Does News Affect Intraday Foreign Exchange Market Volatility? A Genetic Algorithms Approach

by

Morgan Kidd

A Thesis Submitted in Partial Fulfillment of the Requirements for the
Degree of Bachelor of Arts/Sciences, Honours
in the Department of Economics
University of Victoria
April 2017

Supervised by Dr. Marco Cozzi

for

Dr. Chris Auld, Honours co-advisor

Dr. Merwan Engineer, Honours co-advisor

Abstract

In this paper, I use statistical genetic algorithms to try to find patterns in volatility returns caused by news shocks in the foreign exchange markets for the US Dollar, European Euro, and Japanese Yen pairings for the first half of 2016. I find that the genetic algorithms are not successful in identifying any significant correlations, and the methods do not perform significantly better than a purely random selection of news articles. This result appears to be consistent across a wide array of robustness checks including the currency pairing, month, degree of time aggregation, and different algorithm parametrizations. Possible explanations for the results could be that the foreign exchange markets are sufficiently large to be resilient to most news shocks (compared to, for instance, an individual equity), and that the number of significant news events is sufficiently small that algorithms fail to identify any pattern attributable to them. An alternate view may be that the Reuters news wire used in this study is not representative of breaking news, and that traders may be able to gain access to the information reported before it becomes public.

Table of Contents

1. Introduction..1
2. Data………………………………………...2
3. Methodology..6
4. Results..8
5. Robustness Checks...11
6. Conclusion..14
Bibliography……..16

1. Introduction

The efficient market hypothesis (EMH) postulates that asset prices should adjust as new information becomes known (Andersen at al, 2007). The foreign exchange market (FOREX) is no different, and a sizeable amount of literature has been dedicated to better understanding structural breaks arising from a variety of different factors, including time of the day (Andersen et al., 2001, Mizuno et al., 2003, Guillaume et al., 1997) and news releases (Andersen et al., 2007, Balduzzi et al., 2001).

One of the issues with analysing price changes is the need to classify informational shocks as “positive” or “negative” with respect to the asset price, this is particularly challenging with foreign exchange rates due to their pairwise nature (e.g., what's “good” for the Yen is “bad” for the Dollar and vice versa). However, price volatility will appear during price shifts as market participants realign their portfolios. To this extent, it becomes easier to identify volatility shifts (i.e., that the price moving at all) compared to the direction of price changes (Andersen, Bollerslev & Das, 2001).

There have been several forays in to the prediction of asset prices (or price volatility) through external information sources. Examples include general news releases and the US bond market (Balduzzi et al. 2001), Twitter and the stock market (Bollen, Mao, & Zeng, 2011), and macro-economic surprises and their impact on several asset types (Andersen et al., 2007). Although most of these studies have focused on time-series analysis of the neighbourhood around pre-specified events. By comparison, this study focuses on the specification of these events without prior knowledge of which type of events may be relevant. This aspect of the analysis may be controversial, as different researchers may have different assessments on what kind of event may be relevant. To address this issue, in this paper I take a more agnostic approach using genetic algorithms.

Genetic algorithms have numerous applications, including to identify relevant regressors from large factor models (Goldberg, 1989, Givens, Hoeting 2012) or to recognize feature patterns in a data set (Yang, Honavar 1998). This makes genetic algorithms a theoretically viable approach to the problem examined in this paper.

The outcome of this research has several applications. First, it is a verification of the efficient market hypothesis and an exploration of how responsive some events may be under the EMH. Secondly, this research aids with identifying patterns amongst the types of news that cause volatility increase. This may be of practical interest as it allows traders to be better anticipate shifts in the volatility regime arising due to news.

The structure of this paper will be as follows: I will first introduce the data sets employed in the analysis and the relevant transformations to the data in section one. Section three will then explain the experiment design and methodology employed in the analysis, and include a brief overview of the use of genetic algorithms. Following that, section four is a discussion of the results and section five is a discussion some of the robustness checks. Finally, I will summarize the findings and comment on the potential of future research in the conclusion.

2. Data

2.1 News

I employ two sets of data. The first is a collection of all news articles published by Reuters during summer 2016. These were collected via an online repository of all the URLs of news articles, and then using a web-scraping algorithm to collect the headline, timestamp (to the nearest minute), and contents of each article. There are typically around 30,000 news articles per month, occurring at all hours of all days of the week. For comparison with the quantities of interest, the articles are aggregated to the nearest 2-minute interval (rounded down). For instance, a news article appearing at 1:35 will be assigned to the interval beginning at 1:34. This results in several instances where multiple news articles are released simultaneously, but they are treated as a single headline for the analysis in this paper. The resulting aggregation is displayed in Table 1, with each month having approximately 7,000 instances of news article release after aggregation.

Month	n Articles	n After Aggregation
January	27,045	7,248
February	27,315	6,943
March	29,739	7,137
April	35,248	7,240
May	35,766	7,074
June	35,275	7,364

Table 1: Reuters News

2.2 FOREX Data

FOREX data proves to be somewhat more difficult to access. The FOREX market generally exists directly between banks and brokers with no central exchange, and with direct prices visible only to direct market participants. I obtain foreign exchange tick data provided by truefx.com, a website operated by Integral Development Corp. Integral provides technology services for several major currency traders and institutions accounting for roughly 20% of the entire FOREX market, so the data is likely to be representative of the overall market. Due to market close, the data applies to weekdays from 22:00 GMT on Sunday (when the Australian market opens) until 22:00 GMT Friday (when the New York market closes).

The FOREX data itself consists only of a best bid and ask quote, along with a timestamp to the millisecond, the data has a new entry at every instance in which the bid or ask quote changes.

The data unfortunately does not have a direct measure of order book density or market volume, but it is still possible to infer some measure of market activity based on the frequency of quote updates. Furthermore, due to the extreme density of the data and to smooth the observations somewhat, I aggregate data up to the two-minute level for comparison with the news articles. The bid and ask price are defined as the average of dominant bid and ask quotes during the interval, i.e.

$$\mathrm{P}_{\mathrm{Bid,Ask,t_1}} = \frac{1}{k-j}\sum_{i=j}^{k}(Q_i) \quad t_j = \max(t < t_1)\,, t_k = \min(t > t_1 + 1)$$

Where Q_i is the i-th bid or ask quote, t_i is the corresponding i-th timestamp, which lie in the time interval t with minimum and maximum times t_j,t_k. Similarly, frequency is defined as the number of quote updates during the period, i.e. k-j. For volatility, I employ two separate definitions found in the literature (Andersen et al., 2001) . The first is simply a classic measure of standard deviation, similar to the bid and ask price defined above, I define volatility as

$$\mathrm{Vol}^2_{Bid,Ask,t_1} = \frac{1}{k-j}\sum_{i=j}^{k}(Q_i - \overline{Q_t})^2 \quad t_j = \max(t < t_1)\,, t_k = \min(t > t_1 + 1)$$

Where $\overline{Q_t}$ is the sample mean of price quotes for the period, and corresponds to the price P_t derived above. Just as the price is defined as the mean of quoted prices across a time interval, the volatility is analogous to the standard deviation of quoted prices during the same time interval. An alternate volatility measure is referred to as the relative volatility, defined as:

$$\mathrm{Rel.Vol_t} = \frac{|\sum_{i=j}^{k}\Delta Q_i|}{\sum_{i=j}^{k}|\Delta Q_i|} = \frac{|P_t - P_{t-1}|}{\sum_{i=j}^{k}|\Delta Q_i|}$$

This volatility ratio explained by Guillaume et. al (1997) uses the absolute value of the price change as a measure of volatility instead of the variance. The ratio takes values between 0 and 1, corresponding to a purely random and purely trending process respectively.

The summary statistics after these transformations are listed in Table 2. Volatility is significantly lower for the EURUSD pairing due to the magnitude of the price units, but relative volatility is comparable across all currency pairings. The number of intervals should be approximately the same, as they are only a reflection of the hours that the markets are open for. The slight discrepancy between same-month N across different currencies is due to data anomalies where a quote at the very beginning or end of the trading week is erroneously timestamped slightly earlier or later than market open or close. This makes virtually no difference in the analysis, however, so the anomalies are left in.

Month	Price	Volatility	Relative Vol.	Quotes/Interval	N
January	118.1	0.0104	0.062	352.4	14130
February	114.8	0.0121	0.061	400.0	14816
March	113.0	0.0085	0.088	216.5	16180
April	109.6	0.0086	0.096	617.1	14673
May	108.9	0.0077	0.072	849.8	15483
June	105.5	0.0108	0.058	1069	15477

Table 2.A: USDJPY

Month	Price	Volatility	Relative Vol.	Quotes/Interval	N
January	128.3	0.0087	0.044	522.7	14123
February	127.4	0.0118	0.045	554.7	14818
March	125.6	0.0099	0.059	362.7	16194
April	124.3	0.0107	0.071	636.6	14672
May	123.2	0.0088	0.067	883.3	15483
June	118.5	0.0148	0.058	1273	15477

Table 2.B: EURJPY

Month	Price	Volatility	Relative Vol.	Quotes/Interval	N
January	1.086	8.700e-05	0.058	364.5	14131
February	1.110	8.706e-05	0.057	355.2	14821
March	1.112	7.580e-05	0.077	236.6	16180
April	1.134	7.177e-05	0.086	569.5	14673
May	1.131	5.958e-05	0.063	805.5	15483
June	1.124	9.174e-05	0.065	962.1	15477

Table 2.C: EURUSD

Table 2: Summary statistics for each asset, 2-minute intervals

3. Methodology

The experiment design essentially looks to use relevant news as a treatment effect on FOREX volatility. This creates two problems to resolve, first the definition of the treatment effect (news selection) and then quantifying the actual changes resulting or correlating with the treatment effect, which is undertaken from a forecasting perspective.

3.1 News Selection

In this paper, I define “relevant news” to refer to news articles that are likely to have a direct consequence on prices in the FOREX market. This could be due to events directly related to the foreign exchange market (e.g. central bank activity), or secondary effects resulting from economic or political actions in either of paired countries. Due to the volume of news, it is impractical to manually sort the articles by relevancy, so I instead employ the use of two algorithmic approaches using keyword searches. The results are scored based on changes in the conditional mean of the volatility, i.e. comparing the mean volatility 60 minutes before and 90 minutes after a news event. However, before fitting any model requiring coefficient estimations, a feasibility check can be conducted by examining the following hypothesis test:

$$\mathrm{H}_0\text{:}\ \ Vol(P|News) = Vol(P)$$
$$\mathrm{H}_1\text{:}\ \ Vol(P|News) \neq Vol(P)$$

Where *Vol(P|News)* is measured as the ratio of the mean volatility of the 2 minute intervals during the 90 minutes following the selected news to the mean volatility of the 2 minute intervals during the 60 minutes before the same news. *Vol(P)* is similarly measured but across the sample of all news (selected or not), i.e.:

$$Vol(P|News) = \frac{\sum_{i=1}^{k}(Vol_{Post}|News_i)}{\sum_{i=1}^{k}(Vol_{Pre}|News_i)}$$

$$Vol(P) = \frac{\sum_{i=1}^{n}(Vol_{Post}|News_i)}{\sum_{i=1}^{n}(Vol_{Pre}|News_i)}$$

I make the conjecture that the volatility preceding news should be no different than the average volatility across the sample, i.e.:

$$\frac{\sum_{i=1}^{k}(Vol_{Pre}|News_i)}{k} \approx \frac{\sum_{i=1}^{n}(Vol_{Pre}|News_i)}{n}$$

As such, the hypothesis can be tested using the following test statistic.

$$F = \frac{Vol(P|News)}{Vol(P)} \sim F_{k,n}$$

Thus, the news article selection can be scored based on some function of the magnitude of *Vol(P|News)*, and the algorithm introduced in the following section will try to find the specification of news articles that leads to the largest score.

3.2 Genetic Algorithms

Genetic algorithms are a variant of ascent-based optimization algorithms which mimic the process of Darwinian natural selection. Candidate solutions are tested for fitness via some criteria, with high-fitness candidates more likely to breed the next generation. Breeding combines some of the traits of both parent candidates, and ensures diversity of solution while gradually evolving to become increasingly fit. Genetic algorithms are a desirable search algorithm for problems that can easily be categorized via binary design vectors and have clear objective functions, such as the selection of regressors (Reeves, Rowe, 1993, Givens, Hoeting, 2012).

The first parametrization uses combinations of the 500 most frequent words in news articles for a given month iteratively over the sample. The first month is scored based upon a function of the volatility ratio F-statistic defined in (5), and penalized for the number of keywords suggested.

The keyword search itself only uses linear combinations of words, e.g. mentions of the "bank of Japan" may lead to relevant articles, but the algorithm only searches for the appearance of either or both "bank" and "Japan", and does not add additional restrictions to certain combinations of words. For subsequent months, the ideal candidate determined during the previous month is pre-seeded as a candidate in the next month's generation.

The second parametrization is identical to the first, except it uses a curated list of words that are likely to be relevant a priori to running the genetic algorithm. The intention of this is to reduce the sample space, as even the 500 most common words have a sample space of 2^{500} possible specifications. I remove words that are almost definitely irrelevant such as sentence structure words (e.g. is, a, but, etc.) or words that are unlikely to be relevant on a FOREX-level magnitude (e.g. names of individual companies or minor countries). This produces a small list of 80 words that are then passed through the genetic algorithm identically to above.

4. Results

4.1 Keyword Selection

Keyword results are organized based on the two algorithmic approaches discussed in section 4.2, using a complete list of words based on the 500 most commonly used words in the article headlines, and a smaller list of 80 frequent keywords that are likely to be of macroeconomic importance. The algorithms are ran for each month and currency with generation sizes of 30 and 100 iterations. Some of the resulting keyword suggestions are listed below.

Common Words:	"Shares, Market, Attack, Global, Revenue, Capital, Probe, Lawsuit, Agreement, Regulator, Appoints, Stable, Rebound"
Curated Words:	"China, Global, Rate, Yuan, Military, Brexit, Crisis, Europe, Inflation"

Table 3A: Some keyword suggestions under both methodologies

		Jan.	Feb.	Mar.	Apr.	May	Jun.
Common Words (500)	USDJPY	103	148	173	150	175	200
	EURJPY	107	151	189	204	237	227
	EURUSD	92	152	157	170	208	217
Curated Words (80)	USDJPY	15	21	25	19	17	17
	EURJPY	39	37	29	29	33	34
	EURUSD	19	27	30	33	29	29

Table 3B: Quantity of words selected by list and currency

The results show some degree of inconsistency between the methodologies, with significantly more words being chosen when selecting from the unrestricted word list. As seen in Table 3B, the number of words selected is significantly greater when using the larger sample space provided by the common word list. However, both methods generally lead to approximately 20-30% of the available word choices being used to filter news articles. Nonetheless, many of the suggestions seem sensible. Mentions of countries or currencies are explicit, and words that suggest actions or changes (e.g. appoints, wins, etc.) are often from political events that may cause currency fluctuations, or references to macro-economic reports (e.g. June, core, factors). Still, the overwhelming number of words selected by the unrestricted methodology and its frequent inclusion of certainly irrelevant words suggests that the curated methodology may be superior.

The analysis of the summary statistics for the neighbourhood around the news articles selected by the keywords is listed below, the base volatility is the recorded ratio of volatility before and after any news event (relevant or otherwise). The filtered volatility is the ratio of volatility from before and after news events determined to be relevant via the selection algorithm. The percentage change compares the filtered volatility to the base volatility (i.e., how much greater the volatility ratio of the selected news is compared to the full sample). Finally, the percentage of sample refers to how many of the total sample of all news articles was selected by the genetic algorithm.

Month	Base Vol.	All Words			Curated Words		
		Filtered Vol.	% Change	% Sample	Filtered Vol.	% Change	% Sample
January	1.020	1.037	1.69%	14%	1.051	2.97%	5%
February	1.026	1.028	0.12%	43%	1.030	0.39%	7%
March	1.021	1.023	0.25%	46%	1.038	1.66%	7%
April	1.026	1.028	0.15%	45%	1.040	1.29%	5%
May	1.297	1.303	0.42%	29%	1.323	1.98%	7%
June	1.290	1.310	1.49%	18%	1.378	6.79%	8%

Table 4: Genetic Algorithm Results for USDJPY, 2-minute filtering

Month	Base Vol.	All Words			Curated Words		
		Filtered Vol.	% Change	% Sample	Filtered Vol.	% Change	% Sample
January	1.018	1.032	1.41%	14%	1.042	2.35%	9%
February	1.016	1.017	0.09%	43%	1.020	0.39%	9%
March	1.019	1.026	0.68%	13%	1.056	3.57%	6%
April	1.022	1.025	0.29%	45%	1.035	1.32%	8%
May	1.225	1.226	0.04%	29%	1.242	1.37%	10%
June	1.260	1.264	0.30%	30%	1.285	1.96%	11%

Table 5: Genetic Algorithm Results for EURJPY, 2-minute filtering

Month	Base Vol.	All Words			Curated Words		
		Filtered Vol.	% Change	% Sample	Filtered Vol.	% Change	% Sample
January	1.029	1.044	1.42%	14%	1.052	2.23%	7%
February	1.027	1.034	0.67%	14%	1.041	1.41%	9%
March	1.032	1.035	0.31%	13%	1.052	1.96%	8%
April	1.033	1.039	0.59%	13%	1.044	1.02%	8%
May	1.077	1.106	2.64%	16%	1.123	4.19%	10%
June	1.151	1.156	0.48%	22%	1.183	2.79%	12%

Table 6: Genetic Algorithm Results for EURUSD, 2-minute filtering

Notable from the tables is that the lack of significance of the first model. The improvements by selecting from the 500-word list is very marginal, and the algorithm tends to select a very large proportion of the news articles across most asset types and months. This over-selection (e.g. 40-50% of the news article sample) suggests that the algorithm is failing to detect any sort of distinct patterns amongst the words. This reflects the result of the keyword selection, that the un-curated word listing is simply too broad. By comparison, the algorithm performs much better with curated words, and identifies a much smaller proportion of the news sample as relevant (<10% across most samples). The changes still fail to reach statistical significance, but the filtered volatility changes are uniformly higher than when using the common word list. This result is also consistent across all asset types and months, suggesting that the curated list of words is a better model specification.

5 Robustness Checks

To ensure robustness of results, I employ several variations to the genetic algorithm selection discussed above. The null results are maintained consistently across all checks.

5.1 GA Parameters

Re-running the genetic algorithms with varying generation size, iteration number, and mutation rate all produce similar results. Similarly, changing the fitness criteria (i.e. removing the candidate size penalty) worsens the results as the algorithm tends to select too many news articles and the results become too large a portion of the overall sample.

Month	Base Vol.	Filtered Vol.	% Change	% Change (from basic model)	% Sample
January	1.020	1.053	3.19%	0.22%	6%
February	1.026	1.031	0.47%	0.08%	6%
March	1.021	1.038	1.65%	-0.01%	7%
April	1.026	1.042	1.53%	0.23%	6%
May	1.297	1.333	2.74%	0.74%	7%
June	1.290	1.362	5.53%	-1.18%	7%

Table 7: Algorithm Parameter Sensitivities – USDJPY

Month	Base Vol.	Filtered Vol.	% Change	% Change (from basic model)	% Sample
January	1.018	1.045	2.68%	0.32%	6%
February	1.016	1.020	0.42%	0.03%	11%
March	1.019	1.059	3.89%	0.30%	6%
April	1.022	1.040	1.76%	0.43%	7%
May	1.225	1.256	2.55%	1.16%	9%
June	1.260	1.299	3.11%	1.12%	11%

Table 8: Algorithm Parameter Sensitivities - EURJPY

Month	Base Vol.	Filtered Vol.	% Change	% Change (from basic model)	% Sample
January	1.029	1.051	2.09%	-0.15%	5%
February	1.027	1.039	1.20%	-0.21%	10%
March	1.032	1.058	2.47%	0.50%	6%
April	1.033	1.050	1.66%	0.63%	6%
May	1.077	1.116	3.54%	-0.63%	11%
June	1.151	1.172	1.86%	-0.91%	14%

Table 9: Algorithm Parameter Sensitivities - EURUSD

The tables above include one set of replications for the three currency pairings using better performing curated keyword list. The replicates take a larger population size of 40 per generation instead of the original 30, and reduces the mutation rate to 5% from 10%. The number of iterations remains the same at 100. Notable in the tables, the algorithm still successfully filters increases in volatility, but it only improves on the basic model when examining the EURJPY pairing in Table 8. The other two parings, Table 7 & 9, show no distinguishable trend for finding greater or lesser filtered volatility than the original specification, and in fact lead to very similar estimations of filtered volatility. This result suggests that the algorithmic selection process is mostly unresponsive to algorithm parameter specifications.

5.2 Time Intervals

An alternate conjecture is that the volatility is being over-smoothed by examining too wide of a neighbourhood around the news releases. To address that, I run a further replicate where pre-and post-news volatility is calculated from a sample of 30 and 50 minutes respectively (down from 60 and 90 minutes in the original model. The results are tabulated below.

Month	Base Vol.	Filtered Vol.	% Change	% Change (from basic model)	% Sample
January	1.071	1.064	0.67%	-2.30%	28%
February	1.066	1.060	0.61%	0.21%	33%
March	1.042	1.021	2.09%	0.43%	5%
April	1.083	1.075	0.77%	-0.53%	39%
May	1.087	1.067	1.86%	-0.12%	10%
June	1.073	1.072	0.09%	-6.70%	44%

Table 10: Interval Sensitivities - USDJPY

Month	Base Vol.	Filtered Vol.	% Change	% Change (from basic model)	% Sample
January	1.061	1.056	0.46%	-1.88%	27%
February	1.052	1.045	0.68%	0.29%	32%
March	1.088	1.053	3.38%	-0.19%	5%
April	1.066	1.059	0.67%	-0.65%	39%
May	1.070	1.056	1.32%	-0.06%	10%
June	1.073	1.072	0.09%	-1.87%	11%

Table 11: Interval Sensitivities- EURJPY

Month	Base Vol.	Filtered Vol.	% Change	% Change (from basic model)	% Sample
January	1.075	1.069	0.54%	-1.70%	5%
February	1.068	1.062	0.59%	-0.81%	34%
March	1.088	1.077	1.08%	-0.88%	7%
April	1.082	1.079	0.28%	-0.75%	40%
May	1.096	1.073	2.09%	-2.10%	10%
June	1.090	1.086	0.36%	-2.44%	43%

Table 12: Interval Sensitivities - EURUSD

Visible in all three tables is that by narrowing the neighbourhood of interest, the magnitude of change in volatility decreases greatly due to a higher base volatility when compared with the original model specifications. The higher base volatility is likely attributable to less smoothing from the smaller time window around the news releases. Also noteworthy is the inconsistent but large portion of the news sample filtered by the genetic algorithm. This leads to the same problem of pattern detection failure discussed in the results section.

6. Conclusion

In this research, I have used a genetic algorithmic approach to attempt to filter news based on impactfulness on the foreign exchange markets. This is an attempt to extend the existing literature which focuses on event analysis around specific, pre-defined macro-economic events.

Although we may have a strong prior that news impacts asset prices, and proof that certain forms of unexpected news may cause shifts in the volatility structure, it does not seem possible to easily determine which forms of news are the most impactful via simple statistical learning algorithms. Clear shortfalls of this research are the relative simplicity with which the algorithms are ran. Specifically, the selection methods rely only on the existence of certain words, but are unable to qualitatively sort them further. This leads to the possibility of necessarily relevant articles being conflated with entirely irrelevant events by coincidentally having some of the same words. Moreover, it is possible that Reuters is not an adequate source of news and that traders may receive their news from more direct sources, giving them a speed advantage. Additionally, genetic algorithms themselves are imperfect and have only probabilistic, not guaranteed, convergence (Whitley, 1994).

Another possible explanation for the results is due to unquantifiable market surprise. Although unexpected geopolitical events are likely to shift currency rates, the continued discussion of them may not, and may in fact desensitize traders to certain topics. This is consistent with previous literature, where the independent variable if macro-economic surprise, quantified from the future’s market (Andersen et al., 2007). Unfortunately, there is no simple way to quantify the degree of surprise arising from the types of news used in this research, and so this extension is left for future research.

Finally, even under the conditions of news being theoretically impactful and surprising, traders may be rationally inattentive to news. In other words, the benefits to a trader of adjusting his portfolio following a particular news release may be marginal and in fact suboptimal due to the volume of information flowing in at any given point in time (Sims, 2015).

Further areas of research should be focused on more advanced learning techniques. Specifically, I believe that neural networks and multi-layered learning are both likely to produce more tangible results. Wide web-scraping of several news sources may also produce a more comprehensive news data set. Finally, FOREX movements are susceptible to a very large number of factors, and studying a more isolated asset class (medium cap equities, for instance) may also be a useful exercise to determine which form of news selection algorithm performs the best.

References

Andersen, T. G., Bollerslev, T., & Das, A. (2001). Variance-ratio statistics and high-frequency data: testing for changes in intraday volatility patterns. The Journal of Finance, 56(1), 305-327.

Andersen, T.G., Bollerslev, T., Diebold, F.X., & Labys, P. (2003). Modeling and forecasting realized volatility. Econometrica, 71(2), 579-625.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Vega, C. (2007). Real-time price discovery in global stocks, bond and foreign exchange markets. Journal of International Economics, 73(2), 251-277.

Balduzzi, P., Elton, E.J., Green, T.C. (2001). Economic news and bond prices: evidence from the U.S. treasury market. Journal of Financial and Quantitative Analysis, 36, 523-543.

Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. Journal of computational science, 2(1), 1-8.

Bollerslev, T., (1986) Generalized autoregressive conditional heteroskedasticity, Journal of Econometrics, 31, 307-327.

Boyd, J.H., Jagannathan, R., Hu, J. (2005). The stock market's reaction to unemployment news: why bad news usually good for stocks. Journal of Finance, 60, 649-672.

Franses, P.H., & Van Dijk, D. (1996). Forecasting stock market volatility using (nonlinear) GARCH models. Journal of Forecasting, 229-235.

Givens, G.H., & Hoeting, J.A. (2012). Computation statistics (Vol. 710). John Wiley & Sons.

Goldberg, D.E. (1989). Genetic algorithms in search, optimization, and machine learning. Addison-Wesley.

Guillaume, D.M., Dacorogna, M.M., Dave, R.R., Muller, U.A., Olsen, R.B., & Pictet, O.V. (1997). From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. Finance and Stochastics, 1(2), 95-129.

Mizuno, T., Kurihara, S., Takayasu, M., & Takayasu, H. (2003). Analysis of high-resolution foreign exchange data of USD-JPY for 13 years. Physica A: Statistical Mechanics and its Applications, 324(1), 296-302.

Peiers, B. (1997). Informed traders, intervention, and price leadership: A deeper view of the microstructure of the foreign exchange market. The Journal of Finance, 52(4), 1589-1614.

Reeves, C.R., Rowe, J.E. (1993) Modern heuristic techniques for combinatorial problems. Wiley, New York.

Sims, Christopher A. (2015). Rational inattention and monetary economics. In: Handbook of Monetary Economics. Elsevier.

Whitley, D. (1994). A genetic algorithm tutorial. Statistics and Computing, 4, 65-85.

Yang, J., & Honavar, V. (1998). Feature subset selection using a genetic algorithm. In feature: extraction, construction and selection (117-136). Springer US.