(doParallel)
        library(readr)
        library(dplyr)
        library(purrr)
        library(rvest)
        library(dbplyr)
        library(RSQLite)
        })
        cl <- makeCluster(detectCores())
registerDoParallel(cl)

# global constants
        BASE_URL    = "http://uk.reuters.com/news/archive/GCA-EconomyUK?view
        =page&pageSize=10&page="
        DB_PATH     = ifelse(argv$db == "//campus.eur.nl/users/home/
        433638yb/...",
```

```
        tempfile(tmpdir=".", fileext=".sqlite3"), argv$db)
        PAGE_START  = argv$from
        PAGE_END    = argv$to
        CACHE_PAGES = argv$cache
        OUTPUT_CSV  = ifelse(argv$csv == "//campus.eur.nl/users/home/
        433638yb/...", FALSE, TRUE)
        CSV_PATH    = argv$csv

# functions
## create empty stories object
empty_stories <- function(size=0) {
        data.frame(
                headline_url = character(size),
                story_url    = character(size),
                title        = character(size),
                text         = character(size),
                time         = character(size)
        )
        }

## scrape stories in the given page
get_stories <- function(page_url) {
        tryCatch({
          stories <-
                read_html(page_url, encoding = "UTF-8") %>%
                html_nodes("article.story") %>%
                map(function(s) {
                  url   <- s %>% html_node("a") %>% html_attr("href")
                  title <- s %>% html_node(".story-title") %>% html_text()
                  text  <- s %>% html_node("p") %>% html_text()
                  time  <- s %>% html_node(".timestamp") %>% html_text()
                  tibble(headline_url=page_url, story_url=url, title=title,
                  text=text, time=time)
                }) %>% bind_rows()
          stories
        }, error=function(e) {
          print(page_url)
          print(e)
          empty_stories()
        })
}

# setup database
db <- src_sqlite(DB_PATH, create=TRUE)
stories_db <- copy_to(db, empty_stories(), name="stories", temporary = FALSE)
```

```
# start scraping
pb <- progress_bar$new(
        total  = (PAGE_END - PAGE_START + 1) %/% CACHE_PAGES, #PAGE_END
        - PAGE_START - 1,
        format = "scraping [:bar] :percent in :elapsed"
)

for(i in seq(PAGE_START, PAGE_END, CACHE_PAGES)) {
  stories <- foreach(j=seq(i, min(i+CACHE_PAGES-1, PAGE_END)),
        .combine=bind_rows, .packages = c("dplyr", "purrr", "rvest")) %dopar%
        {
          page_url <- paste0(BASE_URL, j)
          get_stories(page_url)
        }
  try({
        db_insert_into(db$con, "stories", stories)
  })
  pb$tick()
}

# save as a csv file
if(OUTPUT_CSV) {
  output_tbl <- tbl(db, "stories") %>% collect()
  write_csv(output_tbl, CSV_PATH)
}
```

B.2 Scraping HTML and Extracting Article

```
#### Packages ####
library("boilerpipeR")
library("RCurl")
library("tm")
library("stringr")
library("wordcloud")
library("SnowballC")

#### Extract html from web ####
art20170612a <- getURL("http://uk.reuters.com/article/
  uk-britain-economy-idUKKBN1930MQ", .opts=curlOptions(followlocation = TRUE))
art20170612b <- getURL("http://uk.reuters.com/article/
  uk-britain-election-s-p-idUKKBN1930R7", .opts=curlOptions(followlocation = TRUE))
...
...
...
art20110405c <- getURL("http://uk.reuters.com/article/
  uk-banks-smes-idUKLNE73400220110405", .opts=curlOptions(followlocation = TRUE))
art20110404a <- getURL("http://uk.reuters.com/article/
  uk-pmi-construction-britain-idUKLNE73301820110404",
   .opts=curlOptions(followlocation = TRUE))

#### Extract article ####
art20170612a <- ArticleSentencesExtractor(art20170612a)
art20170612b <- ArticleSentencesExtractor(art20170612b)
...
...
...
art20110404c <- ArticleSentencesExtractor(art20110404c)
art20110401a <- ArticleSentencesExtractor(art20110401a)

#### Make txt. files ####
write.table(art20170612a, "art20170612a.txt", sep = " ")
write.table(art20170612b, "art20170612b.txt", sep = " ")
...
...
...
write.table(art20110404c, "art20110404c.txt", sep = " ")
write.table(art20110401a, "art20110401a.txt", sep = " ")
```

B.3 Cleaning Texts and Creating News Sentiment Score

```
#### Packages ####
library("RCurl")
library("tm")
library("stringr")
library("wordcloud")
library("SnowballC")
library("xlsx")

#### Create Corpus ####
corp_source <- VCorpus(DirSource("//campus.eur.nl/users/home/433638yb/...")) 
corp <- corp_source

#### Clean Corpus ####

# To remove lines with certain words
for(i in seq(corp_source)) {
corp[[i]]$content <- corp_source[[i]]$content[!grepl("FILE",
corp_source[[i]]$content, ignore.case = FALSE)]
}

for(j in seq(corp)) {
corp[[j]]$content <- corp[[j]]$content[!grepl("BST",
corp[[j]]$content, ignore.case = FALSE)]
}

for(k in seq(corp)) {
corp[[k]]$content <- corp[[k]]$content[!grepl("REUTERS",
corp[[k]]$content, ignore.case = FALSE)]
}

for(l in seq(corp)) {
corp[[l]]$content <- corp[[l]]$content[!grepl("GMT",
corp[[l]]$content, ignore.case = FALSE)]
}

for(m in seq(corp)) {
corp[[m]]$content <- corp[[m]]$content[!grepl("Reporting by",
corp[[m]]$content, ignore.case = FALSE)]
}

for(n in seq(corp)) {
corp[[n]]$content <- corp[[n]]$content[!grepl("Polling by",
corp[[n]]$content, ignore.case = FALSE)]
```

```
}

for(o in seq(corp)) {
corp[[o]]$content <- corp[[o]]$content[!grepl("Analysis by",
corp[[o]]$content,ignore.case = FALSE)]
}

for(p in seq(corp)) {
corp[[p]]$content <- corp[[p]]$content[!grepl("Editing by",
corp[[p]]$content,ignore.case = FALSE)]
}
remove(i,j,k,l,m,n,o,p)

# Other text manipulations
corp <- tm_map(corp, content_transformer(tolower))
corp <- tm_map(corp, removeWords, stopwords("english"))
corp <- tm_map(corp, removePunctuation)
corp <- tm_map(corp, removeNumbers)
corp <- tm_map(corp, stripWhitespace)
corp <- tm_map(corp, stemDocument)

# Inspect the corpus as text
corp_text <- sapply(corp, as.character)
corp_text[["art140501a.txt"]]                          #example

#### Load lexicon/dictionary ####
setwd("//campus.eur.nl/users/home/433638yb/...")
pos <- readLines("Positive_own.txt")
neg <- readLines("Negative_own.txt")

pos <- tolower(unique(pos))
neg <- tolower(unique(neg))

pos <- stemDocument(pos)
pos <- unique(pos)

neg <- stemDocument(neg)
neg <- unique(neg)

#### Bag of words ####
corp_list <- lapply(corp, FUN = paste, collapse=" ")
corp_bag <- str_split(corp_list, pattern = "\\s+")

# match words
posmatch_list <- lapply(corp_bag, function(x) {sum(!is.na(match(x, pos)))})
```

```
negmatch_list <- lapply(corp_bag, function(x) { sum(!is.na(match(x, neg)))})

posmatch <- unlist(lapply(corp_bag, function(x) { sum(!is.na(match(x, pos)))}))
negmatch <- unlist(lapply(corp_bag, function(x) { sum(!is.na(match(x, neg)))}))

# Naive score: substract negative from positive
score_naive_list <- lapply(corp_bag, function(x) { sum(!is.na(match(x, pos)))
- sum(!is.na(match(x, neg)))})

score_naive <- unlist(lapply(corp_bag, function(x) { sum(!is.na(match(x, pos)))
- sum(!is.na(match(x, neg)))}))

# Relative score: (pos-neg)/(pos+neg)
score_list <- lapply(corp_bag, function(x) { (sum(!is.na(match(x, pos)))
- sum(!is.na(match(x, neg)))) / (sum(!is.na(match(x, pos)))
+ sum(!is.na(match(x, neg))))})

score <- unlist(lapply(corp_bag, function(x) { (sum(!is.na(match(x, pos)))
- sum(!is.na(match(x, neg)))) / (sum(!is.na(match(x, pos)))
+ sum(!is.na(match(x, neg))))}))

#### Make xls spreadsheet ####
write.xlsx(negmatch, "//campus.eur.nl/users/home/433638yb/...")
write.xlsx(posmatch, "//campus.eur.nl/users/home/433638yb/...")
write.xlsx(score_naive, "//campus.eur.nl/users/home/433638yb/...")
write.xlsx(score, "//campus.eur.nl/users/home/433638yb/...")
```