The Effects of Donald Trump's Tweets on US Financial and Foreign Exchange Markets

By Constantin Colonescu^{*}

Twitter is the US president's, Donald Trump, preferred media for communicating his thoughts to his followers. This project looks at the effect of the daily flow of Donald Trump's tweets on the US financial and foreign exchange markets, represented by the Dow Jones Industrial Average (DJI) index and some exchange rates, over the period of Donald Trump's ongoing presidency. Using text mining techniques, some correlation is found between various moving average window lengths of tweet content and the DJI index. Some short term and lasting effects are also detected on US-Canada and US composite exchange rates.

Keywords: Donald Trump, financial markets, text mining, tweets

Introduction

Some observers of financial markets point out that President Donald Trump's tweets create short term volatility in the share prices of one or another company. In this regard, FXCM (2018) gives examples of companies such as Toyota, Lockheed Martin, Boeing, and General Dynamics that have been in the past affected by the president's tweets. Otani and Shifflett (2017) compile a "Trump Target Index" to follow the stock market evolution of the companies that Donald Trump has ever mentioned, mostly with a negative connotation. The authors find that companies have quickly recovered after a brief period of volatility in their stock. Another source, Seeking Alpha (2017), maintains a "Trump Index" that includes companies presumably favored by the protectionist policy promoted by President Trump, companies in the steel, banking, and heavy machinery sectors. The authors find that these companies have been performing much better than the market after Donald Trump's accession to the White House on his "America first" platform.

Most of the existing literature seeks to relate financial market variables, such as stock returns, to an aggregate measure of investor sentiment, based on a large number of individuals; there is no attempt in the literature, as far as I know, to investigate the effect of one personality's opinion, action, or rhetoric on financial markets in general. This article seeks to identify a possible correlation between the president of the United States' declarations and the Dow Jones Industrial Average, an aggregate measure of the state of US financial markets. The effect of Donald Trump's rhetoric on some exchange rates of the American dollar is also investigated.

^{*}Associate Professor, MacEwan University, Canada.

Donald Trump's Economic Credo

Donald Trump assumed the presidency of the United States on January 1st, 2017. His vision is one of economic nationalism, which entails mercantilism and protectionism in international relations, limiting immigration, discouraging domestic companies from employing foreign labor, deregulating economic activities to increase the competitiveness of American companies in international markets, and lowering taxes for large entrepreneurs. Along these lines, the Trump administration withdrew the United States from several international free trade agreements, stopped some ongoing negotiations, and started renegotiating other political and economic agreements. Prominent cases are the negotiations for a free trade agreement with the European Union, NAFTA, and several bilateral agreements such as the one with China. A typical act of Trumpian mercantilism is the imposition of a 25% tariff on imported steel and 15% on aluminum, which the Trump administration enacted on March 1st, 2018 and extended to Canada, Mexico, and the European Union in May 2018.

How would financial markets respond to mercantilist and protectionist decisions and in particular what would be the effect of such decisions on the Dow Jones Industrial Average (DJIA) index? The DJIA index is the sum of the stock prices of 30 major American companies, of which very few, such as Caterpillar, Boeing, and General Electric may still have a manufacturing component. Most of the other companies come from various sectors, such as pharmaceuticals, banking, finance, computer technology, and software. Thus, some decisions, as the imposition of a tariff on steel and aluminum imports may not have a significant direct effect on the DJIA. However, in the medium term the effects could spill over other sectors, some included in the DJIA, as trade partner countries retaliate to tariffs. For instance, both Canada and the European Union announced on the same day of the imposition of the tariff retaliatory measures that concerned a wide range of products.

Literature Review

Using Social Media to Forecast Stock Prices

Measuring market sentiment based on social network data and seeking the effect of sentiment on financial markets is a recent, but not a new trend in the literature. Several authors investigate the effect of social media such as Twitter on financial markets, gathering data from not one, but many social media users. For instance, using an event study methodology and data that span a 16-month period, Ranco, Aleksovski, Caldarelli, Grčar, and Mozetič (2015) find a small but significant short-run effect of sentiment spikes on 30 stocks of the DJI index.

In the same vein of short term effects, P. Papaioannou, Russo, Papaioannou, and Siettos (2013) use Twitter information to model and predict high frequency daily fluctuations of the EUR/USD exchange rate. They notice, first, that more and more studies challenge the *efficient market hypothesis* and suggest including

information gathered from social media into time series models to forecast exchange rates. Unlike other studies, though, this one uses specific numbers collected from the tweets, instead of assessing sentiment or opinion. The study finds that such information can improve the capacity of certain models to predict short term, intra-day exchange rates between the euro and the US dollar.

Bollen, Mao, and Zeng (2011) ask a more general question: does society-wide sentiment affect the economy? The authors gather Twitter information using Opinion Finder and Google-Profile of Mood States, two text mining engines for determining mood content in tweets. The former instrument provides only a positive versus negative scale, while the latter measures mood along a few dimensions. Like in the article at hand, Bollen et al. (2011) use the Dow Jones Industrial Average stock index to measure market changes, but they rely on closing values instead of daily returns. The results also show that including mood variables in time series models improves short run prediction of stock prices.

Zhang, Fuehres, and Gloor (2011) use a sample of tweets collected over a period of six months for the same purpose as many other authors such as Bollen et al. (2011) discussed above, namely to determine a correlation between tweet sentiment and some aggregated market variables. Besides the Dow Jones Industrial Average index, though, Zhang et al. (2011) also consider NASDAQ, S&P 500, as well as the VIX index, which is a measure of market volatility. The authors find a negative correlation between emotion in general (whether it is positive or negative) and the market price aggregates, but a positive correlation with the market volatility index. A shortcoming of this study may be the apparently limited lexicon that the authors use for gauging sentiment.

Nguyen, Shirai, and Velcin (2015) use both topics and sentiment to determine the effect of market mood on a set of 18 stocks, many of which are also parts of the DJIA index. For the source of mood data, the authors use Yahoo Finance Message Board.

Literature on Financial Text Mining

The major challenge in text mining for financial modeling is that this kind of data, i.e., text written by a variety of authors, is amorphous, unstructured. Text data often contain very important information, but this information is difficult to separate from the noise of jargon, misspelled words, nonstandard abbreviations, a wealth of symbols, and a large set of irrelevant words and messages. For example, a text that is in general negative may contain a small positive part that is the only relevant part for a financial analyst's purpose. To add to the difficulty of the task, text can come on a wide variety of supports, formats, and languages; perhaps all the existing text analysis tools require exporting content from all these various sources into a structured electronic text format with a unique encoding. Encoding is necessary for transforming human readable characters into machine-specific, binary code.

The simplest method, but surprisingly efficient for text analysis is the "bag of words" method, which tries to classify texts based on counting significant words and sorting these words in meaningful categories. This method, however, has severe limitations when used for more specific purposes; one obvious shortcoming is that a positive word can become negative by adding "not" in front of it, a subtlety that the bag of words methods tends to miss.

There is already a rich collection of works trying to provide better methods of detecting and gauging mood for financial analysis purposes out of written text. For example, noting the shortcomings of the "bag of words" method of text analysis, Chan and Chong (2017) develop a "sentiment analysis engine" (SAE) that tries to account for syntactic substructures (parts of speech) in text data. The bag of words method, though, still remains popular because it can be applied to diverse problems and requires the least human intervention.

Methodology

Briefly, our method consists in constructing a time series of the sentiment content of Donald Trump's tweets and then uses some econometric method (time series analysis) to relate sentiment to a representative market index, exchange rate, or other economic variable of interest.

The data for compiling sentiment in Donald Trump's tweets come from Trump Twitter Archive (2018), a collection of about 33.000 tweets, all the tweets that Donald Trump wrote since 2009, for many years before his advent to the presidency of the United States. This study, though, only concerns the about 3500 tweets Donald Trump produced since January 2017, when he became the president, until May 10, 2018.

Data Description



Figure 1. Number of Tweets over Time and Over the Day

Figure 1 depicts the distribution of the number of tweets over the period since Donald Trump assumed presidency. It shows a noticeable variation, with lows in spring and highs in summer and fall. The second panel shows the average number of tweets by the time of the day, indicating that Donald Trump's preferred time for tweeting is a couple of hours about noon, with a thick tail towards evening and midnight hours. While the daily distribution of twitter activity may be interesting in itself, or important when studying very short term, hourly effects, it may not be relevant for an investigation of longer period interactions.

Building the Sentiment Time Series

We use a text mining method described in Silge and Robinson (2017), based on the **tidyverse** package (Wickham and Grolemund 2017). The purpose of text mining is to assign a sentiment measure (positive or negative) to each tweet and to construct a time series of daily sentiment expressed in tweets. First, each tweet is decomposed in words (the "bag-of-words" method); then, the words are compared to a lexicon, which is a previously annotated dictionary in which each word is assigned a positive or negative sentiment value.

A general-purpose lexicon called AFINN (Nielsen 2011), which has been tested against other similar packages by F. Arup Nielsen (2011), is used to quantify sentiment. AFINN assigns sentiment values in the range of -5 to 5, thus allowing a more accurate quantification of the sentiment content of words than other similar lexicons and is adequate for constructing a time series on a continuous scale.

Figure 2. Positive and Negative Sentiment Values in Tweet Word Cloud



Figure 2 gives a polarized cloud of words, to give a sense of which words are considered positive and negative by the sentiment lexicon. The size of the words in the graph is proportional to the frequency with which they appear in the text. This graph, however, is only shown here for an example, since it has been constructed with the "Bing" lexicon, which uses only two sentiment categories, "positive" and "negative;" the remaining analysis uses the AFINN lexicon described above. A notable miss-assignment in Figure 2 is the word "trump," which in tweets appears extremely frequently and most likely represents the president's name, not the common word. If this is the case, then the word "trump" should probably be considered neutral for sentiment analysis and removed from the tweet list of words. Except for removing a few of such words from the lexicon and adding a few others, no major human intervention has been made on the AFINN lexicon.

This research uses two methods of sentiment aggregation over a day and compares them to determine the robustness of the results. The first method uses the average sentiment of the day and constructs a sentiment measure using the sentiment scores in the AFINN lexicon. The second aggregation method identifies the most extreme sentiment score of the day, which can be either negative or positive, and retains it as the representative sentiment score for the respective day. The next sections report the results separately for each of these two methods.

The Average Sentiment Method

We construct a daily sentiment time series by calculating the average sentiment measure for each day. A time series plot (not shown here) of the daily average sentiment in Donald Trump's tweets displays no trend and its volatility seems to be sizeable but constant; these features suggest that the series is stationary. An augmented Dickey Fuller test is calculated, which indicates no unit root in this series, thus confirming the stationary of the series.

After having built the sentiment time series, the next step is to make available a financial market indicator for the same period. As a measure of financial market behavior, we consider the daily *log* returns of the Dow Jones Industrial Average index, for which the trading symbol is DJI. The data for the DJI index come from Yahoo! Finance, retrieved with the function tq get of the **tidyquant** package (Dancho and Vaughan 2018). A time series plot shows a peculiarity of this series: On 2018-02-05, the DJI index dropped almost 1,175 points, the largest drop in its history; this happened for no apparent reason, while the index was following an increasing trend since the beginning of Donald Trump presidency. This outlier point in the DJI has been removed from our dataset, but the DJI time series has remained noisy and volatile ever since the day of the historic drop. We should expect difficulties in finding a relationship between Donald Trump's tweet sentiment and the DJI over the recent months in the presence of such a persistent noise.

The DJI daily returns series shows some, but very weak autocorrelation as revealed by an autocorrelation function plot (acf), but the Augmented Dickey Fuller test does not detect evidence of nonstationarity. Since none of the two series - DJI and sentiment - is nonstationary (they are both I(0)), we can use linear regression to find a possible relationship between them.



Figure 3. Sentiment and DJI Moving Average Series

Figure 3 shows five-day moving average plots of the two series together, appropriately scaled such that their amplitudes be comparable; the dotted line represents the sentiment, and the solid line represents the market index series. A first observation would be that the two series seem to be contemporaneously negatively correlated: low returns coincide with positive sentiment scores. This observation may suggest either that prices are more stable at times of positive sentiment and thus returns are lower, or that there might be some lag in the two series. One would need to include, in a regression model, autoregressive terms to account for other factors that influence the market but are not in the model, assuming that such factors, which are otherwise important in determining market fluctuations are independent of our other independent variables, which are lags of the sentiment series.

Let us estimate the model described in Equation (1), where (y) is the market variable, daily returns of the Dow Jones index, and (x) is the sentiment at time (t); both sentiment and the market variable are measured as five-day moving averages in an attempt to further reduce noise, but, of course, enhancing the disadvantages of aggregation.

$$y_t = \alpha_0 + \alpha_1 y_{t-1} + \beta_1 x_{t-1} + \beta_2 x_{t-2} + e_t \tag{1}$$

The number of lags in each model has been chosen to minimize the Akaike Information Criterion (AIC) of the model.

	Coefs 1	P-Vals 1	Coefs 2	P-Vals 2	Coefs 3	P-Vals 3
Intercept	0.0002	0.0729	0.0004	0.0381	-0.0006	0.3026
Stock(-1)	0.7620	0.0000	0.8053	0.0000	0.6729	0.0000
Sent(-1)	0.0006	0.0445	-0.0010	0.0149	-0.0003	0.8711
Sent(-2)	-0.0008	0.0078	0.0008	0.0483	0.0024	0.2369

Table 1. Stock on Mean Sentiment over Three Intervals

Findings and Discussion

Table 1 collects the results of the model (1) for three time intervals, which are determined over the whole period by two break points in the stock series. The break points have been found using a method based on Chow tests as described in (Zeileis 2006), by regressing the market index series on a constant.

The results show statistically significant short-term effects of Trump tweet sentiment on the Dow Jones Industrial Average index for the first two time intervals and no significance for the last interval. As Figure 3 shows, the DJI index displays unusually large variability over the last time interval, which makes the detection of the effect of a relatively less important factor such as Trump's tweets more difficult. Although the effects of tweet sentiment on the market come out statistically significant, their importance or magnitude cannot be assessed from the regression results because these results are based on an arbitrary sentiment scale. The results turn out to be very sensitive to the choice of the moving average window; as I have already mentioned, the reported results correspond to a five-day moving average window.

An interesting feature of the sentiment coefficients produced by our regression model is their sign and relative magnitude: during the first interval the first lag is positive and the second negative, while during the second interval their signs are reversed. One may ask if there is an overall lasting effect beyond these short-term fluctuations.

To determine a possible lasting effect, we test the (null) hypothesis that the sum of all sentiment coefficients is equal to zero; the test comes out negative; that is, it does not reject the null hypothesis. This result shows that there is no lasting effect of tweets on the DJI index.

Using the Max and Min Sentiment Scores Instead of Average

Averaging tweet content over a day may not be the best aggregation method because tweets may contain opposing views, some important for the market and some not important. Thus, the effect of an important message may be neutralized by an unimportant one. This section repeats the previous analysis, but using a different aggregating method. Instead of averaging the sentiment score, it retains the most extreme one, either positive or negative, assuming that the strongest message of the day is the one that captures the attention of the public.



Figure 4. Sentiment Extremes and DJI Moving Average Series

 Table 2. Stock on Sentiment with Extreme Sentiment Scores

	Coefs 1	P-Vals 1	Coefs 2	P-Vals 2	Coefs 3	P-Vals 3
Intercept	0.0001	0.3263	0.0003	0.1613	-0.0007	0.2219
Stock(-1)	0.8722	0.0000	0.8683	0.0000	0.7333	0.0000
Sent(-1)	-0.0003	0.0162	-0.0003	0.0853	0.0000	0.9979
Sent(-2)	0.0003	0.0219	0.0003	0.1608	0.0006	0.3517

Table 2 shows the results of the same model and time intervals as before, but with the new sentiment aggregation method. This time, the moving average window that turns out to be significant is of 10 days. As before, a test of the overall significance of the sentiment terms in the regression comes out negative, suggesting no lasting effects of tweet sentiment on the DJI index. This case shows that the results in the previous section are robust to the aggregation method. However, the two methods differ by the length of the moving average window, which may imply that they disagree with regard to the amount of time it takes for the fluctuations induced by the tweets to diminish. Moreover, since the moving average window is wider in the second method, the two methods may not be too different after all in the way they average the sentiment scores.

Tweet Sentiment and Market Variability

Finding so far only a short-term effect of tweets on a market index, one may think that there may be an effect on market variability. The Federal Reserve Bank of St. Louis (Federal Reserves of the U.S. 2018) provides VIXCLS, a market volatility index, which I use here for a measure of market variability in relation with Donald Trump's tweets. The VIXCLS data are based on the volatility index VIX, calculated by Chicago Board Options Exchange (CBOE), which measures fluctuations in S&P 500 options. The VIX index measures, in fact, the expected volatility of the underlying asset. Since the volatility index turns out to have significant autocorrelation, I use its one period difference series. The data is imported from the FRED (Federal Reserves of the U.S., 2018) website, using the R function tq get from the package **tidyquant**.

	Coefs 1	P-Vals 1	Coefs 2	P-Vals 2	Coefs 3	P-Vals 3
Intercept	-0.0002	0.9081	0.0063	0.2063	0.0086	0.2437
Stock(-1)	0.7098	0.0000	0.4657	0.0023	0.6997	0.0000
Sent(-1)	-0.0008	0.6234	0.0077	0.0428	-0.0042	0.4382
Sent(-2)	0.0000	0.9872	-0.0055	0.2275	-0.0053	0.3590
Sent(-3)	-0.0013	0.5790	-0.0021	0.6201	0.0127	0.0376
Sent(-4)	0.0039	0.0869	0.0076	0.1051	-0.0023	0.6969
Sent(-5)	-0.0017	0.3172	-0.0089	0.0251	-0.0109	0.0503

 Table 3. Market Variability on Mean Sentiment over Three Intervals

Table 3 shows the result of regressing the log difference in the market variability index VIXCLS on tweet sentiment, with all variables measured in five day moving averages. The results show, again, some short-term relationship, but a linear hypothesis test does not reveal a lasting effect.

Tweet Sentiment and the Exchange Rate

The previous sections investigated a possible effect of Donald Trump's tweets on an aggregate measure of the US financial market, the DJIA index. Financial markets, though, still benefit from capital mobility across the US border, as opposed to goods and services, which are the objects of an important aspect of Donald Trump's economic vision, international trade. As I have already mentioned, trade is strongly affected by this administration's policies. It is, then natural to ask whether information emerging from the White House under the form of tweets has an effect on the exchange rates between the dollar and other currencies.

The method here is the same as before, just the economic time series are changed. We restrict our attention in this section to the mean sentiment series, since the previous section identified no major change when moving from daily average to daily extreme sentiment. Moreover, when transforming the series into moving averages, the distinction between the two methods becomes less pronounced.



Figure 5. Sentiment and CAD-USD Exchange Rate Moving Averages

I find no evidence of a relationship between tweet sentiment and the US dollar-euro exchange rate, but I find some significant short-term effect of tweets on the US dollar-Canadian dollar exchange rate. Figure 5 shows the five-day moving average plots of the two series: sentiment and exchange rate.

Table 4 gives the regression coefficients, showing statistical significance for two lags in the sentiment series. Again, statistical significance does not necessarily indicate an important effect, given that one lag shows a negative sign and the other a positive one while their absolute values are comparable in magnitude.

The results certainly suggest significant fluctuations, but not necessarily a lasting effect. The linear hypothesis (2), where the notations are as in Equation (1) determines whether a significant lasting effect exists.

$$H_0: \beta_1 + \beta_2 = 0$$
 (2)

	Coef.	Std. Error	t Value	p Value
(Intercept)	0.00000	0.00002	0.20595	0.83692
Lag(y, 1)	0.96808	0.01173	82.53211	0.00000
Lag(x, 1:2)Lag.1	-0.00032	0.00016	-2.04684	0.04127
Lag(x, 1:2)Lag.2	0.00033	0.00016	2.08496	0.03765

Table 4. Regression of Sentiment on USD/CAD Exchange Rate

Using the linear Hypothesis function, we do not find a significant lasting effect.

Investigating Other Series

Another measure of the US dollar exchange rate is "Trade Weighted U.S. Dollar Index: Major Currencies," which has the series code DTWEXM on Federal Reserves' database (Federal Reserves of the U.S. 2018). Table 5 shows the

coefficients and p-Values of regressing this variable on a 15-day moving average of the sentiment variable; the sample is divided in four periods based on a structural change test. A linear hypothesis test has been conducted, which revealed even a rare lasting effect of tweets on this measure of exchange rate for the first two intervals.

	Coef. 1	p-Val. 1	Coef. 2	p-Val. 2	Coef. 3	p-Val. 3	Coef. 4	p-Val. 4
Intercept	0.0001	0.1915	0.0000	0.5917	0.0000	0.3215	0.0000	0.6987
Stock(-1)	0.9378	0.0000	0.9303	0.0000	0.9481	0.0000	0.9487	0.0000
Sent(-1)	-0.0005	0.0775	-0.0004	0.0134	-0.0004	0.0217	-0.0004	0.0175
Sent(-2)	0.0003	0.2736	0.0003	0.0974	0.0003	0.0641	0.0003	0.0672

Table 5. Sentiment on WEXM Exchange Rate, Four Intervals

Conclusion

Advances in social media technology have changed the ways political leaders communicate, a change that has both its enthusiasts and critics. The US president, Donald Trump, has been in the center of this debate since his advent to the White House. Some observers have found short term (daily) effects of Donald Trump's major tweet events on selected individual stock and exchange rates, but no longer have term or lasting effects been investigated. There is also no research, to the best of my knowledge, on the effects of Donald Trump's tweets on an aggregate market variable such a Dow Jones Industrial Average index.

This paper finds some evidence of short term, as well as some persistent effects of twitter announcements by the US president on some financial and foreign exchange aggregates, such as the Dow Jones Industrial Average, the US-Canadian currency exchange rate, and the "Trade Weighted U.S. Dollar Index: Major Currencies," an aggregate US dollar exchange rate index. It also investigates other bilateral exchange rates, but finds no significant effects. These effects are identified in moving average series of various window sizes, since the results seem to be sensitive to the choice of the moving average window. The only lasting effects found are in the case of the US dollar composite exchange rate.

Future research may focus on creating a more specialized lexicon for sentiment quantification, to include words that are normally neutral on a sentiment scale but are not neutral in the current international and economic context. Words like "Syria", "Iran", "immigration", and "tariff" have no particular connotation in a standard sentiment dictionary as the one used in this paper, but may significantly influence the markets. Another potentially fruitful venue of research could be identifying relevant topics, rather than words in the tweet content and using only the tweets related to a certain topic instead of the overall sentiment expressed in tweets, somehow along the lines of Nguyen et al. (2015), who use both topics and sentiment to forecast several individual stocks.

In terms of method, better results may be achieved using machine learning algorithms to identify relevant topics in tweets and to test the prediction power of the model.

Acknowledgements

All calculations have been produced using the statistics application R (R Core Team 2018); for various purposes, the following R packages have been used: **car** (Fox and Weisberg 2011), **tidyquant** (Dancho and Vaughan 2018), **read text** (Benoit and Obeng 2017), **lubridate** (Grolemund and Wickham 2011), **tidytext** (Silge and Robinson 2016), **ggplot2** (Wickham 2009), **word cloud** (Fellows 2014), **quant mod** (Ryan and Ulrich 2018), **strucchange** (Zeileis 2006), and others.

References

- Benoit K, Obeng A (2017) *Readtext: Import and handling for plain and formatted text files.* Retrieved from https://bit.ly/2tJ8RLm.
- Bollen J, Mao H, Zeng X (2011) Twitter mood predicts the stock market. Journal of Computational Science 2(1): 1–8. doi: 10.1016/j.jocs.2010.12.007.
- Chan SW, Chong MW (2017) Sentiment analysis in financial texts. *Decision Support Systems*, 94: 53–64. doi:10.1016/j.dss.2016.10.006.
- Dancho M, Vaughan D (2018) *Tidyquant: Tidy quantitative financial analysis*. Retrieved from https://bit.ly/2KuBP98.
- Federal Reserves of the U.S. (2018) *FRED economic data. FRED, Federal Reserve Bank* of St. Louis. Retrieved from https://bit.ly/2hWtGPP.
- Fellows I (2014) Wordcloud: Word clouds. Retrieved from https://bit.ly/2N7Wy3Y.
- Fox J, Weisberg S (2011) An R companion to applied regression 2nd. Thousand Oaks CA: Sage. Retrieved from https://bit.ly/2yO5mch.
- FXCM (2018) How does president Trump's twitter. Retrieved from https://bit.ly/2tJOUEk.
- Grolemund G, Wickham H (2011) Dates and times made easy with lubridate. *Journal of Statistical Software* 40(3): 1–25. Retrieved from https://bit.ly/2tM5xzf.
- Nguyen TH, Shirai K, Velcin J (2015) Sentiment analysis on social media for stock movement prediction. *Expert Systems with Applications* 42(24): 9603–9611. doi:10. 1016/j.eswa.2015.07.052.
- Nielsen FA (2011) *AFINN*. Denmark: Informatics; Mathematical Modelling, Technical University of Denmark. Retrieved from https://bit.ly/2pZzWL4.
- Nielsen FA (2011) A new ANEW: evaluation of a word list for sentiment analysis in microblogs. CoRR, abs/1103.2903. Retrieved from https://bit.ly/2IA1tHx.
- Otani A, Shifflett S (2017) *Think a negative tweet from trump crushes a stock? Think again.* Retrieved from https://on.wsj.com/2yM09ln.
- Papaioannou P, Russo L, Papaioannou G, Siettos C (2013) Can social microblogging be used to forecast intraday exchange rates? *Netnomics: Economic Research and Electronic Networking* 14: 47–68. doi:10.1007/s11066-013-9079-3.
- R Core Team (2018) *R: A language and environment for statistical computing*. Vienna: R Foundation for Statistical Computing. Retrieved from https://bit.ly/ 1gm1uk2.
- Ranco G, Aleksovski D, Caldarelli G, Grčar M, Mozetič I (2015) The effects of twitter sentiment on stock price returns. *PLOS ONE* 10(9). doi:10.1371/journal.pone.01 38441.
- Ryan JA, Ulrich JM (2018) *Quantmod: Quantitative financial modelling framework*. Retrieved from https://bit.ly/2Mt3WpC.
- Seeking Alpha (2017) The trump index. Retrieved from https://bit.ly/2KqtGT6.

Vol. 4, No. 4 Colonescu: The Effects of Donald Trump's Tweets on US Financial...

Silge J, Robinson D (2016) Tidytext: Text mining and analysis using tidy data principles in r. JOSS 1(3). doi:10.21105/joss.00037.

Silge J, Robinson D (2017) Text mining with R. Retrieved from https://bit.ly/2IybR2y.

- Trump Twitter Archive (2018) *Up to date archive of all Donald Trump's twits*. Retrieved from https://bit.ly/2lCVQiI.
- Wickham H (2009) *Ggplot2: Elegant graphics for data analysis*. New York: Springer-Verlag. Retrieved from https://bit.ly/Tea7Ss.
- Wickham H, Grolemund G (2017) *R for data science*. Retrieved from https://bit.ly/2ax 76Hr.
- Zeileis A (2006) Implementing a class of structural change tests: An econometric computing approach. *Computational Statistics & Data Analysis* 50: 2987–3008.
- Zhang X, Fuehres H, Gloor PA (2011) Predicting stock market indicators through twitter "i hope it is not as bad as i fear". *Procedia - Social and Behavioral Sciences* 26: 55–62. doi:10.1016/j.sbspro.2011.10.562.