Machine news and volatility: The Dow Jones Industrial Average and the TRNA sentiment series

David E Allen[a], Michael McAleer[b] and Abhay K Singh[c]

[a]Adjunct Professor Centre for Applied Financial Studies, UniSA, and Visiting Professor, School of Mathematics and Statistics, University of Sydney

[b]Department of Quantitative Finance, National Tsing Hua University, Taiwan, Econometric Institute, Erasmus School of Economics, Erasmus University Rotterdam, Tinbergen Institute, The Netherlands, and Department of Quantitative Economics, Complutense University of Madrid, Spain

[c]School of Business, Edith Cowan University,Perth, Australia

Abstract

This paper features an analysis of the relationship between the volatility of the Dow Jones Industrial Average (DJIA) Index and a sentiment news series using daily data obtained from the Thomson Reuters News Analytics (TRNA) provided by SIRCA (The Securities Industry Research Centre of the Asia Pacific). The expansion of on-line financial news sources, such as internet news and social media sources, provides instantaneous access to financial news. Commercial agencies have started developing their own filtered financial news feeds, which are used by investors and traders to support their algorithmic trading strategies. In this paper we use a sentiment series, developed by TRNA, to construct a series of daily sentiment scores for Dow Jones Industrial Average (DJIA) stock index component companies. A variety of forms of this measure, namely basic scores, absolute values of the series, squared values of the series, and the first differences of the series, are used to estimate three standard volatility models, namely GARCH, EGARCH and GJR. We use these alternative daily DJIA market sentiment scores to examine the relationship between financial news sentiment scores and the volatility of the DJIA return series. We demonstrate how this calibration of machine filtered news can improve volatility measures.

Keywords: DJIA, Sentiment Scores, TRNA, Conditional Volatility Models.
JEL: C58, G14.

Email address: profallen2007@gmail.com

January 17, 2014

1. Introduction

There has been a revolution in the speed of news transmission over the past century that began with wire services, whose use spans a period from around World War 1 to the 1940s, during which news agencies in the U.S.A. transmitted copy over telephone wires to teletypewriters in newspaper offices. In the late 1940s, things changed again with the introduction of Teletypesetter machines. These permitted the use of perforated paper tape, which was fed into typesetting, or linotype, machines, without human intervention, further reducing processing times. Newspapers subsequently switched from linotype to photocomposition in the late 1960s to 1970s.

A more recent innovation has been the use of the internet. Information is now transmitted by satellite service or the Internet, and newspapers reconstruct the information in their own format. News has always been the lifeblood of financial markets, and being the first to know provides a first mover advantage. However, some parties, such as corporate officers, are likely to be the 'first in the know', and this has attracted the attention of market regulators over the years, who have attempted to ensure that investors face a 'level playing field'. For example, in the USA, sections 16(b) and 10(b) of the Securities Exchange Act of 1934 address insider trading.

There is also the issue that the information has to be pertinent and value-relevant, and other investors also need to be convinced of its value. This brings us to consider Keyne's (1936) famous analogy between choosing investment stocks and a fictional newspaper beauty competition in which entrants are asked to choose from a set of six photographs of women that are the most beautiful. Those who picked the most popular face would then be eligible for a prize.

"It is not a case of choosing those [faces] that, to the best of one's judgment, are really the prettiest, nor even those that average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees." (Keynes (1936), Chapter 12, p.100). At any given moment in time, a security's price must be a weighted average of investor trading strategies.

Clearly, the information embodied in news items is one information source that has the potential to influence investor opinions. This paper features an exploration of the impact of a machine created news series drawn from Thomson Reuters News Analytics (TRNA) which could be termed news sentiment, and which is produced by the application of machine learning techniques to news items.

The paper is a companion paper to two other studies by Allen, McAleer and Singh (2013a, b). The first of these papers examines the influence of the Sentiment measure as a factor in pricing DJIA constituent company stocks in a Capital Asset Pricing Model (CAPM) context. The second paper uses these daily DJIA market sentiment scores to study the relationship between financial news sentiment scores and the DJIA return series using entropy-based measures. Both studies find that the sentiment scores have a significant information

component, which in the former study is priced as a factor in an asset pricing context. The current paper further explores the influence of sentiment scores in the context of their impact on the DowJones30 index's volatility.

The series we use are based on Thomson Reuters News Analytics (TRNA), which takes news items calibrated into either positive, negative or neutral values per news item, and used to construct its Sentiment series. The key issue is the extent to which the series influences investors' investment strategies which, in turn, influence the market and the evolution of stock prices. They are also used as an input to algorithmic trading techniques.

There has been attention recently on the role of market news sentiment, in particular, machine-driven sentiment signals, and their implications for financial market processes. The research on this topic argues that news items from different sources influence investor sentiment, which feeds into asset prices, asset price volatility and risk (Tetlock, 2007; Tetlock, Macskassy and Saar-Tsechansky 2008; Da, Engleberg and Gao, 2011; Barber and Odean, 2008; diBartolomeo and Warrick 2005; Mitra, Mitra and diBartolomeo 2009; Dzielinski, Rieger and Talpsepp 2011). The diversification benefits of the information impounded in news sentiment scores provided by RavenPack has been demonstrated by Cahan, Jussa and Luo (2009) and Hafez and Xie (2012), who examined its benefits in the context of popular asset pricing models.

Another important research question is the extent to which the availability of these machine-driven series contributes to market information and the evolution of security prices. Baker and Wurgler (2006) demonstrated a link between investor sentiment and stock returns. Recent work by Hafez and Xie (2012) examines the effect of investor's sentiment using news-based sentiment, generated from the RavenPack Sentiment Index as a proxy for market sentiment in a multi-factor model. They report a strong impact of market sentiment on stock price predictability over 6- and 12-month time horizons. Allen, McAleer and Singh (2013a) demonstrate, in an analysis of DowJones Index constituent companies, that a Sentiment series can make up a distinct factor that is priced in a CAPM framework.

The issue of the news content of sentiment scores for volatility behaviour is the central focus of this paper. We address it by analysing the relationship between a commercially available series, the Thomson Reuters News Analytics (TRNA) series and the volatility behaviour of a major index, the DJIA. These large US stocks are likely to be among the most heavily traded and analysed securities in the world. Therefore, the issue of the relevance of these news feeds to the volatility behaviour of this major index is an important one.

We take the TRNA news series for the DJIA constituent stocks and aggregate them into a daily time series. This facilitates an analysis of the relationship between the two daily sets of series, TRNA news sentiment on the one hand, and the DJIA volatility behaviour on the other. We analyse the relationship between the two series using three standard univariate volatility models, namely GARCH(1,1), GJR(1,1) and EGARCH(1,1).

The extent to which these news series have relevant information for volatility behaviour is germane for both investors and market regulators. If access to these

particular information feeds provides a trading advantage, then the market is no longer a level playing field for all investors. Institutions and algorithmic traders with access to these analytics will have an advantage. However, this paper does not address the issue of the timing of access to news items, but the more general question of the degree to which these sentiment-based series contain 'relevant information', as revealed by an analysis of the volatility of the DJIA and its links to a daily average of the TRNA series. The paper is a component of three separate analyses of this relationship. Allen, McAleer and Singh (2013a) explore the links between the series in an asset pricing framework, while Allen, McAleer and Singh (2013b) explore the informational relationship between the two series using entropy-based measures.

The paper is organized as follows: Section 1 provides an introduction. Section 2 features an introduction to sentiment analysis and an overview of the TRNA data set, and introduces the research methods adopted. Section 3 discusses the major results, and Section 4 draws some conclusions.

2. Research methods and data

2.1. News Sentiment

In this paper we examine the sentiment scores provided by TRNA as a single information vector which is added to the mean and variance equations for three commonly-used volatility models, GARCH(1,1), GJR(1,1) and EGARCH(1,1), as applied to the volatility behaviour of the DJIA. We use daily DJIA market sentiment scores constructed from high frequency sentiment scores for the various stocks in DJIA. The empirical analysis includes data from the Global Financial Crisis and other periods of market turbulence to assess the effect of financial news sentiment on stock prices in both normal and extreme market conditions. The relationship between stock price movements and news sentiment has recently been examined by Tetlock (2007), Barber and Odean (2008), Mitra, Mitra and diBartolomeo (2009), Leinweber and Sisk (2009), Sinha (2011), and Huynh and Smith (2013).

The scale of competing news sources in the electronic media means that there is scope for the commercial use of sources of pre-processed news. Vendors such as TRNA and RavenPack produce sentiment scores to provide direct indicators to traders and other financial practitioners of changes in news sentiment. They use text mining tools to electronically analyse available textual news items. The analytics engines of these sources apply pattern recognition and identification methods to analyse words and their patterns, and the novelty and relevance of the news items for a particular industry or sector. These news items are converted into quantifiable sentiment scores.

We use sentiment indicators provided by TRNA in our empirical analysis. Thomson Reuters was a pioneer in the implementation of a sophisticated text mining algorithm as an addition to its company and industry-specific news database, starting from January 2003, which resulted in the present TRNA data set. The TRNA data guide states that: "Powered by a unique processing system the Thomson Reuters News Analytics system provides real-time numerical

insight into the events in the news, in a format that can be directly consumed by algorithmic trading systems".

Currently the data set is available for various stocks and commodities until October 2012. The TRNA sentiment scores are produced from text mining news items at a sentence level, which takes into account the context of a particular news item. This kind of news analytics makes the resulting scores more usable as they are mostly relevant to the particular company or sector. Every news item in the TRNA engine is assigned an exact time stamp, and a list of companies and topics it mentions. A total of 89 broad fields are reported in the TRNA data set, which are broadly divided into following 5 main categories:

1. Relevance: A numerical measure of how relevant the news item is to the asset.
2. Sentiment: A measure of the inherent sentiment of the news item, quantifying it as either negative (-1), positive (1) or neutral (0).
3. Novelty: A measure defining how new the news item is; in other words, whether it reports a news item that is related to some previous news stories.
4. Volume: Counts of news items related to the particular asset.
5. Headline Classification: Specific analysis of the headline.

Figure 1 shows a snapshot of the headline text as reported in BCAST_REF field of the TRNA database for General Motors during the year 2007. These are not the sentences which are analysed by TRNA to produce sentiment scores, but are the headlines for the news item used to generate the TRNA sentiment and other relevant scores.

Figure 2 provides another example featuring theAustralian company BHP Billiton, which is reported in TRNA as having generated more than 3000 news items in the year 2011. Figure 2 shows the sentiment scores (-1 to +1) for BHP Billiton during the month of January 2011, where the red line is the moving average of the scores.

Similar to BHP, there are various news stories reported per day for the various DJIA traded stocks. These news stories result in sentiment scores which are either positive, negative or neutral for that particular stock. Figure 3 gives a snapshot of the sentiment scores for the DJIA traded stocks during the year 2007. The bar chart of Figure 3 shows that the most sentiment scores generated during the year 2007, which marked the beginning of the period of Global Financial Crisis, were for the Citi Bank group (C.N), General Motors (GM.N), General Electric (GE.N) and J. P Morgan (JPM.N). This is a reflection of the market sentiment during the GFC period, as these financial institutions were among the most affected during the GFC.

Figure 4 shows the number of positive, negative or neutral sentiment scores stacked against each other in 2007. It is evident that the number of negative and neutral sentiment news was exceeded by the number of positive sentiments for the majority of stocks, as it was only later in the year that the GFC really began to impact. However, Bank of America (BA.N), Citibank (C.N), General Motors

Figure 1: TRNA-Snapshot of News Headlines Generated for General Motors in the Year 2007

Figure 2: TRNA-Sentiment Scores Generated for BHP Billiton in January 2011

Figure 3: Sentiment Score Distribution for DJIA Stocks in 2007
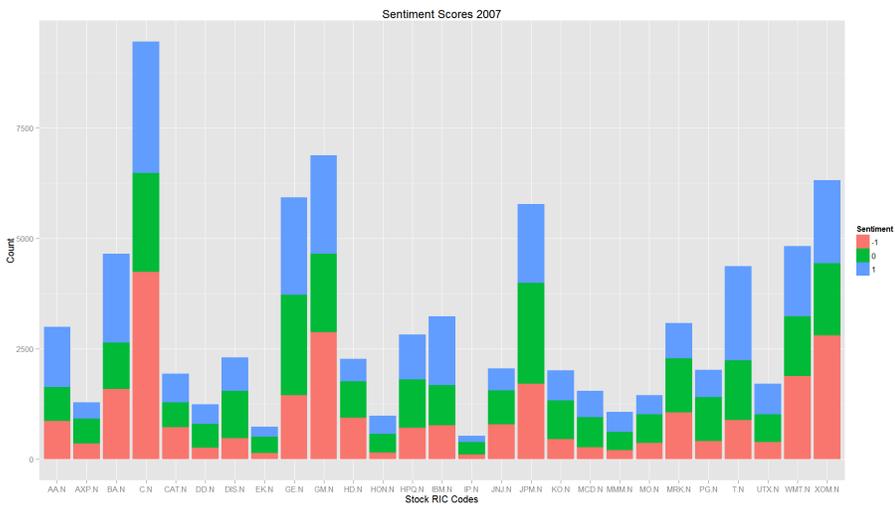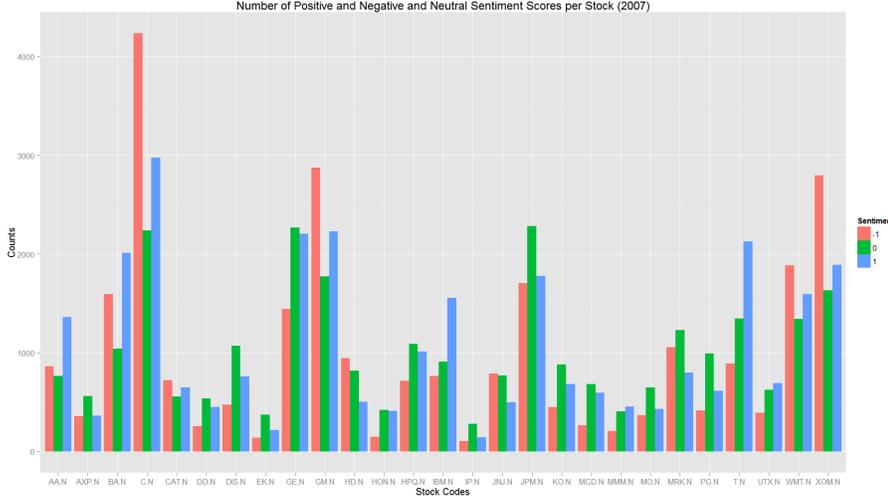
Figure 4: Positive, Negative and Neutral Sentiment Score Distribution for DJIA Stocks in 2007

(GM.N), Wal Mart (WMT.N) and Exxon (XOM.N), have a preponderance of negative sentiment during the year.

Applications of TRNA news data sets in financial research have grown recently. Dzielinski (2012), Groß-Kulßman and Hautsch (2011), Smales (2013), Huynh and Smith (2013), Borovkova and Mahakena (2013). Storkenmaier et al. (2012), and Sinha (2011) have explored the usefulness of the TRNA dataset in stock markets and in commodity markets. In this paper we use the TRNA data set to analyse the effect of news sentiment on the DJIA daily volatility behaviour. We construct daily sentiment index score time series for the empirical exercise based on the high frequency scores reported by TRNA.

The empirical analysis in this paper analyses the effect of news sentiment on stock prices of the DJIA by considering the daily DJIA market sentiment as an additional exogenous factor in volatility models of the DJIA We construct daily sentiment scores for DJIA market by accumulating high frequency sentiment scores of the DJIA constituents obtained from the TRNA dataset. We use data from January 2007 to October 2012 to analyse the senstivity of the DJIA daily volatility to the daily market sentiment scores. The daily stock prices for all the DJIA traded stocks are obtained from the Thomson Tick History database for the same time period.

The TRNA provides high frequency sentiment scores calculated for each news item reported for various stocks and commodities. These TRNA scores for the stocks traded in DJIA can be aggregated to obtain a daily market sentiment score series for the DJIA stock index components. A news item, s_t, received at time t for a stock is classified as positive ($+1$), negative (-1) or neutral (0). $I_{s_t}^+$ is a positive classifier (1) for a news item, s_t, and $I_{s_t}^-$ is the negative (-1) classifier

for a news item, s_t. TRNA reported sentiment scores have a probability level associated with them, $prob_{s_t}^+$, $prob_{s_t}^-$, $prob_{s_t}^0$ for positive, negative and neutral sentiments, respectively, which is reported by TRNA in the Sentiment field. Based on the probability of occurrence, denoted by P_{s_t} for a news item, s_t, all the daily sentiments can be combined to obtain a daily sentiment indicator. We use the following formula to obtain the combined score:

$$S = \frac{\sum_{q=t-1}^{t-Q} I_{s_q}^+ P_{s_q} - \sum_{q=t-1}^{t-Q} I_{s_q}^- P_{s_q}}{n_{prob_{s_q}^+} + n_{prob_{s_q}^-} + n_{prob_{s_q}^0}} \tag{1}$$

The time periods considered are $t - Q, \ldots, t - 1$, which covers all the news stories (and respective scores) for a 24-hour period.

2.2. Our sample characteristics and preliminary analysis

Table 1 lists the various stocks traded in DJIA, along with their RIC (Reuters Instrument Code) and time periods. We use the TRNA sentiment scores related to these stocks to obtain the aggregate daily sentiment for the market. The aggregated daily sentiment score, S, represents the combined score of the sentiment scores reported for the stocks on a particular date. We construct daily sentiment scores for the DJIA market by accumulating high frequency sentiment scores of the DJIA constituents obtained from the TRNA dataset. We use data from January 2006 to October 2012 to examine the sensitivity of the daily DJIA volatility to the daily market sentiment scores. The daily stock prices for all the DJIA traded stocks are obtained from the Thomson Tick History database for the same time period, and are provided by SIRCA (The Securities Industry Research Centre of the Asia Pacific).

The stocks with insufficient data are removed from the analysis and the stocks prices for EK.N and EKDKQ.PK are combined to obtain a uniform time series.

The summary statistics in Table 3 show that the sample of Sentiment scores for the full sample is predominantly negative, with a mean of -0.034532. The minimum score is -0.52787 and the maximum score is 0.28564. It appears that negative news has greater prominence than positive news on the scale running from +1 to -1. The Hurst exponent for the Sentiment score, with a value of 0.925828, suggests that there is long memory or persistence in the scores, which makes intuitive sense, given that items of news may take several days to unfold, as greater scrutiny of a story leads to greater disclosure of information. When an event is classified as positive or negative, this will tend to occupy the media for several days, and is consistent with trending behaviour. The Hurst exponent for DJIA is 0.557638, which suggests that the DJIA shows much less tendency to display memory and, as might be expected, behaves more like a random walk. The significant Jarque-Bera Lagrange multiplier test statistics for both series suggest that both are non-Gaussian.

We also used a number of variants of the sentiment score, squared score, absolute value of the score and the first difference to explore which might better capture the influence of market sentiment scores. The plots of the various series

Table 1: DJIA Stocks with Thomson Tick History RIC Codes

RIC Code	Stocks	First Date	Last Date
.DJI	Dow Jones INDU AVERAGE	1-Jan-96	17-Mar-13
AA.N	ALCOA INC	2-Jan-96	18-Mar-13
GE.N	GENERAL ELEC CO	2-Jan-96	18-Mar-13
JNJ.N	JOHNSON&JOHNSON	2-Jan-96	18-Mar-13
MSFT.OQ	MICROSOFT CP	20-Jul-02	18-Mar-13
AXP.N	AMER EXPRESS CO	2-Jan-96	18-Mar-13
GM.N	GENERAL MOTORS	3-Jan-96	18-Mar-13
GMGMQ.PK	GENERAL MOTORS	2-Jun-09	15-Aug-09
JPM.N	JPMORGAN CHASE	1-Jan-96	18-Mar-13
PG.N	PROCTER & GAMBLE	2-Jan-96	18-Mar-13
BA.N	BOEING CO	2-Jan-96	18-Mar-13
HD.N	HOME DEPOT INC	2-Jan-96	18-Mar-13
KO.N	COCA-COLA CO	2-Jan-96	18-Mar-13
SBC.N	SBC COMMS	2-Jan-96	31-Dec-05
T.N	AT&T	3-Jan-96	18-Mar-13
C.N	CITIGROUP	2-Jan-96	18-Mar-13
HON.N	HONEYWELL INTL	2-Jan-96	18-Mar-13
XOM.N	EXXON MOBIL	1-Dec-99	18-Mar-13
MCDw.N	MCDONLDS CORP	6-Oct-06	4-Nov-06
MCD.N	MCDONALD'S CORP	1-Jan-96	18-Mar-13
EK.N	EASTMAN KODAK	1-Jan-96	18-Feb-12
EKDKQ.PK	EASTMAN KODAK	19-Jan-12	18-Mar-13
IP.N	INTNL PAPER CO	2-Jan-96	18-Mar-13
CAT.N	CATERPILLAR INC	2-Jan-96	18-Mar-13
HPQ.N	HEWLETT-PACKARD	4-May-02	18-Mar-13
MMM_w.N	3M COMPANY WI	18-Sep-03	27-Oct-03
MMM.N	MINNESOTA MINIhNG	1-Jan-96	18-Mar-13
UTX.N	UNITED TECH CP	2-Jan-96	18-Mar-13
DD.N	DU PONT CO	2-Jan-96	18-Mar-13
IBM.N	INTL BUS MACHINE	2-Jan-96	18-Mar-13
MO.N	ALTRIA GROUP	2-Jan-96	18-Mar-13
WMT.N	WAL-MART STORES	2-Jan-96	18-Mar-13
DIS.N	WALT DISNEY CO	2-Jan-96	18-Mar-13
INTC.OQ	INTEL CORP	20-Jul-02	18-Mar-13
MRK.N	MERCK & CO	2-Jan-96	18-Mar-13

Figure 5: Basic Series Plots: DJIA and Sentiment Scores

(a) DJIA Returns and absolute value of sentiment scores
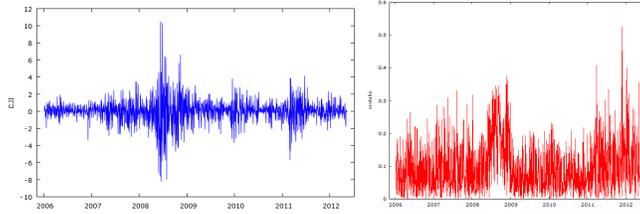
(b) Sentiment series, squared sentiment series and first differences of sentiment series
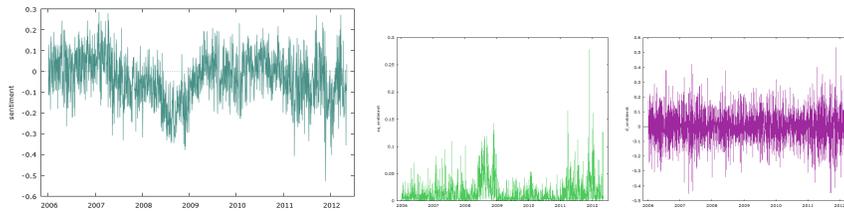

Table 2: Summary statistics, DJIA returns and Sentiment Scores

Jan 4th 2006 to 31st October 2012					
	DJIA return (%)	Sentiment Score	Sentiment Squared	Sentiment Abs	Sentiment difference
Min	-8.2005	-0.52787	5.38240e-010	2.32000e-005	-0.452678
Median	0.053410	-0.031140	0.00623149	0.0789397	-0.00349615
Mean	0.013971	-0.034532	0.0148177	0.0960405	3.70853e-005
Maximum	10.5083	0.28564	5.38240e-010	0.527867	0.534308
St. Deviation	1.3640	0.116762	0.0222546	0.0748150	0.125985
Hurst Exponent	0.557638	0.925828	0.861467	0.853927	0.178098
Jarque-Bera test	5320.84 (0.00)	18.2197 (0.00)	27737.1(0.00)	489.515(0.00)	14.8959(0.00)

are shown in Figure 5. Summary statistics for these series are presented in Table 2. The variants of the sentiment score have quite similar values for their Hurst exponent. All suggest trending behaviour, apart from the first differences of sentiment scores, which have a low Hurst exponent of 0.178, suggesting a tendency to display reversals. The Jarque-Bera Lagrange multiplier tests strongly reject the null hypothesis of a normal distribution for all series.

2.2.1. Volatility models utilised

Engle (1982) developed the Autoregressive Conditional Heteroskedasticity (ARCH) model that incorporates all past error terms. It was generalised to GARCH by Bollerslev (1986) to include lagged conditional volatility. In other words, GARCH predicts that the best indicator of future variance is the weighted average of long-run variance, the predicted variance for the current period, and any new information in this period, as captured by the squared return shocks (Engle (2001)).

The framework is developed as follows: consider a time series $y_t = E_{t-1}(y_t) + \varepsilon_t$, where $E_{t-1}(y_t)$ is the conditional expectation of y_t at time $t-1$ and ε_t is the error term. The GARCH model has the following specification:

$$\varepsilon_t = \sqrt{h_t}\eta_t, \ \ \eta_t \sim N(0,1) \tag{2}$$

$$h_t = \omega + \sum_{j=1}^{p}\alpha_j\varepsilon_{t-j}^2 + \sum_{j=1}^{q}\beta_j h_{t-j} \tag{3}$$

in which $\omega > 0$, $\alpha_j \geq 0$ and $\beta_j \geq 0$ are sufficient conditions to ensure a positive conditional variance, $h_t \geq 0$. The ARCH effect is captured by the parameter α_j, which represents the short run persistence of shocks to returns. β_j captures the GARCH effect, and $\alpha_j + \beta_j$ measures the persistence of the impact of shocks to returns to long-run persistence. A GARCH(1,1) process is weakly stationary if $\alpha_1 + \beta_1 \leq 1$.

We explore the impact of the various sentiment series on both the conditional mean and conditional variance equations.

Engle (2001), Nelson (1991), McAleer (2005), and Harris, Stoja and Tucker (2007) outline some of the disadvantages of the GARCH model as follows: GARCH can be computationally burdensome and can involve simultaneous estimation of a large number of parameters. The standard GARCH model tends to underestimate risk (when applied to Value-at-Risk, VaR), as the normality assumption of the standardized residual does not always hold with the behaviour of financial returns. The specification of the conditional variance equation and the distribution used to construct the log-likelihood may also be incorrect.

The basic symmetric model rules out, by assumption, the negative leverage relationship between current returns and future volatilities, despite empirical evidence to the contrary. GARCH assumes that the magnitude of excess returns determines future volatility, but not the sign (positive or negative returns), as it is a symmetric model. This is a significant problem as research by Nelson (1991)

and Glosten, Jagannathan and Runkle (GJR) (1993) shows that asset returns and volatility do not react in the same way for negative information, or 'bad news', as they do for positive information, or 'good news', of equal magnitude.

An alternative asymmetric model is the GJR model (1993), which is specified as:

$$h_t = \omega + \sum_{j=1}^{r}(\alpha_j + \gamma_j I(\varepsilon_{t-j}^2))\varepsilon_{t-j}^2 + \sum_{j=1}^{s}\beta_j h_{t-j} \qquad (4)$$

where

$$I_{it} = \left\{ \begin{array}{ll} 0, & \varepsilon_{it} \geq 0 \\ 1, & \varepsilon_{it} < 0 \end{array} \right\}$$

and i_{it} is an indicator function that distinguishes between positive and negative shocks of equal magnitude. In this model, when there is only one lag, that is, when $r = s = 1$, the sufficient conditions to ensure that the conditional variance is positive ($h_t > 0$) are that $\omega > 0$, $\alpha_1 \geq 0$, $\alpha_1 + \gamma_1 \geq 0$, and $\beta_1 \geq 0$; where α_1 and ($\alpha_1 + \gamma_1$) measure the short-run persistence of positive and negative shocks, respectively, and the given conditions apply for a GJR(1,1) model.

In the EGARCH model, the conditional variance h_t is an asymmetric standardized function of the lagged disturbances, ϵ_{t-1}:

$$ln(h_t) = \omega + \sum_{j=1}^{p}\beta_j ln(h_{t-j}) + \sum_{i=1}^{g}\alpha_i \left|\frac{\epsilon_{t-1}}{\sqrt{h_{t-i}}}\right| + \sum_{k=1}^{r}\gamma_k \frac{\epsilon_{t-k}}{\sqrt{h_{t-k}}} \qquad (5)$$

The fact that the log of the conditional variance is used in equation (5) implies that the leverage effect may be exponential and guarantees that forecasts of the conditional variance will be non-negative. The presence of asymmetric effects can be tested by the hypothesis that $\gamma_i = 0$, and the impact is asymmetric if $\gamma_i \neq 0$. A sufficient condition for the stationarity of the EGARCH(1,1) model is that $|\beta| < 1$.

In this paper we analyse the impact of the news series on volatility using these three standard variants of the GARCH model and four different measures of the sentiment index, namely weighted sentiment scores, squared values of the sentiment score, absolute values of the sentiment score, and its first difference. The results of our analysis are shown in the next section. We explore the influence of the sentiment scores on both the conditional mean and conditional variance equations using the methods introduced in equations (2) to (5).

3. The significance of the sentiment scores in the GARCH analysis of Dow Jones Index (DJIA) return series

We commence by estimating a standard GARCH(1,1) model, and augment both the conditional mean and conditional variance equations by adding a vector of the variants of the sentiment scores to assess whether they add information to the basic model. The results are shown in Table 3.

Table 3: GARCH(1,1) model of DJIA, with mean and variance equations augmented by sentiment scores

Variable	Sentiment	Sentiment squared	Sentiment absolute val.	Sentiment difference
Constant ω	6.9508598 (0.00)	12.0447668 (0.00)	14.3688799 (0.00)	6.9319048 (0.00)
Sentiment Φ_1	1.7808631 (0.00)	-4.4281874 (0.00)	-0.8853287 (0.01)	-0.0007089 (0.23)
Constant	268.0133312 (0.0))	156.7162601 (0.00)	101.0827012 (0.29)	192.2448593 (0.00)
α	0.1088096 (0.00)	0.1041336 (0.00)	0.1033907 (0.00)	0.1043631 (0.00)
β_1	0.8678536 (0.00)	0.8833507 (0.00)	0.8837965 (0.00)	0.8825165 (0.00)
Sentiment β_2	-17.4582744 (0.03)	27.9358986 (0.51)	10.5553179 (0.35)	0.0435489 (0.37)
Loglikelihood	-9965.6178	-10004.9911	-10008.2604	-10011.2447

Note: probabilities in parentheses.

In Table 3, which represents a standard GARCH(1,1) model under normality, the sentiment series raw scores appear to work the best in that they have the smallest loglikelihood value of the four sentiment measures, and the coefficient is highly significant in both the conditional mean and conditional variance equations. Sentiment squared performs the next best, but it is not significant in the conditional variance equation, though it is highly significant in the conditional mean equation. The least effective sentiment metric is the difference of the sentiment scores, which is insignificant in both the conditional mean and conditional variance equations for the GARCH(1,1) model.

We also estimated the GJR model with the student t distribution and report estimates with robust standard errors. The results are shown in Table 4.

Table 4: GJR(1,1) model of DJIA, with mean and variance equations augmented by sentiment scores

Variable	Sentiment	Sentiment Squared	Sentiment absolute	Sentiment Difference
Conditional mean equation				
Constant ω	0.0448461 (0.01)	0.0701814 (0.00)	0.0871242 (0.00)	0.0349153 (0.06)
Sentiment	1.57960 (0.00)	3.96032 (0.00)	0.728568 (0.02)	1.16198 (0.00)
Conditional variance equation				
Constant	0.0141329 (0.01)	0.00786521 (0.15)	0.00393231 (0.67)	0.0125011 (0.00)
Sentiment	0.171409 (0.00)	0.700363 (0.13)	0.134351 (0.24)	0.354530 (0.05)
Alpha α	0.0468063 (0.00)	0.0500625 (0.00)	0.0493889 (0.00)	0.0471467 (0.00)
Gamma γ	1.01626 (0.00)	1.00935 (0.00)	1.00936 (0.00)	1.00957 (0.00)
Beta β	0.898544 (0.00)	0.892031 (0.00)	0.893386 (0.00)	0.897946 (0.00)
Likelihood	-2264.05965	-2296.52548	-2300.39130	-2269.12295

Note: probabilities in parentheses.

The variants of the Sentiment series score are significant in all four equations in the conditional mean return specification. They are less influential in the conditional variance equation, but the Sentiment score and the Sentiment in differences are significant in their respective conditional variance equations. The log likelihood statistic again suggests that the most useful form of the Sentiment score is the weighted average.

The final set of GARCH models feature Nelson's (1991) EGARCH model. The results are shown in Table 5, and feature a skewed t distribution and robust standard errors. The various Sentiment score measures are always highly significant in the conditional mean equation and the first two Sentiment measures, the weighted average score and the square of the score, are highly significant in the conditional variance equation, while the other two Sentiment metrics are significant at the 10% level. The log likelihood statistic suggests that the weighted average Sentiment score is the most informative for the EGARCH specification.

Table 5: EGARCH model of DJIA, with mean and variance equations augmented by sentiment scores

Variable	Sentiment	Sentiment Squared	Sentiment absolute	Sentiment Difference
Conditional mean equation				
Constant ω	0.0128581 (0.46)	0.0676681 (0.00)	0.0902486 (0.00)	0.0298678 (0.00)
Sentiment	1.65606 (0.00)	4.41787 (0.00)	0.838581 (0.00)	1.14923 (0.00)
Conditional variance equation				
Constant	0.107452 (0.00)	0.110305 (0.00)	0.129897 (0.00)	0.0978228 (0.00)
Sentiment	0.433137 (0.00)	0.936212 (0.02)	0.238725 (0.09)	0.742513 (0.07)
Alpha α	0.128429 (0.00)	0.129616 (0.00)	0.129897 (0.00)	0.126812 (0.00)
Gamma γ	0.191717 (0.00)	0.195877 (0.00)	0.194817 (0.00)	0.179412 (0.00)
Beta β	0.966777 (0.00)	0.973785 (0.00)	0.974168 (0.00)	0.981003 (0.00)
Likelihood	-2245.48488	-2290.97489	-2295.05655	-2263.91824

Note: probabilities in parentheses.

4. Conclusion

In this paper we have analysed the relationship between the TRNA news series for the DJIA constituent stocks after having aggregated them into a daily average Sentiment score time series using all the constituent companies in the DJIA. This was then used in an analysis of the relationship between the two daily sets of series, TRNA news sentiment on the one hand, and DJIA returns on the other. We analysed the relationship between the two series using the basic GARCH, GJR and EGARCH models. The conditional mean and conditional variance equations are augmented for each model by including one of the four variants of the sentiment score.

The results for all three models suggested that the weighted average Sentiment score was the most informative in all cases, with the lowest log likelihood score. Nevertheless, all variants of the score contained useful information about factors impacting on the volatility of the DJIA. These findings support our previous work on the topic (Allen et al. (2013a, b)), which suggested the usefulness of the sentiment series in an asset pricing context and the informativeness of the series, as revealed by entropy-based metrics.

Acknowledgements

References

[1] Allen, D.E., M. McAleer and A.J. Singh (2013a), "Daily Market News Sentiment and Stock Prices", working paper commisioned by QUANTVALLEY/FdR: 'Quantitative Management Initiative', available at http://www.qminitiative.org/quantitative-management-initiative.html

[2] Allen, D.E., M. McAleer and A.J. Singh (2013b), "An entropy based analysis of the relationship between the DOW JONES Index and the TRNA Sentiment series", working paper commisioned by QUANTVALLEY/FdR: 'Quantitative Management Initiative', available at http://www.qminitiative.org/quantitative-management-initiative.html

[3] Baker, M, and J. Wurgler (2006) "Investor Sentiment and the Cross-Section of Stock Returns", *Journal of Finance*, 61, 4, 1645-1680

[4] Barber, B. M., and T. Odean (2008) "All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors". *Review of Financial Studies, 21(2)*, 785–818.

[5] Bollerslev, T. (1986) Generalized autoregressive conditional heteroscedasticity, *Journal of Econometrics,* 31, 307-327.

[6] Borovkova, S., and D. Mahakena (2013) *"News, Volatility and Jumps: The Case of Natural Gas Futures"*. Working Paper. Retrieved From : http://ssrn.com/abstract=2334226

[7] Cahan R., Jussa J., & Luo Y. (2009). Breaking News: How to Use News Sentiment to Pick Stocks: MacQuarie Research Report.

[8] Da, Z. H. I., Engelberg, J., and Gao, P. (2011) "In Search of Attention", *The Journal of Finance,* 66 (5), 1461-1499.

[9] diBartolomeo, D., & Warrick., S. (2005). "Making covariance based portfolio risk models sensitive to the rate at which markets react to new information", in J. Knight & S. Satchell. (Eds.), *Linear Factor Models* : Elsevier Finance.

[10] Dzielinski, M. (2012) "Which news resolves asymmetric information?", Working Paper, nCCR.

[11] Dzielinski, M., Rieger, M. O., & Talpsepp, T. (2011) "Volatility asymmetry, news, and private investors", *The Handbook of News Analytics in Finance* (pp. 255-270): John Wiley & Sons, Ltd.

[12] Engle, R.F. (1982) Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation, *Econometrica,* 50, 987-1007.

[13] Engle, R. F. (2001) "GARCH 101: An Introduction to the Use of Arch/Garch Models in Applied Econometrics", *Journal of Economic Perspectives,*15, 4, 157-168.

[14] Glosten, L.R., R. Jagannathan, D. Runkle (1993) On the relation between the expected value and the volatility of the nominal excess return on stocks, Journal of Finance, 48, 1779-1801.

[15] Groß-Klußmann, A., and N. Hautsch (2011) "When machines read the news: Using automated text analytics to quantify high frequency news-implied market reactions". *Journal of Empirical Finance, 18*(2), 321-340

[16] Hafez, P. & Xie J. (2012). Factoring Sentiment Risk into Quant Models. RavenPack International S.L

[17] Harris, R.D.F., E. Stoja, J. Tucker (2007) "A simplied approach to modeling the co-movement of asset returns", *Journal of Futures Markets,* 27(6), 575-598.

[18] Huynh, T. D., and D.R. Smith (2013). News Sentiment and Momentum. FIRN Research Paper.

[19] Keynes. J.M. (1936) The General Theory of Employment, Interest and Money, Palgrave Macmillan, Basingstoke, UK.

[20] Leinweber, D., and J. Sisk (2011) Relating news analytics to stock returns *The Handbook of News Analytics in Finance* (pp. 147-172): John Wiley & Sons, Ltd.

[21] McAleer, M. (2005) "Automated inference and learning in modeling financial volatility", *Econometric Theory*, 21, 232-261.

[22] Mitra, L., Mitra, G., and D. diBartolomeo (2009) Equity portfolio risk (volatility) estimation using market information and sentiment. *Quantitative Finance, 9*(8), 887–895

[23] Nelson, D.B. (1991) Conditional heteroskedasticity in asset returns: A new approach, Econometrica, 59(2), 347-370.

[24] Sinha, N. (2011) *Underreaction to news in the US stock market.* Working Paper. Retrieved From: http://ssrn.com/abstract=1572614

[25] Smales, L. A. (2013). *News Sentiment in the Gold Futures Market.* Working Paper, Curtin University of Techonology.

[26] Storkenmaier, A., Wagener, M., and C. Weinhardt (2012) "Public information in fragmented markets". *Financial Markets and Portfolio Management, 26*(2), 179-215.

[27] Tetlock, P.C. (2007) "Giving content to investor sentiment: the role of media in the stock market". *Journal of Finance 62*, 1139–1167.

[28] Tetlock, P.C., Macskassy, S.A., and M. Saar-Tsechansky (2008) "More than words: quantifying language to measure firms' fundamentals". *Journal of Finance 63*, 1427–1467