Text Mining for Sentiment Analysis in Bond Portfolio Construction

Thanos Konstantinidis, Maya Golan, Yao Lei Xu, Tony G. Constantinides, Danilo P. Mandic
Department of Electrical and Electronic Engineering
Imperial College London
London, United Kingdom
Email: {a.konstantinidis16, maya.golan18, yao.xu15, a.constantinides, d.mandic}@imperial.ac.uk

***Abstract*—The Efficient Market Hypothesis stipulates that financial instruments within a market perfectly reflect all public information. As the predominant method by which information is disseminated, it is considered that information contained within news influences the value of financial instruments. Thus, if the Efficient Market Hypothesis holds, then understanding the implications of such information is vital in informing portfolio construction. Sentiment analysis is the attempt of a machine to quantify the opinion of natural language. Therefore, this paper visits the Efficient Market Hypothesis by using a novel approach of analysing the relationship between financial news sentiment and corporate bonds in the time frame between 2015 and 2021. Several methods of sentiment analysis are employed to compare the efficacy of traditional lexicon-based knowledge extraction methods and the novel Transformer model approach. Furthermore, the short- and long-term impact of news is evaluated by comparing lagging and decaying signals. The outputs yielded from the sentiment analysis are used as a parameter in portfolio construction ant their returns are examined in light of Sharpe's Ratio to find success in demonstrating the continued viability of the Efficient Market Hypothesis, and the strong relationship between this theory and modern methods of financial sentiment analysis. Ultimately, this paper yields literature leading results with a Sharpe Ratio of 2.1.**

I. Introduction

Modern society is trending towards market transparency as a result of the new capacity for information-sharing brought by the digital age [1]. In particular, financial news sources and social media platforms have bridged people and companies together by providing for the near instantaneous sharing of wide-reaching information to a highly expansive and diverse international audience. This has enabled a channel of communication between corporations and investors which was inconceivable prior to the Internet Age. Technological developments have also afforded retail and institutional investors the opportunity to respond to this information immediately. Such developments therefore present challenges to investors seeking to gain a competitive edge, but can also enable this opportunity. Sentiment analysis has been employed to quantify the opinions presented in textual data, such as a news article, to categorise whether the overall view point is positive, negative, or neutral. When applied to a sufficiently substantial dataset, this can yield an understanding of the overall direction of macroscopic trends. Sentiment analysis of financial information has received growing interest by connecting qualitative and quantitative measures of financial performance [39]. It might therefore be utilized to inform successful portfolio construction, based on news analysis.

This work aims to utilize financial sentiment analysis to effectively determine if corporate bond returns respond to emerging news. Although literature has evaluated whether various instruments change in response to news, this work uses a novel approach by utilizing sentiment analysis to assess the relationship between the opinions expressed in financial news and corporate bond returns, an instrument which has been inadequately explored. Ultimately, the research objectives are: a) Evaluate the effectiveness and performance of four different sentiment analysis methods, which are based on either lexicon-based knowledge extraction or Transformer architectures. b) Determine the immediate impact of news through evaluating the lagging and decaying effect of sentiment over a period of time. This objective distinguishes itself from previous literature, as only sentiment from a particular day was analyzed in existing literature. c) Develop NER models to correctly assign financial textual data to the relevant corpus.

II. Related Work

The application of sentiment analysis in the financial domain has attracted opposing views. Firstly, Fama introduced the notion of the Efficient Market Hypothesis [3] in 1970, which states that stock prices change in response to unexpected fundamental information and news (used as a proxy for information). This hypothesis suggests that text contains opinions and connections that may be utilized to assess trading rules, predict value movement, identify risk, or to corroborate other news [4]. Results were achieved through investigating the impact of news on earning announcements [9], stock splits [10], mergers [11], dividend changes [12], and common stock issuance [13]. News about these factors, whether positive or negative, accorded to the stock price of the corporations included in the studies. Regarding sentiment analysis on textual data, there are two main methods commonly used in literature. The first method requires the representation of lexicon in a domain-specific dictionary. [39][22][23]. The dictionary contains words that are labelled as either positive or negative, and are subsequently used to determine the polarity of the text [18] [25].

The second method utilizes machine learning following a supervised approach [26][27][28]. Furthermore, Support Vector Machines (SVM) and Bayesian regression methods were developed [19] to support Antweiler and Frank's study in

979-8-3503-3959-8/23/$31.00 ©2023 IEEE

the effect of news on trading volume. Moreover, when using social media as a proxy for information, a Twitter specific lexicon [40], in combination with the Dynamic Architecture for Artificial Neural Networks (DAN2) machine learning model [41], produces sentiment classification with higher accuracy when compared to the SVM approach. Deep-learning methods were developed for sentiment analysis using a cascade of multiple layers of non-linear processing units for complex feature extraction and transformation; whereby each successive layer utilizes the output from the preceding layer as an input, consequently extracting complex features [29]. Attention mechanisms and the Transformer were further developed, omitting the use of recurrence and convolutions, traditionally used in deep-learning tasks [30]. Ultimately, Natural Language Processing (NLP) Transformer Architectures stemmed from these methods to develop the Bidirectional Encoder Representations from Transformer (BERT) [42]. Finally, Financial Bidirectional Encoder Representations from Transformer (FinBERT) was developed by pre-training BERT on the Thomson Reuters Text Research Collection (TRC2) and fine-tuning it for the financial text classification task [2].

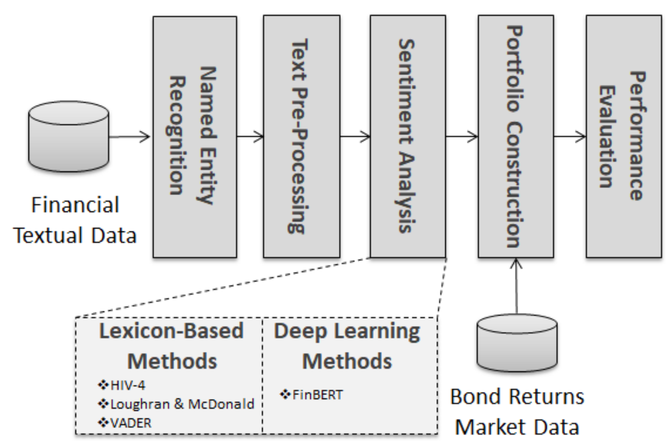

Fig. 1. Framework for sentiment analysis

III. Proposed Framework

This work follows the constructed framework, as observed in Fig. 1, to achieve a thorough understanding of the application of financial sentiment analysis in news articles.

Data Collection: We collected both textual and market data. Textual data (204,017 articles from 2015 to 2021) were collected from Reuters, The Motley Fool, and MarketWatch due to their reputation, reliability, lack of bias, comprehensiveness, and their brand-wide focus on major corporations. Financial market data were collected for the same time period. The collected market data contained bond returns for 1451 companies; thus, the market data pertaining to stocks outside of the Investable Universe were discarded - leaving 500 stocks to be considered. Ultimately, the market data collected consisted of 500 stocks with 1672 days of bond returns data.

Named Entity Recognition (NER): News articles must be identified by directly referencing the organizational entity to be analyzed, and stored accordingly such that each news article is connected to at least one stock [31]. This step reduces the issue of irrelevant news articles being connected to a particular stock [32]. A naive approach for NER defines an entity and associates the particular news article with that entity. In this work, the SpaCy [34] library, was used to count the instances of organizational entities and consequently, the BERT-Base-NER model [35], recognizing four types of entities: location, organization, person, and miscellaneous. If the news article did not meet the requirements defined in the model, the particular news article was removed from the stock's corpus. Thus, the BERT-Base-NER algorithm was applied on the corpus for each company. The algorithm's threshold to ensure that the article related to the company was a 98

Text Pre-Processing: The news article is represented as a "bag-of-words" and we then performed the following steps: a) Tokenization, b) Stop-Word Removal c) Lemmatization d) (Lower) Case Normalization, e) Feature Selection. Regarding Feature Selection, the frequency of each word was used, only for lexicon-based knowledge extraction (namely LMD [25], HIV-4[24], VADER).

Sentiment Analysis: Four sentiment analysis methods were applied. LMD and HIV-4 were implemented using the pysentiment2 Python library, VADER [28] using the NLTK library and FinBERT using the transformers library. Each method was conducted on every article in each company's corpus. If multiple articles were published on a single day, the average sentiment was computed for that day, as defined in (1). St defines the average sentiment for that particular t-th day, Nt is the quantity of news articles published on that particular t-th day, SVit defines the sentiment value of the i-th news article on that particular t-th day.

$$S_t = \frac{1}{N_t} \sum_{i=1}^{N_t} SV_{it} \tag{1}$$

An average of all the sentiment analysis methods was computed to further identify if the average sentiment of the tools was successful in creating a valuable portfolio. Finally, the outputted sentiment for each company was merged to arrive at the final sentiment data that were utilized as a parameter for the portfolio construction.

Portfolio Construction: Once the sentiment for each method has been defined for each company, the long-short portfolio can be constructed. The sentiment is used as a parameter to identify which company should be in a long or short position to maximize the returns from both the positions. The long-short portfolio was constructed as follows:

Define the Investable Universe: Although the S&P 500 contained 500 companies, the financial textual data did not contain articles pertaining to some of the companies between the test time period of January 2015 to May 2021. Consequently, 417 companies were investigated.

Define the parameter to long and short companies: five parameters investigated to produce the optimal portfolio, including the sentiment derived from LMD, HIV-4, VADER, FinBERT, and an average sentiment of the aforementioned methods. Thus, five possible portfolios were constructed. Companies are subsequently ranked daily according to their sentiment. Companies that did not have sentiment data on a particular day were omitted from the ranking; while companies that had a sentiment value of 0 were included, as they suggested neutral sentiment. Companies with a higher sentiment will outrank companies with a lower sentiment.

Allocation: An equal-weighted portfolio strategy was considered when forming the long-short portfolio. Equal weighting is more frequently utilized by hedge funds in portfolio construction and was consequently the chosen strategy in this report [36]. The percentage of companies on a long and short position was fixed at 10%. This entailed that 10% of the companies that had sentiment data on that particular day held either a long position, or 10% held a short position.

Determine daily returns: For each company the daily return that held a long and short position was defined by the market data on that particular day. The total return of companies that held a long position is defined as in (2). Similarly, the total return of companies that held a short position is defined as in (3). (2) and (3), rlong(i) and rshort(i) represents the individual daily returns of the corporate bonds of the top 10% and bottom 10% performing companies, respectively, for each company i. N represents the number of companies that hold a long and short position on that particular day. For each particular day, the number of companies that hold either long or short position are equal. Consequently, the total return of the portfolio on a particular day is the difference between the average long return and average short return, as observed in (4).

$$\text{Average Long Return} = r_{long} = \frac{\sum_{i=1}^{N} r_{long}(i)}{N} \tag{2}$$

$$\text{Average Short Return} = r_{short} = \frac{\sum_{i=1}^{N} r_{short}(i)}{N} \tag{3}$$

$$\text{Total Return} = r_{long} - r_{short} \tag{4}$$

Portfolio Evaluation: The final step is to evaluate the performance of the constructed portfolio using the Sharpe Ratio [43]. A 10% long-short portfolio was fixed for this report. 10% was chosen to ensure that the results across experiments are comparable. Despite this, the control test used a 20% long-short strategy; as the control data obtained utilized a 20% strategy, initial results used a 20% strategy to remain comparable with the control.

IV. Experimental Results

A. Control Test and Benchmark

To evaluate the four previously outlined sentiment analysis methods it was first crucial to evaluate the performance of the constructed portfolio for the specified period of time against a control. In previous literature [37], an annualised Sharpe Ratio of 1.3 was determined for the period between January 2015 and May 2021 while using the LMD dictionary and a 20% long-short portfolio. This control similarly extracts financial textual data from The Motley Fool, Reuters, and MarketWatch. Consequently, this Sharpe Ratio was used as a suitable benchmark to determine whether the computed portfolio was successful. This also enabled consideration of whether the amendments to NER methods utilised in previous literature were useful improvements. An initial test for the sentiment analysis was conducted on the determined relevant news articles using the LMD method and a 20% long-short portfolio, as demonstrated in Fig. 2

Fig. 2 shows striking results during turbulent economic periods and real-world strife, consequently demonstrating that the models developed have shown successful results. For example, there are high levels of standard deviation (volatility) between 2015 and 2016, and the first quarter of 2020. At the beginning of 2016, Donald Trump was inaugurated as President of the United States following a 2015 election, and the results of the 2016 Brexit referendum were announced;

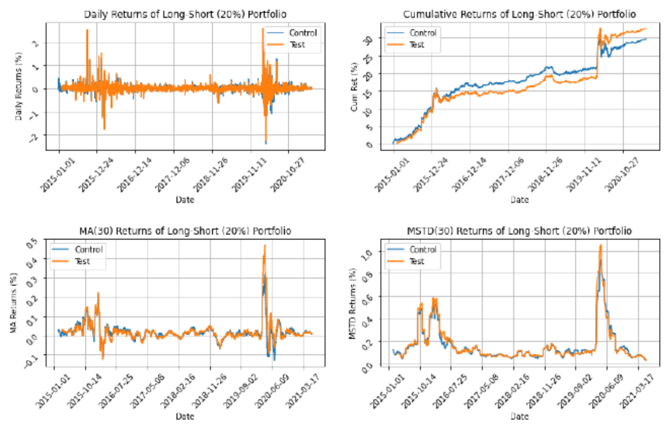

Fig. 2. A comparison of the initial long-short (20%) portfolio constructed using the LMD sentiment analysis method and the control, from the time period of Jan 2015 to May 2021.

It was further crucial to examine the disparities in the characteristics between the control test and obtained results. TABLE I. further demonstrates that the characteristics are similarly matched, with minor disparities.

TABLE I. Comparison between the characteristics of the portfolio between the control test and initial results using the LMD dictionary and a 20% long-short portfolio.

	Control	Initial Test
Sharpe Ratio	1.3	1.4
Cumulative Return	29.53%	32.06%
Avg. Daily Return (%)	0.0179%	0.01936%
Avg. Daily Volatility (%)	0.216%	0.230%

B. Portfolio Construction – Baseline, Lagging & Decaying Experiments

Subsequently to the control test conducted in Section IV-A, a baseline long-short portfolio was created for the four previously outline sentiment analysis method. The baseline experiment was defined as the long-short portfolio reflecting sentiment on that particular day. The baseline portfolio was constructed, as defined in Section III, and its results is observed in Fig. 3. The HIV-4 model performs the least successfully, when compared to other cumulative returns, as the HIV-4 dictionary is a general purpose dictionary, rather than a dictionary proposed for the financial domain. Ultimately, the results observed in TABLE II. suggest that the 10% long-short portfolio constructed using FinBERT method is most successful, with the highest Sharpe Ratio of 2.1 for the baselive experiment. This Sharpe Ratio of 2.1 exceeds the benchmark Sharpe Ratio of 1.3, which is already an improvement on results obtained from previous literature. Furthermore, FinBERT achieves a higher cumulative and average daily returns when compared to the other sentiment analysis methods. This is expected as FinBERT is based on the BERT model, which is pre-trained on financial textual data. However, TABLE II. suggests that the LMD sentiment analysis method has the lowest average daily volatility. This implies that the LMD method computes a portfolio with a lower overall risk. However, this suggests that the portfolio constructed using LMD can expect a lower return over time. Although the Efficient Market Hypothesis states that stock prices will immediately reflect all new publicly available information, the market may still under-react to new publicly available information, and may follow patterns of momentum trading [14][33]. Momentum strategies are bets on past returns predicting the cross-section of future returns [38]. Consequently, the time lag in news sentiment reflects on the impact of news on the market over time.

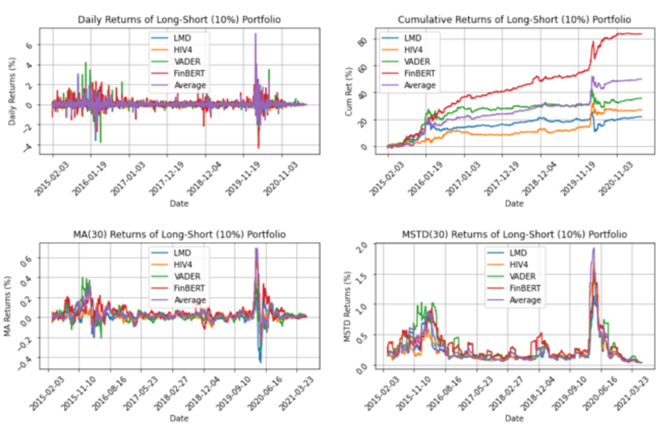

Fig. 3. A comparison of the long-short (10%) portfolio constructed using the various sentiment analysis methods method, from the time period of Jan 2015 to May 2021.

Next, a portfolio was constructed using a lag of one business day to examine the counter hypothesis to the Efficient Market Hypothesis and investigate if the market under-reacts to new publicly available information. News sentiment was lagged according to the following principles. Firstly, all news sentiment was shifted by one day. Secondly, if the news sentiment fell on a non-business day the news sentiment would be shifted to the following business day as this is the only times when the markets will be open to enable an investor to react. The outcomes for the four sentiment methods are similar to the baseline method, however, the Sharpe Ratio of all the evaluated sentiment analysis methods observes a sharp decline of, on average, approximately 42%, a 44% drop of cumulative returns, a 44% drop of average daily returns, and a 3% drop of average daily volatility. This suggests that the market is not misaligned with new publicly available information. In the third experiment, the time series sentiment data accounts for the potential influence of decaying news. The decaying sentiment was found for every no news day. A sample sentiment distribution was computed between January 2021 and May 2021. A long-short portfolio was created using the four previously outlined sentiment analysis methods with a decaying half-life of 2 days. The expected results (TABLE II.), demonstrates that the FinBERT method outperforms the other sentiment analysis method with the highest observed cumulative return.

TABLE II. STATISTICAL COMPARISON BETWEEN THE 4 SENTIMENT ANALYSIS METHODS USING A 10% LONG-SHORT PORTFOLIO WITH AND WITHOUT A LAG OF ONE BUSINESS DAY, AND WITH A DECAYING HALF-LIFE OF 2 DAYS

		LMD	*HIV-4*	*VADER*	*Fin BERT*	*Average*
Baseline	Sharpe Ratio	0.7	0.8	0.9	**2.1**	1.3
	Cumulative Return	21.84	26.97	35.57	**83.66**	49.92
	Avg. Daily Return (%)	0.013	0.016	0.021	**0.050**	0.030
	Avg. Daily Volatility (%)	**0.296**	0.321	0.394	0.388	0.376
Lagging	Sharpe Ratio	0.4	0.3	0.6	**1.2**	0.9
	Cumulative Return	11.51	9.99	23.97	**47.18**	30.00
	Avg. Daily Return (%)	0.007	0.006	0.014	**0.028**	0.018
	Avg. Daily Volatility (%)	**0.307**	0.308	0.393	0.393	0.326
Decay (2days)	Sharpe Ratio	1.2	0.0	0.3	1.7	1.5
	Cumulative Return	24.720	0.080	0.451	**32.890**	28.540
	Avg. Daily Return (%)	0.015	0.000	0.003	**0.020**	0.017
	Avg. Daily Volatility (%)	0.193	0.160	**0.156**	0.190	0.183

V. CONCLUSION

This work identified the relative effectiveness and performance of four different sentiment analysis methods: LMD, HIV-4, VADER, and FinBERT. The efficacy of the Transformer model in the financial domain demonstrated outstanding results in accurately recovering sentiment from a textual dataset. Not only was the objective of evaluating the performance of various sentiment analysis method met, but the constructed portfolio yielded superior results to existing literature. The developed FinBERT model yielded a Sharpe Ratio exceeding previous iterations by approximately 170% in constructed portfolios solely based on sentiment. Ultimately, this report is believed to be the first attempt at providing a comprehensive review of sentiment analysis approaches on corporate bonds. Secondly, this report aimed to successfully evaluate the impact of news on corporate bond returns. While news sentiment holds a short-term decaying signal, it is not able to forecast corporate bond returns in the future. Although this short-term predictive signal of news sentiment has been identified in literature, this report introduces a novel contribution to existing research using corporate bonds as the instrument of focus. Finally, the objective to perform NER was met as articles pertaining to relevant firms were correctly identified through using the BERT-base-NER model.

REFERENCES

[1] C. Biancotti and P. Ciocca, "Financial markets and social media: Lessons from information security," Carnegie Endowment for International Peace, 2022.

[2] D. Araci, "FinBERT: Financial Sentiment Analysis with Pre-Trained Language Models," Computing Research Repository, 2019.

[3] E. F. Fama, "Efficient capital markets: A review of theory and empirical work," The Journal of Finance, vol. 25, no. 2, pp. 383–417, 1970.

[4] S. R. Das, A. Martínez-Jerez, and P. Tufano, "eInformation: A Clinical Study of Investor Discussion and Sentiment," Financial Management, vol. 34, no. 3, pp. 103–137, 2005.

[5] S. R. Das and J. Sisk, "Financial communities," The Journal of Portfolio Management, vol. 31, no. 4, pp. 112–123, 2005.

[6] D. Godes, D. Mayzlin, Y. Chen, S. R. Das, C. Dellarocas, B. Pfeiffer, B. Libai, S. Sen, M. Shi, and P. Verlegh, "The firm's management of social interactions," Marketing Letters, vol. 16, no. 3, pp. 415–428, 2005.

[7] F. Li, "Do stock market investors understand the risk sentiment of corporate annual reports?," SSRN Electronic Journal, Apr 2006.

[8] Y. V. Hochberg, A. Ljungqvist, and Y. Lu, "Whom you know matters: Venture capital networks and investment performance," The Journal of Finance, vol. 62, no. 1, pp. 251–301, 2007.

[9] R. Ball and P. Brown, “An empirical evaluation of accounting income numbers,” Journal of Accounting Research, pp. 159–178, 1968.

[10] E. F. Fama, L. Fisher, M. Jensen, and R. Roll, “The adjustment of stock prices to new information,” International Economic Review, vol. 10, no. 1, 1969.

[11] G. Mandelker, “Risk and return: The case of merging firms,” Journal of Financial Economics, vol. 1, no. 4, pp. 303–335, 1974.

[12] J. Aharony and I. Swary, “Quarterly dividend and earnings announcements and stockholders’ returns: An empirical analysis,” The Journal of Finance, vol. 35, no. 1, pp. 1–12, 1980.

[13] P. Asquith and D. W. Mullins Jr, “Equity issues and offering dilution,” Journal of Financial Economics, vol. 15, no. 1-2, pp. 61–89, 1986.

[14] W. S. Chan, “Stock price reaction to news and no-news: Drift and reversal after headlines,” Journal of Financial Economics, vol. 70, no. 2, pp. 223–260, 2003.

[15] P. C. Tetlock, “Giving content to investor sentiment: The role of media in the stock market,” The Journal of Finance, vol. 62, no. 3, pp. 1139–1168, 2007.

[16] P. C. Tetlock, M. Saar-Tsechansky, and S. Macskassy, “More than words: Quantifying language to measure firms’ fundamentals,” The Journal of Finance, vol. 63, no. 3, pp. 1437–1467, 2008.

[17] A. R. Admati and P. C. Pfleiderer, “Noisytalk.com: Broadcasting opinions in a noisy environment,” Research Papers 1670, Stanford University, Graduate School of Business, Aug 2001.

[18] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, “Persuasion bias, social influence, and unidimensional opinions,” The Quarterly Journal of Economics, vol. 118, no. 3, pp. 909–968, 2003.

[19] W. Antweiler and M. Z. Frank, “Is all that talk just noise? the information content of internet stock message boards,” The Journal of Finance, vol. 59, no. 3, pp. 1259–1294, 2004.

[27] S. R. Das and M. Y. Chen, “Yahoo! for Amazon: Sentiment extraction from small talk on the web,” Management Science, vol. 53, no. 9, pp. 1375–1388, 2007.

[20] L. Mitra, G. Mitra, and D. Dibartolomeo, “Equity portfolio risk estimation using market information and sentiment,” Quantitative Finance, vol. 9, no. 8, pp. 887–895, 2009.

[21] D. Leinweber and J. Sisk, “Event-driven trading and the “new news”,” The Journal of Portfolio Management, vol. 38, pp. 110–124, Oct 2011.

[22] J. L. Rogers, A. Van Buskirk, and S. L. Zechman, “Disclosure tone and shareholder litigation,” The Accounting Review, vol. 86, no. 6, pp. 2155–2183, 2011.

[23] A. K. Davis, W. Ge, D. Matsumoto, and J. L. Zhang, “The effect of manager-specific optimism on the tone of earnings conference calls,” Review of Accounting Studies, vol. 20, no. 2, pp. 639–673, 2015.

[24] P. J. Stone, D. C. Dunphy, and M. S. Smith, “The General Inquirer: A Computer Approach to Content Analysis.,” The American Journal of Sociology, 1966.

[25] T. Loughran and B. McDonald, “When is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks,” The Journal of Finance, vol. 66, no. 1, pp. 35–65, 2011.

[26] S. F. Crone and C. Koeppel, “Predicting exchange rates with sentiment indicators: An empirical evaluation using text mining and multilayer perceptrons,” in 2014 IEEE Conference on Computational Intelligence for Financial Engineering and Economics (CIFEr), pp. 114–121, IEEE, 2014.

[27] M. Y. Kaya and M. E. Karsligil, “Stock price prediction using financial news articles,” in 2010 2nd IEEE International Conference on Information and Financial Engineering, pp. 478–482, IEEE, 2010.

[28] C. Hutto and E. Gilbert, “VADER: A parsimonious rule-based model for sentiment analysis of social media text,” in Proceedings of the International AAAI Conference on Web and Social Media, vol. 8, 2014.

[48] M. Ghiassi and S. Lee, “A domain transferable lexicon set for twitter sentiment analysis using a supervised machine learning approach,” Expert Systems with Applications, vol. 106, pp. 197–216, 2018.

[29] O. Araque, I. Corcuera-Platas, J. F. Sánchez-Rada, and C. A. Iglesias, “Enhancing deep learning sentiment analysis with ensemble techniques in social applications,” Expert Systems with Applications, vol. 77, pp. 236–246, 2017.

[30] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, and E. Hovy, “Hierarchical attention networks for document classification,” in Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pp. 1480–1489, 2016.

[31] R. Feldman, B. Rosenfeld, R. Bar-Haim, and M. Fresko, “The Stock Sonar – Sentiment Analysis of Stocks Based on a Hybrid Approach,” Proceedings of the National Conference on Artificial Intelligence, vol. 2, Jan 2011.

[32] J. Boudoukh, R. Feldman, S. Kogan, and M. Richardson, “Which news moves stock prices? a textual analysis,” Working Paper 18725, National Bureau of Economic Research, Jan 2013.

[33] M. Baker and J. Wurgler, “Investor sentiment in the stock market,” Journal of Economic Perspectives, vol. 21, no. 2, pp. 129–152, 2007.

[34] M. Honnibal and I. Montani, “spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing.” To appear, 2017.

[35] D. S. Lim, “BERT-base-NER.” https://huggingface.co/dslim/bert-base-NER, 2021.

[36] Z. T. Ke, B. T. Kelly, and D. Xiu, “Predicting returns with text data,” NBER Working Papers 26186, National Bureau of Economic Research, Inc, 2019.

[37] R. K. Jain, “Natural language based recommender systems for financial modelling,” Master’s thesis, Imperial College London, 2021.

[38] K. Daniel and T. J. Moskowitz, “Momentum crashes,” Journal of Financial Economics, vol. 122, no. 2, pp. 221–247, 2016.

[39] K. Mishev, A. Gjorgjevikj, I. Vodenska, L. T. Chitkushev, and D. Trajanov, “Evaluation of sentiment analysis in finance: From lexicons to transformers,” IEEE Access, vol. 8, pp. 131662–131682, 2020.

[40] M. Ghiassi and S. Lee, “A domain transferable lexicon set for twitter sentiment analysis using a supervised machine learning approach,” Expert Systems with Applications, vol. 106, pp. 197–216, 2018.

[41] M. Ghiassi and H. Saidane, “A dynamic architecture for artificial neural networks,” Neurocomputing, vol. 63, pp. 397–413, 2005. New Aspects in Neurocomputing: 11th European Symposium on Artificial Neural Networks.

[42] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” Computing Research Repository, 2018.

[43] J. B. Berk and P. M. DeMarzo, Corporate Finance. Pearson, 2019.