

Predicting stock price and spread movements from news

Pontus Wistbacka
Hanken School of Economics
pontus@siloi.ai

Katia Vozian
Hanken School of Economics
Helsinki Graduate School of Economics
katia.vozian@hanken.fi

Samuel Rönqvist
Hanken School of Economics
TurkuNLP, University of Turku
saanro@utu.fi

Satchit Sagade
Goethe University Frankfurt
Leibnitz Institute for Financial Research SAFE
satchit.sagade@gmail.com

Abstract—We explore several ways of using news articles and financial data to train neural network machine learning models to predict shock events in high-frequency market data, and aggregated shock episodes. We investigate the use of price movements in this context, and separately at a daily interval as well. We describe in detail how training sets are created from our data sources and how our machine learning models are trained. We find that pairing company-related news text with events or movements in financial time series proves less straight-forward than the literature would indicate. We discuss possible reasons for negative results, especially relating to the combination of minute-level news and millisecond-level market data.

I. INTRODUCTION

The market adjusts the price of a company's stock as a response to every known piece of information about the company as well as any other related information available, public or private, according to the Efficient Market Hypothesis [1]. It is not known if this effect is instantaneous. Brogaard [2] identified four components that affect the share price; public, as well as private, company specific information, market-wide information and noise. In this report, we are specifically interested in public company-specific information, and describe modeling experiments that seek to identify how such information from news may affect markets in terms of liquidity shocks and price movements. We perform this study using Reuters news and FTSE-100 market data. We encounter fundamental challenges in modeling high-frequency phenomena based on comparatively infrequent news.

The tasks performed by financial analysts and informed investors, such as reading and analyzing financial news reporting and regulation-mandated reports (e.g. 8K reports), are prime targets for automation because of their large economic value. Automated textual analysis as a component in algorithmic trading means that the price impact of e.g. news articles on a stock can be anticipated and traded upon earlier than would otherwise be possible [3]. Compared to the study of structured data such as historical share price, which is the domain of technical analysis, textual analysis is more challenging because it deals with human language that

is inherently variable and ambiguous, and often subjective and abstract, but also highly expressive.

Handcrafting a rule-based computer program to predict how any particular stock will be affected by news articles is very time consuming, error prone and poorly scalable, compared to using machine learning. Machine learning, in particular supervised machine learning, can be described as using a parameterized model, and varying its parameters to fit a training set of example inputs and outputs [4].

Several types of parameterized models and ways of tuning their parameters have been devised for problems in the field of written text, or natural language processing (NLP). The most universally successful ones are neural networks of various design, referred to as deep learning models, with their parameters varied using the technique known as error back-propagation [4].

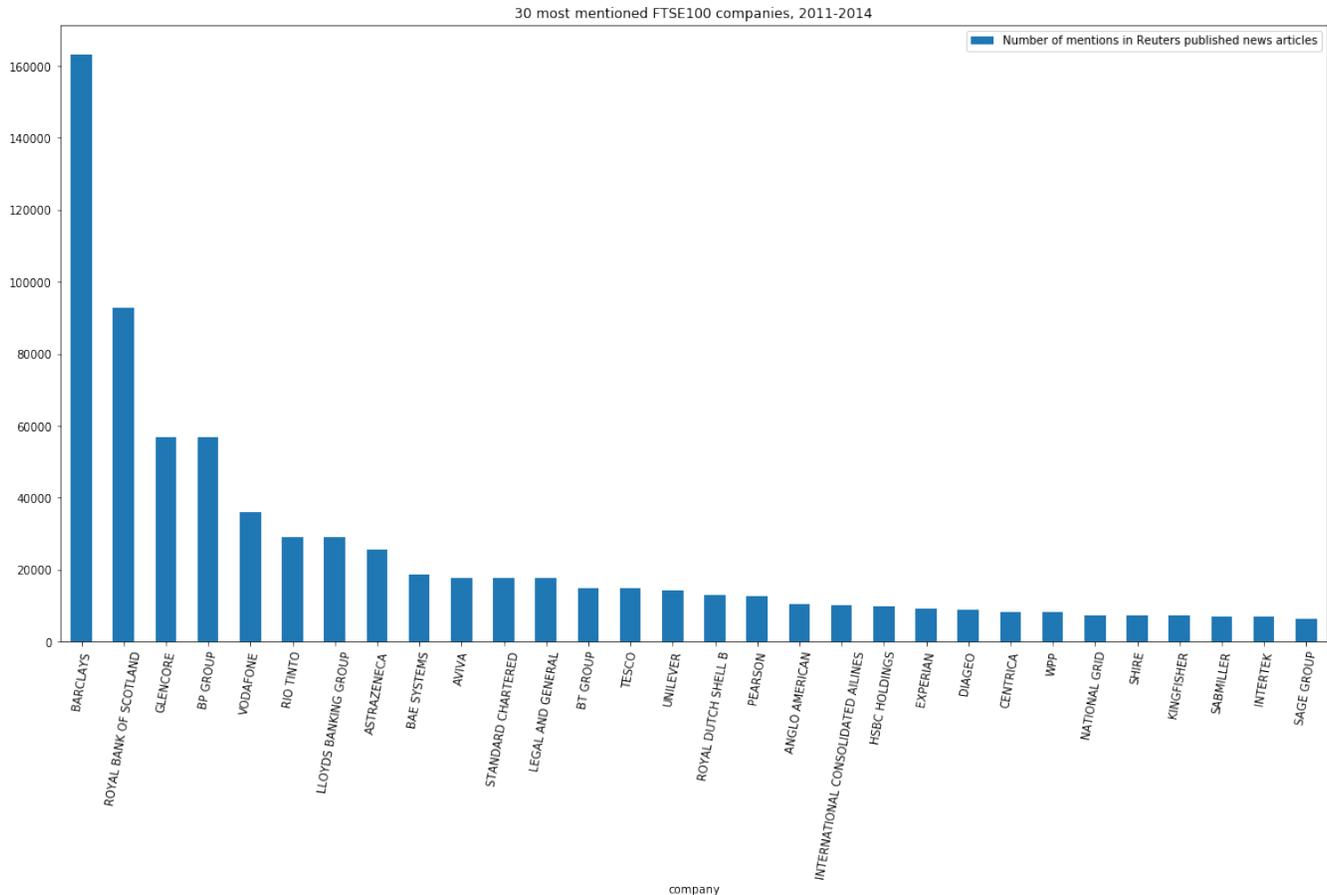
Rönqvist and Sarlin [5] successfully trained a neural network on a large corpus of news articles, downloaded from Reuters, to predict bank distress based on events that have resulted in government interventions (e.g. bailouts) in the past. This report details our work on adapting and extending their methodology towards the prediction of variations in financial time series, particularly, liquidity shocks observed in high-frequency data. We explore both inter-day and intra-day impacts of news in order to account for different dynamics in low-frequency versus high-frequency financial time series, and compare against related work.

In the following sections, we outline this related literature, the data sets that we use in our experiments, as well as the selection of machine learning methods and preprocessing steps. Finally, we present the results from our experiments and reflect on the study.

II. RELATED LITERATURE

The identification of financially relevant information from text, such as news, has been of considerable interest both in academia and industry for more than a decade (cf., e.g., [6]). In early work [7], text data is usually represented numerically using Bag of Words (BoW) representations that indicate the presence of words individually, while ignoring word order and syntactic structure, or represented through the extraction

Fig. 1: 30 most mentioned companies in the FTSE-100, 2011–2014



of linguistic features such as part-of-speech or named entity tags. Ding et. al [8] used structured event representations of the headlines of news articles, i.e. “Microsoft”, “sues”, “Barnes”, “Noble” represented as (Actor = Microsoft, Action = sues, Object = Barnes & Noble). In total, they find some 60,000 events from news headlines. To create their training set, these events are then aligned with a financial time series of the stock price [9] and labelled as either positive or negative according to the difference in the stock price between the day that the article was published and the subsequent day, week and month. They then trained neural networks which were able to predict the stock price movement of companies from these events with an accuracy upwards of 65% [8, 10]. The authors used Bag of Words features with a feed-forward neural network as a baseline [10] reaching 57% accuracy.

Ding et al. note several findings, including that the accuracy of their models decrease with longer time windows, and that predicting the price change of individual stocks seems to outperform predicting it for the S&P 500 index as a whole, something which was later verified by others [8, 11, 12]. Ding et al. released their data set which has subsequently been used by other authors [11, 12]. The data set consists of 109,110 news articles related to companies

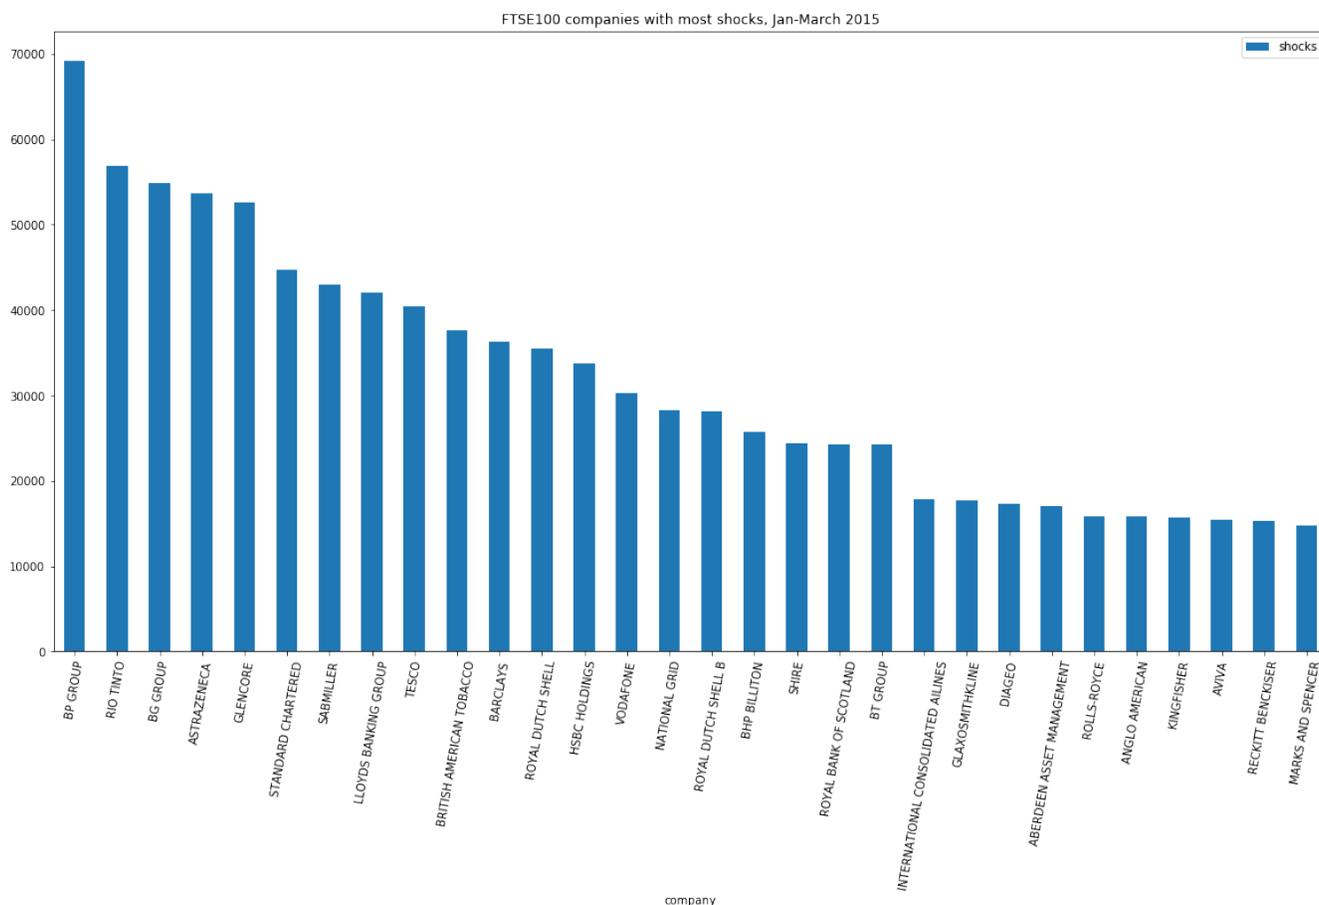
in the the S&P 500 in the time period October 2006 to November 2013, which had been scraped from Reuters and Bloomberg. The articles publication date are noted on a minute resolution timestamp.

Price movements in-between trading days are most commonly used as the target for prediction, as intra-day price data is not readily available. The articles [9, 11, 12] all use financial data for the S&P 500 acquired from Yahoo Finance.

More recently, several authors [11]–[13] have been working on improving the results with deep learning techniques for this same data set using unstructured text with no linguistically informed methods. They use a data-driven approach for *representation learning* similar to the popular *word2vec* method by Mikolov et al. [14], in order to compute word embeddings, i.e., dense vector representations. These vectors represent the semantics of the words in the sense that words that are similar in meaning will have word vectors that are close to each other.

Rönnqvist and Sarlin [5] investigated the use of deep learning for detecting discrete events without annotated textual resources. Particularly, they address events of bank distress, related to 101 European banks. They find 262,000 news articles from Reuters that mentioned these banks on 716,000 occasions. They train vectors representations from

Fig. 2: The 30 companies with the highest number of liquidity shocks in the FTSE-100, Jan-March 2015



scratch using the news text. Furthermore, they split label mentions as either co-occurring or not co-occurring with the distress events, using a variable time window around the articles. With this distant supervision approach, they train a three-layer feed-forward neural network to distinguish news text that is indicative of bank distress. They show that it is possible to predict bank distress in individual banks based on the news reporting, and use the predicted distress probabilities as means to explore the evolution of distress over time and across countries, as well as retrieving descriptions of events causing the elevated levels to be signaled.

III. DATA SETS

The text collection used in this study consists of 3,329,397 news articles from Reuters. From these, we filtered out 261,842 unique sentences mentioning companies in the FTSE-100 in the time period 2011-2014, as well as 28,312 sentences in the first three months of 2015.

A typical sentence in the resulting data set is found below:

* ASHTEAD GROUP : The industrial equipment hire firm said on Tuesday it reported a record pre-tax profit of 141 million pounds in the first

half of the year and proposed an interim dividend of 1.5 pence per share.

Tue Dec 11, 2012 2:31am EST

The distribution of mentions of the 30 most frequent companies is illustrated in Figure 1. For instance, we see that Barclays and Royal Bank of Scotland were the most dominant in the reporting.

At intra-day level, a proprietary limit order book data set, consisting of transactions at millisecond resolution was used in the project. This data set, which covers the first quarter of 2015, was coupled with classifications of liquidity shock occurrences following Danielsson et al. [15]. The definition of the liquidity shock follows a peak-over-threshold approach. A liquidity shock occurs in the point in time at which the quoted spread deviates from a threshold in the direction of less liquidity and may last until the quoted spread returns below the threshold level. The threshold is the 95th percentile of the hourly empirical distribution of the quoted spread of the respective day for a selected stock e.g. 8 a.m. 15-01-2019. We purposely chose the 95th percentile as the threshold since our interest lies in capturing severe liquidity shocks in contrast with any departure from the median or mean. The distribution of shocks by companies is seen in Figure 2.

For inter-day financial data, we downloaded data related to the FTSE-100 from Yahoo Finance. This data set includes market open and close stock prices, going back until 2011.

IV. METHODS

For the modeling of liquidity shocks we follow the labelling approach of Rönnqvist & Sarlin [5], in order to obtain labels for each mention of FTSE-100 companies in the news as to whether they co-occur with a shock event or not. Later, for modeling price data, we label mentions according to the price movement the next day.

1) *Bag of Words – Neural Network*: As an elementary model in our search for a signal, we use a feed-forward neural network that implements a bag-of-words approach, i.e., the identities of words present in a text are feed as independent features. This type of model is typically used as a simple but strong baseline, even though it does not account for word order or sentence structure.

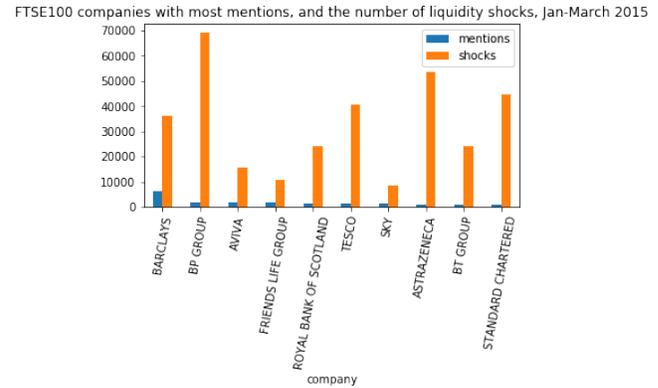
For feature extraction, we used the CountVectorizer from scikit-learn to turn the sentences into a sparse matrix representation. To construct the matrix, each unique term in the corpus, known as the dictionary, is first assigned a unique integer. Each sentence is represented by a column in the matrix, where the indices in this column denote the number of times each word in the dictionary occurred in the sentence. The label of the sentence (co-occurring/non-co-occurring or positive/negative/neutral), is turned into a one-hot vector representation. These are vectors of length equal to the number of classes, with one element having the entry one and the rest zero.

We use a feed-forward neural network architecture where each word in the dictionary has its own node in the input layer. The network further consists of a hidden layer with 300 nodes and a tanh activation function, an optional dropout layer and a dense output layer with softmax activation function for classification. We use categorical cross-entropy as loss function and optimize with stochastic gradient descent. The models are trained for 80 epochs, all the while monitoring the validation set accuracy and saving the best model.

2) *Recurrent Neural Network*: We also train a recurrent neural network architecture, namely a Long Short-Term Memory (LSTM) network, in order to test if accounting for word order and sentence structure may provide better performance compared to the bag-of-words model. The LSTM network models sentences as a sequence of words by traversing the sequence and keeping an internal state representation, which can model long-range dependencies and is therefore suitable for text data. The model consisted of an input layer, an embedding layer that learns word encodings, an LSTM layer with 300 nodes, the tanh activation function and dropout, and a final dense output layer with softmax activation function for classification. We used categorical cross-entropy for the loss function, and the Adam optimizer.

We evaluate the methods based on accuracy on a held-out test set, and compare performance against the majority class

Fig. 3: Shocks and mentions for the 10 most reported on companies in the FTSE100



baseline, i.e., the frequency of the largest class, which represents the accuracy achieved following a best guess based on the prior class distributions.

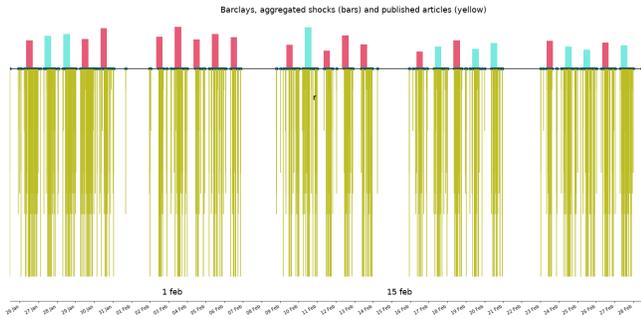
A. Data Preprocessing

The Reuters dataset contains articles of varying type; news articles, earnings reports, bond emissions and other press announcements. To reduce noise, we filtered articles to find only those that mentioned companies in the FTSE 100. For each company, we accounted for common abbreviations and spelling variations using regular expressions. In some experiments, we performed further filtering in order to find only news articles, or to find documents that mentioned certain keywords, such as *liquidity*. In processing the articles, due to the size of the raw data set, the code was written to be executed in parallel on several cores.

We decided to use the sentence as the unit for training examples. Every article that mentioned a company was split into sentences using NLTK (Natural Language Toolkit). We then added only those sentences that contained the company name to our training set. The company name was replaced by a placeholder token, in order to avoid overfitting on company names. The sentences were further filtered by their length and by the number of digits and special characters, as we believed that documents which contained a lot of special formatting, e.g. tables, would not be useful for the machine learning models as opposed to descriptive text. Reuters will often publish the same text on multiple occasions with only minimal changes. Because of this, our data set would contain the same sentences several times. To avoid the bias this could cause on the machine learning model accuracy, we removed duplicate examples from the training set. The articles have their publication date noted on a resolution of one minute. This timestamp is saved along with each sentence to construct the training sets.

1) *Liquidity Shock Training Set*: The liquidity shocks that we study are discrete events, and we label sentences using a co-occurrence heuristic (in accordance with [5]), whereby we define a time window around the time of publication for each mentioning sentence/article. We then label a sentence

Fig. 4: Liquidity shocks (aggregated episodes, above line) and news articles (below line) during 5 weeks



as co-occurring or not co-occurring with a shock event if it falls within this time window for any event. We test different window sizes starting from news published up to, e.g., 6h, 3h, 1h, 30min, 10min, 5min or 1min prior to the event until 1min past (in order to account for possible inaccuracies in the timestamps). Due to the number of the events found by the shock identification approach, using this labelling scheme alone, produces orders of magnitude more shocks than news articles, as seen in Figure 3. On average, there are 200 shocks per mention in the period we were researching (January-March 2015).

This is problematic for the predictive modeling, as the model would be expected to infer labels for several different shocks based on any one mention. To mitigate this, we chose to aggregate shocks into shock episodes on an hourly or daily basis, as well as apply some filtering. In the liquidity shock data, a severity measure describing the magnitude of the shock is defined for each shock event as the maximum spread $max(S)$ during the shock relative the percentile threshold θ :

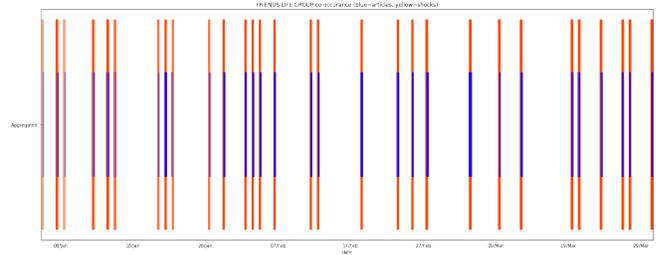
$$\frac{max(S) - \theta}{\theta}$$

This measure provides a means for filtering shock events and episodes. In Figure 4, the aggregated liquidity shocks for Barclays (red/cyan bars) are visualized along with news articles (yellow bars), with the height of the shock bars scaled according to their severity. By tuning the time for aggregation and the filtering applied, we could create training sets with equal number or fewer liquidity shocks than articles. Figure 5 shows a visualization of this kind of balanced training set with the co-occurring articles shown in blue along with the aggregated liquidity shocks in red.

2) *Intra-Day Mid-Price and Spread Training Sets:* In order to model how the market responds to news, we consider price and spread movements following the publication of an article. These movements can provide additional means for identifying significant events that may be possible to predict based on news.

In order to automatically label sentences, we used the UTC-converted timestamp of each article and compared the price or the spread for the relevant company along with the same variable the following day. We would label the sentence

Fig. 5: Aggregated liquidity shocks, overlap with articles, TESCO



as either positive, negative or neutral. The process is shown in Algorithm 1. We tune the prevalence of the neutral class in training sets with the use of a threshold value. This value is calculated as set multiples of the standard deviation δ in the financial time series.

Algorithm 1 Sentence labelling for a company using the stock price P , a threshold value and the forward and backwards time windows

Result: Labelled sentences

```

for each sentence at time  $T$  do
  if  $Abs(P_{T+w_f} - P_{T-w_b}) > \delta$  then
    if  $P_{T+f} > P_{T-b}$  then
      | label = positive
    else
      | label = negative
    end
  else
    | label = neutral
  end
end

```

For either the mid-price or the spread, we vary the time window within which we calculate the change in the relevant variable in ranges from minutes to hours. In each case, we set the beginning of the time window at one minute before the article publication timestamp. If we did not have a value inside the time window, we discarded the sentence from the data set.

The effect of increasing the threshold can be seen in Figure 6, where each of the panels show the labelling between positive (blue), negative (red) and neutral (yellow), according to price movements at three different threshold levels: 0, 1 and 2 times the standard deviation.

3) *Inter-Day Mid-Price Training Sets:* Finally, we conduct experiments on inter-day price data, in order to establish performance on our news data set with our text extraction procedures, and compare against the more prevalent literature on this topic.

In the case of the inter-day training sets, the time window in Algorithm 1 stretches forward one day and backwards zero days.

Fig. 6: Threshold settings (0, 1 and 2 times σ) for labeling Barclays stock price movements as positive (blue) / negative (red) / neutral (yellow)

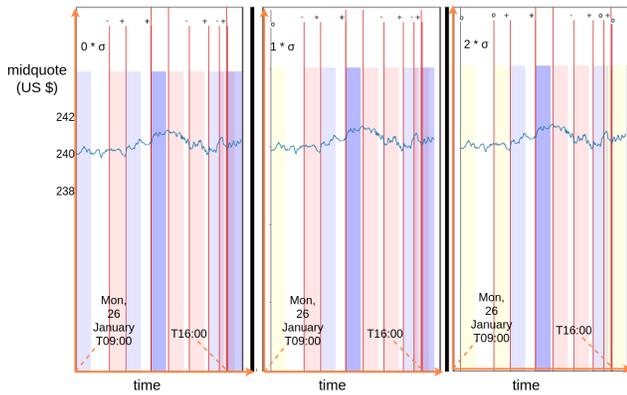
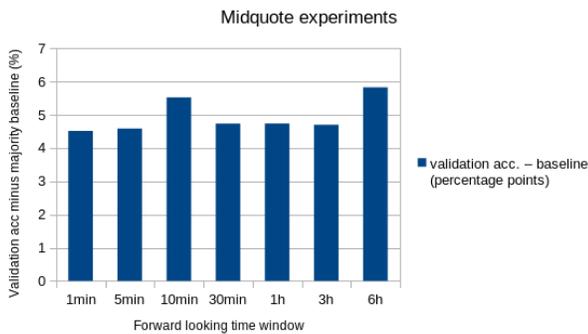


Fig. 7: Varying the time window parameter, mid-quote training set



V. RESULTS

We began the set of experiments by focusing on using the liquidity shock co-occurrence training set. Different versions of the training set were created by varying the time window as well as the time used for aggregating individual shock events. The initial liquidity shock data set consisted of 1.5 million events in total. We believed that the orders of magnitude difference compared to the number of sentences was causing our models to not pick up a signal. In our liquidity shock training sets, we varied the number of shocks from hundreds up to a few thousands. These experiments did not allow for the training of a model that would exceed the majority baseline.

It is possible that because the shock events, being identified at a resolution of milliseconds, would be caused by many factors that are not explained by news events, and thus that the signal might be too noisy to predict with the methods we were using.

When switching our focus to the intra-day mid-price and spread training sets, our methods would yield models that would modestly exceed the majority baseline, typically by one-two percentage points. We decided to probe the signal we were seeing by varying the time window and threshold

Fig. 8: Varying the time window parameter, spread training set

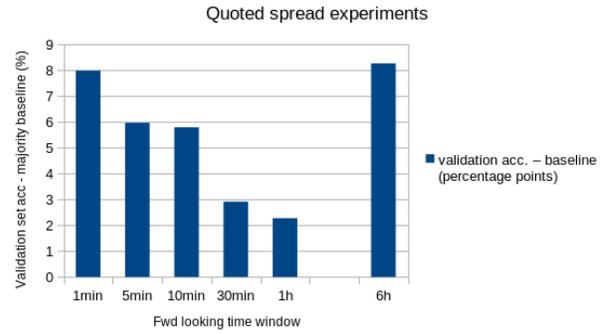
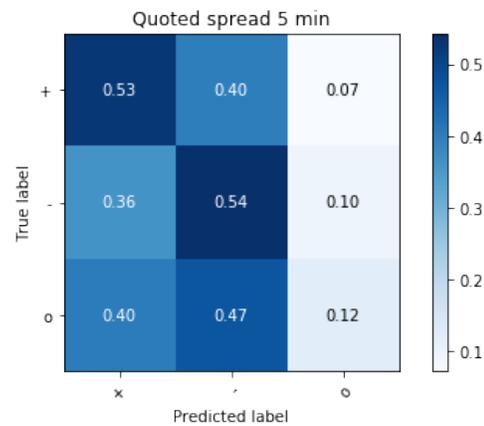


Fig. 9: Confusion matrix for model trained on spread training set



parameters. Shown in Figures 7 and 8 are the results of varying the time window for the articles in the intra-day mid-price and spread training sets, measured in terms of the difference between the training set validation accuracy and the majority class baseline. Based on these experiments, there would seem to be quite a large effect of the time window used when trying to predict the quoted spread.

Following this, we also investigated classification errors with a confusion matrix, as seen in Figure 9. As can be seen, the model has a strong tendency to not predict the neutral class, which hurts overall performance. We investigated the effect of varying the class distribution by using a threshold parameter, as we believe that the models would be better at predicting large movements in the financial time series.

A. Sanity Checking

Before moving on to further experiments, we performed sanity checks on our methods. We considered that the weak signal we observe in Figure 10 should not be unrealistic given published results in the literature, for instance Ding et al. [9], granted that our time resolution was higher than theirs was. We wanted to identify any cause of potentially overoptimistic results. This included refining the data cleaning already mentioned in section IV, by making sure our data sets did not include any duplicate articles or

Fig. 10: False signal, training-validation split at random

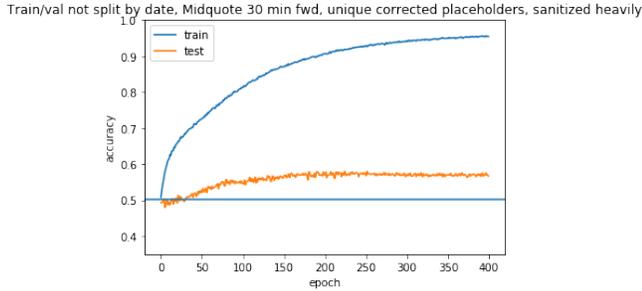
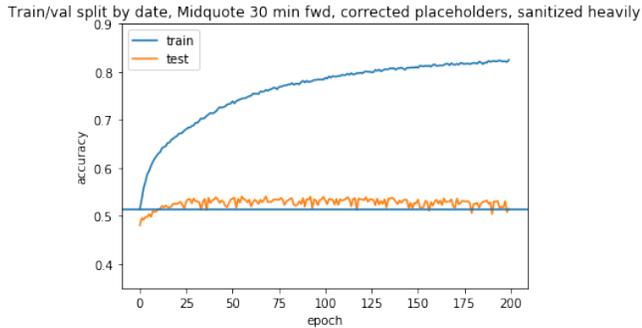


Fig. 11: Splitting training and validation sets by date



other repetition of information that may spill over between train and test set.

We also performed error analysis by inspecting examples from the test set that the model assigned a high confidence to, to try and establish if the model was learning anything that was sensible to a human. We found many sentences from documents that were not news articles, and seemingly did not have any information about the markets. As noted earlier, the data set was very heterogeneous. This exercise indicated that we needed to further clean the data sets. The most important thing we identified, however, was an error in the initial experiments where we had been splitting the labelled sentences into training and validation sets randomly, with 85% of examples used for training and 15% for validation. We corrected this by splitting the time-dependent data set according to publication date, such that the training set strictly preceded the test data. As shown in the Figure 11, we were unable to obtain a signal that would be significantly above the majority class baseline when we corrected the split (keeping the same 85–15 % ratios). Following this realization, probing the earlier signal by for instance varying the threshold parameter would now not yield any results.

B. Market-Close Data Sets

We decided that in order to make progress in finding a new signal, we would more closely follow the methods and data sets used in the literature. For this, we were required to switch to the task of intra-day price modeling, for which we could find relevant points of reference in the literature.

We expanded the time period beyond the three months in 2015 that the proprietary high frequency data set covered to encompass the years 2011-2014. This gave us better opportunity to heavily clean the data sets while still having more training data. In the literature, most authors use readily available daily market close data to label the sentences, thus we scraped data from the same commonly mentioned source, namely Yahoo Finance. The labelling of the sentences was also very straight-forward as the financial time series is on a daily resolution, so our data set would be constructed in a manner closely matching what was already published by others. The new data set consisted of 261,842 sentences in total, which when filtered for news articles exclusively gave us 47,091 sentences. We also investigated filtering by finding only articles which contained the keyword "liquidity". Both the complete 2011-2014 training set and the more filtered training sets were used to train a bag-of-words neural network model as well as an LSTM model. The results of these experiments, however, did not yielded any new signal above the majority class baseline.

VI. DISCUSSION AND CONCLUSIONS

In this report, we have outlined a study in which we set out to explore the relationship between financial news from Reuters mentioning FTSE-100 companies and liquidity shocks involving their stocks at a high-frequency level. We set out to adapt the methodology of Rönqvist and Sarlin [5] to the prediction of liquidity shocks, in order to link mentions in news with shocks events as defined by Danielsson et al. [15]. In the experiments, we encountered a fundamental challenge in the mismatch between minute-resolution news and millisecond-resolution market events. On the one hand, this mismatch makes it difficult to reliably define a time span in which a piece of news could be expected to have an effect on the market. On the other hand, due to the relative prevalence and brevity of the shocks, the matching yields orders of magnitudes more labeled events than news, which makes it impossible to estimate any function of a mention to distinguish multiple events.

In order to address this problem, we conducted a number of experiments directed at reducing the number of events in order to successfully model the relationship from mention to shock event. These experiments included aggregating individual shocks into continuous shock episodes based on their temporal clustering, and filtering these episodes based on estimated severity and associated price movements. These investigations were nevertheless unsuccessful in yielding an experimental setup that would be able to identify a significant predictive signal. This negative result can be compared to that of [16] who were also investigating the impact of news articles on market liquidity, and likewise did not find any signal.

We carried out experiments on intra- and inter-day price and spread movements prediction based on the same news data and text processing methods. In an attempt to validate this part of the setup, we begun comparing our methods against the available literature on this topic. We were unable

to identify a significant signal in this setup. Comparing our inter-day training set and machine learning models with those of Ding et al. [9], it is worth mentioning that we use a similar baseline model (Bag of words and feed-forward neural network). Any more advanced models, such as the LSTM we trained, or the Convolutional Neural Network used by Ding et al., will not learn anything if a baseline method fails to learn. Thus, future work should rather focus on improving the general setup than on more advanced machine learning models. Furthermore, comparing with Ding et al., we are using comparable text data sets as well as financial data sets, from similar sources, and with the same time resolution. The only difference being that Ding et al. studied the S&P 500. Future experiments in the vein of trying to validate our methods could benefit from performing aggregation on the input side, which would follow the related work more closely than we have done coming from the high-frequency data setting. Ding et al. and others [11, 12] found the best results from using only the headlines, rather than sentences taken from the content of the articles. This could be attempted. However, the authors did also report a (weaker) signal when using the content. Validating the setup using the same textual training set as Ding et al. is also possible, as it is available as an open source project on Github.

As future directions, we see that further investigation into methods for filtering of shock episodes and analysis of the effects of such methods could yield positive results. A necessary parallel line of inquiry would be to reexamine the shock event from a domain perspective and possibly redefine them, in order to embed more meaningful means of filtering at an earlier stage based on refined notions of severity, intensity, etc. Such an inquiry could also take into account possible ways to a priori distinguish between exogenous, event-driven shocks and endogenous or self-reinforcing shocks.

REFERENCES

- [1] E. Fama, "Efficient capital markets: A review of theory and empirical work.," 1970.
- [2] J. Brogaard, T. Nguyen, T. Putnins, and E. Wu, "What moves stock prices? the role of news, noise, and information," *SSRN Electronic Journal*, 01 2018.
- [3] B. Huang, Y. Huan, L. Xu, L. Zheng, and Z. Zou, "Automated trading systems statistical and machine learning methods and hardware implementation: a survey," *Enterprise Information Systems*, vol. 13, pp. 1–13, 07 2018.
- [4] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [5] S. Rönnqvist and P. Sarlin, "Bank distress in the news: Describing events through deep learning," *Neurocomputing*, 2017.
- [6] F. Hogenboom, *Automated Detection of Financial Events in News Text*. PhD thesis, Erasmus Research Institute of Management, December 2014. No. EPS-2014-326-LIS. ERIM Ph.D. Series Research in Management.
- [7] R. Schumaker and H.-c. Chen, "Textual analysis of stock market prediction using breaking financial news: The azfin text system," *ACM Trans. Inf. Syst.*, vol. 27, 02 2009.
- [8] X. Ding, Y. Zhang, T. Liu, and J. Duan, "Deep learning for event-driven stock prediction," in *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI'15*, pp. 2327–2333, AAAI Press, 2015.

- [9] X. Ding, Y. Zhang, T. Liu, and J. Duan, "Using structured events to predict stock price movement: An empirical investigation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, (Doha, Qatar), pp. 1415–1425, Association for Computational Linguistics, Oct. 2014.
- [10] X. Ding, Y. Zhang, T. Liu, and J. Duan, "Using structured events to predict stock price movement: An empirical investigation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, (Doha, Qatar), pp. 1415–1425, Association for Computational Linguistics, Oct. 2014.
- [11] M. Vargas, B. Lima, and A. Evsukoff, "Deep learning for stock market prediction from financial news articles," pp. 60–65, 06 2017.
- [12] "Stock market prediction with deep learning: A character-based neural language model for event-based trading," in *Proceedings of Australasian Language Technology Association Workshop*, pp. 6–15, 2017.
- [13] H. Liu, "Leveraging financial news for stock trend prediction with attention-based recurrent neural network," *ArXiv*, vol. abs/1811.06173, 2018.
- [14] Q. Le and T. Mikolov, "Distributed representations of sentences and documents," in *International conference on machine learning*, pp. 1188–1196, 2014.
- [15] J. Danielsson, E. Panayi, G. Peters, and J.-P. Zigrand, "Market resilience," *SSRN Electronic Journal*, 2018.
- [16] P. Gomber, U. Schweickert, and E. Theissen, "Liquidity dynamics in an electronic open limit order book: an event study approach," *European Financial Management*, vol. 21, no. 1, pp. 52–78, 2015.

VII. ACKNOWLEDGEMENTS

We gratefully acknowledge financial support from the "Digging into High Frequency Data" project funded by the Trans-Atlantic Platform for the Social Sciences and Humanities. The authors also thank EUROFIDAI for providing the BEDOFIH database.